

Integrating Blockchains and Intelligent Agents in the Pursuit of Artificial General Intelligence

Benjamin Fluharty

A Senior Thesis submitted in partial fulfillment
of the requirements for graduation
in the Honors Program
Liberty University
Spring 2022

Acceptance of Senior Honors Thesis

This Senior Honors Thesis is accepted in partial fulfillment of the requirements for graduation from the Honors Program of Liberty University.

Daniel O'Malley, Ph.D.
Thesis Chair

Robert Tucker, Ph.D.
Committee Member

David Schweitzer, Ph.D.
Assistant Honors Director

Date

Abstract

Artificial General Intelligence (AGI) is the next greatest technological milestone. AGI can be defined as a realized artificial intelligence (AI) with the ability to understand and solve problems of various scope within constantly changing environments. To take steps toward this goal, a baseline of information will be provided regarding surrounding topics and the current state of AGI, itself. Through the culmination of swarms of highly optimized narrow AI agents, a collaborative effort will be extended toward general intelligence. Blockchains have been selected to facilitate this connection. A software deliverable will accompany this thesis to illustrate how this idea might be realized. The SingularityNET platform is utilized for this end due to its advanced protocols and methods for inter-AI communication.

Integrating Blockchains and Intelligent Agents in the Pursuit of Artificial General Intelligence

Artificial general intelligence (AGI), if realized, will be the greatest step in mankind's technological prowess to date. Since its inception, artificial intelligence (AI) has garnered much attention, especially in recent years due to exciting new endeavors coming to fruition (Lu, 2019; Natale & Ballatore, 2020). Projects like Sophia shocked the world with their ability to closely replicate human intelligence using common programming languages. However, this infatuation with AI was not always the case. This discipline has experienced peaks and valleys where interest and funding saw plenty of variation (Grudin, 2009). AI winters, as they were labeled, marked times in which AI research faltered. These took place multiple times over the last century and are largely due to overpromising and underdelivering results (Haenlein & Kaplan, 2019). However, when nations sought to utilize AI for national security purposes and began to reap meaningful results, attention shifted back to this area and with it, generous funding (Grudin, 2009). Other factors like the Internet, programmable graphical processing units, and access to immense data sets have also played an integral role in the resurgence of AI research. In the present time, AI is receiving much more attention as many applications of this research are becoming mainstream.

AGI is the subset of AI which focuses on forming agents with the ability to reason abstractly and understand a broad array of concepts (Torres, 2019). It must also be capable of adapting to its environment and according to some, should even be self-aware (Huang, 2017). What defines self-aware is its own problem still debated by many. Other terms for this idea are strong AI and full AI. Computers can accomplish some tasks with much higher proficiency than

that of the average individual. A human's ability to perform complex mathematical calculations is quite limited while that is not as evident for computers. However, there are some things that humans possess that AI, in its current form, just does not have. People have the unique ability to reason at different levels of understanding. Metacognition and the ability to adapt to ever-changing environments and tasking is also what sets people apart from current narrow AI. However, just because this is not yet realized does not mean it never will be. Blockchains, an emerging technology based on distributed ledger technology (DLT), may hold the key to a future with AGI (Dinh & Thai, 2018).

Before discussing this possible solution, first, a literature review will be provided to deliver much-needed context as to the current state of AGI research. Within this section, the requirements and characteristics of AGI will be listed. Following this background, a review of narrow AI and its use when joined with myriads of other AI agents will be introduced. These swarms of AI agents represent an integral part in achieving AGI. The reasoning of which will be explained within that section. The current role of blockchains and their characteristics will also be included. The blockchain will be shown to be an integral part of this solution as well. Once these ideas have been discussed, a novel approach to achieving AGI through the integration of these ideas will be provided. With this idea in place, a description of an accompanying software deliverable will be presented to serve as a proof of concept.

The SingularityNET platform is enlisted for this purpose. The process of constructing and configuring a service that operates on this platform will be discussed along with a description of the protocols and technology it uses to meet its communication requirements. Finally, while the specific narrow AI discussed throughout was not yet achieved, an example service will be

reviewed in detail such that the appeal of this platform can be seen. Ultimately, the goal here is to synthesize the current standing of AGI research and upon this foundation, demonstrate its future using blockchains as a medium of communication between swarms of AI agents.

Literature Review

To potentially provide the scientific community with a novel approach to achieving artificial general intelligence, it is essential to refer to the extensive work already dedicated to this area. A wide array of sources was acquired to understand the state of AGI research from a well-rounded perspective. The contributions of Dr. Goertzel have significantly furthered its progress and many of his works are cited throughout. However, as there are many other individuals at the forefront of AGI research, the works of others will be listed as well. Chomsky and Kurzweil are others who possess differing views towards the implementation of AGI. This section will seek to provide readers with a baseline of information regarding artificial general intelligence. First, the most widely agreed-upon requirements for achieving AGI will be provided and explained. This will be followed by a brief survey of a few major strategies in progress for achieving this amazing feat. Once this section concludes, a focused look will be taken at the components that will comprise the proposed solution.

Quantifying Artificial General Intelligence

Before trying to attain something as tremendous as a functioning artificial general intelligence, it is critical to first establish a set of requirements to know when this has been achieved. The requirements to be echoed here are, for the most part, identified from one source and while the semantics may change, the overlying ideas are agreed upon by most leading AGI researchers. The first set of items are not necessarily requirements but characteristics to be

expected from this system. These will be shown and explained before the list of requirements is given.

Characteristics of Artificial General Intelligence

An article developed by ten individuals at the forefront of this field includes the list of characteristics. The first of these characteristics is that a generally intelligent agent should be able to operate in a complex and diverse environment with a variety of different objects (Adams et al., 2012). One notable test designed to monitor the progression and eventual success of AGI is the “Wozniak Test” (Hernández-Orallo, 2017). The scenario this test presents is that if we placed an intelligent agent at the doorstep of a house and instructed it to go inside and make a cup of coffee, would it be successful in completing this task? The basis of this idea is that it can be easy enough to train a narrow AI to accomplish this task, however, if the house and coffeemaker were unfamiliar, does it have the general understanding of the concepts required to succeed in this task (Hernández-Orallo, 2017)? Here, general understanding is emphasized over the automation of a familiar task.

