

# Evolving definitions of powerful AI

Please comment or email ([jaswsunny at gmail dot com](mailto:jaswsunny@gmail.com)) if something is wrong/missing!

Subtypes of definitions:

- Processes: Does AI think in “human-like” ways?
- Generality: Can AI learn new tasks it hasn’t seen?
- Performance: Can AI complete various tasks?
- Impacts: Will the AI have a big impact on society?

Year	Person	Definition and description	Link and context	Subtype
1947	Alan Turing	<p>“A good working rule is that the ACE can be made to do <b>any job that could be done by a human</b> computer, and will do it in one ten-thousandth of the time.”</p> <p>“I was researching on what might now be described as an [3] investigation of the theoretical possibilities and limitations of digital computing machines. I considered a type of machine which had a central mechanism, and an infinite memory which was contained on an infinite tape. This type of machine appeared to be sufficiently general.”</p>	<p>“Lecture to the London Mathematical Society” from the <i>Collected Works of Alan M. Turing, Volume 1: Mechanical Intelligence</i></p> <p>First known concept of a machine intelligent enough to complete human-like tasks, learn from experience, etc. Also discusses automation, training, alignment, hallucination.</p>	Processes, Performance
1955	John McCarthy, Marvin Minsky, Nathaniel Rochester, Claude Shannon	<p>“We propose that a 2 month, 10 man study of <b>artificial intelligence</b> be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.”</p>	<p><a href="#">Proposal for the Dartmouth Summer Research Project on AI</a></p> <p>First published usage of “artificial intelligence”</p>	
1980	John Searle (philosopher)	<p>“According to <b>strong AI</b>, the computer is not merely a tool in the study of the mind; rather, the appropriately programmed computer really is a mind, in the sense that computers given the right programs can be literally said to understand and have other cognitive states.”</p>	<p><a href="#">Minds, brains, and programs</a></p> <p>Searle’s “Strong AI” hypothesis, which he refutes, suggests that a computer program simulating a mind is also conscious and has its own mental states. (This is not how most AI researchers or theorists would define AI.)</p>	Processes
1997	Mark Avrum Gubrud (physicist)	<p>“By <b>advanced artificial general intelligence</b>, I mean AI systems that rival or surpass the human brain in complexity and speed, that can acquire, manipulate and reason with general knowledge, and that are usable in essentially any phase of ... operations where a human intelligence would</p>	<p>Used in context of advanced military technology in article <a href="#">“Nanotechnology and National Security”</a></p>	Generality, Performance

		otherwise be needed.”	First recorded usage of AGI but went unnoticed by AI community	
2002	Ben Goertzel & Shane Legg	<p>“What is meant by <b>AGI</b> is, loosely speaking, AI systems that possess a reasonable degree of self-understanding and autonomous self-control, and have the ability to solve a variety of complex problems in a variety of contexts, and to learn to solve new problems that they didn’t know about at the time of their creation.</p> <p>A marked distinction exists between practical AGI work and, on the other hand:</p> <ul style="list-style-type: none"> <li>• Pragmatic but specialized “<b>narrow AI</b>” research which is aimed at creating programs carrying out specific tasks like playing chess, diagnosing diseases, driving cars and so forth (most contemporary AI work falls into this category.)</li> <li>• Purely theoretical AI research, which is aimed at clarifying issues regarding the nature of intelligence and cognition, but doesn’t involve technical details regarding actually realizing artificially intelligent software.”</li> </ul>	<p>Goertzel's <a href="#">book</a> (pub. 2005) needed a title. He didn’t like “strong AI,” and Shane came up with “AGI.”</p> <p><a href="https://goertzel.org/who-coined-the-term-agi">https://goertzel.org/who-coined-the-term-agi</a></p> <p>Goertzel and Legg are credited with popular diffusion of the term AGI.</p>	Processes, Generality, Performance
