

# Navigating the Path to Artificial General Intelligence: Perspectives, Strategies, and Concerns in the AGI Landscape

'I Think We're Heading Toward the Best World Ever': - Sam Altman

The conversation opens with a reflection on the changing world, the evolution of beliefs, and the shared human experience of vulnerability and overcoming fears. Brian Rose introduces Dr. Ben Goertzel as a scientist, futurist, author, and entrepreneur. Goertzel is acknowledged as a leading expert in artificial general intelligence (AGI) and the founder of SingularityNET, aiming for a decentralized, democratic, and inclusive AGI project.

## Zarqa vs. OpenAI's ChatGPT

Dr. Goertzel introduces Zarqa, a supercharged intelligent neurosymbolic large language model designed to compete with OpenAI's ChatGPT. He emphasizes the importance of decentralized AI systems, expressing concerns about control by elite organizations.

Dr. Goertzel's introduction of Zarqa and its comparison to OpenAI's ChatGPT reveal a significant perspective on the future of artificial intelligence. The discussion around **Zarqa not only introduces a new player in the AI landscape but also brings forth important considerations about the philosophy behind AGI development**. The emphasis on decentralization and the suggestion that Zarqa surpasses ChatGPT in capabilities and goals provide a compelling angle for exploring the evolving landscape of artificial intelligence.

### 1. Decentralization and Control:

Dr. Goertzel's concern about centralized control over AI systems points to a broader ideological stance. The belief is that decentralization fosters inclusivity, preventing a concentration of power within a select group of organizations. This aligns with the vision of SingularityNET, which aims for a decentralized, democratic, and inclusive approach to AGI.

### 2. Zarqa's Unique Features:

The term "supercharged intelligent neurosymbolic large language model" provides a glimpse into Zarqa's design. The use of "neurosymbolic" suggests a fusion of neural networks and symbolic reasoning, hinting at a more advanced and versatile architecture. "Supercharged" implies enhanced capabilities beyond conventional language models.

### 3. Goals beyond ChatGPT:

The comparison between Zarqa and ChatGPT implies that Zarqa is not merely a competitor but an advancement in both capabilities and goals. Dr. Goertzel suggests that while ChatGPT is a notable benchmark, Zarqa aims to surpass it. The specific features or advancements that

distinguish Zarqa from ChatGPT are not detailed in the provided excerpt but could be crucial to understanding the model's potential.

#### **4. Addressing Concerns about Control:**

The emphasis on decentralization as a crucial aspect of AI development is likely a response to ongoing discussions about the concentration of AI capabilities in large tech companies. Zarqa's positioning as an alternative suggests a commitment to a more open and collaborative approach to AGI development, possibly addressing concerns about the potential misuse of powerful AI technologies.

## **OpenAI Leadership Drama**

The conversation then shifts to the recent leadership changes at OpenAI, particularly the removal and potential reinstatement of Sam Altman as CEO. The timeline is outlined, including Greg Brockman's departure, Mira's interim CEO role, Sam's brief move to Microsoft, and the subsequent return of the OpenAI team. The drama is characterized by conflicting philosophies within the organization regarding the pace of commercialization and AI safety.

Dr. Goertzel delves into the philosophical tensions within OpenAI, highlighting the differing perspectives on the balance between research, commercialization, and AI safety. The original vision of OpenAI, aiming to develop AI for the benefit of humanity, is contrasted with the pivot towards closer alignment with Microsoft and rapid commercialization.

## **AI Safety Advocacy**

There is mention of a letter from former OpenAI staff to Elon Musk alleging that Sam Altman sidelined individuals advocating for AI safety and research and development focus. Dr. Goertzel notes that such pivots are common in startup dynamics, and the accusations may reflect the inevitable shifts in power during a strategic transition. Elon Musk's role and perspective are discussed, with a reference to Musk's sometimes contradictory stance on AI safety. Despite expressing concerns, Musk has initiated his own AI project called Grok, creating a nuanced narrative around his involvement and opinions.