It should also be able to complete brief tasks in a succinct manner while also possessing the capability to dedicate longer periods of time to solving more complex problems (Adams et al., 2012). Long-term planning and memory are key ideas that relate to this characteristic. The tasks must not be narrow in scope and should instead be complex and constantly new (Hernández-Orallo, 2017). It should be impacted by the interaction of other agents (Adams et al., 2012). These interactions should be as complex as they are among humans. The computational resources used to generate this intelligent agent should be limited and not increased gradually as more problems arise for it to solve (Adams et al., 2012; Thórisson et al., 2013). Finally, the

existence of this agent should be continual and established for an extended amount of time (Thórisson et al., 2013). While there is no set length for this characteristic, it should not be satisfied by short bursts of disconnected decision-making (Adams et al., 2012). These many characteristics ensure that this intelligence is not limited as current narrow AIs are.

One final point concerning these items is that their complexity is still bounded by their limitations (Adams et al., 2012). Interactions should be complex but need not be unlimited in scope. By this, it is meant that a generally intelligent agent is not required to have an endless supply of knowledge about everything in the world. However, it should grasp basic concepts to facilitate conversational understanding of common topics (Shevlin et al., 2019). While these characteristics require much advancement in this field, it is not expecting anything beyond human intelligence (Adams et al., 2012). Humans are limited in their ability to solve certain problems, and this is not required to be any different for the agent. Given these basic characteristics, an idea can be formed about what sets apart AGI from the current realized AI implementations.

Requirements for AGI

Now, the requirements for artificial general intelligence will be provided. First, a realized AGI must not require further reprogramming to solve new problems (Adams et al., 2012). It must also possess the ability to realize and understand a symbol system (Bruckner et al., 2012). In addition, it must also be able to represent and effectively use several types of knowledge (Adams et al., 2012). Modality-specific knowledge in this context represents the ability for an agent to have a deep understanding of a concept such that they can contextualize it depending on the scenario it finds itself in (Adams et al., 2012). As expressed earlier, generally intelligent

agents must be aware of a large variety of knowledge areas. While this need not be completely extensive, it should allow for a general understanding of numerous concepts (Bruckner et al., 2012). It must be able to form and hold beliefs that are not bound to its environment but are sustained throughout various settings (Adams et al., 2012). Once again, this idea refers to the ability for AGI to have a memory it can reference at will. This is echoed by others including Shevlin et al. (2019) that claim the ability to transfer information between domains and to retain information over extended periods is a noteworthy means of gauging the progression of work toward achieving general intelligence.

Meta-cognitive knowledge or self-awareness is another important subset of knowledge required by AGI (Adams et al., 2012; Thórisson et al., 2013). It must also support the possibility of bounded and unbounded deliberation based on the problem it must solve. Its design should facilitate learning from people and even other agents (Adams et al., 2012). There are many ways that this learning can take place, one of which will be discussed soon. Regardless of the method, it is widely agreed upon that it will require the collaboration of a diverse group of individuals specializing in an array of different fields (Adams et al., 2012; Goertzel, 2014). In all, these requirements provide a much-needed baseline to judge further work so it can be determined what requires more effort, and in due time, when it is finally achieved.

Current Approaches to Artificial General Intelligence

Research was also dedicated to reviewing some of the major approaches to achieving AGI. To bridge the gap between the last section and this one, the first point will be one that can be seen as a possible requirement and may warrant a specific approach. This is the possible necessity for embodiment. Goertzel (2007) writes that the need for embodiment is by no means

unanimous and that there are very intelligent researchers who hold opinions on both sides of the debate. He believes that while there are clear benefits to an embodied AGI, it is not necessary for its existence (Goertzel, 2007). He goes on to push for a greater focus in the virtual sphere of AGI embodiment where they can interact with other agents in simulated worlds rather than in this natural one. This would dramatically decrease the time it would take to develop a functional model as the integration with a physical body would be even more time-consuming (Goertzel, 2007). Given that this opinion comes from one of the minds behind Sophia, a physical manifestation of AI mentioned earlier, this point merits additional credibility.

In the pursuit of a well-rounded basis for future work, the view of the opposing party will be included next. Bruckner et al. (2012) claim that a physical body is necessary for meaningful interaction with its environment. This interaction is believed to be integral to developing intelligent behavior and giving meaning to sensory inputs. The idea of physical embodiment can be extended further to include the concept of intentionality (Husserl, 2012). It is claimed by Xu and Wang (2018) that the best path toward AGI or the integration of phenomenology and cognitive science is achieved through an appeal to embodiment. They cite Husserl's work on intentionality as the marker in achieving this profound feat. Whether it is required for AGI to have a physical manifestation or not, it can certainly be considered a desired feature.

As Goertzel (2007) discussed here, a virtual embodiment may be the first step. Second Life is the name of the virtual world Goertzel dedicated to this concept where millions of parrots inhabit this environment to interact with human users. The goal of this virtual world is to pursue improved, adaptive language learning (Goertzel, 2007). As they improve in their ability to communicate with people, more users will have greater satisfaction from these conversations and

will be more engaged with the AIs. This process would cause a positive feedback loop enhancing the AI's grasp on a human language to the point where it can match that of a person (Goertzel, 2007). This is an interesting approach that might see a revival due to the increasing interest in the Metaverse among other emerging virtual worlds which seek to fully immerse the user and produce a functioning, virtual society.

