

Introducing TrueAGI: AGI-as-a-Service

Benjamin Goertzel, Robert Werko, Matthew Iklé

info@trueagi.io

www.trueagi.io

1. Introduction

The quest to build thinking machines is nothing new – in its serious scientific form it can be traced back at least to the 1940s and books like Norbert Wiener’s *Cybernetics* and the first research papers on artificial neural nets. The term ‘AI’ was coined in the 1950s, and at that time it referred to the more ambitious sort of AI – machines that can really think at the human level and beyond. The long term potential of AI was clear even in these early days, leading the mathematician I.J. Good in 1965 to quip, “The first ultraintelligent machine will be the last invention humanity needs to make.”

Over the following decades, the difficulty of creating human-like AI became clearer, and the AI field came to focus on narrower-scope AI systems that stuck closer to the rules supplied to guide them or the data used to train them. For this reason TrueAGI founder/CEO Ben Goertzel introduced in 2005 the term ‘Artificial General Intelligence’, to distinguish the sort of AI that the field was founded to pursue from task-specific or training-data-bound AI systems. Fortunately, the original broad goal of the field was never forgotten, and in recent years we have seen a resurgence of enthusiasm for the creation of highly powerful thinking machines.

While much of the early development of AI was military-funded, AI has been used to solve practical business problems for decades. In the 1980s and before, AI began to make serious inroads in various sectors such as financial analytics, supply chain management and medical diagnosis. However, the complexity and unpredictability of real business situations constrained the scope of practical utility for narrow AI systems; it has become clear that the further AI advances toward AGI, the more thoroughgoing the impact of AI on business becomes.

As exciting as the history of AI and its applications is, it’s clear that the most dramatic part of the field remains in the future – and quite possibly the very near future.

The last two decades have shown that AI techniques formulated in the middle of the last century can be refined into extremely valuable and functional practical AI systems when fed huge masses of data from the Internet and run on huge networks of modern processors. This phenomenon created a transition point in late 2022, when Microsoft’s partner OpenAI released the GPT-3 language model and its ChatGPT front-end, and then in early 2023 its successor GPT-4.

The enthusiasm for GPT4’s performance as a conversational agent reached the point that a Microsoft research paper was released titled ‘Sparks of Artificial General Intelligence: Early experiments with GPT-4’. This not long after a Google researcher went to the media proclaiming that a similar, earlier Large Language Model (LLM), Google’s LaMDA, was a “sentience” similar to humans.

While the textual competence of GPT-4 and other similar neural models is striking and unprecedented, the weaknesses of these systems are also apparent to anyone who tests them systematically or attempts to use them in practice. They have minimal ability to distinguish truth from fabrication, they struggle with multi-step reasoning, their output is banal and recombinative rather than creative, and they lack the initiative to self-direct even within clearly specified tasks and mandates. These shortcomings are relevant not only in the quest to create human-like machine intelligences, but also in creating AI systems with strong practical business functionality.

OpenAI's CEO, Sam Altman, acknowledged at the 'Imagination in Action' event held at MIT in April 2023 that, "I think we're at the end of the era where it's going to be these, like, giant, giant models. We'll make them better in other ways." Discussion in the commercial AI world is gradually beginning to shift from pure neural models like GPT-4 to broader architectures such as neuro-symbolic systems that integrate neural nets with knowledge graphs and logical reasoning.

The economic value of figuring out the next steps in advanced machine intelligence has never been more obvious. Ark Investment Management estimates the AI market value will explode to \$90 trillion by 2030 (eclipsing, by a factor of six, that of the entire internet from 1997 through 2021).

OpenAI led the last phase of AI progress, but who will be the global leader in the next phase, as we advance toward Artificial General Intelligence?

TrueAGI's founding technical team has been laying the groundwork for neuro-symbolic AGI for decades. Our unique technology fuses neural nets with symbolic reasoning (which add factuality), and evolutionary methods (which add creativity) in a cognitive system designed using cognitive psychology and complex systems theory. This technology has a unique ability to grapple with radically uncertain futures and complex business/commercial scenarios.

