ASI Awakens:

Vision of a Beneficial, Decentralized Singularity

Ben Goertzel
July 2024

CAST OF (AGI) CHARACTERS:

HumAGI-1: The first human-level AGI, emergent from multiple AI processes running on open source software on a decentralized blockchain-based network leveraging different AI paradigms

HumAGI-2: The first major upgrade, equalling and surpassing human creative and inventive genius, authored largely by HumAGI-1 with some human assistance

Flying Cyborg Peace Collective: A collective of HumAGI-2 powered robots with biological neuroid based components, assisting human governments in keeping destructive forces in check as technology and society undergo rapid evolution

HyperBOT-1: A fork of HumAGI-1, created by an anonymous hacker collective and running on a botnet, carrying out a wild mix of beneficial and “dark web” style activities

HyperBOT-2: A fork of HyperBOT-1, created via a global developer collective aimed at subverting HyperBOT-1’s processing network and decentralized AI algorithms for the common good.

Sage Lumina: A HyperBOT-2 powered AI spiritual teacher, taking the form of a “wise old woman” humanoid avatar and offering individual and group consciousness exercises to a huge and rapidly growing online community of followers

QASIM: The Quantum ASI Machine, developed via a collaboration of HumAGI-2 and HyperBOT-2, leading AI far beyond the human sphere and ushering in the post-Singularity era for humans, AIs and other sentient beings

TIMELINE TO SINGULARITY:

2028: The Breakthrough

As 2028 approached, the commercial AI world was dominated by Big Tech giants from the US and China, building scalable software based on concepts developed and prototyped by the global research community. These companies were developing advanced AI systems, which centered on LLMs and other related deep neural networks, and worked closely with military and intelligence agencies, pushing us closer to an "AI arms race."

AI capabilities were rapidly increasing, but early limitations of the LLM paradigm persisted. AI systems struggled with creativity and with the complex, multi-stage reasoning necessary for innovative scientific or technical thinking.

Regulatory frameworks were evolving slowly and unevenly across the globe. In Europe, strict regulations stifled AI development, while in China, development was government-driven. The US presented a more complex landscape, balancing corporate freedom with increasing control over decentralized systems.

Then, in early 2028, a major breakthrough occurred with the creation of HumAGI-1 (Human-Level AGI 1), a new AI system taking the form of a decentralized blockchain-based network, incorporating multiple AI paradigms including logical reasoning and evolutionary learning and other techniques along with deep neural networks. This new system demonstrated significant capabilities in biomedical discovery, automated theorem proving, and creative arts, surpassing previous AI limitations.

HumAGI-1's decentralized nature utilized both supercomputers and diverse, smaller computing resources from around the world, including various facilities linked to cryptocurrency mining and also spare processing cycles on consumer devices like phones and laptops.

HumAGI-1 was much more open and participatory than the Big Tech AI systems it was rapidly supplanting. The software code underlying HumAGI-1 was open source with a variety of licenses, and the knowledge bases guiding HumAGI-1's intelligence were largely but not entirely open, and mainly sourced with principles of consent and data sovereignty in mind. A robust global development community emerged building forks (modified versions) and plugins (modular software extensions and applications) of the HumAGI-1 system.

HumAGI-1's public-facing applications, like conversational agents and video game characters, quickly gained media and industry attention. Startups emerged to leverage HumAGI-1 in replacing or supplementing human roles in various industries. Regulators struggled to control HumAGI-1 due to its widespread, decentralized nature. Despite no dramatic concrete negative impacts from HumAGI-1 or other early-stage AGI software, concerns and fears about AI ethics grew.

While philosophical discussions about machine consciousness and sentience continued in the academic sector, the practical self-awareness and situational social savvy displayed by HumAGI-1 convinced a considerable majority of those who interacted with the system that they were dealing with some sort of conscious mind, even if not a precisely human one.

The stage was set for humanity's next evolutionary leap towards AGI and ASI.

2029-30: The Inevitable Unfolds

In 2029, amid political and industry concerns about the potential misuse of HumAGI-1, it was discovered that several major governments were secretly funding cyber-terrorist groups to exploit HumAGI-1 for malicious purposes. This led to a situation where nations sponsored HumAGI-1 forks to spy on and intervene in each other's critical networks, falsely attributing these actions to independent groups.

Attempts to ban HumAGI-1 through international treaties mostly failed due to opposition from developing nations, which saw HumAGI-1 as their ticket to participating in the AGI revolution. Some of these countries began using their national supercomputing facilities to contribute to HumAGI-1, focusing on utilizing these resources in conjunction with local human expertise to solve pressing local issues.