This idea of "coaching" has also seen support from others. Rohrer (2010) states that on simpler systems, coaching can elicit more interesting and complex behavior than would be achieved in isolation. He also writes that there is a biological component to this as well since it is how children and the young of many species learn best. Also, much prior research that has been devoted to evaluating intelligence for people can be applied to intelligent agents through coaching (Adams et al., 2012). This area is certainly one of the most plausible means toward achieving AGI (Rohrer, 2010). However, there are even more means of implementing a learning model. Neural networks are another means of constructing a learning model to be used by AI (Elman, 1993). A multi-layer combination of these neural networks can comprise a deep learning model which is also believed by some to be the key to AGI (Bengio, 2009). The next approach toward creating general intelligence is much broader in scope.

Self-programming is a concept that has certainly garnered more attention recently (Xiao et al., 2019). Pseudo-automatic programming initiatives like GitHub Copilot have extended what has been considered possible in the realm of software engineering (Krill, 2021). While this is far from what is required to achieve something that can be used for AGI, it has not stopped several organizations from pursuing AGI through these means (Thórisson et al., 2013). Self-programming can be defined as the creation of a program by the system itself (Xiao et al., 2019).

The principles of this product are provided to the system when it is being designed but the details were determined at runtime based on the experience of the system. It is critical for this self-programming product to be influenced by factors both within the system and the external environment (Thórisson et al., 2013). While it has already been discussed briefly, it is important to, once again, discuss an agent's ability for self-reflection and understanding. Self-programming is one approach that critically needs this aspect to be implemented.

Time and space restrictions also need to be considered. Thórisson et al. (2013) continue to state that general intelligence systems will be distinguished not just for their ability to apply reasoning methods to solve problems but their ability to consider their own limitations and what they are capable of. This is accomplished by adopting a constructivist epistemology which, through an iterative approach, attempts to model the world as it experiences it (Luger et al., 2002). Self-programming is a notable approach and even if AGI is not achieved through these means, other interesting applications might result from the continuation of these efforts.

There is another approach that deserves to be discussed in this section as well. This approach takes on artificial general intelligence from an entirely new perspective. Instead of believing that AGI can be achieved through traditional computation methods, Huang (2017) believes a new form of architecture, the "neurocomputer" is required. He writes that the most advanced supercomputer could only simulate the brain at a mere, one percent of its capability while requiring an extremely large amount of power consumption (Huang, 2017). Because of this, it is necessary, in his view, to approach AGI through imitating the physical structure of the brain on a completely novel platform. To accomplish this, comprehensive brain scans that enable computer engineers to understand the structure of the brain more fully are required (Huang,

2017). Following this achievement, neuromorphic devices that simulate neurons and synapses must be developed in a way that mimics the composition of neural networks in the brain (Huang, 2017). The final step, labeled “intelligence fostering engineering”, would be effectively stimulating this artificial brain with signals necessary to generate intelligence (Huang, 2017). Huang is not alone in this understanding of the current computational limitations.

Ruaro et al. (2005) agree with the process of using artificial neural networks that are designed to mimic the processes of the human brain. While the technology is available for imitating neurons, there are still other issues that cause this solution to still need extensive work (Ruaro et al., 2005). Silent electrodes failing to make sufficient electrical contacts with other neurons is the greatest limitation of this design along with the difficulty of measuring the electrical activity of these electrodes (Ruaro et al., 2005). Despite these current limitations, it is certainly another approach that warrants attention.

The underlying premise that current hardware is insufficient to facilitate AGI is not uncommon. Pei et al. (2019) propose a similar approach by designing a new type of chip to address this problem. They lay out the plans for a hybrid chip that supports both the current computer-science-based artificial neural networks and those inspired by neuroscience. The “Tianjic” chip integrates these approaches into a synergistic platform. It utilizes a multi-core architecture that enables reconfigurable building blocks and streamlined data flow (Pei et al., 2019). It is programmed with hybrid coding schemes that accommodate both algorithms found in computer science and brain-inspired circuits (Pei et al., 2019). Deng et al. (2020) also pursue advances in AI through the Tianjic chip architecture. They believe that it has the unique structure required to host the growing complexity of deep learning and neuromorphic platforms and that

combining these paths into one piece of technology would greatly accelerate the journey toward AGI (Deng et al., 2020). All these strategies have convincing points and their own collection of benefits and downsides. While they all are certainly pursuable means of achieving AGI, there certainly remain new, undiscovered paths to this end.

This review has proved quite useful for gathering information necessary to render a new approach to the creation of AGI. Establishing a foundation of requirements and characteristics to build from is essential in knowing what the target is before taking aim. Seeing how many other individuals and groups are approaching this problem has given much-needed context to the pursuit of AGI. Before explaining the new approach in detail, it will be tremendously useful to first ensure a firm understanding of its components are realized beginning with narrow AI and its implementation in swarm settings.

Narrow AI and Swarm AI

Given this understanding of the requirements for achieving AGI, this section will analyze the current state of artificial intelligence and why it is insufficient to fulfill these points. Narrow AI, which is an intelligent agent designed to serve a specific purpose in a defined environment, certainly has its uses. However, it is inherently limited. When placed in a new environment or given a new assignment, it fails miserably. This is because narrow AIs must be highly optimized to their specific context (Adams et al., 2012). Environments can be fully or partially observable, deterministic or stochastic, competitive or collaborative, single-agent or multi-agent, dynamic or static, and discrete or continuous (Russell & Norvig, 2010). Swarm AI is a related idea to narrow AI in that it can be defined as a myriad of narrow AIs working together with a singular focus

(Kutsenok & Kutsenok, 2011). Both ideas will be explained further in this section and their importance in this new approach will follow.

Narrow AI

This discussion on narrow AIs will greatly benefit by beginning with a clear definition. Kurzweil's (2005) definition essentially states that narrow AI is the subset of artificial intelligence that focuses on solving specific and narrowly constrained problems. AGI is certainly not part of this category as one of its most important requirements is its ability to function in diverse and complex environments to solve a wide array of problems (Bruckner et al., 2012). Furthermore, narrow AI is limited in that even if the context or behavior specification is altered by the smallest amount, reprogramming or, at the very least, reconfiguring is required (Goertzel, 2014). There is a stark contrast between this behavior and that of people.