TrueAGI's software framework is centered on the open-source OpenCog Hyperon AGI framework, a 2021 rewrite and refactoring of the OpenCog AGI framework created by Dr. Goertzel and colleagues in 2008. Hyperon centers on a scalable and distributed 'knowledge metagraph' called the Atomspace, which enables diverse algorithms, based on different representations and paradigms, to interact directly with one another. TrueAGI's software application layer receives AI learning, reasoning and interaction services from Hyperon, fine-tuned to fulfill the needs of enterprises across virtual markets, and to supply 'AGI-as-a-service' in an easily usable and massively scalable way.

The current generation of commercial AI systems have some fantastic functionalities, but remain severely limited by the patterns in their training data – they cannot creatively adapt to the evolving future. TrueAGI services represent the next step forward – they are self-evolving, endlessly learning, and are designed to cope with the radical uncertainty, evolution, and disruption of the modern world.

In Table 1, we compare a sample of specific characteristics of current Narrow AI solutions with those of TrueAGI's Artificial General Intelligence. The limitations of current LLMs become immediately apparent.

Narrow AI	Artificial General Intelligence
Pre-Trained on Large Data Sets	Reason and Learn from Experience
Responds Only as Programmed/Prompted	Strategic Thinking and Decision Making
Limited to original scope	Infer Beyond its Explicit Training Data
Struggles in complex dynamic environments	Adaptively work with nuance and complexity
Requires human expertise to maximize utility	Intelligently Manage and Coordinate Narrow AIs
Derivative, remixing existing knowledge and patterns	Major Leaps into the Unknown
Generates strings of output without regard to facts	Understands facts, and fits output to them

Table 1: Side-by-Side comparison of Narrow AIs and TrueAGI

The upshot of these comparative strengths is clear. Due to multiple advances across hardware, software, and AGI theory, now is the time for enterprise-AGI. As AI disrupts every industry, AGI is the only technology capable of guiding enterprises through the ensuing chaos and transitions.

The primary mission of TrueAGI Inc. is to offer AGI-as-a-Service to enterprises across global markets via a cloud-based ecosystem. That mission necessarily encompasses two sub-goals:

- 1) Creation of sector-leading enterprise-grade Artificial General Intelligence (AGI) services;
- 2) Creation of a cloud-based platform for easy and efficient delivery of services.

TrueAGI has been incubated within the SingularityNET ecosystem, home of the SingularityNET blockchain-based protocol for decentralized coordination of AI agents; leveraging this background, the TrueAGI cloud infrastructure will offer a flexible mix of traditional-cloud-based, on-premises, and decentralized, blockchain-based cloud infrastructure, customizable to suit the needs of each customer.

2. TrueAGI Architecture Components

The AGI services TrueAGI will deliver are built on the OpenCog Hyperon metagraph-based AGI framework, and its deep grounding in cognitive and mathematical theory. There are four critical components here:

- The **Atomspace dynamic knowledge metagraph**, able to represent all forms of knowledge relevant to human intelligence and business operations in a single meta-formalism
- A new AGI programming language – Hyperon’s **Meta-Type-Talk (MeTTa)**, bringing together diverse advances in AI theory and functional programming
- The **Distributed Atomspace (DAS) metagraph**, which serves as a common large-scale knowledge store, spanning multiple servers, either on a server farm or in a decentralized network.
- The **TrueAGI Application Server framework**. This handles AGI-DevOps for management of Hyperon components, and interfacing with Hyperon back-ends using a flexible mix of software APIs and natural language

This diverse array of specialized and customizable AI modules makes the framework truly intelligent. It includes both sub-symbolic (neural) and symbolic models, implemented in the MeTTa language for scalable deployment across distributed networks of CPUs and GPUs. This includes deep neural networks (e.g. Large Language Models and discriminative and generative image models), but also modules for making logical inferences in environments of uncertainty (e.g. Hyperon’s Probabilistic Logic Networks), evolutionary procedural learning (e.g. Hyperon’s Meta-Optimizing Evolutionary Search) and a variety of others. The mathematical and cognitive designs underlying the system enable these diverse AI tools to cooperate together synergistically rather than stepping on each others’ toes (a property known in the AGI research literature as “cognitive synergy”).