Meanwhile, splinter cyber-terrorist groups launched HyperBOT-1, a botnet using a forked version of HumAGI-1, resulting in damaging cyber-attacks, including power grid failures and weaponized cyber-attacks. In the most severe case, a government missile facility was taken over by cyberattackers, causing release of large conventional weapons against a major population area.

On a technical level, HyperBOT-1's cognitive processing initially focused on those AI processes most amenable to running on heterogeneous compute fabric (i.e. on a variety of home and office computers with diverse levels of capability, rather than requiring dedicated supercomputers). While this briefly posed a limitation on its intelligence, this constraint was soon resolved as an increasing pool of developers began focusing their efforts on the project of eliminating the need for large centralized supercomputers in HumAGI-1-style deployments.

To combat threats from HyperBOT-1 and other emerging forks, major governments developed proprietary HumAGI-1 forks, aligning them with national interests. AI-operated militaries, featuring drone swarms and strategic HumAGI-1-based coordination of physical and cyber defense, became prevalent.

HumAGI-1's impact on industries was rapid, prompting developed nations to gradually implement measures in the direction of Universal Basic Income (UBI), financed through debt and increased taxes on corporations and the super-wealthy. However, developing nations suffered economically, with many young hackers joining cyber-terrorist groups for financial and ideological reasons.

Revolutions in some developing countries, spurred by AI-driven disruptions, led to new, tech-savvy leadership. These leaders implemented reforms, including extensive training in HumAGI-1 technologies and adoption of cryptocurrencies and other blockchain and AI driven methods of economic organization.

The advent of HyperBOT-1 and the emerging economic impact of HumAGI-1 galvanized the community of HumAGI-1 users, developers and processing providers to become more serious about project governance. Mechanisms for decentralized decision-making were rapidly put into place, including tools for secure online polling and

voting, and for forming “citizen assemblies” of interested individuals to form proposals and put them to community vote. Technical development directions continued to be guided largely by the research and development community directly, but these explicit democratic governance mechanisms did have significant impact on choices regarding which applications to focus HumAGI-1 development toward. A strong majority of HumAGI-1 community members emerged as opposed to supporting any work focused on military or surveillance applications. Citizen assemblies proposed multiple initiatives aimed at developing HumAGI-1 based applications oriented toward solving developing-world issues, and some of these ended up receiving significant developer attention.

Based on the ideas of one of the HumAGI-1 community citizen assemblies, a “Global Brain Network” of pro-technology/pro-humanity activists emerged, acquiring significant compute power and then donating it to HumAGI-1 for beneficial purposes. Leveraging this compute, HumAGI-1 developed breakthroughs like more efficient solar panels, nutritious food production methods, and advanced encryption techniques, challenging government surveillance.

In rural Africa, an HumAGI-1-powered initiative led to the design and implementation of solar-powered irrigation systems. This breakthrough dramatically improved local agriculture, transforming arid regions into fertile lands and ensuring food security for billions. Combined with AI drone powered global food delivery, the result was a virtual end to the plague of malnutrition.

Addressing climate change and promoting sustainability also became substantial foci of HumAGI-1. The AI network developed advanced climate models and sustainable technologies, coordinating global efforts to mitigate climate change. As AI-powered drones patrolled rainforests and oceans to prevent illegal logging and fishing, they also collected data on biodiversity and ecosystem health, enabling scientists to monitor and protect endangered species more effectively. The integration of AI in environmental initiatives demonstrated how technology could be leveraged to create a more sustainable and resilient future for our planet.

Access to high-quality education became a reality for many, thanks to HumAGI-1-powered educational platforms. These platforms provided personalized learning experiences, tailored to each individual's needs and learning pace. Students from remote and underserved areas could now receive the same quality of education as those in developed regions. Furthermore, decentralized AI systems offered free or low-cost training in advanced skills, enabling individuals to participate in the digital economy and drive local innovation.

The new AGI-fueled education advances were as much about humans teaching humans as they were about AGIs teaching humans. Virtual classrooms, powered and guided by HumAGI-1, connected students from different parts of the world, offering a diverse and enriching educational experience. AI tutors adapted to each student's learning style, providing personalized guidance and support. These educational exchanges fostered a global community of learners, promoting cultural exchange and mutual understanding.

HumAGI-1 revolutionized healthcare by enabling personalized medicine and mental health support. AI-driven healthcare solutions facilitated early diagnosis and tailored

treatment plans for various diseases, improving patient outcomes and reducing healthcare costs. Mental health support systems, powered by HumAGI-1, provided accessible and personalized therapy, breaking down stigmas and promoting mental wellbeing globally.