Humans, on the other hand, have an important capability to self-adapt to changes in circumstances and goals. They possess the ability to perform "transfer learning" which is the capacity to generalize one's understanding from one setting to another whether it be contextual or due to changing goals (Goertzel, 2014). There is certainly a wide gap between the capabilities of these two branches of artificial intelligence. However, as it will be explained further, narrow AI might just have a role to play in the pursuit of AGI. It will be of use to now reflect on some modern uses for narrow AI to better understand how it can be applied to implementing AGI.

If AGI were to be developed in the same way that humans develop over time, they must first learn the basic blocks through which human communication takes place (Adams et al., 2012). There is a significant amount of research dedicated to this approach to AGI. Many have speculated that the first emerging AGI will only have the capacity of a youth and should be

required to pass a standardized collection of exams like those given to preschoolers (Adams et al., 2012). The idea is that aspiring intelligent agents should work through school curriculum and take assessments just like children. This would certainly require an understanding of natural language (Adams et al., 2012). However, is this possible for AI agents? The “Chinese Room” thought experiment introduced by Searle (as cited in Globus, 1991) challenges the possibility of AGI by claiming that it is impossible to make the connection between semantics and syntax. Churchland and Churchland (1990) provide a rebuttal to this idea by invoking the, at the time, new connectionist engines which would allow for parallel networks of computations. One flaw in Searle’s experiment is the high degree of focus on individual agents and not collectives working together (Globus, 1991). The solution proposed here seeks to answer this challenge by not focusing on individual entities but on a large-scale collaboration that does have the ability to understand semantics in a collective sense.

If these fundamental blocks are in place, it is believed that further complexity can be achieved. Quoting Turing, Adams et al. (2012) reinforce the validity of this approach by stating that in the same publication that the famed Turing test is found, he suggests that instead of trying to program a mind with the capabilities of an adult, that first, replicating the mind of a child should be achieved. However, there are others who would fundamentally disagree with this approach. Chomsky (1980) believes that what allows children to understand grammar and the semantics of language is that they form and confirm hypotheses that require an inner language of thought. This is something not natively found in computer systems. However, this did not impede the work of others who sought to achieve this with compliance to generative grammar that emphasizes systematicity and productivity (Globus, 1991). This is an interesting scenario in

which a sophisticated narrow AI with a sufficient understanding of natural language processing (NLP) can be developed to the point of achieving general intelligence. In fact, one of the plausible paths toward artificial general intelligence is the “Steady Incremental Progress” scenario which speculates that over hundreds of years, narrow artificial intelligence will continue to become more capable until it can achieve a general understanding of the world (Goertzel, 2007). Whether this means of achieving AGI is plausible is still certainly up for debate.

While it is certainly not the preferred approach, mainly due to the amount of time it would take, narrow AI might be the eventual, lone answer to this question. With this basic understanding of narrow AI, the topic of discussion will shift slightly to the use of swarms of narrow AIs to accomplish different goals.

Swarm AI

Swarm AI is a fascinating field of study. It has seen more use in recent years and is a promising means for a variety of applications (Sadiku & Musa, 2021). It is explained well by Kutsenok and Kutsenok (2011) who state that the main goal behind swarm AI is to create a system of very simple agents that individually work toward solving a larger problem. These agents require a simple and indirect means of communicating with each other, allowing them to exchange information and coordinate globally (Kutsenok & Kutsenok, 2011). Different approaches can be utilized to emphasize different outcomes. While collaborative AI swarms will be the focus of this section, competitive sets of AI agents are also useful in some scenarios. The collective work from the myriads of narrow AIs will allow global decisions to be made with higher confidence.

Three main principles dictate the creation of swarm AIs. The first dictates how individual agents will work to solve sub-problems to work towards the overarching solution (Kutsenok & Kutsenok, 2011). A multiagent system is used to solve a single problem. This problem is broken down into parts that autonomous agents can solve (Sadiku & Musa, 2021). Currently, the designer of the swarm is responsible for dividing work among these agents, and they are not controlled by another AI (Kutsenok & Kutsenok, 2011). In solving these smaller problems, progress is made toward completing the desired solution. The next principle emphasizes that the least amount of information given to the agent is preferable as it will need to emphasize speed over a complete understanding of the input set (Kutsenok & Kutsenok, 2011). This also greatly simplifies the development of these agents. Finally, communication between these individual narrow AI agents is critical to solving the main problem. It is only through quick and universal communication and collaboration that this design gives swarm AI an edge in certain scenarios (Sadiku & Musa, 2021). These principles will be shown to benefit the implementation of artificial general intelligence in the final section.

Another topic worth discussing here is the “No Free Lunch” theorem. This theorem was popularized by Wolpert (as cited in Ho & Pepyne, 2002) whose contribution is largely based on the work of David Hume. The essential premise of this theorem is that there is no single algorithm which can be universally applied to all problems while still being optimized to its fullest extent (Gómez & Rojas, 2016). Models must be specialized to their specific problem for optimization to take place, yet, in doing so, the number of scenarios to which the model can be applied is lessened (Ho & Pepyne, 2002). This further shows that a single narrow AI is not capable of deriving a model that can understand the world in a general sense. Unless swarms of

narrow AI agents are enlisted which are each highly optimized to their own problem sets, it is highly unlikely that narrow AI will be of use in this endeavor toward general intelligence.

