The Tong Test

By John Benson

May 2025

1. Introduction

It is agreed that "artificial general intelligence" will be a cybernetic tool with many human behaviors. But is the "human" more like Albert Einstein, or say, a "digital Chinese girl that looks to be about 3 to 4 years old." Although the former will be a very powerful tool, the latter may be a more powerful tool that will help humans understand how we develop, learn and mature.

I'm thinking the latter may be more useful in the long term.

2. DeepSeek

Sometimes, less is more. In January, DeepSeek released the latest version of its chatbot, upending the artificial intelligence (AI) world. A training AI built for under $6 million, DeepSeek seems to rival the technical capabilities of other large language model (LLM) AIs, including ChatGPT, with only a fraction of the processing power. The breakthrough was a welcome development for Song-Chun Zhu, dean of the Institute for Artificial Intelligence at Peking University in Beijing, who has been challenging the current LLM-dominated AI paradigm in his efforts to create artificial general intelligence (AGI).[1]

Zhu, a trailblazer in the AI field, graduated from Harvard University in 1996 and has published more than 400 papers covering computer vision, cognitive science, robot autonomy, and commonsense reasoning, among other topics. Now, he is the founder and director of the non-profit Beijing Institute for General Artificial Intelligence (BIGAI).

"We as a society may have misunderstood the term 'AI'," says Zhu. "Just like how we call a multifunctional cellphone 'smart', the popular AI models we use today are not truly intelligent." That's because today's AI, he explains, is driven by big data built upon massive computing power. Zhu pioneered data-driven statistical approaches and created the world's first large-scale annotated image dataset at the Lotus Hill Institute in 2005. However, he realized that big data sets and specific machine learning models alone are not enough to make true intelligence. "One of the major Chinese philosophical schools, the Yangmingism or the 'Teachings of the Heart', argues 'the reality we see comes from how our minds perceive" Zhu says." To make AI more like humans it needs to have a framework that emulates the top-down mechanisms in the brain."

According to Zhu, the future of AGI should be a kind of autonomous AI that doesn't require vast datasets. In 2020, Zhu returned to China to establish and lead the BIGAI. Its mission: To pursue a unified theory of artificial intelligence in order to create general intelligent agents for lifting humanity.

[1] Custom Publishing Office, Science Magazine, March 28, 2025 Issue, "AI gets a mind of its own," https://www.science.org/content/article/ai-gets-mind-its-own

2.1. Defining AGI Agents in CUV-Space

Zhu and his team's focus at BIGAI is on creating value-driven human-like cognition that goes beyond data-driven imitation. "The difference between artificial general intelligence (AGI) and current LLM-based AI is just like the difference between a crow and a parrot," he said. While parrots can mimic many words, he says, crows can achieve their goals autonomously in the real world. In an article published in 2017, Zhu discusses how statistical models, which modern LLMs are based upon, function like randomized parrots. While leading two Multidisciplinary University Research Initiatives at UCLA, Zhu pursued research to make machines more crow-like, exploring the brain mechanisms that make it possible for crows—and humans—to understand the physical and social world and act accordingly.

Human intelligence evolves over time, as the body changes and experiences accrue. AGI also matures over time. To help define, evaluate, and improve AGI development, Zhu proposed to define AGI in the mathematical space of the "CUV framework" In this framework, C is the AGI's "cognitive architecture" to think, or its simulation of the decision-making processes in the brain. U is a set of "potential functions" that represent an AGI's ability to understand and interact with its environment. V is a set of hierarchical internal "value functions" that supply the AGI's motivation. With this formulation, Zhu and colleagues can define AGI agents as points in this CUV space and characterize their learning and self-reflection processes.

2.2. The Tong Test

In Chinese, the word "general" is translated as Tong (the tong test), a character that is also the logo of BIGAI. (通)

Artistically arranged, the character also holds the English letters "AGI." Tong Tong is the name Zhu gave to world's first AGI agent born at BIGAI, a digital Chinese girl that looks to be about 3 to 4 years old. Tong Tong is a step forward in AGI research, and researchers really want to know, "What is she thinking?" and "How is she learning and making decisions?" Researchers have long relied on tests to assess AI models. The Turing test was developed to determine whether a machine could mimic human intelligence through dialogue. ChatGPT and other AI built on big data can pass the Turing test, but Zhu wanted a test that could assess broad human intelligence. Thus, the Tong Test was born, which relies on the CUV framework.

What sets Tong Tong apart from ChatGPT is that she doesn't exist in a vacuum, but is rather embodied in a virtual world that emulates the complexity of the real physical social world. The Tong test examines an AGI's understanding of this world— its abilities—as well as the AGI's internal motivations for behaviors—its values. For example, how an AGI responds to a crying baby sitting on a floor can say a lot about its commonsense reasoning, inference of social interactions, and self-awareness. "Those natural abilities such as emotions and languages are true embodiment of human intelligence," Zhu says. "Tong Tong may be an AGI agent, but she is just like a real human child, able to understand and behave according to her own environment even if it changes. The goal of the Tong test is to build a systematic evaluation system to promote standardized, quantitative, and objective benchmarks and evaluation for AGI." And Tong Tong is just the beginning; researchers at BIGAI are developing diverse AGI agents that may someday enter the physical world through robotics and other mediums to serve society in meaningful ways.

2.3. AGI Safety

As Tong Tong and the Tong test continue to grow and mature, AGI safety is front of mind for Zhu. Because AGI behavior is human-like, and not all humans are benevolent, there are risks that AGI will take actions that are not in humanity's best interests. On the other hand, AGI's cognitive architecture may be able to incorporate a mutual theory of mind—in other words, the golden rule: do unto others as you would have them do unto you.