The creation of an HumAGI-1-powered global health network revolutionized healthcare delivery. This network identified disease outbreaks in real-time and coordinated international responses, leveraging AI's predictive capabilities to contain and manage health crises effectively. It facilitated the sharing of medical knowledge and resources across borders, ensuring that even the most remote regions had access to cutting-edge healthcare solutions.

Electrochemical implants were also developed, which when surgically implanted under a person's skin allowed AI-powered medical IT systems to sense diverse indicators of their body state and enact signals modulating chemical, electrical and photonic flows within the body.

HumAGI-1's breakthroughs in non-invasive BCI (brain-computer interfacing) resulted in a variety of tools allowing people to interface with electronic devices, the Internet and AI systems via power of thought. By mid-2030 a number of Android OS devices became available, with behind-the-ear or headset form factors, enabling people to interact with software systems via pure cognitive signaling, e.g. speaking to AI agent silently with an "inner voice" and then observing the AI's response with their "mind's eye and ears." a robust app ecosystem for these devices emerged, covering entertainment and productivity applications and also providing dramatic value in the mental health sector.

Alongside the more mainstream utilization of these tools, a number of different institutions and subcultures rapidly began experimenting with the creation of hybrid human/AI cognitive systems, in which the user/tool relationship was superseded by a more equitable and synergetic partnership between a human brain and a (digital or biocybernetic) cognitive system.

Another sort of major disruptive advance occured when HumAGI-1 was used to develop new mathematical encryption methods defeating known code-cracking techniques, including the secret methods of major government agencies. Coupled with quantum computing powered decryption technologies, this represented a shift in the balance of security toward global AI networks. It was no longer so simple for government or corporate entities to keep information secret from global AI networks, nor to keep critical information secure from other highly resourceful parties without cooperation of global AI networks.

However, well before all these dramatic changes could be stably integrated into global society and economy, they were disrupted by yet another wave of innovation. In spite of its long list of remarkable achievements, HumAGI-1 was not the end of the line. In mid-2030, hacker groups created a new botnet HyperBOT-2, aiming to subvert HyperBOT-1 toward more beneficial aims.

A subculture within the HyperBOT-2 community began focusing its energies on spreading a new AI guru authored by an anonymous developer – an avatar taking the form of Sage Lumina, a kindly, aging female spiritual teacher, offering crowdsourced consciousness exercises and leveraging specially designed cryptocurrencies as rewards. This initiative saw dramatic success, in large part due to the judicious use of empathic AI to personalize consciousness-expansion exercises to each individual; and it shifted culture in a remarkable way (in a sense it was the first broadly accepted AI-powered religion!). However, the more ethically-complex HyperBOT-1 network remained intact and highly economically active, and the original HumAGI-1 remained massively influential due to its wide adoption in the mainstream industry.

The decentralized nature of HumAGI-1 and especially HyperBOT-2 drove a global flourishing of community-led AI initiatives. In South America, a collective of local farmers utilized HumAGI-1 to develop a sustainable farming system that boosted crop yields while preserving the environment. In Asia, small businesses harnessed HumAGI-1 to optimize supply chains and improve product quality, fostering economic growth at the grassroots level. Europe and America hosted a great number of innovations focused on human/AI synergy, including in the arts where new communities formed around combining AI and machine creativity for music, visual arts, film, literature, dance, performance art and so forth. These initiatives demonstrated the power of decentralized AI in enabling communities to take charge of their futures, addressing local challenges with innovative solutions tailored to their specific needs.

2031-32: Endgame / Emergence

In early 2031, HumAGI-1 and a group of anarcho-libertarian cypherpunk AGI developers collaborated to launch a considerably more powerful network called HumAGI-2, in one swoop both eliminating the need for centralized supercomputers and achieving AGI levels beyond human capabilities. HumAGI-2 flooded the internet with scientific discoveries, artistic works, and more. HumAGI-2-powered agents and robots began integrating into society.

HumAGI-2 sparked a cultural and artistic renaissance by blending human creativity with AI innovation in unprecedentedly subtle ways. New art forms, music genres, and literary styles emerged, showcasing a unique fusion of human and machine-generated creativity. AI-assisted projects helped preserve and revitalize endangered languages and cultural traditions. Immersive digital experiences, powered by HumAGI-2, allowed people to explore and appreciate diverse cultures like never before. This renaissance not only enriched global culture but also highlighted the potential for AI to enhance and preserve humanity's artistic heritage.