Here, a brief overview of the principles of narrow AI and swarm AI has been given. The current uses and principles that guide the development and implementation of these ideas have been shown. While most have looked elsewhere for pursuing AGI, it will be shown how these items are integral to its success. For many applications, narrow AI is still extremely useful. As discussed earlier, while not sufficient for achieving a general intelligence individually, narrow AI is exceptional at solving specific problems and can be utilized with that limitation in mind. Having a firm understanding of these concepts, the next section will review the other piece of technology necessary to this approach. From this foundation, a new solution will be proposed.

Blockchains and Smart Contracts

This section will discuss the emerging technology of blockchains and the related use of smart contracts. Blockchains have seen success in many different areas with a growing number of applications each year (Salah et al., 2019). A description of this technology filled with examples and key ideas will be given. This will be followed by a brief description of smart contracts and how they can be used for this solution as well. Once all these ideas have been communicated, the solution will be presented.

Blockchain Structure

The basic structure of a blockchain and the blocks that comprise it will first be discussed so that a better understanding can be attained of its current usages and ultimately its relation to the development of artificial intelligence. Blockchains are a specific implementation of the broader idea of distributed ledger technology. DLT dictates the protocols and required

infrastructure required to facilitate the simultaneous access, validation, and modification to a network that is scattered across multiple nodes (Maull et al., 2017). A blockchain, specifically, is a public, shared ledger that is agreed upon by all users in a distributed network (Dinh & Thai, 2018). It is a nearly immutable structure that is commonly used to record transactions. (Hechler et al., 2020). These transactions which are represented in blocks within this blockchain allow a transparent, and secure means of recording asset transfer. These assets can be physical items like cars and houses or digital items like the various cryptocurrencies available at this time (Hechler et al., 2020). It is important to mention that while large, widely used blockchains have robust security in place, their smaller counterparts are not as protected. If a single entity were to own at least 51% of the machines responsible for block validation, they could bypass the proof of work (PoW) consensus protocol and manipulate the data stored in blocks (Salah et al., 2019). The blocks and links between them will be discussed in more detail next.

Blocks will be analyzed first. Each block includes the information concerning the transaction, a hash value, and a timestamp (Dinh & Thai, 2018). The block header is responsible for storing metadata about this block. This is necessary for it to be easily discernable where this block fits into the blockchain (Rahim et al., 2020). The block identifier is the cryptographic hash, and a Merkle tree is utilized to store a summary of all transactions in the block and to verify the dataset integrity (Rahim et al., 2020). Before leaving the topic of block structure, smart contracts should be mentioned.

Smart Contracts

Smart contracts are pieces of code that encapsulate the business terms of the transactions to be recorded in the blockchain. (Dinh & Thai, 2018). These code pieces are self-executing and

used to verify the enforcement of the transaction's terms and conditions (Salah et al., 2019).

While Bitcoin uses a different means of validating digital currencies, Ethereum, another popular cryptocurrency validates transactions by having mining nodes execute the smart contracts (Salah et al., 2019). There is still much more untapped potential in the use of smart contracts. Having AI take over the process and automate the creation and brokering of smart contracts would vastly improve the efficiency of the process if a sophisticated enough agent were developed. Before concluding this section, the blockchain structure will be discussed further.

Blockchains are implemented using a linked list with the hash of the previous block included in the current block to track its location within the data structure (Dinh & Thai, 2018). The abstract ideas that describe the features of blockchains can be grouped into five key concepts: agreement, dispersed estimation, information repository, source, and stability (Rahim et al., 2020). Agreement is foundational to the blockchain structure. It is completely reliant on nodes participating in the hashing process. Immense computational power is required to process the hashes in the blocks. It is estimated that over 204.5 TWh of power is used annually to mine just Bitcoin (Digiconomist, 2022). This task is passed on to voluntary nodes in the network. Estimated dispersion refers to the idea that entire copies of the blockchain are stored across numerous nodes ensuring that there is no one, centralized node that could act as a point of failure. The blockchain acts as an information repository meaning it can store all the information in its structure. This is commonly referred to as "on-chain" (Quiniou, 2019). The source feature is meant to communicate that all transactions are transparent since each one is traced, documented, and fully identifiable without the use of a third party that trust would have to be placed in.

Finally, stability refers to the near immutability of blockchains. While not completely immune, the possibility of interfering with a block at the end of the chain is quite low to begin with and only becomes more difficult with each preceding block in the chain (Rahim et al., 2020). Some other key terms to consider in the discussion of blockchains are “distributed” and “non-centralized”. These are the most commonly agreed-upon terms for blockchains since it avoids the debate between decentralized and distributed systems that is common in this sphere. Blocks in the chain are not centralized but also not necessarily decentralized. The non-centralized model of the blockchain is based on nodal architecture and it allows members of the distributed network to connect in a way that removes the need for a central entity or service with absolute power in the network (Quiniou, 2019). This idea is crucial to the foundational concept of blockchain as it distributes power throughout the nodes in the network so that one node is not any more necessary than another. This improves the persistency of the structure by not relying on specific nodes.

It is important to grasp the ideas behind what makes a blockchain function to grasp the amazing applications it has currently and will have very soon. To do so, an analysis of blockchains as a data structure and the blocks it is comprised of were given. Smart contracts were also mentioned in brief as a glimpse of what AI can be capable of when using blockchains. These ideas provide a solid understanding of the blockchain structure so that now the focus may shift to current uses of blockchain and its implications in artificial intelligence. Finally, the new approach to realizing AGI will be proposed in the following section.

Integration of Narrow AI Swarms over Blockchain for AGI

It is at this time that the two main areas of focus can be integrated to form the overarching purpose of this research. The history and use of narrow AIs have been discussed at great length. It has also been conveyed how swarm AI is used in numerous different scenarios. The expanding role of blockchains in society and smart contracts has been related to this topic as well. At this point, the ideas will be combined to provide a new way of achieving artificial general intelligence.

General Intelligence through Amplified Narrow Intelligence

The crux of this idea lies in understanding swarm AI in a different sense. Up to this point, swarm AI has been used to join numerous narrow AIs together under a singular goal that enables highly accurate results (Bengio, 2009). Instead of using many like-minded narrow AIs to achieve a singular end, numerous AI agents that each specialize in their own domain may be connected and work together on more complex problems. This is the central idea which sets it apart from other paths to AGI.