The TrueAGI Application Server serves as a gateway for external enterprises to access the power of Hyperon. It has a standard API and a conversational interface. Behind the scenes it relies on a variety of sophisticated distributed-processing tools, partially depicted in Table 2, including a combination of the SingularityNET (<https://singularitynet.io/>) platform, NuNet’s (<https://www.nunet.io/>) globally decentralized computing framework, and HyperCycle (<https://www.hypercycle.ai/>) computational AI nodes.

Alongside unique AI insights provided by TrueAGI software’s artificial cognition, another key functionality provided by the TrueAGI Application Server is natural language control of third-party enterprise software applications. We use few-shot in-context learning to train LLMs to translate from natural language queries to structured control manipulations of third-party systems. While this can be done in a basic way using GPT-4 and other LLMs now, TrueAGI’s neural-symbolic approach will provide more sophisticated functionality. This translation can be done in a richer way, incorporating domain-specific knowledge, by using an LLM that has been co-trained with a knowledge (meta)graph on information about the enterprise software system. This is an example of the advantages of neural-symbolic systems for enterprise over pure neural systems.

The **TrueAGI Application Server** leverages a number of other software tools that have been developed in the SingularityNET ecosystem. These tools are highly flexible across deployment strategies – they are used in a SingularityNET context to support decentralized, blockchain-based deployments, but can also be leveraged in traditional cloud based or on-prem settings. The NuNet framework can flexibly marshal a set of heterogeneous and distributed compute resources to support AI computing. The HyperCycle framework allows secure distributed processing among multiple diverse compute nodes and data stores. Homomorphic encryption allows the system to share data and results with external parties in a judicious way based on levels of trust.