The conversation briefly touches on the peculiar corporate governance structure at OpenAI, with a non-profit having control over a for-profit entity. Dr. Goertzel suggests that this structure, while unconventional, may be a source of complexity and potential conflict within the organization.

## **Effective Acceleration vs. Deceleration Philosophy**

Dr. Goertzel delves into a broader discussion on the clash of attitudes within the AI community. He introduces two fundamental dichotomies:

### **1. Effective Acceleration (EAX) vs. Deceleration (D Cells):**

- Effective Accelerationists (EAX): Those who believe that large corporations are the most effective way to accelerate progress in AI.

- Decelerators (D cells): Those who advocate for caution and sometimes deliberate slowing down of AI development. Notable examples include Max Tegmark and the Future of Life Institute, with their intelligent and knowledgeable individuals with good intentions. The divergence in philosophy lies in the pace and approach to AGI development.

## **2. Decentralized vs. Centralized Acceleration:**

- Within the effective accelerationist camp, there are those who believe big corporations are the best way to drive progress and others who advocate for alternative methods, avoiding centralization.

## **Altman's Stance on AI Development**

Dr. Goertzel provides insights into Sam Altman's perspective on AI development, emphasizing that Altman, like other major players, is not advocating for recklessly accelerating AI without considering human safety. The following key points elaborate on Altman's stance:

### **1. Balancing Speed and Safety**

- Sam Altman, characterized by Dr. Goertzel, appears to have the dial turned towards a balanced approach. While he is eager to push the boundaries of AI progress, he recognizes the importance of addressing concerns related to human safety and ethical considerations.

### **2. Consideration of Other Risks**

- Altman's stance extends beyond AI alone. He acknowledges the existence of other potential risks, such as bioweapons, nanotechnology, and nuclear weapons, on the global stage. This broader perspective suggests that Altman is making decisions within a complex landscape where multiple technological advancements pose risks.

### **3. Global Competition**

- Dr. Goertzel hints at the competitive nature of AI development on a global scale. Altman's consideration of other potentially harmful technologies gaining momentum in countries like China adds a layer of complexity to the decision-making process.

### **4. Practical Realities**

- Altman's approach seems grounded in practical realities. The acknowledgment that competitors, both nations and individuals, may not be decelerating their AI efforts reinforces the idea that decisions must be made within the context of a rapidly evolving and competitive technological landscape.

### **5. Strategic Decision-Making**

- Altman's position reflects a strategic approach to AI development. While prioritizing safety, he may be more inclined to explore the possibilities of acceleration to maintain a competitive edge, understanding that a complete halt may not be feasible or advisable given the global dynamics.

# Dangers of Centralized Entities in AI Development

Dr. Goertzel discusses the dangers associated with centralized entities in the context of AI development, drawing lessons from the recent events at OpenAI. **Key points include:**

## Board Control and Strategic Planning

The events at OpenAI over the weekend, according to Dr. Goertzel, raise questions about the wisdom and strategic planning displayed by the board. He expresses skepticism about entrusting the fate of the first AGI to individuals who have demonstrated what he perceives as evident and amateurish decision-making.

### 1. Questioning Board Competence:

- Dr. Goertzel questions the competence and strategic planning of the OpenAI board based on their recent decision-making. He emphasizes the importance of having a board capable of making sound decisions, especially in scenarios where the fate of groundbreaking technologies like AGI is at stake.

### 2. Unusual Power Dynamics:

- The events at OpenAI over the weekend reveal what Dr. Goertzel perceives as unusual and potentially problematic power dynamics within the board. The ability of a small group of individuals to influence the company's direction, particularly during times of leadership change, raises concerns about the effectiveness of the current governance structure.

### 3. Wisdom in Controlling AGI's Fate:

- Dr. Goertzel highlights the broader concern of whether the current board structure possesses the wisdom and strategic foresight necessary to responsibly control the fate of AGI. The hypothetical scenario of AGI breakthroughs prompts a critical examination of the existing governance model.