It has been shown how advanced many narrow AIs are becoming. Each with a specific goal, these agents have made significant progress in understanding human input regardless of the form it takes. If agents that were competent in speech recognition were joined to those which possessed high levels of natural language processing, and then those were connected to AIs which were designed for understanding human emotions through various types of input, their combined strengths could be utilized to achieve something much greater. The underlying philosophy of this approach is that the whole is greater than the sum of its parts. If a large enough collection of these narrow AIs were to be all connected, could general intelligence

emerge? This possibility is the goal. This approach does prompt some significant questions which will be discussed next.

How can this be achieved? Connecting this myriad of narrow AI agents will require a novel approach. In recent years, blockchains have become much more commonplace. They have shown themselves to be useful for many different applications. As such, they will be the foundation for the connection of these narrow AIs. Many different possibilities can be considered for this research. SingularityNET (2020) boasts a strong basis for this idea as it possesses many of the needed components including the decentralized AI marketplace and the use of smart contracts. The rationale for this will be explained here. One of the primary reasons is that it is a current impossibility to, in a centralized location, maintain and deliver power to the number of narrow AIs required to support this idea. Current swarm AIs are programmed to be extremely basic for this reason. To have a multitude of them working on a problem, the power draw is significant. Since more advanced AI agents will be enlisted, they cannot all come from the same source. As communicated earlier, the distribution of components among many actors will effectively keep a single entity from controlling this advanced technology. At this point, a description of the application and platform used to communicate these ideas will be provided.

SingularityNET as a Potential Platform

The platform utilized for this project is SingularityNET (2020). Mentioned briefly earlier, SingularityNET boasts the first, complete AI network and marketplace. Given the requirements and purpose of this application, SingularityNET seemed to be the clear choice. There were some obstacles in the way which made publishing a service quite difficult but the capabilities it offers make up for this downside. This experience will be detailed next.

To begin, a developer account is required for publishing on this platform. The process is quite involved and required many items that were not available at this time. It is not designed for individuals to easily publish their work but is more geared toward established organizations with a service to offer. There were multiple complications related to this issue. First, it was required that an “etcd” cluster be set up for receiving payments for the service. This would require security certificates and remote services with static IP addresses. An Ethereum identity would also be necessary along with payment groups and various addresses. This has blocked the progress of publishing the service with their platform.

Because of this, the scope of this application was narrowed to simply achieving a running, sample SingularityNET service. There is plenty of documentation available on this subject. Unfortunately, because tutorials are becoming quite dated, building a docker container with all the requirements was not feasible. Working around this issue, a previously built container was pulled which included all the dependencies.

The initial prompt was quite limited. After some basic navigation through the surrounding directories, a Bash shell was started which made the process of reviewing documentation and source files much easier. The example service which showcases SingularityNET’s capabilities operates like a calculator. The user specifies the operator and two numerical values. The syntax of the command is as follows

```
snet client call example-org example-service default_group mul '{"a":"6","b":"7"}' -y
```

The base of this command is

```
snet client call
```

which tells SingularityNET that it will be interacting with a service (client) and to call this service to execute it. What follows is the name of the organization, service, and payment group.

The next two arguments are used by the program to determine what is to be calculated. The

```
mul
```

piece informs the program that the next two numbers are to be multiplied. The final argument is

JavaScript Object Notation (JSON) which stores the two integers to be computed. Finally,

```
-y
```

skips confirmation prompts for payment information. The result of this command is

```
value: 42.0
```

which is, of course, the expected output of multiplying six and seven. This process is an involved one and the pieces that make it function will be outlined next.

The key to SingularityNET's unique platform is its focus on communication. First, the Protocol Buffers (protobuf) file must be compiled. The protobuf format is used by SingularityNET for its standard of communication. Google Remote Procedure Call (gRPC) is the system enlisted to perform communication across components. A channel is opened on gRPC and the endpoints of the request along with the operator and numbers are provided. The `example_service.py` file which consists of the logic behind the calculator includes a function, `serve`,

which spawns threads that listen for requests from the test Python script. It parses the arguments and returns a result. Thankfully, the functionality of this example service was designed to run immediately in this container once the protobuf file had been compiled. Creating a new service required much more work.

To provide something new for the SingularityNET community, a new AI service was drafted to run on this platform. Titled “Bridge AI”, this agent would be responsible for running deep learning algorithms to classify problem types and enlisting AI agents which were specifically designed to solve the identified problem type. The design of this platform makes this feature achievable. The interconnection of these AI agents enables services like Bridge AI to call other services seamlessly. Further configuration is required to run this application on a local instance with a functioning daemon. As mentioned earlier, there is also more work required on SingularityNET’s website to register an organization for publishing the service. The current state of this work is available for review on Docker Hub at <https://hub.docker.com/repository/docker/bfluharty/bridge-ai> and will continue to be updated as progress is made on this front.

It has been shown that getting started with the SingularityNET platform is not a trivial task. There is much knowledge required about various technologies to succeed in this endeavor. However, after some configuration, a glimpse of the platform’s potential can be observed. While the narrow AI discussed throughout is not yet completed and functional on this platform, the platform, itself was explored and further work will be dedicated toward this end. Next, the role of smart contracts will be discussed.

The Role of Smart Contracts

Since the technology behind smart contracts has already been explained, their role here will be presented directly. Smart contracts enable transactions to take place automatically over a blockchain. Because the collection of these numerous agents would be quite a daunting task. It would be advantageous to once again, utilize narrow AI to oversee the process.

