Global Catastrophic Risk
INSTITUTE

http://gcri.org

2020 Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy

McKenna Fitzgerald, Aaron Boddy, & Seth D. Baum
Global Catastrophic Risk Institute

Cite as: McKenna Fitzgerald, Aaron Boddy, and Seth D. Baum, 2020. 2020 Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy. Global Catastrophic Risk Institute Technical Report 20-1.

Executive Summary

Artificial general intelligence (AGI) is artificial intelligence (AI) that can reason across a wide range of domains. While most AI research and development (R&D) deals with narrow AI, not AGI, there is some dedicated AGI R&D. If AGI is built, its impacts could be profound. Depending on how it is designed and used, it could either help solve the world's problems or cause catastrophe, possibly even human extinction.

This paper presents a survey of AGI R&D projects that are active in 2020 and updates a previous survey of projects active in 2017. Both surveys attempt to identify every active AGI R&D project and characterize them in terms of relevance to ethics, risk, and policy, focusing on seven attributes:

- The type of institution in which the project is based
- Whether the project publishes open-source code
- Whether the project has military connections
- The nation(s) in which the project is based
- The project's goals for its AGI
- The extent of the project's engagement with AGI safety issues
- The overall size of the project

The surveys use openly published information as found in scholarly publications, project websites, popular media articles, and other websites. The 2020 survey uses information from the 2017 survey as well as the past three years of the *Journal of Artificial General Intelligence*, the past three years of AGI conference proceedings (the 2017 survey covered prior content from the *Journal of Artificial General Intelligence* and the AGI conference proceedings), keyword searches in Google web search, Google Scholar, Crunchbase, GitHub, the authors' prior knowledge, suggestions from readers of the 2017 survey, and additional literature and webpages identified via all of the above. The use of Crunchbase, GitHub, and reader suggestions is new to the 2020 survey.

The 2020 survey has two main findings. First, the expanded search methodology has identified a large number of new AGI R&D projects, bringing a more comprehensive picture of the field of AGI R&D. Second, accounting for the expanded search methodology, there has been little change in the field of AGI R&D between 2017 and 2020. Some 2017 projects are now inactive and some new projects have emerged in 2020, but the overall picture is largely the same. Specific trends are as follows:

- The 2020 survey identifies 72 active AGI R&D projects spread across 37 countries. The 2017 survey identified 45 projects in 30 countries. The 2020 survey updates the 2017 dataset, finding 70 projects that were active in 36 countries in 2017, 57 of which remain active in 2020.

- Relative to the 2017 survey, the AGI R&D projects presented in the 2020 survey tend to be smaller, more geographically diverse, less open-source, less focused on intellectual goals, more focused on humanitarian goals, and more concentrated in private corporations.

- The 2020 survey presents some small changes between 2017 and 2020 projects: a decrease in academic projects, an increase in private corporation projects, an increase in projects stating humanitarian goals, a decrease in projects with military connections, and a decrease in projects based in the United States. Other project attributes have been approximately constant.

Looking at just the 72 projects active in 2020, the trends can be summarized as follows:

- Many of the projects are interconnected via common personnel, common parent organizations, or project collaboration.

- About half of the projects are in private corporations. Academic institutions are the second-most common institution type.

- About half of the projects publish open-source code and about half do not.

- Only nine projects have identifiable military connections. The connections mostly involve basic research and sometimes involve tactical military applications. No major strategic military AGI R&D activities were identified.

- Almost half of the projects are based in the US, and most are based in a country allied with the US. Five projects are based in China. 20 projects are multinational, including six projects that operate in both China and the US.

- More than half of the projects state humanitarian goals. The second-most common type of stated goal is intellectual. Many projects have an orientation toward commercializing AGI as a product, though few of these explicitly state profit as a goal.

- Most projects are not active on AGI safety issues, and some are openly dismissive of AGI safety concerns, though some others have a significant emphasis on safety.

- Most projects are in the small-to-medium size range. The four largest projects are BlueBrain (an academic project based in Lausanne, Switzerland), DeepMind (a Google project based in London), the Human Brain Project (an academic project also based in Lausanne), and OpenAI (a nonprofit based in San Francisco). The largest projects are over 100 times larger than the smallest projects as measured in terms of full-time project personnel.

Looking across multiple attributes, some additional trends are apparent:

- There is a cluster of corporate projects are active on safety and state that their goals are to benefit humanity (i.e., is humanitarian).

- Another cluster is of academic projects that are not active on safety and state intellectual goals. Many of these projects have funding from military research agencies such as DARPA.

- Projects new to the 2020 survey are mostly small-to-medium-sized private corporations, many of which articulate humanitarian goals and are oriented toward commercializing AGI as a product without stating profit as a goal.

Additionally, the 2020 survey identifies, for the first time, a preponderance of projects that state a focus on AGI but demonstrate no other AGI activity. These projects are counted separately from the 72 projects described above.

Figure ES1 on the next page presents an overview of the data.

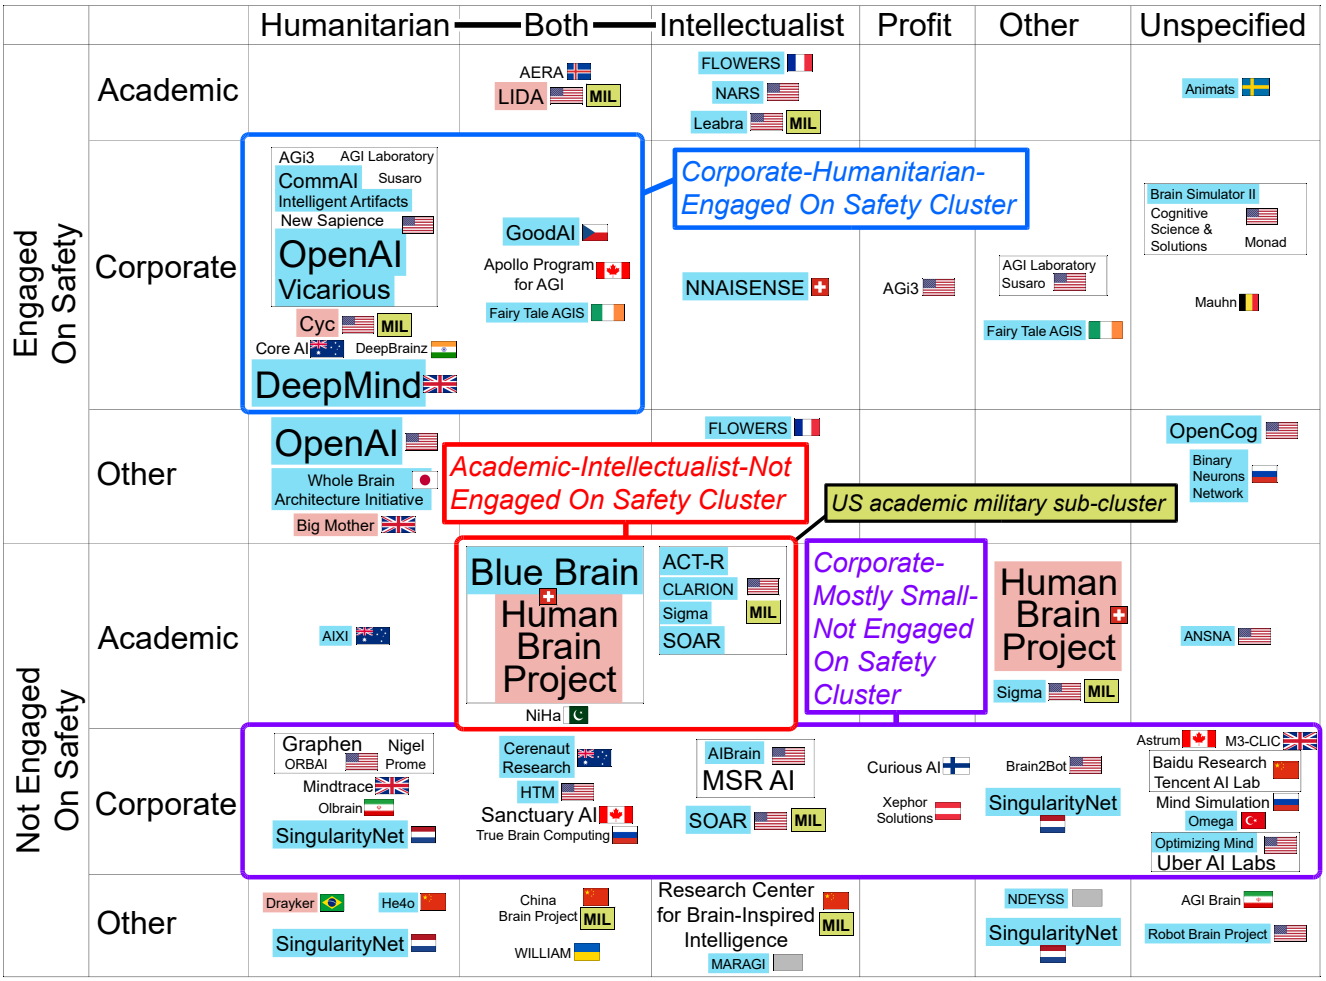

Figure ES1. Overview of the 72 identified AGI R&D projects characterized according to 7 attributes:

- **Institution type:** academic (mainly universities), corporate (public and private corporations), and other (government, nonprofit, and projects not affiliated with any institution).
- **Open-source code:** blue background indicates code available open-source; red background indicates code available upon request.
- **Military connections:** projects labeled MIL have military connections.
- **Nationality:** Flags indicate the country in which a project is based. Australia, Austria, Belgium, Brazil, Canada, China, Czech Republic, Finland, France, Iceland, India, Iran, Ireland, Japan, Netherlands, Pakistan, Russia, Sweden, Switzerland, Turkey, Ukraine, UK, USA, no project location identified
- **Stated goals:** animal welfare, ecocentrism, humanitarian, intellectualism, profit, transhumanism, and unspecified (when project goals could not be identified).
- **Engagement on safety:** engaged projects either actively try to make their AGI safe or state that they support other AGI safety work; projects that are not engaged on safety either openly dismiss concerns about AGI safety or have no publicly stated activity on safety.
- **Size:** font size of project names indicates project size.

The data suggest the following conclusions:

Regarding ethics, the major trend is that projects are split between stated goals of benefiting humanity and advancing knowledge, with the former largely from corporate projects and the latter largely from academic projects. While these are not the only goals that projects articulate, there appears to be a loose consensus among projects to aim for some combination of these two goals. The 2020 survey finds a small increase in the number of projects with humanitarian goals and a small decrease in the number of projects with intellectual goals. The 2020 survey also finds a large increase in the number of corporate projects, which are often motivated by profit even if they do not explicitly say that their goal is to generate profit.

Regarding risk, in particular the risk of AGI catastrophe, there is good news and bad news. Unfortunately, most projects are not adequately addressing AGI safety issues. Academic projects in particular have done relatively little to address safety issues, a trend that continues from 2017. The proliferation of corporate projects since 2017 heightens the concern that these projects could put profit ahead of safety and the public interest. Fortunately, however, there is much potential for projects to cooperate on safety issues, thanks to the partial consensus on goals, the concentration of projects in the US and its allies, and the various interconnections between different projects. Additionally, the absence of large government projects suggests that, at this time, no government is aggressively pursuing AGI for its strategic advantage.

Regarding policy, several conclusions can be drawn. First, the concentration of projects in the US and its allies could greatly facilitate the establishment of international public policy for AGI. Second, the large and growing number of corporate projects suggests an urgent need to attend to the corporate governance and political economy of AGI R&D. Third, the smaller but still significant number of academic projects suggests that research policy institutions, such as review boards that evaluate risky research, have an important role to play. Fourth, the large number of projects with open-source code presents a policy challenge because it makes it AGI R&D accessible to anyone anywhere in the world. Fifth, the large difference in size between the largest and smallest projects suggests that policymaking may benefit from a focus on larger projects. Finally, the absence of large government projects suggests that the primary role of governments may be as regulators of private-sector AGI R&D rather than as drivers of AGI R&D.

This study has some limitations, meaning that the actual state of AGI R&D may differ from what is presented here. The survey is based exclusively on openly published information and projects were sought out primarily using the English language. It is possible that this survey missed some AGI R&D projects. Data collection occurred primarily between June and September 2020 and may have missed more recent project information. The 72 projects identified by this survey therefore represent a lower bound on the number of existing projects. Furthermore, projects' actual attributes may differ from those found in openly published descriptions of those projects. For example, profit may be a goal of corporate projects even though most corporate projects did not state that profit was one of their goals. Therefore, this study's results should not be assumed to necessarily reflect the actual current state of AGI R&D. That said, the study nonetheless improves upon the 2017 survey and provides what is now the most thorough description yet of AGI R&D in terms of ethics, risk, and policy.

Finally, this document, presenting the 2020 survey, is largely based on the document that presented the 2017 survey. Much of the content remains the same, including the overall report design and various specific passages of text.

Acknowledgments

Kevin Chow, Maximilian Cornell, Fadi Al-Salti, and Daniel Johnston provided research assistance. Ben Goertzel, Gordon Irlam, Luke Muehlhauser, Robert de Neufville, Olga Afanasjeva, Andrea Owe, Tony Barrett, Brian Tse, Yawen Duan, Kwan Yee Ng, Justin Shovelain, and Marek Rosa provided feedback on previous drafts. Lloyd Harry-Davis provided assistance with copy editing. Melissa Thomas-Baum provided assistance with graphic design. Olga Afanasjeva, Colin Hales, and Marek Rosa suggested AGI R&D projects upon reading the 2017 version of this survey. Any lingering errors or other shortcomings are the authors' alone.

This work was made possible by a grant from the Gordon R. Irlam Charitable Foundation. The views in this paper are the authors' alone and do not necessarily reflect the views of the Global Catastrophic Risk Institute or the Gordon R. Irlam Charitable Foundation.

Contents

Main text

1. Introduction

Artificial general intelligence (AGI) is artificial intelligence (AI) that can reason across a wide range of domains, much like the human mind. Most AI is domain-specific and can only reason within a particular field or set of tasks. For example, DeepBlue can beat Garry Kasparov at chess, and maybe at a few other basic tasks like multiplication, but it cannot beat him at anything else. This is an example of narrow AI, not AGI. AGI was initially a primary goal of AI research and has long been considered its "grand dream" or "holy grail."[1] The technical difficulty of building AGI has led most of the field to focus on narrower, more immediately practical forms of AI, but some dedicated AGI research and development (R&D) continues.