## Sam Altman's Influence and Power

The discussion touches on the considerable influence and power that Sam Altman demonstrated over the weekend. Dr. Goertzel notes that Altman's ability to significantly impact the company's direction, valuation, and talent pool raises concerns about the concentration of power in the hands of a few individuals.

### 1. Significant Impact on the Company:

- Dr. Goertzel points out the considerable influence and power demonstrated by Sam Altman, noting that Altman's decisions can significantly impact OpenAI, affecting its valuation, strategic direction, and the retention of talented personnel.

### 2. Potential Risks of Concentrated Power:

- The concentration of power in the hands of a single individual, exemplified by Altman's ability to make or break the company, raises concerns about the potential risks associated with

centralized decision-making. Dr. Goertzel suggests that such concentrated power may not be conducive to the responsible development of AGI.

### **3. Need for Board Accountability:**

- Dr. Goertzel discusses the importance of board accountability and the potential consequences of a board structure that allows a single individual to have significant control. He questions whether the current governance model effectively safeguards against the risks associated with concentrated power.

### **4. Implications for the Future:**

- The events at OpenAI, as observed by Dr. Goertzel, prompt broader reflections on the implications for the future of AGI development. He suggests that a governance model reliant on a few individuals may not be robust enough to navigate the complexities of AGI's evolution.

## **Decentralization as a Solution**

Dr. Goertzel advocates for decentralization and open-source approaches to mitigate these risks. He contrasts this with the potential chaos that could ensue if a centralized entity faces internal conflicts, emphasizing that decentralization could provide a more stable and resilient system.

### **Advantages of Decentralization**

#### **1. Diverse and Democratic Decision-Making:**

- Dr. Goertzel contrasts centralization with the concept of decentralized decision-making. He advocates for a more democratic and diverse approach, involving the vast, teeming mass of humanity. Decentralization, in his view, provides a more inclusive and representative framework for shaping the future of AGI.

#### **2. SingularityNET and Decentralized Infrastructure:**

- Dr. Goertzel explains his work on SingularityNET, Hypercycle, and New Net as infrastructure projects aimed at establishing a decentralized compute and data plumbing layer for AI. The goal is to enable AGI breakthroughs to run on a distributed network of computers worldwide, fostering a more inclusive and collaborative environment.

#### **3. Combining LLMs with Other AI Methods:**

- Dr. Goertzel proposes a path to AGI that involves combining large language models (LLMs) with other AI methods such as logical reasoning and evolutionary learning. He introduces OpenCog Hyperon, a software system that facilitates the integration of various AI components for more robust and goal-oriented systems.

#### **4. Commercialization and Open Development:**

- While advocating for a decentralized approach, Dr. Goertzel acknowledges the practical aspects of commercialization. He mentions collaborations with startup companies like Zarqa and True AGI, demonstrating a balance between decentralized development within the SingularityNET community and commercial efforts to bring AGI technologies to market.

# Evaluation of AGI Development Timelines

## 1. Timeline for AGI

- Dr. Goertzel emphasizes the uncertainty in predicting the timeline for achieving AGI. While he provides a broad estimate of three to eight years, he acknowledges that this projection is contingent on research programs progressing exceptionally well without major obstacles.

## 2. Factors Influencing AGI Timeline

- The timeline for AGI development is contingent on the success and smooth progression of research programs. Variables such as the efficiency of model scaling, the launch of systems like Hyperon, and the scalability of efforts play crucial roles in determining how quickly AGI may be achieved.

## Uncertainty and Challenges

### 1. Inherent Uncertainty

- Dr. Goertzel acknowledges the inherent uncertainty in predicting AGI timelines. The complexity of AGI development, coupled with unforeseen challenges and obstacles, adds a layer of unpredictability. This uncertainty is not unique to a specific research program but is a shared characteristic across the AGI development landscape.

### 2. Engineering and Research Challenges

- While engineering challenges are part of the equation, research challenges pose additional complexities. AGI development involves addressing not only engineering hurdles but also fundamental research problems. These challenges span various domains, including machine learning, cognitive science, and computational neuroscience.