2005	Ray Kurzweil	<p>“Artificial intelligence permeates our economy. It’s what I define as “<b>narrow</b>” AI: machine intelligence that equals or exceeds human intelligence for specific tasks... So what are the prospects for “<b>strong</b>” AI, which I describe as machine intelligence with the full range of human intelligence?”</p>	<a href="#">Long Live AI / The Singularity Is Near</a>	Generality
2005	Nils Nilsson	<p>“Machines exhibiting true <b>human-level intelligence</b> should be able to do many of the things humans are able to do. Among these activities are the tasks or “jobs” at which people are employed.</p> <p>I suggest we replace the Turing test by something I will call the “employment test.” To pass the employment test, AI programs must be able to perform the jobs ordinarily performed by humans. Progress toward human-level AI could then be measured by the fraction of these jobs that can be acceptably performed by machines.”</p> <p>“Rather than work toward this goal of automation by building special-purpose systems, I argue for the development of general-purpose, educable systems that can learn and be taught to perform any of the thousands of jobs that humans can perform.”</p>	<p>Nilsson says “general-purpose, educable” systems are the right approach to achieve HLMI</p> <p>Proposes an “employment test” to measure AI capabilities</p> <p><a href="#">Human-Level Artificial Intelligence? Be Serious!</a></p>	Performance, Impacts
2007	Shane Legg & Mark Hutter	<p>“<b>Intelligence</b> measures an agent’s ability to achieve goals in a wide range of environments.”</p> <p><b>universal intelligence</b></p>	<p>First formal definition of “universal intelligence”</p> <p><a href="https://arxiv.org/pdf/0712.3329">https://arxiv.org/pdf/0712.3329</a></p>	Performance

		<p>Bringing all these pieces together, we can now define our formal measure of intelligence for arbitrary systems. Let <math>E</math> be the space of all computable reward summable environmental measures with respect to the reference machine <math>\mathcal{U}</math>, and let <math>K</math> be the Kolmogorov complexity function. The expected performance of agent <math>\pi</math> with respect to the universal distribution <math>2^{-K(\mu)}</math> over the space of all environments <math>E</math> is given by,</p> $T(\pi) := \sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}.$ <p>We call this the <i>universal intelligence</i> of agent <math>\pi</math>.</p>		
2011	Nick Bostrom & Eliezer Yudkowsky	<p><b>“Artificial General Intelligence”</b> (hereafter, AGI) is the emerging term of art used to denote “real” AI (see, e.g., the edited volume Goertzel and Pennachin 2006). As the name implies, the emerging consensus is that the missing characteristic is generality. Current AI algorithms with human-equivalent or -superior performance are characterized by a deliberately-programmed competence only in a single, restricted domain... It is a qualitatively different class of problem to handle an AGI operating across many novel contexts that cannot be predicted in advance.”</p>	<a href="#">The Ethics of Artificial Intelligence</a>	Generality, Performance
2014	Vincent C. Müller & Nick Bostrom	<p>“Define a <b>‘high-level machine intelligence’ (HLMI)</b> as one that can carry out most human professions at least as well as a typical human.”</p> <p>“We need to avoid using terms that are already in circulation and would thus associate the questionnaire with certain groups or opinions, like “artificial intelligence”, “singularity”, “artificial general intelligence” or “cognitive system”. For these reasons, we settled for a definition that a) is based on behavioral ability, b) avoids the notion of a general ‘human-level’ and c) uses a newly coined term.”</p>	<p>2014 survey of AI expert opinion — HLMI is used to be debiasing</p> <p><a href="#">Future Progress in Artificial Intelligence: A Survey of Expert Opinion</a></p>	Performance