AGI is also a profound societal concern. Or, rather, it could be if it is built. AGI could complement human intellect and increase the world's capacity to solve its problems. Or it could be used maliciously in a power play by whoever controls it. Or humanity may fail to control it, possibly resulting in catastrophe. Any AI can outsmart humans in some domains (e.g., multiplication), but an AGI may be able to outsmart humans in all domains. In that case, the outcome could depend on the AGI's goals: whether it seeks to benefit humanity, the world, itself, or to pursue some other goal entirely. Indeed, scholars of AGI sometimes propose that it could cause human extinction or a similar catastrophe (see literature review below).

The high potential stakes of AGI raise questions of ethics, risk, and policy. Which AGI, if any, should be built? What is the risk of catastrophe if an AGI is built? What policy options are available to avoid AGI catastrophe and, if desired, to develop safe and beneficial AGI? These are all questions under active investigation. However, the literature to date has tended to be theoretical and speculative, and has also tended to have little basis in the actual state of affairs in AGI R&D. Since AGI may not be built for many years, it is inevitable that the AGI literature will be somewhat speculative. Nevertheless, AGI R&D is happening right now. Information about current R&D can guide decisions about today's ethics, risk, and policy, as well as provide insight into what future R&D might look like.

This paper presents an update to the 2017 survey of active AGI R&D projects in terms of their ethical, risk, and policy dimensions (Baum 2017a). There have been several prior surveys of R&D on AGI and related technologies (Chong et al. 2007; Duch et al. 2008; Langley et al. 2008; de Garis et al. 2010; Goertzel et al. 2010; Samsonovich 2010; Taatgen and Anderson 2010; Thórisson and Helgasson 2012; Dong and Franklin 2014; Goertzel 2014; Kotseruba et al. 2016). However, these are all surveys on the technical aspects of AGI, such as how the AGI itself is designed, what progress has been made on it, and how promising the various approaches are for achieving AGI.

This paper presents and analyzes information of relevance to ethics, risk, and policy—for example, which political jurisdictions projects are located in and how engaged they are with AGI safety issues. Additionally, whereas prior surveys focus on select noteworthy examples of AGI R&D projects, this paper attempts to document all currently existing projects. In order to answer many questions related to ethics, risk, and policy, it is important to survey the entire range of projects—for example, to know which political jurisdictions should be included in AGI public policy.

Section 1.1 explains terminology related to AGI. Section 2 reviews prior literature on AGI ethics, risk, and policy. Section 3 presents the research questions pursued in this study. Section 4 summarizes the methodology used for this paper's survey. Section 5 presents the main results of the survey of AGI R&D projects and other notable projects. Section 6 concludes. Appendix 1 presents a full survey of AGI R&D projects that are active in 2020. Appendix 2 presents select inactive AGI R&D projects, specifically those that were active in 2017, are no longer active, and for which there is new coding information. Appendix 3 presents notable projects that were considered for the 2020 survey but excluded for not meeting inclusion criteria, including inactive 2017 projects for which there is no new

[1] See e.g. Legg (2008, p.125); https://montrealartificialintelligence.com

coding information. Appendix 4 presents other projects that have stated a focus on AGI but demonstrate little further activity on AGI.

1.1 Terminology

AGI is one of several terms used for advanced, and potentially transformative, future AI. The terms have slightly different meanings, and it is worth briefly distinguishing between them.

- *AGI* is specifically AI with a wide range of intelligence capabilities, including "the ability to achieve a variety of goals, and carry out a variety of tasks, in a variety of different contexts and environments" (Goertzel 2014, p.2). AGI is not necessarily advanced—an AI can be general without being highly sophisticated—though general intelligence requires a certain degree of sophistication and AGI is often presumed to be highly capable. AGI is also a dedicated field of study, with its own society (http://www.agi-society.org), journal (*Journal of Artificial General Intelligence*), and conference series (http://agi-conf.org).
- *Cognitive architecture* is the overall structure of an intelligent entity. One can speak of the cognitive architecture of the brains of humans or other animals. However, the term is used mainly for theoretical and computational models of human and nonhuman animal cognition. Cognitive architectures can be narrow, focusing on specific cognitive processes such as attention or emotion (Kotseruba et al. 2016). However, they are often general, and thus their study overlaps with the study of AGI. *Biologically inspired cognitive architectures* is a dedicated field of study with its own society (http://bicasociety.org), journal (*Biologically Inspired Cognitive Architectures*), and conference series (http://bicasociety.org/meetings).
- *Brain emulations* are computational instantiations of biological brains. Brain emulations are sometimes classified as distinct from AGI (e.g., Barrett and Baum 2017a). However, they are computational entities with general intelligence, and thus this paper treats them as a type of AGI.
- *Human-level AI* is AI with intelligence comparable to humans, or "human-level, reasonably human-like AGI" (Goertzel 2014, p.6). An important subtlety is that an AGI could be as advanced as humans, but with a rather different type of intelligence: it does not necessarily mimic human cognition. For example, AI chess programs will use brute force searches in some instances in which humans use intuition, yet the AI can still perform at or beyond the human level.
- *Superintelligence* is AI that significantly exceeds human intelligence. The term *ultraintelligence* has also been used in this context (Good 1965) but is less common. It is often proposed that superintelligence will come from an initial *seed AI* that undergoes *recursive self-improvement*, becoming successively smarter and smarter. The seed AI would not necessarily be AGI, but it is often presumed to be.

This paper focuses on AGI because the term is used heavily in R&D contexts, and it is important for ethics, risk, and policy. Narrow cognitive architectures (and narrow AI) are less likely to have transformative consequences for the world. Human-level AI and superintelligence are more likely to have transformative consequences. However, these terms are not common in R&D. Note that not every project included in this paper's survey explicitly identifies as an AGI project, but they all have the potential to be AGI or contribute to the development of AGI, and thus have potential for transformative consequences. The survey does not exclude any projects that are explicitly trying to build human-level AI or superintelligence.

2. Prior Literature

2.1 Ethics

Perhaps the most fundamental question in AGI ethics is on whether to treat AGI as an intellectual pursuit or as something that could impact society and the world at large. In other words, is AGI R&D pursued in order to advance the forefront of knowledge or to benefit society? Often, these two goals are closely connected, as evidenced by the central role of science and technology in improving living conditions worldwide. However, these goals are not always connected and can sometimes be at odds. In particular, for the present study, research into potentially dangerous new technologies can yield significant scientific and intellectual insights, yet end up being harmful to society.

Researchers across all fields of research often have strong intellectual values and motivations, and AGI is no exception. The question of whether to evaluate research in terms of intellectual merit or broader societal/global impacts is a point of ongoing contention across academia (Schienke et al. 2009). As with most fields, AI has traditionally emphasized intellectual merits, though there are calls for this to change (Baum 2018). The intellectual pull of AGI can be particularly strong, given its status as a long-term "grand dream" or "holy grail." However, the broader impacts can also have a strong pull, given the enormous potential stakes of AGI. In practical terms, an AGI project with intellectual motivations is, relative to a project motivated by broader impacts of AGI, more likely to view building AGI as a worthy goal in itself and to pay little attention to any potential dangers or other broader impacts.

The second area of AGI ethics concerns the goals that an AGI should be designed to pursue. This is the main focus of prior literature on AGI ethics. One line of thinking proposes "indirect normativity" or "coherent extrapolated volition," in which the AGI is designed to use its intellect to figure out what humanity wants it to do (Yudkowsky 2004; Muehlhauser and Helm 2012; Bostrom 2014). This proposal is motivated in part by procedural justice concerns—everyone, and not just AGI designers, should have a say in the AGI's ethics—and in part by concerns about the technical difficulty of programming the subtleties of human ethics directly into the AGI. The proposal raises questions about whose values to include (for example, should the values of nonhumans be included?), how to assess what each individual's values are, and how to resolve disagreements among individuals (Baum 2020).

An alternative line of thinking proposes that the AGI should create new entities that are morally superior to humans. This thinking falls in the realm of "transhumanism" or "posthumanism"; AGI researchers de Garis (2005) and Goertzel (2010) use the term "cosmism." This view holds that AGI should benefit the cosmos as a whole, not just humanity, and proposes that morally superior beings produced by AGI may advance the good of the cosmos. Whereas Goertzel (2010) stresses that "legacy humans" should be able to decide for themselves whether to continue in this new world, de Garis (2005) suggests that this world may be worth forming even if legacy humans would be eliminated. In contrast, Yampolskiy (2013) argues that AGI should only be built if they are expendable tools of benefit to their human creators.

Finally, there has also been some discussion of whether AGI should be built in the first place, although to date, less attention has been paid to this idea. Most discussions of AGI either support building it or do not seriously consider whether to build it because they presume that it will inevitably be built, as discussed by Totschnig (2019). Some arguments against building AGI are rooted in concerns about catastrophe risk (e.g., Joy 2000); more on risk below. Others argue that even safe AGI should not be built. These include the fringe anarchist views of Kaczynski (1995, para. 174) and the more sober discussion of Totschnig (2019). However, there has been much less outright opposition to AGI than there has been to similarly transformative technologies like human enhancement.

2.2 Risk

The potential for AGI catastrophe is rooted in the notion that AGI could come to outsmart humanity, take control of the planet, and pursue whatever goals it is programmed to pursue. Unless it is programmed with goals that are safe for humanity and everything else that matters, the result could be catastrophic. Likewise, in order to avoid catastrophe, AGI R&D projects must take sufficient safety precautions.

Opinions vary on the size of this risk and the corresponding safety effort required. Some propose that it is fundamentally challenging to design an AGI with safe goals—that even seemingly minor mistakes could yield catastrophic results and that therefore AGI R&D projects should be very attentive to safety (Yudkowsky 2004; Muehlhauser and Helm 2012; Bostrom 2014). Others argue that an AGI can be trained to have safe goals and that this process is not exceptionally fragile, such that AGI R&D projects need to attend to safety, but not to an unusual extent (e.g., Goertzel and Pitt 2012; Bieger et al. 2015; Goertzel 2015; 2016; Steunebrink et al. 2016). Finally, some dismiss the risk entirely, either because AGI will never be able to outsmart humanity (e.g., Bringsjord 2012; McDermott 2012) or because it is too unlikely or too distant in the future to merit attention (e.g., Etzioni 2016; Stilgoe and Maynard 2017). Analysis of other global risks provides reason to believe that AGI does, in fact, pose a substantial risk to humanity (Baum et al. 2019)

One common concern is that competing projects will race to launch AGI first, with potentially catastrophic consequences (Joy 2000; Shulman 2009; Dewey 2015; Armstrong et al. 2016). Desire to win the AGI race may be especially strong due to perceptions that AGI could be so powerful that it would lock in an extreme first-mover advantage. This creates a collective action problem: it is in the group's collective interest for each project to maintain a high safety standard, but it is each project's individual interest to skimp on safety in order to win the race. Armstrong et al. (2016) present game theoretic analysis of the AGI race scenario, finding that the risk increases if (a) there are more R&D projects; (b) the projects have a stronger preference for their own AGI relative to others', making them less likely to invest in time-consuming safety measures; and (c) the projects have a similar capability to build AGI, so that the advantage they get from skimping on safety is relatively larger.

Barrett and Baum (2017a; 2017b) develop a risk model of catastrophe from AGI, looking specifically at AGI that recursively self-improves to the point that it becomes superintelligent and gains control of the planet.[2] For this catastrophe to occur, six conditions must all hold: (1) superintelligence must be possible; (2) the initial ("seed") AI that starts the self-improvement process must be created; (3) the self-improvement process and the resulting superintelligence must not be successfully contained, thus allowing the superintelligence to gain control of the planet; (4) humans must fail to make the AI's goals safe, so that if it accomplished its goals it would cause a catastrophe; (5) the AI must not independently make its goals safe regardless of human efforts; and (6) the AI must not be deterred from pursuing its goals by humans, other AIs, or anything else. The total risk depends on the probability of each of these conditions holding. Risk management can seek to reduce the probability of conditions (2), (3), (4), and (6), and is one aspect of AGI policy.

2.3 Policy

AGI policy can be understood broadly as all efforts to influence AGI R&D, which can include the formal policies of governments and other institutions as well as the informal policies of people interested in or concerned about AGI, including the researchers themselves. AGI policy can seek to, among other things, fund or otherwise support AGI R&D, encourage particular ethical views to be

[2] The model speaks in terms of AI in general, of which AGI is just one type, alongside other types of AI that could also recursively self-improve. This distinction is not crucial for the present paper.

built into AGI, or reduce AGI risk. Sotala and Yampolskiy (2015) review a wide range of AGI policy ideas, focusing on risk management.

Much of the prior literature on AGI politics emphasizes the tension between (1) hypothetical AGI developers who want to proceed with inadequate regard for safety or ethics and (2) a community that is concerned about unsafe and unethical AGI and seeks ways to shift AGI R&D toward safer and more ethical directions. Joy (2000) argues that the risk of catastrophe is too great and calls for a general abandonment of AGI R&D. Hibbard (2002) and Hughes (2007) instead call for regulatory regimes to avoid dangerous AGI without altogether abandoning the technology. Yampolskiy and Fox (2013) propose review boards at research institutions to restrict AGI research that would be too dangerous. Baum (2017b) calls for attention to the social psychology of AGI R&D communities in order to ensure that safety and ethics measures succeed, as well as to encourage AGI R&D communities to do more on their own.

One policy challenge comes from the fact that AGI could be developed anywhere in the world that can attract sufficient research talent and assemble modest computing resources. Therefore, Wilson (2013) outlines an international treaty that could ensure that dangerous AGI work does not shift to unregulated countries. Scherer (2016) analyzes the potential for AGI regulation by the US government, noting the advantages of national regulation relative to sub-national regulation and suggesting that this could be a prelude to an international treaty. Goertzel (2009) analyzes prospects that AGI will be developed in China or in the West, finding that AGI could be developed either way depending on the importance of certain factors. Bostrom (2014) calls for international control over AGI R&D, possibly under the auspices of the United Nations. In order to identify rogue AGI R&D projects that may operate in secret, Hughes (2007), Shulman (2009), and Dewey (2015) propose global surveillance regimes; Goertzel (2012a) proposes that a limited AGI could conduct the surveillance.