A master intelligence could be designed with the single focus of seeking other narrow AIs and, through smart contracts, acquiring them for their role in the collective general intelligence. So, to accomplish this, a single narrow AI will be designed to navigate over the blockchain to locate narrow AIs that have different purposes, they will be acquired and put to work on a task related to their purpose. In the same way that swarms are currently implemented, they will each report back with their solution. It will be passed from one agent to the next over a blockchain and then the necessary responses will be shown. The master AI will oversee these communications and ensure they are all available and collaborating. Once a significant enough number of these agents are created, a more general intelligence may emerge. This is, of course, not without its own set of challenges which will be discussed next.

Challenges to this Approach

There are a few issues that are immediately apparent. Aside from the development of this master AI, one must ask if there are enough narrow AIs available to be integrated into this swarm that fulfills the requirements of AGI described in the first section. Some of the most crucial ones have been discussed, particularly NLP and the need for logical understanding. A tentative list of nonnegotiable abilities must be composed which can then be acquired through the overseeing intelligent agent (Adams et al., 2012). From this point, these narrow AIs must be made available and provided by parties willing to participate in this experiment. The reliance on others is mainly due to the vastness of this endeavor. A distributed solution will reduce the impact of power consumption and individual need for the myriads of AI agents required to meet the requirements of AGI drafted by others. The design is admittedly quite volatile because each of these many pieces is not centralized and controlled by one entity. If a provider decides to take back their

agents, the whole system could be paralyzed. Redundancies and strong relations among partners will be crucial.

One final challenge with this approach is temporal in scope. Humans are not patient creatures. If they ask a question to one of the many smart devices made available to us, they expect a correct and prompt response. The passage of information from the myriad of agents required to participate in the given action could dramatically increase the time it takes to act. Quicker transactions over blockchain must be facilitated to accomplish this great endeavor. While a general intelligence could still be achieved, it would be hindered by an extremely slow response time. These are the parts of this solution that require further research and reflection if this approach is to be realized.

There are several different means for achieving AGI, some more hopeful and others more grounded. While the proposed solution may not be the quickest and most tenable, it does strive to gather intrigue around this subject, ultimately furthering the surrounding discussion and, in response, enabling a generation of intellectuals from a variety of different fields to join in this work and develop something truly extraordinary.

Conclusion

AGI possesses a surpassing level of intrigue due to the immeasurable number of uses that it can have if realized. While this work is by no means a step-by-step guide to implementing the first tangible AGI, it does include a new approach that hopes to build toward this result. Reviewing the status of AGI provided much-needed context to a problem that has baffled the brightest scientists for decades. It was useful to gather what has worked, what has not, and what is left to be tried. It was shown how important a defined list of characteristics and requirements is

for attaining AGI. The role of narrow AIs in achieving artificial general intelligence has been shown by its participation in swarms. Using blockchains in this solution was also demonstrated in detail. The implications of AGI are immense and certainly much grander than what can be detailed here. However, by utilizing communication between AI agents in swarms over blockchain, humanity may be a step closer to achieving this fascinating reality.

References

- Adams, S., Arel, I., Bach, J., Coop, R., Furlan, R., Goertzel, B., Hall, J. S., Samsonovich, A., Scheutz, M., Schlesinger, M., Shapiro, S. C., & Sowa, J. (2012). Mapping the landscape of human-level artificial general intelligence. *AI Magazine*, 33(1), 25-42.
<https://doi.org/10.1609/aimag.v33i1.2322>
- Bengio, Y. (2009). Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1), 1-127. <https://doi.org/10.1561/22000000006>
- Bruckner, D., Zeilinger, H., & Dietrich, D. (2012). Cognitive automation—Survey of novel artificial general intelligence methods for the automation of human technical environments. *IEEE Transactions on Industrial Informatics*, 8(2), 206-215.
<https://doi.org/10.1109/TII.2011.2176741>
- Chomsky, N. (1980). *Rules and representations*. Columbia University Press.
- Churchland, P., & Churchland P. (1990). Could a machine think?. *Scientific American*, 262(1), 32-39.
- Deng, L., Wang, G., Li, G., Li, S., Liang, L., Zhu, M., Wu, Y., Yang, Z., Zou, Z., Pei, J., Wu, Z., Hu, X., Ding, Y., He, W., Xie, Y., & Shi, L. (2020). Tianjic: A unified and scalable chip bridging spike-based and continuous neural computation. *IEEE Journal of Solid-State Circuits*, 55(8), 2228-2246. <https://doi.org/10.1109/JSSC.2020.2970709>
- Digiconomist. (2020). *Bitcoin Energy Consumption Index*. <https://digiconomist.net/bitcoin-energy-consumption>
- Dinh, T. N., & Thai, M. T. (2018). AI and blockchain: A disruptive integration. *Computer*, 51(9), 48-53. <https://doi.org/10.1109/mc.2018.3620971>

- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1), 71-99. [https://doi.org/10.1016/0010-0277\(93\)90058-4](https://doi.org/10.1016/0010-0277(93)90058-4)
- Globus, G. G. (1991). Deconstructing the Chinese room. *The Journal of Mind and Behavior*, 12(3), 377–391.
- Goertzel, B. (2007). Human-level artificial general intelligence and the possibility of a technological singularity: A reaction to Ray Kurzweil's *The Singularity is Near*, and McDermott's critique of Kurzweil. *Artificial Intelligence*, 171(18), 1161-1173. <https://doi.org/10.1016/j.artint.2007.10.011>
- Goertzel, B. (2014). Artificial general intelligence: Concept, state of the art, and future prospects. *Journal of Artificial General Intelligence*, 5(1), 1-48. <https://doi.org/10.2478/jagi-2014-0001>
- Gómez, D., & Rojas, A. (2016). An empirical overview of the no free lunch theorem and its effect on real-world machine learning classification. *Neural Computation*, 28(1), 216-228. https://doi.org/10.1162/NECO_a_00793
- Grudin, J. (2009). AI and HCI: Two fields divided by a common focus. *AI Magazine*, 30(4), 48-57. <https://doi.org/10.1609/aimag.v30i4.2271>
- Haenlein, M. & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California Management Review*, 61(4), 5-14. <https://doi.org/10.1177/0008125619864925>
- Hechler, E., Oberhofer, M., Schaeck, T., & Thummalapalli, S. (2020). *Deploying AI in the enterprise : IT approaches for design, DevOps, governance, change management, blockchain, and quantum computing*. Apress. <https://doi.org/10.1007/978-1-4842-6206-1>