2014	Nick Bostrom	<p><b>Superintelligence</b> is defined as “any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest.”</p>	From <i>Superintelligence</i> book	Performance, Impacts
2016	Holden Karnofsky (OpenPhil)	<p>“Roughly and conceptually, <b>transformative AI</b> refers to potential future AI that precipitates a transition comparable to (or more significant than) the agricultural or industrial revolution. I also provide (below) a more detailed definition.</p> <p>The concept of “transformative AI” has some overlap with concepts put forth by others, such as “superintelligence” and “artificial general intelligence.” However, “transformative AI” is intended to be a more inclusive term, leaving open the possibility of AI systems that count as “transformative” despite lacking many abilities humans have.”</p>	<a href="#">Some Background on Our Views Regarding Advanced Artificial Intelligence   Open Philanthropy</a>	Impacts
2016	Katja Grace et al. (FHI)	<p>“<b>“High-level machine intelligence”</b> (HLMi) is achieved when unaided machines can accomplish every task better and more cheaply than human workers.”</p>	Regular survey of AI experts on views on AGI, first in <a href="#">2016</a>	Performance, Impacts
2018	OpenAI (Sam Altman, et al?)	<p>“OpenAI’s mission is to ensure that artificial general intelligence (<b>AGI</b>) — by which we mean highly autonomous systems that outperform humans at most economically valuable work — benefits all of humanity.”</p>	Defined in the OpenAI charter: <a href="https://web.archive.org/web/20180409161852/https://blog.openai.com/openai-charter/">https://web.archive.org/web/20180409161852/https://blog.openai.com/openai-charter/</a>	Performance, Impacts

2019	Francois Chollet	<p>“The <b>intelligence</b> of a system is a measure of its skill-acquisition efficiency over a scope of tasks, with respect to priors, experience, and generalization difficulty.”</p> <p>““General intelligence” is not a binary property which a system either possesses or lacks. It is a spectrum, tied to 1) a scope of application, which may be more or less broad, and 2) the degree of efficiency with which the system translate its priors and experience into new skills over the scope considered, 3) the degree of generalization difficulty represented by different points in the scope considered. It is conceptually unsound to set “artificial general intelligence” in an absolute sense (i.e. “universal intelligence”) as a goal.”</p> <p>“The consensus definition of <b>AGI</b>, “a system that can automate the majority of economically valuable work,” while a useful goal, is an incorrect measure of intelligence. Measuring task-specific skill is not a good proxy for intelligence.”</p>	<p><a href="#">On the Measure of Intelligence</a> <a href="#">What is ARC-AGI?</a></p> <p>Chollet has been a critic of other definitions of AGI, arguing that they measure memorization rather than generality / new-skill acquisition</p>	Generality
2022	Yann LeCun (Meta)	<p>“I think the phrase <b>AGI</b> should be retired and replaced by “<b>human-level AI</b>”. There is no such thing as AGI. Even human intelligence is very specialized. We do not realize that human intelligence is specialized because all the intelligent tasks we can think of are task that we can apprehend. But that is a tiny subset of all tasks. The overwhelming majority of tasks are completely out of reach of un-augmented human intelligence.”</p>	<p><a href="#">LinkedIn post</a></p>	Performance