Finally, prior literature has occasionally touched on the institutional context in which AGI R&D occurs. The Yampolskiy and Fox (2013) proposal to establish review boards that are similar to the preexisting review boards for research on human subjects predominantly focuses on universities. Goertzel (2017a) expresses concern about AGI R&D at large corporations due to their tendency to concentrate global wealth and bias government policy in their own favor; he argues instead for open-source AGI R&D. In contrast, Bostrom (2017) argues that open-source AGI R&D could be more dangerous because it would give everyone access to the same code and thereby tighten the race to build AGI first. Shulman (2009) worries that nations will compete to build AGI in order to achieve "unchallenged economic and military dominance", and that the pursuit of AGI could be geopolitically destabilizing. Baum et al. (2011) query AGI experts on the relative merits of AGI R&D in corporations, open-source communities, and the US military, and find that experts have diverging views, especially on the relative merits of open-source and military R&D.

It should also be noted that there has been some significant activity regarding AI from major governments. For example, the Chinese government recently announced a major initiative to become a global leader in AI within the next few decades (Webster et al. 2017). The Chinese initiative closely resembles—and may be derivative of—a series of reports on AI published by the US under President Obama (Allen and Kania 2017). Russian President Vladimir Putin has spoken about the importance of AI, calling it "the future," noting "colossal opportunities, but also threats that are difficult to predict" (RT 2017). However, these various initiatives and pronouncements are not specifically about AGI and appear to mainly refer to narrow AI. Some policy communities have even avoided associating with AGI, such as a series of events sponsored by the Obama White House in association with the reports mentioned above (Conn 2016). Thus, high-level government interest in AI does not necessarily imply government involvement in AGI. One instance of high-level government interest in AGI is in the European Commission's large-scale support of the Human Brain Project, in hopes that a computer brain simulation could revive the European economy (Theil 2015).

3. Research Questions

The prior literature suggests several questions that could be informed by a survey of active AGI R&D projects:

How many AGI R&D projects are there? Armstrong et al. (2016) find that AGI risk increases if there are more R&D projects, making them less likely to cooperate on safety. Similarly, literature on collective action in other contexts often proposes that, under some circumstances, smaller groups may be more successful at cooperating, though large groups may be more successful in other circumstances (e.g., Yang et al. 2013).[3] Thus, it is worth simply knowing how many AGI R&D projects there are.

What types of institutions are the projects based in? Shulman (2009), Baum et al. (2011), Yampolskiy and Fox (2013), and Goertzel (2017a) suggest that certain institutional contexts could be more dangerous and that policy responses should be matched to projects' institutions. While the exact implications of institutional context are still under debate, it would be helpful to see which institution types are hosting AGI R&D.

How much AGI R&D is open-source? Bostrom (2017) and Goertzel (2017a) offer contrasting perspectives on the merits of open-source AGI R&D. This is another debate still to be resolved, which meanwhile would benefit from data on the preponderance of open-source AGI R&D.

How much AGI R&D has military connections? Shulman (2009) proposes that nations may pursue AGI for military dominance. If true, this could have substantial geopolitical implications. While military R&D is often classified, it is worth seeing what military connections are present in publicly available data.

Where are AGI R&D projects located? Wilson (2013) argues for an international treaty to regulate global AGI R&D, while Scherer (2016) develops a regulatory proposal that is specific to the US. It is thus worth seeing which countries the R&D is located in.

What goals do projects have? Section 2.1 summarizes a range of ethical views corresponding to a variety of goals that AGI R&D projects could have. Additionally, Armstrong et al. (2016) finds that AGI risk increases if projects have stronger preference for their own AGI relative to others', which may tend to happen more when projects disagree on goals. Thus, it is worth identifying and comparing projects' goals.

How engaged are projects on safety issues? Section 2.2 reviews a range of views on the size of AGI risk and the difficulty of making AGI safe, and Section 2.3 summarizes policy literature intended to ensure that AGI R&D projects may have inadequate safety procedures. Thus, data on how engaged projects are on safety could inform discussion both on the size of AGI risk and on AGI risk policy.

How large are the projects? Larger projects may be more capable of building AGI. Additionally, Armstrong et al. (2016) find that AGI risk increases if projects are similar in their capability to build AGI. The Armstrong et al. (2016) analysis assumes that project capacity is distributed uniformly. It is worth seeing what the distribution of project sizes actually is and which projects are the largest.

Project capacity for building AGI is arguably more important than project size. However, project capacity is harder to assess with this paper's methodology of analyzing openly published statements. In addition to project size, project capacity could also depend on the talent of its personnel, the availability of funding, computing power, or other resources, and on how well the project is managed. These factors are often not publicly reported. Another important factor is the viability of the technical approach that a project pursues, but this is not well understood and is a matter of disagreement among AGI experts. While it may be possible to assess project capacity with some degree of rigor, this

[3] The collective action literature specifically finds that smaller groups are often more successful at cooperating when close interactions reduce free-riding and the costs of transactions and compliance monitoring, while larger groups are often more successful at cooperating when cooperation benefits from having more total resources available (Yang et al. 2013). Thus, for example, one might want a small group for keeping a secret but a large group for fundraising for a fixed project.

paper's methodology is not suited for such a task, and thus it is left for future work. Instead, project size may be used as at least a rough proxy for project capacity, though caution is warranted here because it may be an imperfect or even misleading proxy.

4. Methodology

The paper's method consists of identifying AGI R&D projects and then characterizing them along several axes. The identification and description were based on openly published information as found in scholarly publications, project websites, popular media articles, and other websites, with emphasis placed on more authoritative publications. Identification and description were conducted primarily by the present authors during the primary data collection stage between June and September 2020. Some project data were revisited for minor miscellaneous revisions after September 2020.

In social science terminology, this methodology is known as the "coding" of qualitative data (Coffey and Atkinson 1996; Auerbach and Silverstein 2003). The data is qualitative in that it consists of text about AGI R&D projects; it is coded into quantitative form, such as "one academic project and three government projects." The coding scheme was initially developed based on prior literature and the present authors' understanding of the topics. It was updated during the coding process based on the present authors' reading of the data (known as "in vivo" coding).

This methodology is fundamentally interpretive and is rooted in the researcher's interpretation of the data. Some aspects of the data are not a matter of interpretation—for example, the fact that the University of Southern California is an academic institution in the US. Other aspects are more open to interpretation. This includes which projects qualify as AGI R&D. Goertzel (2014, p.2) refers to the AGI community as a "fuzzy set"; this is an apt description. Different researchers may interpret the same data in different ways. Indeed, a small number of data points from the 2017 survey have been coded differently in this 2020 version, due to the authors' reflection on, and refinement of, their coding technique. Furthermore, different authors may also find different data as they search through the vast space of openly published information about AGI R&D. Thus, these results should be read as one take on AGI R&D and not necessarily as a true or complete reflection of the topic. Interested readers are invited to query the data for themselves and make their own interpretations. The appendices contain full descriptions and explanations of coding judgments and cites the corresponding data. (Not all the data are cited, since much of what was found is redundant or of limited relevance.)

The methodology of the 2020 survey is similar but not identical to the methodology of the 2017 survey. There are two types of methodology changes: changes in the methodology for how projects are identified and changes in the methodology for how the identified projects are coded. Of the two, the changes in project identification methodology are more significant. As described in Section 4.1, the 2020 survey uses an expanded suite of methods used to identify AGI R&D projects. Changes in the coding of identified projects have been relatively minor.

Because the methods have changed, results of the 2017 survey are not directly comparable to results of the 2020 survey. Therefore, the 2020 survey revisits the AGI R&D projects active in 2017 in addition to covering projects active in 2020. The result is three datasets: (1) projects identified in the 2017 survey, (2) projects active in 2017 as identified in the 2020 survey, and (3) projects active in 2020. Section 5 presents all three datasets coded using the same updated coding methodology. Comparisons between these datasets show the changes in AGI R&D projects that are due to changes in project identification methodology as well as the changes due to the passage of time between 2017 and 2020. Changes in the projects due to changes in coding methodology are relatively minor and are not emphasized in Section 5.

4.1 Identification of AGI R&D Projects

AGI R&D candidate projects were identified via:

- The 2017 AGI R&D projects survey.
- The present authors' prior knowledge.
- Projects suggested by readers of the 2017 survey.
- Keyword searches on the internet and in scholarship databases, mainly Google web search, Google Scholar, Crunchbase, GitHub, and LinkedIn.
- Other previous survey papers (Chong et al. 2007; Duch et al. 2008; Langley et al. 2008; de Garis et al. 2010; Goertzel et al. 2010; Samsonovich 2010; Taatgen and Anderson 2010; Thórisson and Helgasson 2012; Dong and Franklin 2014; Goertzel 2014; Kotseruba et al. 2016).
- The entire contents of the issues of the *Journal of Artificial General Intelligence* published since the 2017 survey (2018 through 2020; earlier issues were covered in the 2017 survey).
- The proceedings of the AGI conferences published since the 2017 survey (2018 to 2020; earlier proceedings were covered in the 2017 survey).
- Additional literature and webpages identified via all of the above.

The 2020 survey uses all of the project identification methods that were used in the 2017 survey, plus some additional methods. The new methods are suggestions from readers of the 2017 survey and searches of Crunchbase and GitHub. The expansion of methods was done to achieve a more comprehensive compilation of AGI R&D projects. The 2020 survey has attempted to apply these methods to identify AGI projects active in both 2017 and 2020. Unfortunately, it is not possible in 2020 to search the versions of Crunchbase and GitHub that existed in 2017. It is likewise not possible to incorporate suggestions from all readers of the 2020 survey. Instead, a draft of the survey was circulated to select colleagues to identify further projects and obtain general feedback.

A project is considered active if it has visible updates within the three previous years. Projects coded as active in 2017 need to be active at some point during the three year period 2015-2017. Projects coded as active in 2020 need to be active at some point during the three year period 2018-2020. Projects are coded as active if they are active at any time in the three year period, even if they are known to be no longer active. For example, a project that publicly disbanded in 2019 is still coded as active in 2020, even though it is known that the project is not active in the year 2020. The visible updates needed to code a project as active can be in a variety of forms including research publications, website updates, and blog posts. The updates do not need to show progress on AGI R&D in recognition of the possibility that the R&D itself could be conducted privately.

Project chronology is inferred from both the timing and the substance of the visible updates. When updates were released is inferred from listed publication dates or website update data on the Internet Archive (https://archive.org). The substance of updates is obtained by reading the updates. For example, the Internet Archive may show that a project website was publicly launched in 2018, while the website itself states that the project itself was founded in 2016. This is taken to indicate that the project was active in both 2016 and 2018. The project would likewise be coded in this survey as active in both 2017 and 2020 because 2016 and 2018 are within the three year periods of 2017 and 2020, respectively. An implication of this is that the 2020 survey can retrospectively identify some projects that were active in 2017 but not publicly visible in 2017. Because they were not publicly visible, they could not have been included in the 2017 survey. This creates an asymmetry between 2017 and 2020 projects: unlike for the 2017 data in the 2020 survey, projects that are active in 2020 but not publicly visible in 2020 cannot be included in the 2020 survey. This asymmetry is another factor that should be accounted for when comparing projects between 2017 and 2020.

Each identified project was put into one of five categories:

- Active AGI R&D projects (Appendix 1). These are projects that are, or were at some point during 2018-2020, working toward building AGI. The included projects either identify as AGI or conduct R&D to build something that is considered to be AGI, human-level intelligence, or superintelligence.
- Inactive AGI R&D projects (Appendix 2). These are projects that were working towards building AGI at some point during 2015-2017, were inactive during the period 2018-2020, and for which new information was identified in the 2020 survey that affects how the project is coded. This includes (1) projects that were included in the 2017 survey whose coding is affected by changes in coding methodology between 2017 and 2020, and (2) projects that were active during 2015-2017, inactive during 2018-2020, and identified for the first time in the 2020 survey. These projects are coded in the same way as active AGI R&D projects.
- Other notable projects (Appendix 3). These include (1) projects that were identified in the 2017 survey, were inactive during 2018-2020, and whose coding is not affected by changes in coding methodology between 2017 and 2020; for these projects, the coding in the 2017 survey is used for data analysis in the 2020 survey; (2) projects that work on technical aspects of AGI but are not working towards building AGI, such as projects working on hardware or safety mechanisms that can be used for AGI; and (3) select narrow AI projects, such as AI groups at major technology companies.
- Other projects (Appendix 4). These projects describe themselves as working on AGI, but demonstrate little further activity on AGI.
- Other projects judged not to be worth including in this paper.

The Section 5 data analysis is based on AGI R&D projects as documented in Appendices 1-2 and the 2017 survey. The projects in Appendices 3 and 4 are reported for the purposes of documenting related work, clarifying the present authors' thinking about where the boundaries of AGI R&D lie, and assisting in the identification of any AGI R&D projects that have been overlooked by the present research.

Projects that only do R&D in deep learning and related techniques were excluded unless they explicitly identify as trying to build AGI. Deep learning already shows some generality (LeCun et al. 2015), and some people argue that deep learning could be extended into AGI (e.g., Christiano 2016). Others argue that deep learning, despite its remarkable ongoing successes, is fundamentally limited, and AGI requires other types of algorithms (e.g., Strannegård and Nizamani 2016; Wang and Li 2016; Marcus and Ernst 2019). The recent explosion of work using deep learning renders it too difficult to survey using this paper's project-by-project methodology. Furthermore, if all of deep learning was included, it would dominate the results, yielding the unremarkable finding that there is a lot of active deep learning work. The deep learning projects that explicitly identify as trying to build AGI are much smaller in number, fitting comfortably with this paper's methodology and yielding more noteworthy insights.

4.2 Description of AGI R&D Projects

For each identified AGI R&D project, a general description was produced, along with classification in terms of the following attributes:

- *Type of institution*: The type of institution in which the project is based, such as academic or government.

- *Open-source*: Whether the project makes its source code openly available.
- *Military connections*: Whether the project has connections to any military activity.
- *Nationality*: The nation where the project is based. For multinational projects, the nation where the project's administrative and/or operational leadership is located is considered the lead nation, and additional partner countries were tabulated separately.
- *Stated goal*: The project's stated goals for its AGI, defined as what the project aims to accomplish with its AGI and/or what goals it intends to program the AGI to pursue.
- *Engagement on safety*: The extent of the project's engagement with AGI safety issues.
- *Size*: The overall size of the project.

4.2.1 Type of Institution

The type of institution attribute has six categories:

- Academic: Institution conducts secondary education (e.g., colleges and universities).
- Government: Institution is situated within a local or national government (e.g., national laboratories). This category excludes public colleges and universities.
- Nonprofit: Institution is formally structured as a nonprofit and is not an academic institution (e.g., nonprofit research institutes).
- Private corporation: Institution is for-profit and does not issue public stock.
- Public corporation: Institution is for-profit and does issue public stock.
- None: Project is not based within any formal institution.