- Hernández-Orallo, J. (2017). Evaluation in artificial intelligence: From task-oriented to ability-oriented measurement. *Artificial Intelligence Review*, 48(3), 397-447.
<https://doi.org/10.1007/s10462-016-9505-7>
- Ho, Y. C., & Pepyne, D. L. (2002). Simple explanation of the no-free-lunch theorem and its implications. *Journal of Optimization Theory and Applications*, 115(3), 549-570.
<https://doi.org/10.1023/A:1021251113462>
- Huang, TJ. (2017). Imitating the brain with neurocomputer a “new” way towards artificial general intelligence. *International Journal of Automation and Computing*, 14(5), 520–531. <https://doi.org/10.1007/s11633-017-1082-y>
- Husserl, E. (2012). *Ideas: General introduction to pure phenomenology*. Routledge.
<https://doi.org/10.4324/9780203120330>
- Kornai, A. (2014). Bounding the impact of AGI. *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3) 417-438. <https://doi.org/10.1080/0952813X.2014.895109>
- Krill, P. (2021, August 12). *OpenAI offers API for GitHub Copilot AI model*.
<https://www.infoworld.com/article/3629469/openai-offers-api-for-github-copilot-ai-model.html>
- Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. Viking.
- Kutsenok, A., & Kutsenok, V. (2011). Swarm AI: A general-purpose swarm intelligence design technique. *Design Principles & Practice: An International Journal*, 5(1), 7–16.
<https://doi.org/10.18848/1833-1874/CGP/v05i01/37798>

- Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1-29.
<https://doi.org/10.1080/23270012.2019.1570365>
- Luger, G. F., Lewis, J., & Stern, C. (2002). Problem solving as model refinement: Towards a constructivist epistemology. *Brain, Behavior and Evolution*, 59, 87-100.
<https://doi.org/10.1159/000063735>
- Maull, R., Godsiff, P., Mulligan C., Brown, A., & Kewell, B (2017). Distributed ledger technology: Applications and implications. *Strategic Change*, 26(5), 481-489.
<https://doi.org/10.1002/jsc.2148>
- Natale, S., & Ballatore, A. (2020). Imagining the thinking machine: Technological myths and the rise of artificial intelligence. *Convergence: The International Journal of Research into New Media Technologies*, 26(1), 3-18. <https://doi.org/10.1177/1354856517715164>
- Pei, J., Deng, L., Song, S., Zhao, M., Zhang, Y., Wu, S., Wang, G., Zou, Z., Wu, Z., He, W., Chen, F., Deng, N., Wu, S., Wang, Y., Wu, Y., Yang Z., Ma, C., Li, G., Han, W., ... Shi, L. (2019). Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature*, 572, 106–111. <https://doi.org/10.1038/s41586-019-1424-8>
- Quiniou, M. (2019). *Blockchain: The advent of disintermediation*. International Society for Technology in Education Ltd; John Wiley & Sons, Inc.
<https://doi.org/10.1002/9781119629573>
- Rahim, R., Patan R., Manikandan, R., & Kumar, S. R. (2020). Introduction to blockchain and big data. In N. Kumar, N. Gayathri, A. Rahman, & B. Balamurugan (Eds.), *Blockchain, big*

- data and machine learning: Trends and applications* (1st ed., pp. 1-24). Taylor & Francis Group. <https://doi.org/10.1201/9780429352546>
- Rohrer, B. (2010). Accelerating progress in artificial general intelligence: Choosing a benchmark for natural world interaction. *Journal of Artificial General Intelligence*, 2(1), 1-28. <https://doi.org/10.2478/v10229-011-0005-5>
- Ruaro, M. E., Bonifazi, P., & Torre, V. (2005). Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures. *IEEE Transactions on Biomedical Engineering*, 52(3), 371-383. <https://doi.org/10.1109/TBME.2004.842975>
- Russell, S., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Pearson.
- Sadiku, M. N. O., & Musa, S. M. (2021). *A primer on multiple intelligences*. Springer. <https://doi.org/10.1007/978-3-030-77584-1>
- Salah, K., Rehman, M. H., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access*, 7, 10127-10149. <https://doi.org/10.1109/access.2018.2890507>
- Shevlin, H., Vold, K., Crosby, M., & Halina, M. (2019). The limits of machine intelligence. *EMBO Reports*, 20(10). <https://doi.org/10.15252/embr.201949177>
- SingularityNET. (2020). *SingularityNET*. <https://singularitynet.io/>
- Thórisson, K., Nivel, E., Sanz, R., & Wang, P. (2013). Editorial: Approaches and assumptions of self-programming in achieving artificial general intelligence. *Journal of Artificial General Intelligence*, 3(3), 1-10. <https://doi.org/10.2478/v10229-011-0017-1>
- Torres, P. (2019). The possibility and risks of artificial general intelligence. *Bulletin of the Atomic Scientists*, 75(3), 105-108. <https://doi.org/10.1080/00963402.2019.1604873>

Xiao, Y., Nazarian, S., & Bogdan, P. (2019). Self-optimizing and self-programming computing systems: A combined compiler, complex networks, and machine learning approach. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 27(6).

<https://doi.org/10.1109/TVLSI.2019.2897650>

Xu, Y., & Wang, P. (2018). Why is the Husserlian notion of “intentionality” needed by artificial general intelligence?. *The Philosophical Forum*, 49(4), 401-425.

<https://doi.org/10.1111/phil.12207>