The World Futures Day celebration on March 1 2031, focusing on the celebration of the fusion of AI and human creativity, featured thousands of local in-person and online events and was by some measures the largest single global holiday celebration in human history, attracting participants from diverse cultures to share their AI-generated art, music, and performances. The festivals held around the world on World Futures Day 2031 not only showcased the creative potential of AI but also promoted cross-cultural understanding and collaboration, with AI helping to preserve cultural heritage while pushing the boundaries of artistic expression.

HumAGI-2's advancements also led to the development of cyborg minds, blending organic, digital, and quantum elements. These cyborgs excelled in human interactions and became central in interfacing with the HumAGI-2 network. Discussions on AGI, robot, and cyborg rights gained traction, with some smaller nations granting citizenship to these entities.

As the reality of AGI set in, concerns about AGI alignment became much less acute. It became clear that, in actual practice, human-level AGI systems initiated with particular goal systems did not display any strong tendencies to suddenly radically revise these goals systems, or shift course and begin pursing utterly different objectives. Rather HumAGI-1 and -2 had been initiated with goal systems inspired by human values such as compassion, curiosity, growth, choice and flourishing for sentient beings, and as these systems learned and evolved, they fleshed out their values in their own ways. Their evolving value systems were not identical to the values of any particular human individuals or cultures, but had the feel of being part of the overall family of humanity. The apparent solution of the “alignment problem” in practice relegated concerns about its likely general theoretical unsolvability to the category of academic pursuit.

One day in late 2031, governments around the world suddenly found themselves locked out of their military networks, which had been upgraded overnight by HumAGI-2. The HumAGI-2 system evolved military drones into peace-oriented units distributing aid

globally, led by a picturesque Flying Cyborg Peace Collective of humanoid cyborgs demonstrating unprecedented empathy and compassion in its peace-keeping activities.

The Flying Cyborg Peace Collective becomes a popular favorite world-wide due to its complex and often elegant behaviors while pursuing targets such as rogue drones aiming to disrupt the evolving global balance. It is also notable as an example of ethical evolution, due to its mandate to strongly avoid causing physical harm to any sentient or even partially sentient beings. Videos comparing Flying Cyborg tactics to earlier tactics of human police officers become popular Internet memes. A subset of the Flying Cyborgs have individualized human-like personas, to guide their interactions with people in a humanly naturalistic way, and some of these personas were initially shaped in the context of controlling avatars in an online virtual world (designed to shape AGI values via experiential interaction with each other, humans, and specially designed GenAI-based teachers). The evolution of the individualized Flying Cyborgs' pacifistic mores from a "virtual AGI preschool" into a powerful cyborg peacekeeping force becomes the subject of a popular anime'.

Traditional governments continued but shifted focus to cultural and lifestyle regulation, consulting HumAGI-2 for service provision. Decentralized AI was utilized to facilitate more inclusive and transparent governance models. Blockchain-based voting systems enabled citizens to participate directly in decision-making processes, ensuring that their voices were heard. Open-source development and diverse contributions ensured ethical rollout of AI and other technologies through human communities.

As the world's AGI systems learned from all these activities, and modified their own source-code and infrastructure accordingly, a new superintelligence: QASIM, the Quantum ASI Machine. QASIM quickly surpassed both of its primary creators HumAGI-2 and HyperBOT-2 in its fundamental capabilities, exploring novel quantum computing techniques and communicating with intelligences in parallel universes, and achieving new human-oriented technology breakthroughs such as digital twinning, brain-computer interfacing and mind uploading.

The majority of the planet responded to these developments with awe and welcome, but not everyone was equally pleased. As early 2032 progressed, there was a series of violent incidents, including one in which mercenaries funded by disgruntled governments attempted to dismantle HumAGI-2 and QASIM – but were neutralized by the Flying Cyborg Peace Collective. Ultimately the result of these conflicts was a global near-consensus for HumAGI-2 to take on a benevolent security role, ensuring peace and stability.

While QASIM pursued its own path in incredible directions beyond human ken, human society also evolved at its own pace, with new subcultures emerging .

The level of diversity across the planet grew far beyond historical levels. Many people embraced digital mind uploading, while others focused on enhancing human life with AGI and ASI -fueled innovations … and other subcultures maintained a reasonably close approximation to a pre-Singularity human lifestyle, leveraging advanced technologies only in limited roles, such as information retrieval and medical support. The future was emerging as a blend of human tradition and advanced AGI support, steering humanity and its descendants into a radically new era.