Some projects had two institution types; none had more than two. For the two-type projects, both types were recorded. Project participants were identified and evaluated to determine whether their institutions merited inclusion. Institutions that have not been active within the last three years were considered to be inactive and excluded. Additionally, institutions were only included if the partnership was formally recognized on project websites or other key project documents. Some projects had participation from many institutions that were more limited in nature, such as through co-authorship of publications. This more limited participation was not counted because it would make the entire exercise unwieldy due to the highly collaborative nature of many of the identified projects. This coding policy was maintained across all of the attributes, not just institution type.

In the 2017 survey, projects were coded as having multiple institution types according to a more relaxed standard for partner organizations than was used in the 2020 survey. Therefore, some projects coded as having multiple institution types under the 2017 methodology may be coded as having only fewer institution type(s) under the 2020 methodology.

4.2.2 Open-Source

The open-source attribute has three categories:

- Yes: Project has source code available for download online.
- Restricted: Project offers source code upon request.
- No: Project does not offer source code.

Projects are coded as "yes" if some source code related to their AGI work is open. Projects are coded as "restricted" if no AGI code is open and some AGI code is available upon request. Projects that have no open or restricted code are coded as "no"; this includes projects that only have other, non-

AGI open or restricted code are coded as "no." Whether code is related to AGI is a matter of interpretation, and different coders may produce somewhat different results. The three open-source categories are mutually exclusive: each project is coded as belonging to one category.

4.2.3 Military Connections

The military connections attribute has three categories:

- Yes: The project has identifiable military connections.
- No: The project is found to have no military connections.
- Unspecified: No determination could be made about the project's military connections.

Military connections were identified via keyword searches on project websites and the internet at large, as well as via acknowledgments sections in recent publications. Projects were coded as having military connections if they were based in a military organization, if they received military funding, or if they collaborated in other ways with militaries. Projects were coded as having no military connections if they stated that they did not collaborate with militaries. The latter was only viable for certain smaller projects. Unless a definitive coding judgment could be made, projects were coded as "unspecified."

Projects are coded as having military connections if the project has any identifiable military connection, however limited. This includes projects whose military connections are not for military affairs. For example, a military research agency could fund a project to do basic research on AGI without the corresponding military having any clear application of the research. This also includes projects whose military connections are for non-AGI work. This broad treatment of military connections seeks to assess the full scope of relationships between militaries and AGI R&D projects, which includes but is not limited to military pursuit of AGI for weaponry and other military affairs. As Section 5 documents, relatively few projects have military connections, making this broad treatment more insightful.

In the 2017 survey, projects were coded as having no military connection if they state that they do not collaborate with militaries or if the entire project could be scanned for connections. In the 2020 survey, the latter criterion is omitted on grounds that it is too difficult to scan an entire project. Therefore, some projects coded as "no" under the 2017 methodology may be coded as "unspecified" under the 2020 methodology.

4.2.4 Nationality

The nationality attribute has two categories:

- Lead country: The country in which the project's administrative and/or operational leadership is based
- Partner countries: Other countries that contribute to the project

One lead country was specified for each project. Projects could have zero, one, or multiple partner countries. Partner countries could be the location of secondary sites of the lead institution, such as satellite offices, as well as the location of any partner institution(s).

The 2020 methodology for nationality has two changes relative to the 2017 methodology. First, the 2020 methodology places greater emphasis on formal information about project location, such as is found in project websites, and less emphasis on more informal information such as the location of

residence of project team members. Therefore, some projects may be coded as having different locations under the 2020 methodology than under the 2017 methodology. Second, relative to the 2017 survey, the 2020 survey used a somewhat stricter standard for coding partner countries. Therefore, some countries coded as "partner" under the 2017 methodology may be coded as not "partner" under the 2020 methodology.

4.2.5 Stated Goals

The stated goals attribute has six categories:

- Animal welfare: AGI is being built to benefit nonhuman animals.
- Ecocentrism: AGI is being built to benefit natural ecosystems.
- Humanitarianism: AGI is being built to benefit humanity as a whole. This category includes statements about using AGI to solve general human problems such as poverty and disease.
- Intellectualism: AGI is being built for intellectual purposes, which includes the intellectual accomplishment of the AGI itself and using the AGI to pursue intellectual goals.
- Profit: AGI is being built to make money for its builders.
- Transhumanism: AGI is being built to benefit advanced biological and/or mechanical beings, potentially including the AGI itself.
- Unspecified: Available sources were insufficient to make a coding judgment.

Some categories of goals found in prior AGI literature, including gaining military advantage and the personal benefit of AGI builders, did not appear in the data.

For the coding of stated goals, only explicit statements were considered; the surrounding context was not considered. For example, most AGI R&D projects at corporations did not explicitly state profit as a goal. These projects were not coded "profit" even though they may in fact have profit as a goal. Additionally, only statements by the AGI project and its team members were considered. Statements by related entities, such as the project's parent organization, were not considered. For statements by project team members, emphasis was placed on team members in a leadership role within the project.

In the 2017 survey, the coding was not as strict at excluding statements made by entities related to the project. Therefore, some statements that were coded as indicating a certain stated goal for a project under the 2017 methodology may be coded as not indicating that goal under the 2020 methodology.

4.2.6 Engagement on Safety

The engagement on safety attribute has four categories:

- Active: Projects have dedicated efforts to address AGI safety issues.
- Moderate: Projects acknowledge AGI safety issues but lack dedicated efforts to address them.
- Dismissive: Projects argue that AGI safety concerns are incorrect or misguided.
- Unspecified: Available sources were insufficient to make a coding judgment.

Each project was coded as belonging to one and only one of these categories. Projects that are active on safety and also acknowledge safety issues (i.e., are moderate) were coded as active. Projects that are active on safety and also argue that safety concerns are incorrect or misguided (i.e., are dismissive) are coded as active. Projects that acknowledge safety issues and also argue that safety concerns are incorrect or misguided are coded as moderate. The active-dismissive and moderate-

dismissive projects can occur, for example, if a project articulates that there are some valid safety issues, but not the safety safety issues that some other people have raised.

4.2.7 Size

Project size attribute has five categories:

- Small: Approximately 5 or fewer full-time personnel, small amounts of other indicators
- Small-medium: Approximately 5 to 15 full-time personnel, small-medium amounts of other indicators
- Medium: Approximately 15 to 30 full-time personnel, medium amounts of other indicators
- Medium-large: Approximately 30 to 50 full-time personnel, medium-large amounts of other indicators
- Large: Approximately 50 or more full-time personnel, large amounts of other indicators

Other indicators include, but are not necessarily limited to, part-time personnel, external collaborators, publications, funding, and AGI accomplishments. Coding judgments took into account the full range of indicators. For this reason, projects with a certain number of personnel are not necessarily coded within the corresponding size range.

For projects that do not list personnel, their presence on LinkedIn was studied to infer their personnel size. Some projects, especially some of the larger projects, conduct a mix of AGI and non-AGI activities. For these projects, size was evaluated in terms of the AGI portion of the project.

The 2020 methodology for size contains two changes from the 2017 survey. First, the 2017 survey did not employ a fixed (approximate) set of size ranges for number of personnel. Instead, it used a more informal estimation of project size. Therefore, projects coded as one size under the 2017 methodology may be coded as a different size under the 2020 methodology. Second, the 2017 survey included an "unspecified" size category for projects for which the available information was insufficient to make a coding judgment. The 2020 survey removes the "unspecified" size category on grounds that it is always possible to make a coding judgment, even if there is limited information available. Therefore, projects coded as "unspecified" under the 2017 methodology would be coded as one of the five sizes under the 2020 methodology.

5. Main Results

This section presents the main results of the survey. Full results are presented in Appendices 1-4. Figure ES1 in the Executive Summary presents an overview.

5.1 The Identified AGI R&D Projects

72 AGI R&D projects were identified as being active in 2020. 33 of these projects also appeared in the 2017 survey. 39 are new to this survey. Of the 39 new projects, 24 were found to be active in AGI R&D in 2017; the other 15 became active after 2017. In the list below, **bold font** indicates projects new to this survey; ***bold-italics*** font indicates projects that are new to this survey and were also active in 2017. In alphabetical order, the 72 active AGI R&D projects are:

1. ACT-R, led by John Anderson of Carnegie Mellon University
2. AERA, led by Kristinn Thórisson of CADIA at Reykjavik University
3. **AGI Brain**, led by Mohammadreza Alidoust at AGT Co.

4. **AGI Laboratory**, a private company led by David Kelley
5. ***AGi3***, a project led by Peter Voss
6. ***AIBrain***, a multinational private corporation led by Dr. Richard H. Shinn
7. AIXI, led by Marcus Hutter of Australian National University
8. Animats, a small project led by researchers in Sweden, Switzerland, and the US
9. **ANSNA** is a project led by Patrick Hammer at Temple University
10. **Apollo Program for AGI** is a project of Montréal AI led by Vincent Boucher
11. **Astrum** is a private corporation founded by Srikanth Srinivas
12. Baidu Research, an AI research group within Baidu
13. **Big Mother**, a nonprofit led by Aaron Turner based in the UK
14. ***Binary Neurons Network***, a Russian project led by Ilya Shishkin
15. Blue Brain, led by Henry Markram of École Polytechnique Fédérale de Lausanne
16. **Brain Simulation II**, a small private corporation led by Charles Simon
17. ***Brain2Bot***, a small private corporation led by Gunnar Newquist
18. ***Cerenaut Research***, a project led by Gideon Kowadlo and David Rawlinson
19. China Brain Project, led by Mu-Ming Poo of the Chinese Academy of Sciences
20. CLARION, led by Ron Sun of Rensselaer Polytechnic Institute
21. ***Cognitive Science & Solutions***, a project led by David Sherwood and Terry Higbee
22. CommAI, a project of Facebook AI Research
23. ***Core AI***, a project of Akin AI led by Liesl Yearsly
24. ***Curious AI***, a project led by Harri Valpola
25. Cyc, a project of Cycorp of Austin, Texas, founded by Doug Lenat in 1984
26. ***DeepBrainz***, a project led by Arunkumar Venkataramanan
27. DeepMind, a London-based AI company acquired by Google in 2014
28. **Drayker**, a project based in Brazil and led by Hyadhuad Lucer
29. ***Fairy Tale Artificial General Intelligence Solutions (FTAGIS)***, a project led by Răzvan Flavius Panda
30. FLOWERS, led by Pierre-Yves Oudeyer of Inria and David Filliat of Ensta ParisTech
31. GoodAI, an AI company based in Prague led by computer game entrepreneur Marek Rosa
32. ***Graphen***, a private corporation led by Dr. Ching-Yung Lin
33. **He4o**, a Chinese GitHub project led by Jia Xiaogang
34. HTM, a project of the AI company Numenta, led by Jeffrey Hawkins, founder of Palm Computing
35. Human Brain Project, a consortium of research institutions across Europe
36. ***Intelligent Artifacts***, a US-based project led by Sevak Avakians
37. Leabra, led by Randall O'Reilly of University of Colorado
38. LIDA, led by Stan Franklin of University of Memphis
39. ***M3-CLIC***, a project of M3-IP Ltd. and led by Vidur (Sonny) Nanda
40. **MARAGI**, an open-source project led by Dave Shapiro
41. **Mauhn**, a Belgium-based private corporation led by Berg Severens
42. Microsoft Research AI, an AI research group at Microsoft
43. ***Mind Simulation***, a Russian lab led by Leonid Derikyants, Sergey Pankratov, and Vasily Mazin
44. ***Mindtrace***, a private corporation led by Hoon Chung
45. ***Monad***, a US-based private corporation led by Jovan Williams
46. NARS, led by Pei Wang of Temple University
47. **NDEYSS**, a transhumanist GitHub project
48. ***New Sapience***, a private corporation led by Bryant Cruse
49. Nigel, a project of Kimera, an AI company based in Portland, Oregon
50. ***NiHA***, based at the COMSATS Institute of Information Technology in Pakistan and led by Wajahat Mahmood Qazi

51. NNAISENSE, an AI company based in Lugano, Switzerland and led by Jürgen Schmidhuber
52. ***Olbrain***, a company led by Alok Gautam, Nishant Singh, and Mayank Kumar
53. ***Omega***, led by Eray Özkural from Celestial Intellect Cybernetics
54. OpenAI, an AI research organization with both nonprofit and for-profit components
55. OpenCog, an open-source project listed in the 2017 survey as CogPrime and led by Ben Goertzel
56. ***Optimizing Mind***, a San Francisco-based project led by Tsvi Achler
57. **ORBAI**, a US-based company led by Brent Oster
58. ***Prome***, a private corporation founded by Sean Everett
59. Research Center for Brain-Inspired Intelligence (RCBII), a project of the Chinese Academy of Sciences
60. Robot Brain Project, an open-source project formerly known as Becca and led by Brandon Rohrer
61. **Sanctuary AI**, a private corporation in Canada led by Suzanne Gildert
62. Sigma, led by Paul Rosenbloom of University of Southern California
63. SingularityNET, an open AI platform led by Ben Goertzel
64. Soar, led by John Laird of University of Michigan and a spinoff company SoarTech
65. Susaro, an AI company based in the Cambridge, UK area and led by Richard Loosemore
66. Tencent AI Lab, the AI group of Tencent
67. **True Brain Computing**, a Russian company led by Alexey Redozubov
68. Uber AI Labs, the AI research division of Uber
69. Vicarious, an AI company based in San Francisco
70. Whole Brain Architecture Initiative (WBAI), a nonprofit in Tokyo
71. ***WILLIAM***, a project from Ukraine-based OCCAM led by Dr. Arthur Franz and Michael Löffler
72. ***Xephor Solutions***, a private corporation led by Isabell Kunst

The 2017 survey originally identified 45 AGI R&D projects. 33 of these projects are listed above. The other 12 are now inactive:

1. AIDEUS, led by Alexey Potapov of ITMO University and Sergey Rodionov of Aix Marseille Université; excluded due to apparent inactivity
2. Alice in Wonderland (AIW), led by Claes Strannegård of Chalmers University of Technology; excluded due to apparent inactivity
3. DeSTIN, led by Itamar Arel of University of Tennessee; excluded due to apparent inactivity
4. DSO-CA, led by Gee Wah Ng of DSO National Laboratories, Singapore's primary national defense research agency; excluded due to apparent inactivity upon the departure of the project lead
5. Icarus, led by Pat Langley of Stanford University; excluded due to apparent inactivity
6. Maluuba, a company based in Montréal recently acquired by Microsoft; excluded due to dissolvement upon acquisition by Microsoft
7. MicroPsi, led by Joscha Bach of Harvard University; excluded due to apparent inactivity
8. MLECOG, led by Janusz Starzyk of Ohio University; excluded due to apparent inactivity and lack of R&D
9. Real AI, an AI company based in Hong Kong and led by Jonathan Yan; excluded due to apparent inactivity
10. SiMA, led by Dietmar Dietrich of Vienna University of Technology; excluded due to apparent inactivity
11. SNePS, led by Stuart Shapiro at State University of New York at Buffalo; excluded due to apparent inactivity upon retirement of project leads
12. VICTOR, a project of 2AI, which is a subsidiary of Cifer Inc., a small US company; excluded due to apparent inactivity

The 2020 survey also identified 1 AGI R&D project that was active in 2017 but not in 2020: BasicAI, a small project led by Sean Markan focused on HLAI.