During a panel discussion at SafeAI 2023, Zhu and Stuart Russell from the University of California Berkeley, two leading figures in AGI, had an in-depth discussion on the risks and ethics of AGI.

When Russell raised a question about how humans could keep AGI agents in check, Zhu replied, "To prevent potential threats from future AGI agents to humanity, we can gradually loosen the capability and value space of agents. It's similar to how we approach robots: initially, we confine them in a 'cage' and slowly increase their permission. Now, we already have autonomous vehicles operating on specific roads." Zhu added that once AGI agents are proven safe and controllable, they can have more freedom, with the safeguard of understanding and transparency. "If we can explicitly represent the cognitive architecture of AGI agents, understanding how they work, we will be better equipped to control them."

For Zhu, now is the beginning of a new era for AI to evolve into AGI. Zhu's doctoral advisor at Harvard, mathematician and Fields medalist David Mumford, is also an advocate of creating AIs with the top-down neural architecture of the human brain. He gave Zhu a trophy to recognize his perseverance at AGI innovation. "The future of AGI will be a combination of science and philosophy," Zhu says. "Chinese teachings of the heart are crucial to guiding AGI to obtain true beneficial human behavior."

Final author's comment: If Artificial Intelligence is the leading edge of computer science, the artificial general intelligence (AGI) is the sharpest part of that edge, and might lead to many amazing developments. Song-Chun Zhu's new wrinkle on AGI should lead to much better understanding of how the human mind learns and otherwise reasons.

3. Song-Chun Zhu Biographical Information

Although the prior article contains some biographical information on Song-Chun Zhu, I thought that would be best to include his official biographical page from Institute for Artificial Intelligence at Peking University. This is below.

Song-Chun Zhu

Title:	Chair Professor
Department	Institute for Artificial Intelligence
Research Areas	Computer Vision, Statistical Modeling & Computing, Cognition, Machine Learning, Natural Language and Dialogue, AI, Robot Autonomy.
E-mail	s.c.zhu@pku.edu.cn
Homepage	https://zhusongchun.net

Degrees:

1996 Ph.D., Harvard University, Cambridge, MA

1994 M.S., Harvard University, Cambridge, MA

1991 B.S., University of Science and Technology of China, at Hefei, China

Appointments:

2020 Director, Beijing Institute for General Artificial Intelligence (BIGAI)

2020 Chair Professor, Peking University, Director of Institute for Artificial Intelligence

2006 Professor, University of California at Los Angeles, Depts. of Statistics, Computer Science

2002 Associate Professor, University of California at Los Angeles, Depts. of Statistics, Computer Science

1998 Assistant Professor, Ohio State University, Depts. of Computer Science, Cognitive Science

1997 Lecturer, Stanford University, Dept. of Computer Science

1996 Post-doc, Brown University, Division of Applied Math.

Academic Honors:

2020 Best Paper Award, ICML Workshop on Bridge Perception and Reasoning, with Qing Li et al.

2019 Best Paper Award, ACM TURC , with Xu Xie et al.

2017 Computational Modeling Prize, Cognitive Science Society, with Tianmin Shu et al.

2013 Helmholtz Test-of-Time Award***, 14th Int'l Conf. on Computer Vision at Sydney, Australia, for a region competition paper published in 1995.

2011 Fellow, IEEE Computer Society

2008 Aggarwal Prize**, the Int'l Association of Pattern Recognition. Citation

2007 Marr Prize honorary nomination, 11th Int'l Conf. on Computer Vision at Rio, Brazil, for Object Modeling with Y. Wu et al.

2003 Marr Prize, 9th Int'l Conf. on Computer Vision at Nice, France, for MCMC Inference for Image Parsing with Z. Tu et al.

2001 Young Investigator Award, Office of Navy Research.

2001 Sloan Fellow, Alfred P. Sloan Foundation.

2001 Career Award, National Science Foundation.

1999 Marr Prize honorary nomination*, 7th Int'l Conf. on Computer Vision at Corfu, Greece, for Texture Modeling with Y. Wu.

1995 Jury Prize, Harvard University.

1992 Harvard Fellowship, Harvard Graduate School of Art and Sciences.

1986-1991 Numerous undergraduate student awards in China

*The Marr prize used to be the highest honor in computer vision as it was the only award given in the 1980s-1990s, and early 2000s every two years. But in recent years, the community has created a few other categories. ECCV and CVPR also introduced their best paper prizes.

** The Aggarwal prize started from 2006, and is given by the International Association of Pattern Recognition for one person under 40 years old in every two years.

***This new award started in ICCV 2009. It is for a paper published more than 10 years ago. This was actually my first International Conference on Computer Vision (ICCV) paper.

Professional Activities

Int'l Association of Pattern Recognition, committee for the J.K.Aggarwal Prize (2008--2013) , Chair (2011-13).

IEEE Computer Society, Fellow committee, Vice Chair, 2013.

General Chair, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Providence, RI, 2012.

Editing Board, International Journal of Computer Vision (2004 -- Present)

Editing Board, Foundations and Trends in Computer Graphics and Vision (2004 -- Present)

Editing Board, IEEE Transactions on Pattern Analysis and Machine Intelligence (2005 -- 2009)

General Chair, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, 2019.

Principal Investigator, MURI (Multidisciplinary University Research Initiative)[2] program on Scene Understanding 2010-2015.

Principal Investigator, MURI program on Visual Commonsense Reasoning, 2016-2021

[2] Sponsored by UCLA and supported by the Office of Naval Research grant, https://vcla.stat.ucla.edu/MURI_Scene_Understanding/index.html