2023	Microsoft Research (Sebastien Bubeck et al)	<p>“We use <b>AGI</b> to refer to systems that demonstrate broad capabilities of intelligence, including reasoning, planning, and the ability to learn from experience, and with these capabilities at or above human-level.”</p>	<p>Sparks of AGI paper <a href="https://arxiv.org/abs/2303.12712">https://arxiv.org/abs/2303.12712</a></p>	Processes, Performance
2024	Meredith Morris et al (DeepMind)	<p>“Artificial General Intelligence (<b>AGI</b>) is an important and sometimes controversial concept in computing research, used to describe an AI system that is at least as capable as a human at most tasks.”</p>	<p><a href="#">Levels of AGI for Operationalizing Progress on the Path to AGI</a></p> <p>Emphasizes generality <i>and</i> performance – both exceeding human performance on tasks, and adapting to new tasks via learning</p>	Generality, Performance

		<table border="1"> <tr> <td><b>Performance (rows) x Generality (columns)</b></td> <td><b>Narrow</b> <i>clearly scoped task or set of tasks</i></td> <td><b>General</b> <i>wide range of non-physical tasks, including metacognitive tasks like learning new skills</i></td> </tr> <tr> <td><b>Level 0: No AI</b></td> <td><b>Narrow Non-AI</b> calculator software; compiler</td> <td><b>General Non-AI</b> human-in-the-loop computing, e.g., Amazon Mechanical Turk</td> </tr> <tr> <td><b>Level 1: Emerging</b> <i>equal to or somewhat better than an unskilled human</i></td> <td><b>Emerging Narrow AI</b> GOPAI (Boden, 2014); simple rule-based systems, e.g., SHRDLU (Winograd, 1971)</td> <td><b>Emerging AGI</b> ChatGPT (OpenAI, 2023), Bard (Anil et al., 2023), Llama 2 (Touvron et al., 2023), Gemini (Pichai &amp; Hassabis, 2023)</td> </tr> <tr> <td><b>Level 2: Competent</b> <i>at least 50th percentile of skilled adults</i></td> <td><b>Competent Narrow AI</b> toxicity detectors such as Jigsaw (Das et al., 2022); Smart Speakers such as Siri (Apple), Alexa (Amazon), or Google Assistant (Google); VQA systems such as PaLI (Chen et al., 2023); Watson (IBM); SOTA LLMs for a subset of tasks (e.g., short essay writing, simple coding)</td> <td><b>Competent AGI</b> not yet achieved</td> </tr> <tr> <td><b>Level 3: Expert</b> <i>at least 90th percentile of skilled adults</i></td> <td><b>Expert Narrow AI</b> spelling &amp; grammar checkers such as Grammarly (Grammarly, 2023); generative image models such as Imagen (Saharia et al., 2022) or Dall-E 2 (Ramesh et al., 2022)</td> <td><b>Expert AGI</b> not yet achieved</td> </tr> <tr> <td><b>Level 4: Virtuoso</b> <i>at least 99th percentile of skilled adults</i></td> <td><b>Virtuoso Narrow AI</b> Deep Blue (Campbell et al., 2002), AlphaGo (Silver et al., 2016; 2017)</td> <td><b>Virtuoso AGI</b> not yet achieved</td> </tr> <tr> <td><b>Level 5: Superhuman</b> <i>outperforms 100% of humans</i></td> <td><b>Superhuman Narrow AI</b> AlphaFold (Jumper et al., 2021; Varadi et al., 2021), AlphaZero (Silver et al., 2018), Stockfish (Stockfish, 2023)</td> <td><b>Artificial Superintelligence (ASI)</b> not yet achieved</td> </tr> </table>	<b>Performance (rows) x Generality (columns)</b>	<b>Narrow</b> <i>clearly scoped task or set of tasks</i>	<b>General</b> <i>wide range of non-physical tasks, including metacognitive tasks like learning new skills</i>	<b>Level 0: No AI</b>	<b>Narrow Non-AI</b> calculator software; compiler	<b>General Non-AI</b> human-in-the-loop computing, e.g., Amazon Mechanical Turk	<b>Level 1: Emerging</b> <i>equal to or somewhat better than an unskilled human</i>	<b>Emerging Narrow AI</b> GOPAI (Boden, 2014); simple rule-based systems, e.g., SHRDLU (Winograd, 1971)	<b>Emerging AGI</b> ChatGPT (OpenAI, 2023), Bard (Anil et al., 2023), Llama 2 (Touvron et al., 2023), Gemini (Pichai & Hassabis, 2023)	<b>Level 2: Competent</b> <i>at least 50th percentile of skilled adults</i>	<b>Competent Narrow AI</b> toxicity detectors such as Jigsaw (Das et al., 2022); Smart Speakers such as Siri (Apple), Alexa (Amazon), or Google Assistant (Google); VQA systems such as PaLI (Chen et al., 2023); Watson (IBM); SOTA LLMs for a subset of tasks (e.g., short essay writing, simple coding)	<b>Competent AGI</b> not yet achieved	<b>Level 3: Expert</b> <i>at least 90th percentile of skilled adults</i>	<b>Expert Narrow AI</b> spelling & grammar checkers such as Grammarly (Grammarly, 2023); generative image models such as Imagen (Saharia et al., 2022) or Dall-E 2 (Ramesh et al., 2022)	<b>Expert AGI</b> not yet achieved	<b>Level 4: Virtuoso</b> <i>at least 99th percentile of skilled adults</i>	<b>Virtuoso Narrow AI</b> Deep Blue (Campbell et al., 2002), AlphaGo (Silver et al., 2016; 2017)	<b>Virtuoso AGI</b> not yet achieved	<b>Level 5: Superhuman</b> <i>outperforms 100% of humans</i>	<b>Superhuman Narrow AI</b> AlphaFold (Jumper et al., 2021; Varadi et al., 2021), AlphaZero (Silver et al., 2018), Stockfish (Stockfish, 2023)	<b>Artificial Superintelligence (ASI)</b> not yet achieved		
<b>Performance (rows) x Generality (columns)</b>	<b>Narrow</b> <i>clearly scoped task or set of tasks</i>	<b>General</b> <i>wide range of non-physical tasks, including metacognitive tasks like learning new skills</i>																							
<b>Level 0: No AI</b>	<b>Narrow Non-AI</b> calculator software; compiler	<b>General Non-AI</b> human-in-the-loop computing, e.g., Amazon Mechanical Turk																							
<b>Level 1: Emerging</b> <i>equal to or somewhat better than an unskilled human</i>	<b>Emerging Narrow AI</b> GOPAI (Boden, 2014); simple rule-based systems, e.g., SHRDLU (Winograd, 1971)	<b>Emerging AGI</b> ChatGPT (OpenAI, 2023), Bard (Anil et al., 2023), Llama 2 (Touvron et al., 2023), Gemini (Pichai & Hassabis, 2023)																							
<b>Level 2: Competent</b> <i>at least 50th percentile of skilled adults</i>	<b>Competent Narrow AI</b> toxicity detectors such as Jigsaw (Das et al., 2022); Smart Speakers such as Siri (Apple), Alexa (Amazon), or Google Assistant (Google); VQA systems such as PaLI (Chen et al., 2023); Watson (IBM); SOTA LLMs for a subset of tasks (e.g., short essay writing, simple coding)	<b>Competent AGI</b> not yet achieved																							
<b>Level 3: Expert</b> <i>at least 90th percentile of skilled adults</i>	<b>Expert Narrow AI</b> spelling & grammar checkers such as Grammarly (Grammarly, 2023); generative image models such as Imagen (Saharia et al., 2022) or Dall-E 2 (Ramesh et al., 2022)	<b>Expert AGI</b> not yet achieved																							
<b>Level 4: Virtuoso</b> <i>at least 99th percentile of skilled adults</i>	<b>Virtuoso Narrow AI</b> Deep Blue (Campbell et al., 2002), AlphaGo (Silver et al., 2016; 2017)	<b>Virtuoso AGI</b> not yet achieved																							
<b>Level 5: Superhuman</b> <i>outperforms 100% of humans</i>	<b>Superhuman Narrow AI</b> AlphaFold (Jumper et al., 2021; Varadi et al., 2021), AlphaZero (Silver et al., 2018), Stockfish (Stockfish, 2023)	<b>Artificial Superintelligence (ASI)</b> not yet achieved																							
2024	OpenAI	<p><a href="#">Bloomberg</a>: OpenAI has internally defined five “levels” of AGI</p> <p><b>OpenAI Imagines Our AI Future</b></p> <p><b>Stages of Artificial Intelligence</b></p> <table border="1"> <tr> <td>Level 1</td> <td>Chatbots, AI with conversational language</td> </tr> <tr> <td>Level 2</td> <td>Reasoners, human-level problem solving</td> </tr> <tr> <td>Level 3</td> <td>Agents, systems that can take actions</td> </tr> <tr> <td>Level 4</td> <td>Innovators, AI that can aid in invention</td> </tr> <tr> <td>Level 5</td> <td>Organizations, AI that can do the work of an organization</td> </tr> </table> <p>Source: Bloomberg reporting</p>	Level 1	Chatbots, AI with conversational language	Level 2	Reasoners, human-level problem solving	Level 3	Agents, systems that can take actions	Level 4	Innovators, AI that can aid in invention	Level 5	Organizations, AI that can do the work of an organization	<p>Bloomberg: <a href="#">OpenAI Sets Levels to Track Progress Toward Superintelligent AI (Archive link)</a></p> <p>(In a 2023 <a href="#">blog post</a>, Sam Altman similarly distinguishes between “initial AGI” and “successor systems”)</p>	Performance, Impact											