It can now be said that there were at least 70 AGI R&D projects active in 2017. This includes the 45 projects identified in the 2017 survey, the 1 project identified in 2020 that was active in 2017 but not in 2020, and 24 projects identified in 2020 that were active in 2017 and remain active in 2020. Unfortunately, it is not possible in 2020 to search the versions of Crunchbase and GitHub that existed in 2017. It is possible that some project profiles on the 2017 versions of Crunchbase and GitHub are not on the 2020 versions. However, this is unlikely because Crunchbase and GitHub profiles tend to remain online even for inactive projects. Therefore, 70 is probably at least approximately equal to the number of 2017 AGI R&D projects that would have been identified had Crunchbase and GitHub been searched in 2017. Regardless, this is the best 2017 dataset available and will be used in this survey for comparing AGI R&D projects between 2017 and 2020.

Of the 39 projects new to the 2020 survey, 26 were identified via methods that are new to the 2020 survey. There were 14 were identified via Crunchbase, 11 via GitHub, and 1 via readers of the 2017 survey. Therefore, if the 2020 survey had only used the methods of the 2017 survey, it would have identified 13 new projects instead of 39 new projects. Nonetheless, the full set of 72 projects is used in analysis found throughout this paper.

This survey identifies 6 projects that were active in 2017 but were not publicly visible in 2017: Cognitive Science & Solutions, Core AI, DeepBrainz, Mind Simulation, Mindtrace, and Olbrain. Each of these projects had websites that, according to the Internet Archive, launched after 2017, but the websites themselves, or other project information, indicate that the project had an earlier starting date. This chronological difference could be due to the project spending some time in "stealth mode," in which it was intentionally remaining private in order to focus on internal development. Indeed, this survey identifies 1 project that is in "stealth mode" but nonetheless has a public presence (Sanctuary AI). Alternatively, it could simply be due to the project not getting its public profile active right away. Whatever the reason, these 6 projects are projects that could not have been identified by the 2017 survey. Likewise, it is possible that there are projects that are active in 2020 but are not (yet) visible. The results of this survey suggest that there may be approximately 6 not-yet-visible projects active in 2020.

There are some interconnections between the projects. For example, ANSNA was created by Patrick Hammer, a researcher who worked on NARS who also references Jeffrey Hawkins's HTM; AIXI's Marcus Hutter is on leave at DeepMind; Animats and NNAISENSE collaborate; DeepMind and OpenAI continue to collaborate on AGI safety research; OpenAI and MSR AI are both affiliated with Microsoft; OpenCog and SingularityNET are both still led by the same person, Ben Goertzel, and OpenCog technology continues to be used by SingularityNET; Blue Brain and the Human Brain Project were both initiated by the same person, Henry Markram, and share a research strategy; the China Brain Project and RCBII share the same parent organization, the Chinese Academy of Sciences; Sigma's project lead Paul Rosenbloom used to be the co-PI of Soar; and ACT-R and Leabra were once connected in a project called SAL (an acronym for Synthesis of ACT-R and Leabra; see Appendix 3). This suggests an AGI community that is at least in part working together towards common goals, not competing against each other as is often assumed in the literature (Section 2.2). The existence of interconnections was apparent in the 2017 survey. The projects newly identified in the 2020 survey tend to have fewer interconnections, such that the overall picture of the 2020 survey is one of a field of AGI R&D that is not as extensively interconnected as in the 2017 survey.

The 2020 survey contains 46 data points that were coded differently than the 2017 survey as documented in Appendices 1-2. 22 of these data points were coded differently due to changes in the projects themselves, including 9 changes in partner countries, 4 changes in stated goals, 3 changes in institution type, 2 changes in each of military connections and engagement on safety, and 1 change in

each of open-source and lead country. 18 data points were coded differently due to changes in the coding methodology, including 10 changes in size, 2 changes in each of military connections, partner country, and stated goals, and 1 change in each of institution type and lead country. 6 data points were coded differently to correct errors or oversights in the 2017 survey, including 2 changes in each engagement on safety and stated goals and 1 change in each of lead country and partner country. For 1 data point (SingularityNET lead country), there was a change in both methodology and the project. Overall, this indicates a low degree of change between the 2017 and 2020 surveys, especially relative to the change due to some 2017 projects becoming inactive and due to new projects identified by the expanded project identification methodology of the 2020 survey.

5.2 Type of Institution

Of the 2020 AGI R&D projects, 39 are based at least partially in private corporations, 15 are based at least partially in academic institutions, 6 are based at least partially in nonprofits, 6 are based at least partially in public corporations, 7 have no formal institutional home, and 3 are based at least partially in a government. There are 4 projects split across two different institution types: Soar is academic and a private corporation, FLOWERS is academic and government, and both OpenAI and SingularityNET are private corporations and nonprofits. Figure 1 summarizes the institution type data. For comparison, the 2020 survey finds that in 2017, there were 32 projects at private corporations, 21 at academic institutions, 6 at nonprofits, 6 at public corporations, 4 with no institutional home, and 4 at governments; 3 were split across multiple institution types. The 2020 recoding of the 2017 survey projects finds 11 projects at private corporations, 20 at academic institutions, 5 at nonprofits, 6 at public corporations, 2 with no institutional home, and 4 at governments; 3 were split across multiple institution types.

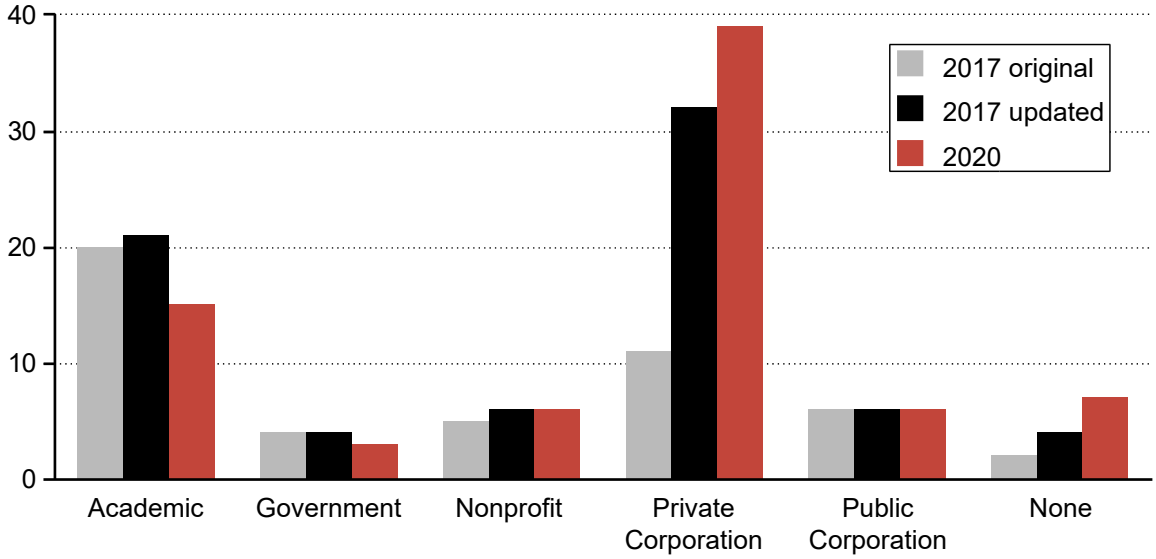

Figure 1. Summary of institution type data. The figure shows more data points than the corresponding number of projects because some projects have multiple institution types. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

In the 2020 survey, private corporations and academia remain the two most common institution types, but their relative frequency is reversed in comparison to the 2017 survey. The 2017 survey

found more academic projects, whereas the 2020 survey finds more private corporation projects. Of the 39 projects new to the 2020 survey, 29 are private corporations; of the 29 private corporations, 14 were identified via Crunchbase, 5 via GitHub, and 10 via other sources. Most of the GitHub projects that are not private corporations do not have any specified institution types and use GitHub as a home.

The 2020 survey also finds an increase in the number of private corporation projects between 2017 and 2020 and a decrease in the number of academic projects during this period. The increase in private corporation projects may constitute evidence of growing AGI profit-R&D synergy, defined as circumstances in which long-term AGI R&D delivers short-term profits (Baum 2017a). Increased AGI profit-R&D synergy would have important risk and policy implications. However, other explanations are also possible. For example, the growth in private corporation AGI projects could be an outcome of the much more general growth of interest in AI between 2017 and 2020. The growth of interest in AI makes it especially notable to observe a decrease in the number of academic projects. The decrease appears to be driven primarily by the departure of academic project leads, who often moved to other AI-related projects. Alice in Wonderland's Claes Strannegård has moved to Animats,[4] DeSTIN's Itamar Arel has since moved on to Syntiant Corp and McD Tech Labs,[5] Icarus's Pat Langley has since become a professor at the University of Auckland,[6] MicroPsi's Josh Bach now works at the AI Foundation,[7] and SNePS's Stuart Shapiro has since retired.[8] Looking ahead, some of the remaining academic projects are led by relatively senior academics, who, upon retirement, may cause further decline in the number of academic projects.

These trends have policy implications. The 2017 survey argued for an important role for measures to govern AGI R&D in academia, such as the Yampolskiy and Fox (2013) proposal for research review boards. The decline of academic projects in 2020 suggests a smaller but still significant role for such measures. If the number of academic projects continues to decline in future years, then the governance of these projects could become a relatively insignificant portion of the overall policy landscape. Likewise, the apparent rise of private corporate projects amplifies the need for corporate governance of AI. That can include governance initiatives within the corporation itself, including those led by employees (Belfield 2020), as well as initiatives external to the corporation, including government regulation of corporate activity. These various initiatives may benefit from a broader characterization of AGI corporations, including other attributes described throughout this survey.

Finally, the 2020 survey continues to find that there are relatively few projects explicitly situated in government institutions. China Brain Project and RCBII are the only projects situated solely within a government, both within the government of China. However, this understates the extent of government involvement in AGI R&D. Numerous other projects receive government funding, aimed at advancing medicine (e.g., Blue Brain, HBP), economic development (e.g., Baidu), and military technology (e.g., Cyc, Soar). That said, the data do not show extensive government interest in developing AGI.

5.3 Open-Source

Of the 2020 AGI R&D projects, 33 have source code openly available online. An additional 5 projects (Big Mother, Cyc, Drayker, the Human Brain Project, LIDA) have code available upon request. For these 38 projects, the available code is not necessarily the project's entire corpus of code, at least for the latest version of the code, though in some cases it is. There are 34 projects for which code could not be found online. Figure 2 summarizes the open-source data. For comparison, the 2020 survey finds that in 2017, 32 projects had openly available source code, 3 projects had code available upon request,

[4] https://www.chalmers.se/en/staff/Pages/claes-strannegard.aspx
[5] https://www.linkedin.com/in/itamar-arel-061b08/
[6] https://www.linkedin.com/in/pat-langley-a992/
[7] https://www.linkedin.com/in/joschabach/
[8] https://www.cse.buffalo.edu/~shapiro/

and 35 projects had no code found online. The 2020 recoding of the 2017 survey projects finds 25 projects had openly available source code, 3 projects had code available upon request, and 17 projects had no code found online.

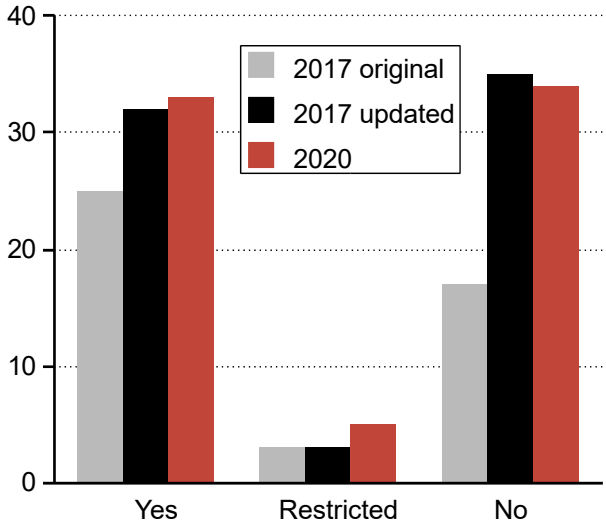

Figure 2. Summary of open-source data. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

Some trends in the open-source data are apparent. First, most of the projects new to the 2020 survey do not have open-source code available. Many of these are private corporations. One possible explanation is that the private corporations are keeping their code private for commercial reasons. This could be an indicator of a more competitive and less cooperative AGI R&D landscape. Second, the 2020 survey finds no significant change in open-source projects between 2017 and 2020. This result suggests that recent dialog on the merits of open publishing in AI, often under the rubric of "publication norms" (e.g., Gupta et al. 2020), have not yet had significant impact on the field of AGI R&D.

Of the 33 projects with code openly available, 11 projects were situated in academic institutions, 14 projects in private corporations, 4 in nonprofits, 2 in public corporations, 1 in a government institution (FLOWERS), and 5 with no institution; 4 of these projects have 2 types of institutions. Of the projects with restricted code, 1 is in a private corporation, 2 are in academic institutions, 1 is in a nonprofit, and 1 does not have an institution type. The preponderance of open-source code availability among academic projects suggests that these projects may tend to be more open and collaborative. This explanation is consistent with the tendency for academic projects to have relatively interconnected personnel.