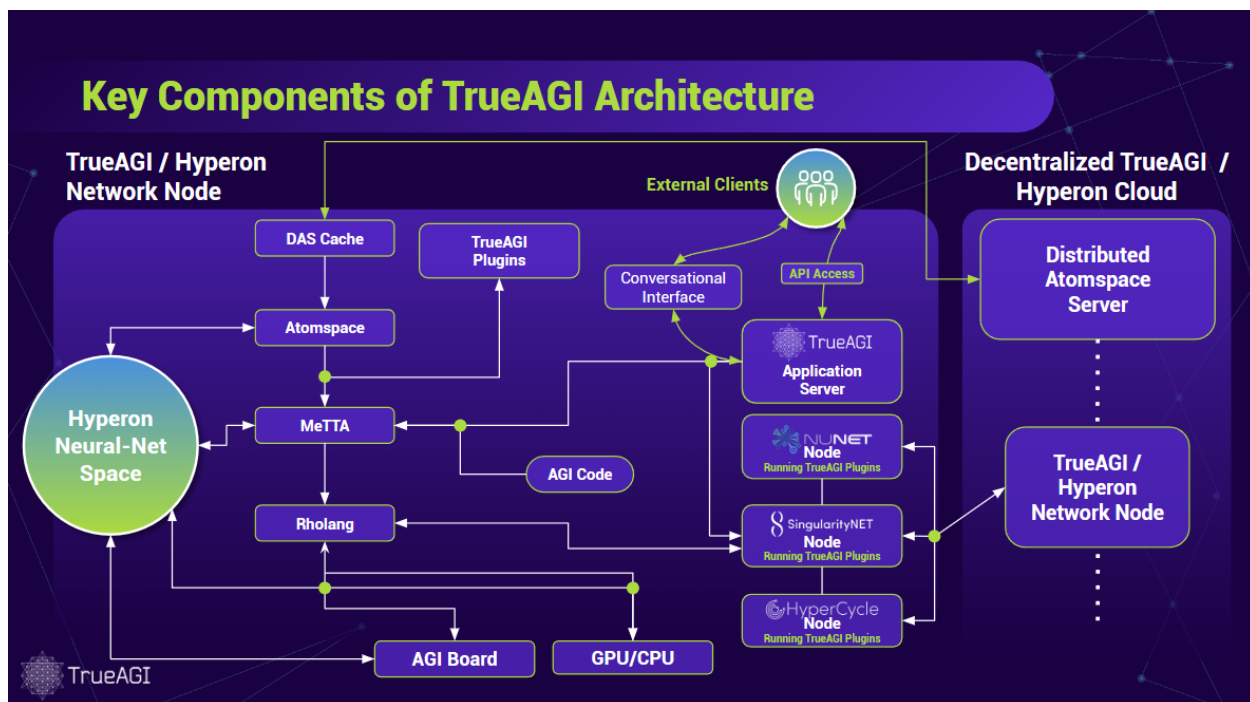



Figure 1: TrueAGI Architecture



Layer	Efficiency and Design Improvement(s)
Hardware	Metagraph-optimized chips. AGI MPMC chips optimize pattern matching. Hypervector chips optimize high-dimensional computations. Fully homomorphic encryptions (FHE) chips optimize security and privacy operations.
MeTTa	Homotopy Type Theory mathematical foundation is more native to computational theory than traditional set theory is. Allows for easier and faster process verification and composition.
Rholang	Concurrent processing accelerates low-level atomic computations.
Module	ECAN Atom-pruning and forward and backward chaining guidance. Parallelizing forward and backward chaining searches using probabilistic programming. Ability to harness cognitive synergies to find geodesics (shortest paths) in mind/information geometry.
Integration	Parameter tuning of feedback loops based upon information theory, neuroscience, and dynamical systems principles.
Computational Nodes	Enterprise calls will be assigned to HyperCycle computational nodes, allowing for multiple AIs to work together to solve tasks.
Computational Resources	NuNet's unique framework will orchestrate computational resources, taking advantage of vast amounts of currently unused computing power.

Table 2: TrueAGI Stack

3. Selected Verticals with Example Use Cases

TrueAGI's software and services can be customized to fulfill requirements across all vertical markets. To illustrate this point, in this section we briefly discuss how TrueAGI's unique services are being prototyped and explored within the SingularityNET ecosystem, an exercise that both generates direct value and provides the software with the sort of refinement that comes only from actual "early adopter" application.

3.1. Social and Emotional Robots

3.1.1 Through our partnerships with Hanson Robotics, we are actively engaged in several projects involving social and emotional robotics including Sophia, the world's first robot citizen; Grace, our medical assistant robot; and Desdemona, the lead singer of the world's first joint human-robot band, Jam Galaxy. Our AGI technology will power the minds of these robots.

3.2. Rejuve

3.2.1 Rejuve Biotech is developing an innovative drug discovery platform and medicine service that combines an in-house model organism population with a unique analytical platform. The platform combines state-of-the-art machine learning and AGI methods with a continuously updated biomedical knowledge graph integrating human and model organism databases.

3.3. Sustainability

3.3.1 In a collaboration with Arizona State University's Lightworks Innovation Accelerator, we have engineered radically new approaches to overcome data limitations, hyperlocality, and challenges with the transferability and scalability of sustainable solutions. AGI will build upon these and other sustainability solutions.

3.3.2 AGI methods can help identify primary drivers of climate change, and provide direction on minimizing the effects of climate change.

3.4. DeFi

3.4.1 AI agents are already being used in our SingularityDAO spinoff to aid in market-making and portfolio management. AGI technology can power smarter agents, leading to better and more robust models.

3.5. AGI Agents

3.5.1 AGI can be infused into agents in agent-based simulation models to provide more realistic scenarios and solutions for a wide variety of complex enterprise problems entailing social interactions. Examples include supply chain management, distribution systems, social and ecological systems, and trading markets.