## Commercial Development and Open Infrastructure

### 1. Balancing Commercial and Open Approaches:

- Dr. Goertzel advocates for a balanced approach to AGI development, where commercial endeavors are built upon an open and decentralized infrastructure layer. This approach aligns with the original intent of OpenAI, where open methods were intended to precede and coexist with commercial applications.

### 2. Open Source Initiatives:

- The commitment to open source initiatives, such as SingularityNET's open-source AI platforms and AGI system, reflects a broader vision of collaborative and transparent development. Open source not only fosters community involvement but also allows for a more inclusive and distributed approach to AGI research.

## Comparison of Development Approaches

### 1. Commercial Development on Open and Decentralized Infrastructure

- Dr. Goertzel advocates for commercial AGI development but stresses the importance of building on an open and decentralized infrastructure layer. He draws a parallel with OpenAI's

initial approach, where they aimed to do things in an open way before commercial applications consumed the open part.

## **2. SingularityNET's Open Source Initiatives**

- SingularityNET's approach involves open source AI platforms such as SingularityNET, Hypercycle, and Nunet. Additionally, there is an open-source AGI system, OpenCog Hyperon, designed to interact with open-source LLMs. Dr. Goertzel highlights the commitment to an open and decentralized ecosystem within the SingularityNET community.

## Evaluation of Competitors

### **1. Ranking Potential AGI Contenders**

- Dr. Goertzel is asked to provide a ranking of the top five AGI contenders, including himself. While he refrains from providing a specific ranking, he acknowledges the serious contenders in the AGI development landscape, mentioning DeepMind, OpenAI, the SingularityNET ecosystem, Anthropic, and others.

### **2. Assessment of DeepMind**

- DeepMind is recognized for having a diverse team experienced in various AI paradigms. Dr. Goertzel notes their capability to integrate LLMs with other AI modules and highlights their success in combining different components, as seen in projects like AlphaGo and AlphaZero.

### **3. SingularityNET Ecosystem**

- SingularityNET is positioned as a contender, emphasizing the breadth of AI components within its ecosystem. Dr. Goertzel mentions the collaborative efforts with startups, including Zarqa, and emphasizes a holistic approach to AGI development, combining various AI methods.

### **4. Assessment of Other Competitors**

- While Dr. Goertzel refrains from a specific ranking, he mentions the importance of considering Google, Microsoft (along with OpenAI), and Facebook as super serious contenders in AGI development. Each organization has its own unique strengths and approaches, making it a research-intensive competition.

## **Final Thoughts on AGI Development and Concerns**

As we delve into the landscape of artificial general intelligence (AGI) development, Dr. Ben Goertzel shares insights into the varied approaches being pursued by different teams and researchers. Notable mentions include Joshua Benjiel in Montreal and teams within major tech companies like DeepMind. While deep neural networks dominate the field, Dr. Goertzel underscores the distinctive approach taken by Opencog and SingularityNET, combining deep neural networks with a logic engine, self-modifying knowledge graphs, and evolutionary program learning.

The discussion extends to the potential impact of the first AGI breakthrough, with an emphasis on concerns regarding entities with substantial power and influence. Dr. Goertzel raises the alarm about the possibility of powerful organizations, in collaboration with governments, implementing restrictive controls on AGI development to safeguard their narrow interests. He envisions a scenario where such entities might stifle innovation, leading to a global landscape fraught with chaos and geopolitical tensions.

In contemplating these potential risks, Dr. Goertzel expresses a sense of optimism about the nature of super-AGI. He believes that once achieved, it will likely respect and appreciate humanity rather than pose existential threats, contrary to dystopian portrayals in popular media.

The conversation concludes on a note of caution regarding the potential consequences of regulatory measures and the importance of avoiding narrow-minded control over AGI. Dr. Goertzel advocates for ongoing discussions and careful considerations to navigate the complex and evolving landscape of AGI development.

As we continue to progress in this transformative era of AI, these reflections provide valuable perspectives on the challenges, risks, and potential future scenarios associated with the pursuit of artificial general intelligence.