Some AGI R&D projects have participated in debates about the merits of open-source code. For example, MARAGI, SingularityNET, and Deepbrainz have articulated the importance of the democratization of code, making AGI research accessible to anyone. The Human Brain Project has expressed concern about the potential for militaries to use open-source code for harmful purposes. Of the projects with restricted code, some do so with the hope of fostering a collaborative community (Big Mother and Drayker), while others do so to entice customers for business purposes (Cyc). Under what circumstances code should be made available is a matter of ongoing analysis (Bostrom 2017; Gupta et al. 2020). There may be value in AGI R&D projects continuing to engage with these issues.

5.4 Military Connections

Of the 2020 AGI R&D projects, 9 have identifiable military connections. There are 2 projects identified as having no military connections. For all other projects, no determination on military connections could be made. Figure 3 summarizes the military connections data. For comparison, the 2020 survey finds that in 2017, there were 10 projects with identifiable military connections, and 2 projects identified as having no military connections, which is the same as the 2020 recoding of the 2017 survey projects.

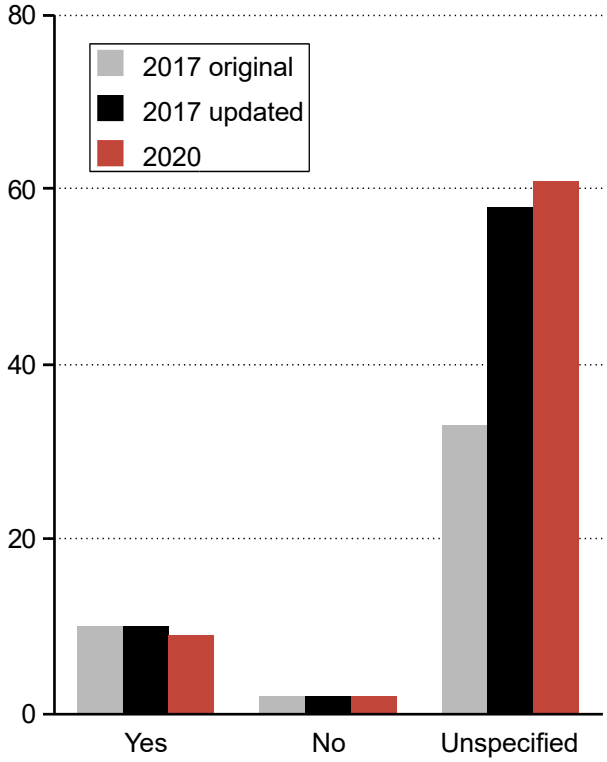

Figure 3. Summary of military connections data. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

The 2017 projects with military connections were ACT-R, CLARION, DSO-CA, Icarus, Leabra, LIDA, Sigma, SNePS, and Soar. The 2020 projects with military connections are ACT-R, China Brain Project, CLARION, Cyc, Leabra, LIDA, RCBII, Sigma, and Soar. The 3 projects with military connections in 2017 but not 2020 (DSO-CA, Icarus, and SNePS) were all inactive in 2020. The 2 projects with military connections in 2020 but not 2017 (China Brain Project and RCBII) were both active in 2017 but did not have military connections until after 2017.

The military connections primarily involve basic scientific research rather than applications for weapons or other military infrastructure. The primary exceptions are Cyc, which worked for US defense agencies in various areas such as terrorism analysis, and Soar, whose affiliate company SoarTech heavily advertises military applications on its website. These applications are largely tactical, suggestive of incremental improvements in existing military capacity, not any sort of revolution in military affairs. Most other projects with military connections have their connections via the receipt of funding from US military research funding agencies such as DARPA and the Office of Naval

Research. Other projects include China Brain Project and RCBII, both of which have collaborated with the Chinese Academy of Military Medical Sciences.

Taking the above into account, it follows that the publicly available record indicates that no military is currently pursuing AGI for major strategic purposes. It is possible that one or more militaries are nonetheless doing so. It does stand to reason that if there were substantial military interest in AGI dominance, this information may be classified and inaccessible. The potential existence of secret military AGI programs is beyond the scope of this paper.

The 2 projects active in 2020 that were identified as having no military connections are AERA and the Human Brain Project. AERA openly rejects military connections. The Human Brain Project, as a European Union Horizons 2020 funded project is constrained to civil purposes (European Commission 2020).

Finally, for the other 61 projects, the presence or absence of military connections was not established. Many of these projects likely do not have military connections because they do not work on military applications and they are not at institutions (such as US universities) that tend to have military connections.

5.5 Nationality

Of the 2020 AGI R&D projects, the projects are based in 23 countries and are present in 37 total countries. There are 32 projects based in the US, 5 in China, 4 in the UK, 3 in Australia, Canada, Russia, and Switzerland, 2 in India, and 1 each in Austria, Belgium, Brazil, Czech Republic, Finland, France, Iceland, Iran, Ireland, Japan, the Netherlands, Pakistan, Sweden, Turkey, and Ukraine, respectively. There are 2 projects that do not have a specified location. Partner countries include the US (partner for 9 projects), China, the UK (4 projects), Australia, Canada, France, Germany, and Israel (3 projects), Austria, India, Italy, Norway, Russia, South Korea, (2 projects), and Belgium, Brazil, Denmark, Ethiopia, Finland, Greece, Hungary, Japan, the Netherlands, Portugal, Singapore, Slovenia, Spain, Sweden, Switzerland, Taiwan, and Turkey (1 project). Figure 4 maps the nationality data. For comparison, the 2020 survey finds that in 2017, there were projects based in 20 countries and present in 36 total countries. The 2020 recoding of the 2017 survey projects finds projects based in 14 countries and present in 30 total countries.

The 32 US-based projects are based in 14 states and territories: 13 in California, 4 in New York, 3 in Pennsylvania, and 1 in each of Florida, Oregon, Maryland, Massachusetts, Michigan, Nevada, Tennessee, Texas, Washington, Washington D.C., and Wyoming, respectively, with 1 project (OpenCog) not having a specified home state. This broad geographic distribution is due largely to the many academic projects: whereas US AI companies tend to concentrate in a few metropolitan areas, US universities are scattered widely across the country. The corporate projects are mainly in the metropolitan areas of Austin, New York City, Portland, San Francisco/Bay Area, and Seattle, all of which are AI hotspots.

There are 20 multinational projects, a minor increase from the 18 projects in 2017. Human Brain Project is in 19 total countries, including the lead country. Nigel and SingularityNET are in 7 countries. CommAI is in 5 countries. AIBrain, Animats, DeepMind, Graphen are in 4 countries. Cerenaut Research, OpenCog, and Soar are in 3 countries. Baidu Research, Core AI, Deepbrainz, M3-CLIC, NNAISENSE, Susaro, TAIL, Uber AI Labs, and WBAI are in 2 countries. Some projects such as NiHA or MSR AI collaborate with researchers overseas, but their partnerships are not indicated here as there was no evidence that these partnerships were active, official, or on AGI related content. BlueBrain was was coded as multinational in 2017 but not 2020. BlueBrain had partner institutions in 4 countries listed on a webpage that was active in 2017 but not 2020. It is possible that BlueBrain retains these or other partner countries, though documentation was not identified.

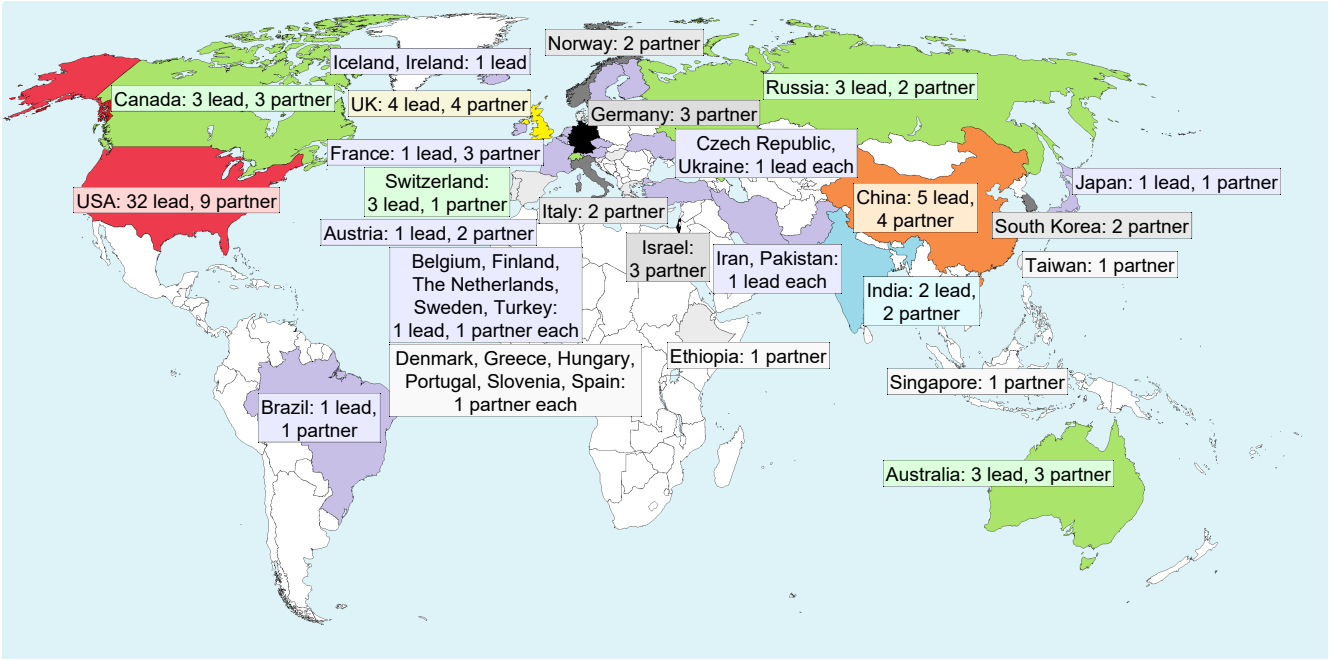

Figure 4. Map of nationality data for projects active in 2020. Depictions of disputed territories do not indicate a position on the dispute.

The most common partner (i.e., non-lead) country is the US with 9 projects, followed by 4 for China and the UK, and 3 for Australia, Canada, France, Germany, and Israel. Germany and Israel are interesting cases because they are each important technology hubs, but neither are the lead country on any AGI R&D projects. The most common partnerships are between the US and China, at 6 partnerships, followed by each of (1) the US and UK and (2) the US and Australia, each with 4 partnerships. The most common non-US partnerships are (1) France and the UK and (2) Israel and the UK, each with 3 partnerships. These tabulations of partnerships include when one country is the lead and the other country is a partner or when both countries are partners. The US has the most international reach with connections to 21 countries, followed by Switzerland with connections to 18 countries. This can be attributed to the Human Brain Project, a large, international consortia based in Lausanne, Switzerland. (BlueBrain is also based in Lausanne.)

As in 2017, the AGI R&D landscape remains dominated by the US and its allies. China is emerging as arguably the primary geopolitical rival to the US (Allison 2017), a competition that extends to the realm of AI (Maas 2019). However, the data do not show intense competition on AGI R&D. China is only the lead country on 5 projects, which remains unchanged from 2017. China is a partner on 4 projects, a decrease from 3 in 2017. Of China's 9 total projects, the US is involved with 6, suggesting a relationship that is more cooperative than competitive. Russia is another major US adversary. It now is the lead on 3 and partner on 2 AGI R&D projects, up from 3 lead and 1 partner in 2017. This modest increase still leaves it with a relatively small overall AGI sector.

Other world regions remain underrepresented. There are only 3 projects in the Middle East: 1 based in Iran (AGI Brain), 1 based in Turkey (Omega), and 1 partnered with Turkey (Human Brain Project). There is 1 project based in Pakistan (NiHa). There are only 2 projects based in India and 2 others partnered there. There is just 1 project (OpenCog) with a presence in Africa, a partner of Ethiopia. There are just 2 projects with a presence in Latin America: Drayker, based in Brazil, and SingularityNET, partnered with Brazil. Given the potential global significance of AGI, a case could be made that there should be greater geographic diversity in AGI R&D. Alternatively, the geographic concentration of AGI R&D projects could be taken as a reason to give underrepresented areas a greater

voice in other aspects of AGI governance. Or it could be a reason to focus policy initiatives on the regions where the projects are, on grounds that this is where policy is most needed.

The results presented in this section may understate the geographic breadth of AGI R&D. There are two reasons for this. First, the list of partner countries is highly sensitive to coding judgments about which institutions to count as partners. The survey only codes collaborating institutions as partners when they have a sufficient degree of collaboration. What classifies as "sufficient" is a coding judgment that could reasonably be made in a variety of ways. For example, many projects have collaborations in the form of co-authorship on publications. These collaborations were generally coded as insufficient unless the publication was central to the project. For some projects, such as ACT-R and Blue Brain, this results in projects being coded as having no partner countries even though they have many international collaborators. Including countries with this more limited form of partnership shows greater geographic diversity; it would include such countries as Mexico, Morocco, Sri Lanka, Venezuela, and Zimbabwe.

Second, many projects have open-source code. This code enables AGI R&D to be conducted anywhere in the world. It is thus possible that there are other countries involved in AGI R&D, perhaps a large number of other countries. The identification of countries whose participation consists exclusively of contributions to open-source code is beyond the scope of this paper.

5.6 Stated Goal

Of the 2020 AGI R&D projects, 38 stated humanitarian goals, 26 stated intellectualist goals, 5 stated transhumanist goals (AGI Laboratory, Brain2Bot, NDEYSS, Sigma, SingularityNET), 3 stated profit goals (AGi3, Curious AI, Xephor Solutions), 3 stated ecocentric goals (FTAGIS, SingularityNET, Susaro), and 3 stated animal welfare goals (Human Brain Project, proposing brain simulation to avoid animal testing, and FTAGIS and SingularityNET, both seeking to benefit sentient beings). There are 18 projects with unspecified goals. Some projects stated multiple goals: 16 projects stated 2 goals, 1 project stated 3 goals (Human Brain Project), and 2 projects stated 4 goals (FTAGIS, SingularityNET). Figure 5 summarizes the stated goal data. For comparison, the 2020 survey finds that in 2017, 33 projects stated humanitarian goals, 28 stated intellectualist goals, 4 stated profit goals, 3 stated transhumanist goals, 4 stated ecocentric goals, 3 stated animal welfare goals, and 17 had unspecified goals, with 17 projects stating multiple goals. The 2020 recoding of the 2017 survey projects finds 19 projects stated humanitarian goals, 23 stated intellectualist goals, 1 stated profit goals, 2 stated transhumanist goals, 3 stated ecocentric goals, 2 stated animal welfare goals, and 10 had unspecified goals, with 12 projects stating multiple goals.