Level 1	Chatbots, AI with conversational language																								
Level 2	Reasoners, human-level problem solving																								
Level 3	Agents, systems that can take actions																								
Level 4	Innovators, AI that can aid in invention																								
Level 5	Organizations, AI that can do the work of an organization																								
2024	Dario Amodei (Anthropic)	<p>“I find <b>AGI</b> to be an imprecise term that has gathered a lot of sci-fi baggage and hype. I prefer “powerful AI” or “Expert-Level Science and Engineering” which get at what I mean without the hype.”</p> <p>“By <b>powerful AI</b>, I have in mind an AI model—likely similar to today’s LLM’s in form, though it might be based on a different architecture, might involve several interacting models, and might be trained differently—with the following properties:</p> <ul style="list-style-type: none"> <li>In terms of pure intelligence, it is smarter than a Nobel Prize winner across most relevant fields – biology, programming, math, engineering, writing, etc. This means it can prove unsolved mathematical theorems, write extremely good novels, write difficult codebases from scratch, etc.</li> <li>In addition to just being a “smart thing you talk to”, it has all the</li> </ul>	<p><a href="#">Machines of Loving Grace</a></p>	Performance																					

		<p>“interfaces” available to a human working virtually, including text, audio, video, mouse and keyboard control, and internet access. It can engage in any actions, communications, or remote operations enabled by this interface, including taking actions on the internet, taking or giving directions to humans, ordering materials, directing experiments, watching videos, making videos, and so on. It does all of these tasks with, again, a skill exceeding that of the most capable humans in the world.</p> <ul style="list-style-type: none"> <li>• It does not just passively answer questions; instead, it can be given tasks that take hours, days, or weeks to complete, and then goes off and does those tasks autonomously, in the way a smart employee would, asking for clarification as necessary.</li> <li>• It does not have a physical embodiment (other than living on a computer screen), but it can control existing physical tools, robots, or laboratory equipment through a computer; in theory it could even design robots or equipment for itself to use.</li> <li>• The resources used to train the model can be repurposed to run millions of instances of it (this matches projected cluster sizes by ~2027), and the model can absorb information and generate actions at roughly 10x-100x human speed<sup>5</sup>. It may however be limited by the response time of the physical world or of software it interacts with.</li> <li>• Each of these million copies can act independently on unrelated tasks, or if needed can all work together in the same way humans would collaborate, perhaps with different subpopulations fine-tuned to be especially good at particular tasks.</li> </ul> <p>We could summarize this as a “country of geniuses in a datacenter”.”</p>		
2025	Sam Altman (OpenAI)	<p>“Systems that start to point to AGI* are coming into view, and so we think it’s important to understand the moment we are in.</p> <p><b>AGI</b> is a weakly defined term, but generally speaking we mean it to be a system that can tackle increasingly complex problems, at human level, in many fields.”</p>	<p><a href="#">Three Observations</a></p> <p>Sam blog post is notably different from the OpenAI site definition, “highly autonomous systems that outperform humans at most economically valuable work”</p> <p>“At human level” and “Many fields” also seem quite vague!</p>	Performance