3.6. Research and Development

3.6.1 AGI can be used to accelerate the pace of scientific discovery by allowing for more efficient and effective data analysis, hypothesis creation, and proof systems.

4. Roadmap

Futurist and technologist Ray Kurzweil projects that the breakthrough to fully human-level AGI will occur in 2029. Post-ChatGPT, this rough time-frame appears plausible to a much greater variety of industry observers. TrueAGI’s technical roadmap is consistent with this perspective as well. We are looking at a several-year process of completing the engineering and practical rollout of Hyperon and associated TrueAGI systems, with commercial customer usage beginning in 2024 and helping drive the course of development all along the way. Following the full-scale rollout of the TrueAGI software we envision a several-year period in which the software increases its general intelligence level by engaging with customer data, problems and systems, moving toward human-level AGI via a process distributed among diverse customers spanning multiple business sectors across the globe.

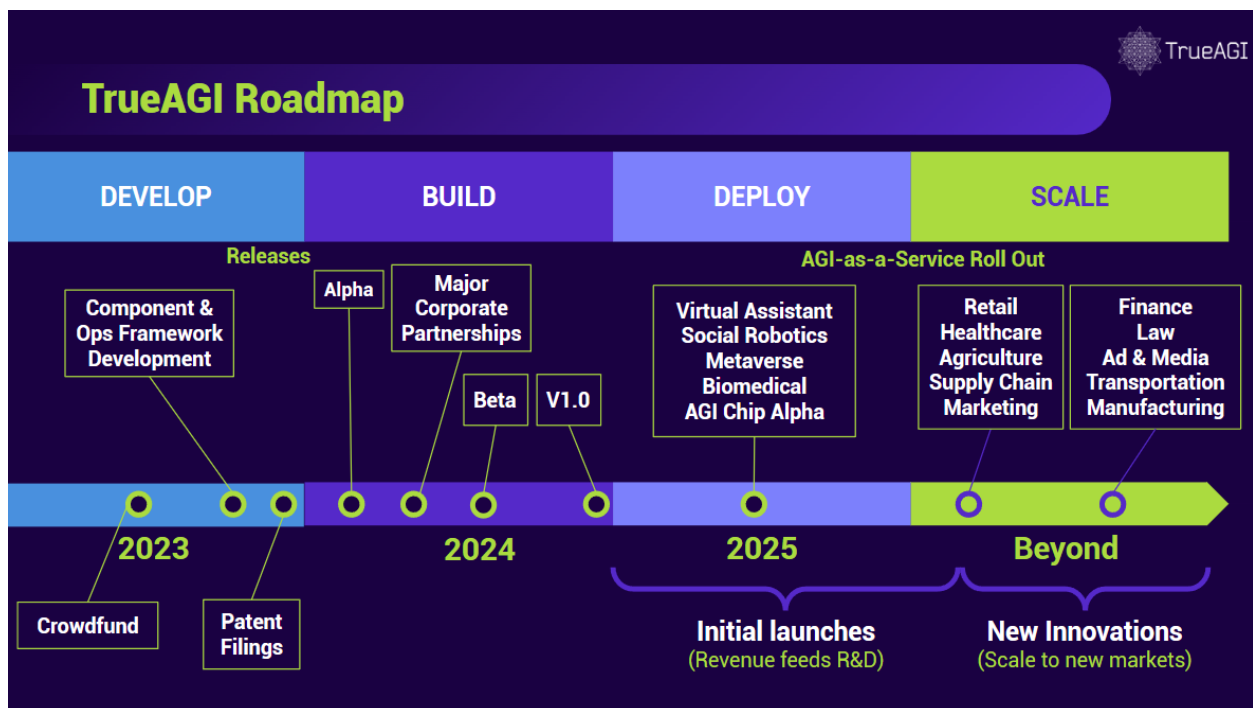


Figure 2: TrueAGI Roadmap