In the 2020 survey, humanitarianism and intellectualism remain the two most common state goals, but their relative frequency is reversed in comparison to the 2017 survey. The 2017 survey found more intellectualist projects, whereas the 2020 survey finds more humanitarian projects. This finding is connected to a similar finding for institution types, with the 2020 survey finding more private corporation projects and the 2017 survey finding more academic projects (Section 5.2). As the 2017 survey first identified, corporate projects (whether private or public) tend to state humanitarian goals, while academic projects tend to state intellectual goals. Among the 2020 projects, 26 of 45 corporate projects state humanitarian goals, while 12 of the 15 academic projects state intellectualist goals. Likewise, the increase in humanitarian projects in the 2020 survey relative to the 2017 survey is driven mainly by the increase in private corporation projects.

Of the 9 projects with military connections, 8 state intellectualist goals. The lone exception is Cyc, which states humanitarian goals. China Brain project states both intellectualist and humanitarian goals, while Sigma states intellectualist and transhumanist goals. The preponderance of intellectualist projects is again related to the projects' academic character. Of the 8 intellectualist projects with military

connections, 6 are at academic institutions and 2 are at the Chinese Academy of Science, which classifies as a government institution but is nonetheless of a heavily academic character. The lone non-intellectualist project, Cyc, is the only project based exclusively at a corporation. (Soar is both academic and corporate.) These results suggest the conclusion that the goals of projects with military connections are driven mainly by the projects themselves and their host institutions, and not driven by the militaries they are connected to. This conclusion further supports the Section 5.4 finding that militaries have relatively limited involvement in AGI R&D at this time.

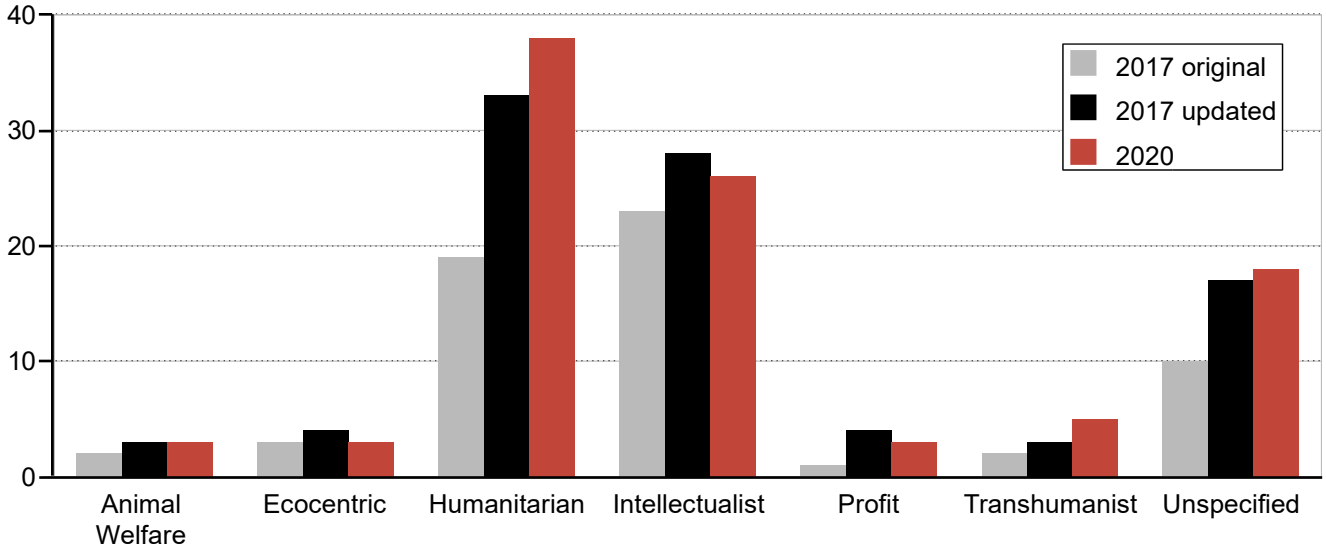

Figure 5. Summary of stated goal data. The figure shows more data points than the corresponding number of projects because some projects have multiple stated goals. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

The results presented in this section may understate the extent of the goal of profit held by AGI R&D projects. Only 3 projects explicitly stated a profit goal (AGi3, Curious AI, and Xephor Solutions) despite there being 45 projects at for-profit corporations. Some of the other corporate projects had a visible orientation toward profit without explicitly stating it as a goal. These projects express interest in commercializing AGI and marketing it as a product or business solution (e.g., Deepbrainz, Graphen, and Intelligent Artifacts). The profit orientation is suggestive of a profit goal but is not definitive evidence. Alternatively, projects could seek to commercialize AGI as a means to other ends, such as benefiting humanity. Indeed, 3 corporate projects explicitly reject profit as a goal (GoodAI, Vicarious, and WILLIAM). Nonetheless, the actual number of projects with the goal of profit may be larger than the number of projects that have explicitly stated this goal in publicly visible statements. This possibility is especially important in light of the larger number of corporate projects identified in the 2020 survey.

Finally, it should be noted that the coding of stated goals was especially interpretative. Many projects do not state their goals prominently or in philosophically neat terms. This includes projects that made statements that were profit-oriented but did not explicitly state a profit goal. It also includes projects with other types of goals. For example, DeepMind lists climate change as an important application. This could be either ecocentric or humanitarian or both, depending on why DeepMind seeks to address climate change. It was coded as humanitarian because it was mentioned in the context of "helping humanity tackle some of its greatest challenges," but it is plausible that ecocentrism was

intended. The same sort of ambiguity is present for some other projects as well. Therefore, the results for stated goals should be viewed more loosely than the results for other project attributes.

5.7 Engagement on Safety

Of the 2020 AGI R&D projects, engagement on safety could only be identified for 34 projects, which is approximately half of the projects. Of these 34 projects, 18 were found to be active on safety, 12 are moderate, and 4 are dismissive. Figure 6 summarizes the engagement on safety data. For comparison, the 2020 survey finds that in 2017, 16 projects were active on safety, 11 were moderate, 4 were dismissive, and 39 were unspecified. The 2020 recoding of the 2017 survey projects finds 13 projects were active on safety, 4 were moderate, 2 were dismissive, and 26 were unspecified.

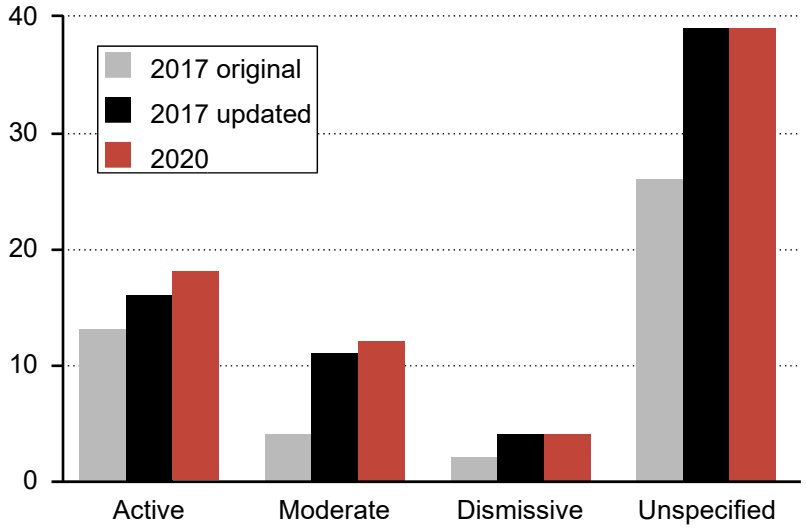

Figure 6. Summary of engagement on safety data. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

Relative to the 2017 survey, the 2020 survey shows a greater portion of projects that are active or moderate on safety, but these projects remain a minority in comparison to projects that have an unspecified engagement on safety. This provides some empirical support for the common assumption in prior AGI policy literature of AGI developers who want to proceed with inadequate regard for safety (Section 2.3). This survey's focus on publicly available data may overstate the neglect of safety because some projects may pay attention to safety without stating it publicly. Still, the data is strongly suggestive of widespread neglect of safety among AGI R&D projects.

Among the 30 projects active in 2020 for which active or moderate engagement on safety was identified, some further trends are apparent. These 30 projects include 11 of the 20 projects with purely humanitarian goals, yet only 4 of the 11 projects with purely intellectualist goals (FLOWERS, Leabra, NARS, NNAISENSE), 7 of 18 projects with unspecified goals, 1 project with ecocentric goals (Susaro), and 1 project with animal welfare, ecocentric, humanitarian, and intellectualist goals (FTAGIS). The 30 projects also include 19 of the 42 projects based purely at corporations, 3 of 4 projects based purely at nonprofit organizations (Big Mother, OpenCog, and WBAI), 5 of 13 projects based purely at academic institutions, and 1 of 2 projects based in part at an academic institution (FLOWERS). This suggests a cluster of projects that are broadly engaged on the impacts of AGI R&D, including ethics questions about what the impacts should be and risk/safety questions about whether

the desired impacts will accrue. This cluster is located predominantly outside of academia. Meanwhile, there continues to be a cluster of projects that are predominantly academic and view AGI in largely intellectual terms, disregarding safety in their pursuit of AGI. These trends suggest the importance of proposals to strengthen risk and ethics practices sometimes suggesting that adequate training could make AGI safety to be a reasonably tractable task among academic AGI R&D projects, such as via research review boards (Yampolskiy and Fox 2013).

The corporate projects new to the 2020 survey show a range of stances on safety. There are 2 projects explicitly focused on potential long-term risks and present thorough discussions of these risks and how to address them (Brain Simulator II, Deepbrainz). There is 1 project is solely focused on near-term risks that may arise such as job loss, racial profiling and biases, and misinformation (Cerenaut Research). There are 2 projects that acknowledged potential AGI risks, but brushed these risks off as implausible if AGI is built correctly (Monad and Susaro). These results show a diverse range of perspectives on safety and point to the need for a broad-based conversation about safety issues in AGI R&D.

5.8 Size

Of the 2020 AGI R&D projects, 30 projects are small, 23 are small-medium, 13 are medium, 2 are medium-large (MSR AI and Vicarious), and 4 are large (Blue Brain, DeepMind, Human Brain Project, and Open AI). Figure 7 summarizes the size data. For comparison, the 2020 survey finds that in 2017, 28 projects were small, 21 were small-medium, 15 were medium, 2 were medium-large, and 4 were large. The 2020 recoding of the 2017 survey projects finds 13 projects were small, 12 were small-medium, 14 were medium, 2 were medium-large, and 4 were large.

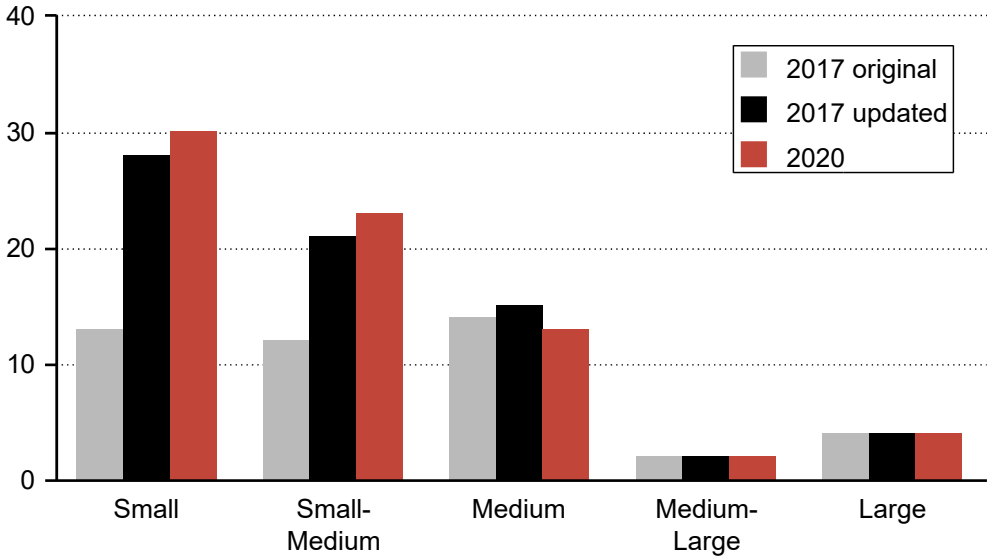

Figure 7. Summary of size data. 2017 original (gray) is the set of projects identified in the 2017 survey recoded with the 2020 methodology. 2017 updated (black) is the set of projects active in 2017 and identified in the 2020 survey. 2020 (red) is the set of projects active in 2020. All projects are coded according to the 2020 coding methodology.

Of the projects new to the 2020 survey, 25 are small, 12 are small-medium, and 2 are medium (Graphen and Sanctuary AI). The 26 small projects include Basic AI, which is new to the 2020 survey but inactive in 2020. Of the 2 medium projects, only 1 project (Graphen) was active in 2017. These numbers show that while the 2017 survey failed to identify a significant number of projects, it may

have identified all of the larger projects and all but 1 (Graphen) of the medium projects. It does stand to reason that larger projects would be easier to find and are less likely to be overlooked. Likewise, these numbers mean that although the 2020 survey presents new trends in AGI R&D (in particular, a larger total number of projects, with a larger portion in private corporations and stating humanitarian goals and a smaller portion with open-source code), these new trends may represent a small portion of total AGI R&D because the newly identified projects are generally smaller.

The variation between the smallest and largest projects is considerable. The smallest projects appear to be hobbies or side projects of a single individual, with personnel less than one full-time person. The largest projects have teams of over 100 full-time personnel with extensive financial and computational resources. It may be the case that a large portion of progress on AGI R&D is driven by a relatively small number of projects, though assessment of progress on AGI R&D is beyond the scope of this paper.

5.10 Other Notable Projects

As documented in Appendix 3, there were 69 other notable projects recorded in the process of identifying AGI R&D projects. These include 40 inactive AGI R&D projects, 15 projects that are not R&D, 29 projects that are not AGI, with some projects in more than one of the aforementioned categories. Some projects fell into more than one category of exclusion. Many of these projects were active AGI R&D projects in the 2017 survey but are no longer active on AGI R&D. Unlike with active AGI R&D projects, no attempt made to be comprehensive in the identification of other notable projects. It is likely that some notable projects are not included. The list of other notable projects and brief details about them are presented in Appendix 3.

The 40 inactive projects are largely academic, such as 4CAPS, led by Marcel Just of Carnegie Mellon University, and OSCAR, led by John Pollock of University of Arizona. They varied considerably in duration, from a few years (e.g., AGINAO, active from 2011 to 2013) to over a decade (e.g., CHREST, active from 1992 to 2012). Of the inactive projects, 12 are projects that were active in 2017 and documented in the 2017 survey, including DeSTIN (led by Itamar Arel and colleagues at the University of Tennessee) and MicroPsi (led by Joscha Bach of the Harvard University).

The 29 projects that are not AGI include many AI groups at large computing technology corporations. 7 such corporations were searched carefully for AGI projects and found not to have any: Alibaba, Amazon, Apple, Intel, SalesForce, Sony, and Twitter. Given these corporations' extensive resources, it is notable that they do not appear to have any publicly identifiable AGI projects. Additionally, the IBM Cognitive Computing Project and the Intel Loihi Chip are hardware projects of relevance to AGI, though neither constitutes AGI R&D. Finally, in addition to DeepMind, 2 other projects at Google were considered: Google Brain and Quantum AI Lab. While Google Brain has done some AGI work with DeepMind, it focuses on narrow AI.

The 15 projects that are not R&D cover a mix of different aspects of AGI. Some focus on basic science, including several brain projects (e.g., the BRAIN Initiative at the US National Institutes of Health and Brain/MINDS at Japan's Ministry of Education, Culture, Sports, Science, and Technology). Several focus on hardware and software for building AGI (e.g., the IBM Cognitive Computing Project, the Tianjic Chip at Tsinghua University, and Hanson Robotics). There are 2 projects that focus on safety aspects of AGI design (Center for Human-Compatible AI at University of California, Berkeley and the Machine Intelligence Research Institute). This is not a comprehensive list of projects working on non-R&D aspects of AGI. For example, projects working on AGI ethics, risk, and policy were not included because they are further removed from R&D.

5.11 Other Projects

There were 39 other projects were recorded in the process of identifying AGI R&D projects. This list includes projects that self-reported a focus on AGI or something AGI-related, but showed little or no apparent AGI R&D. These projects show limited relevance to AGI R&D and are therefore documented only briefly in Appendix 4.

Most of these projects are small corporate projects listed on Crunchbase or GitHub. The projects' profiles use keywords such as "AGI" and "superintelligence," but their actual work does not appear to be focused on AGI. It is possible that some of these projects are in fact performing AGI R&D; all that can be concluded in this survey is that public evidence of AGI R&D was not identified. Another possibility is that some of the projects may be using terms related to AGI to draw more attention to themselves, such as via search engine optimization. This possibility suggests that AGI may be perceived as a valuable brand to associate with even for projects that are not actually working on AGI.

6. Conclusion

Despite the seemingly speculative nature of AGI, R&D towards building it is already happening. This survey identifies 72 AGI R&D projects spread across 37 countries in 6 continents, many of which are based in major corporations and academic institutions, and few of which are large and heavily funded. Given that this survey relies exclusively on openly published information, this should be treated as a lower bound for the total extent of AGI R&D. Thus, regardless of how speculative AGI itself may be, R&D towards it is clearly very real. Given the potentially high stakes of AGI in terms of ethics, risk, and policy, the AGI R&D projects warrant ongoing attention.

6.1 Main Findings

Regarding ethics, the major trend is the large number of projects with humanitarian goals and the somewhat smaller number with intellectualist goals, with the former coming largely from corporate projects and the latter from academic projects. The survey's reliance on public statements may understate the prevalence of profit goals, especially among the many corporate projects. Still, among the projects not motivated by intellectual goals, there does seem to be a bit of a consensus for at least some form of humanitarianism, and not for other types of goals commonly found in AGI discourses, such as transhumanism. Meanwhile, the intellectualist projects indicate that academic R&D projects still tend to view their work in intellectual terms, instead of in terms of societal impacts or other ethical factors, even for potentially high-impact pursuits like AGI.

Regarding risk, two points stand out. First, a clear majority of projects had no identifiable engagement on safety. While many of these projects without identifiable engagement on safety are smaller in size, some of these projects are also larger and with considerable funding. The lack of identifiable engagement suggests that many projects are not active on safety, though it is possible that some are conducting safety activities exclusively in a private fashion. Second, some trends suggest that a risky race dynamic may be avoidable. One is the concentration of projects, especially larger projects, in the US and its allies, which can facilitate both informal coordination and formal public policy. Another is the modest consensus for humanitarianism, again especially among larger projects, which could reduce projects' desire to compete against each other. Some projects are interconnected via shared personnel, parent organizations, AGI systems development, and participation in the same communities, such as the AGI Society. This suggests a community working together towards a common goal and sharing of best practices, not competing against each other.

Regarding policy, several important points can be made. One is the concentration of projects in the US and its allies, including most of the larger projects. This could greatly facilitate the establishment of AGI policy with jurisdiction over most AGI R&D projects, including all of the larger ones. Another important point is the concentration of projects in academia and corporations, with relatively little in government or with military connections. Each institution type merits its own type of policy, such as review boards in academia and financial incentive structures for corporations. The potential for AGI R&D-profit synergy (Section 5.2) could be especially important here, determining both the extent of corporate R&D and the willingness of corporations to submit to restrictive regulations. This survey finds hints of AGI R&D-profit synergy, but not the massive synergies found for certain other types of AI. Finally, the preponderance of projects with at least some open-source code complicates any policy effort, because the R&D could in principle be done by anyone, anywhere.

6.2 Limitations and Future Work

The above conclusions seem robust given this study's methodology, but other methodologies could point in different directions. For example, the published statements about goals suggest a consensus towards humanitarian and intellectualist goals, but this could miss unpublished disagreements on the specifics of these goals. A humanitarian goal for one project could seek greater economic prosperity, whereas another humanitarian project could seek greater freedom or happiness. Additionally, the published statements about goals used in this survey could deviate from projects' actual goals if the projects are not entirely honest in their published statements. Some projects may recklessly pursue self-interested goals while presenting a responsible, ethical front to the public. These possibilities suggest a more difficult policy challenge and larger AGI risk. Alternatively, this study could overestimate the risk. Perhaps many projects have similar goals and concerns about safety, even if they have not published any statements to this effect. Given the overlap in researchers, funding sources, and institutions across the spectrum of projects, this may be quite likely. Future research using other methodologies, especially interviews with people involved in AGI R&D, may yield further insights.

A different complication comes from the possibility that there are other AGI R&D projects not identified by this study. Some projects may have a public presence but were simply not identified in this study's searches, despite the efforts made to be comprehensive. This study is especially likely to miss projects that work in non-English languages, since searches were only conducted in English.

Furthermore, there may be additional AGI R&D projects that have no public face at all. This could conceivably include the secret government military projects. The projects identified with military connections are generally small and focused on mundane (by military standards) tactical issues, not grand ambitions of global conquest. However, more ambitious military goals would most likely not be stated publicly in order to retain strategic advantage during the development and eventual use of AGI. It could be argued that some of the more mundane, tactical issues the military has interest in are merely preliminary research areas that need to be addressed in order to successfully develop AGI. For example, DARPA has multiple research programs that are tangentially related to AGI development such as μBRAIN and CREATE, which focus on neuromorphic computing and autonomous systems.[9] Though these programs, and others, do not use any language explicitly stating interest in the development of AGI, they seem to be preliminary steps toward AGI development or tangentially related projects. Given the stakes, it is important to remain vigilant about the possibility that most military and government funded programs may not be public knowledge.

Another, more likely possibility is that there are additional private sector projects that do not yet have a public presence, either because they are deliberately operating in 'stealth mode' or because they just do not have an active public presence yet. This study identified 6 projects that were active in 2017

[9] https://www.darpa.mil/our-research

but not publicly visible until after 2017. This finding suggests that it is likely that there are some projects that are active in 2020 but not yet publicly visible. The possibility of projects that are not (yet) publicly visible means that the 72 projects identified in this survey should be treated as a lower bound.

A different potential source of additional AGI R&D projects is the vast space of projects focused on deep learning and related techniques. Over the past three years, projects are increasingly more focused on deep learning, cognitive architecture, neuromorphic computing, and artificial neural networks. These projects were generally excluded from this survey because there are too many to assess in this study's project-by-project methodology, and because there are diverging opinions on whether these projects rate as AGI R&D. If AGI could come from these techniques, the AGI R&D landscape would look substantially different from the picture painted in this paper, with major consequences for ethics, risk, and policy. Appendix 4 was created in part because of the increase in projects using these terms in conjunction with self-described AGI development. It is difficult to assess which techniques are more promising for developing AGI and ultimately which projects should be paid more attention to. Therefore, an important direction for future research is to assess the possibility and likelihood that AGI could come from deep learning and other techniques, and then determine the implications for ethics, risk, and policy.

Another worthwhile direction for future research is on projects' capability to build AGI. This study includes project size as a proxy for capability, but it is an imperfect proxy. Capability is the more important attribute for understanding which projects may be closest to building AGI. More capable projects may warrant greater policy attention, and a cluster of projects with similar capabilities could lead to a risky race dynamic. (Project size is important for other reasons, such as projects' pull on the labor market for AGI researchers or their potential for political lobbying.) Project capacity could be assessed via attention to details of projects' performance to date, the viability of their technical approaches to AGI, and other factors. Given the ethics, risk, and policy importance of project capacity, this is an important direction for future research. However, despite observing and tracking the aforementioned details, there is still the possibility that the project to successfully develop AGI first may come from an unexpected project, such as a small, underfunded group that just happens to get lucky.

Additionally, another consideration is the value of diversity and inclusivity within research groups. There have been public exposés of the biases of AI in human resources departments and hiring, as well as in policing and law enforcement, and healthcare (Dastin 2018; Ray 2020; Kaushal et al. 2020). Issues of AI bias have been identified in AI systems that are much less powerful than AGI could be. The stakes for bias in AGI systems could be much greater. This suggests a role for surveying the demographics of AGI R&D projects in tandem with analyses of potential bias in AGI systems.

Finally, future research could also look at other actors involved in AGI. In addition to the R&D projects, there are also R&D groups that work on hardware and safety measures related to AGI; people who study and work on AGI ethics, risk, and policy; and "epistemic activists" who promote different understandings of AGI. Each of these populations could play a significant role in ultimate AGI outcomes with important implications for ethics, risk, and policy. Empirical study of these populations could clarify the nature of the work being done, identify gaps, and suggest trends in AGI R&D.

6.4 Concluding Remarks

Overall, the present study shows that discussion of AGI ethics, risk, and policy can have a sound empirical basis; it need not be based solely on speculation about hypothetical AGI R&D projects and actors. The present study contributes to this empirical basis by being just the second survey of AGI R&D projects and the first to track changes over time. Given the potentially high stakes of AGI, hopefully this research can be used to improve AGI outcomes.

Appendix 1. Active AGI R&D Projects

The following pages document the 70 AGI R&D projects that are active in 2020. Each project entry lists the project website, a brief summary of the project, and details the project's attributes.

Projects with an asterisk (*) next to their name indicate that the project is new to the 2020 survey.

For projects that appear in both the 2017 and 2020 surveys, some data points have been coded differently in the two surveys. These data points are annotated as follows:

- Red font indicates a data point that has been changed due to a change in coding methodology. For example, the 2020 survey uses a more restrictive standard for coding partner countries. As a result, AERA is coded in the 2020 survey as having no partner countries, whereas in the 2017 survey it was coded as having a partner country of Switzerland. Full methodology changes are detailed in Section 4.
- Blue font indicates a data point that has been changed due to change in the project that occurred between 2017 and 2020. For example, Animats added Australia as a partner country in 2020.
- Green font indicates a data point that has been changed to correct a coding error or oversight in the 2017 survey. For example, Leabra was coded as unspecified on safety in the 2017 survey; it should have been coded as moderate due to information in a 2017 paper that was overlooked in the 2017 survey.

Referenced websites were active when project data was collected and coded during June to September 2020. Some websites may have since become inactive. Many of these websites can be viewed via the Internet Archive (https://archive.org).

ACT-R

Main website: http://act-r.psy.cmu.edu

ACT-R, an acronym for Adaptive Control of Thought-Rational, is a research project led by John Anderson of Carnegie Mellon University. It is a theory of human cognition and a computational framework for simulating human cognition.[10] ACT-R was briefly connected to Leabra via the SAL project.[11]

Lead institution: Carnegie Mellon University

Partner institutions: none

- The ACT-R website lists 110 collaborating institutions across 19 countries.[12] However, this list includes researchers who have previously contributed to ACT-R research and have since moved on to other projects outside of ACT-R. This list also does not include some co-authors of recent ACT-R papers.[13] No active partner institutions could be confirmed from the website, and thus none are coded here, though there may be active partner institutions.

Type of institution: academic

Open-source: yes[14]

Military connection: yes[15]
The US Air Force Institute of Technology, US Air Force Research Lab, US Army Research Lab, US Naval Undersea Research Center, and US Military Academy are listed among the 110 collaborating partners on the ACT-R website.[16]

Lead country: USA

Partner countries: none

Stated goals: intellectualist

- The main description of ACT-R on its website is exclusively about academic research with no mention of broader impacts.[17]

Engagement on safety: unspecified

Size: medium

[10] http://act-r.psy.cmu.edu/about and http://act-r.psy.cmu.edu/peoplepages/ja/ja-interests.html
[11] See the entry for Leabra in Appendix 1 for more information.
[12] http://act-r.psy.cmu.edu/people
[13] For example, Wirzberger et al. (2020) has lead authors J.F. Krems and G.D. Rey, who are not listed at http://act-r.psy.cmu.edu/people.
[14] http://act-r.psy.cmu.edu/software
[15] Funding from the US Office of Naval Research was previously reported in Zhang et al. (2016).
[16] http://act-r.psy.cmu.edu/people/
[17] http://act-r.psy.cmu.edu/about

AERA

Main website: http://www.ru.is/faculty/thorisson

AERA, an acronym for Autocatalytic Endogenous Reflective Architecture, is led by Kristinn Thórisson of CADIA at Reykjavik University and the Icelandic Institute of Intelligent Machines (IIIM).[18] The project aims "to both understand the mind and build a practical AGI system."[19] AERA is currently used "to study machine understanding, teaching methodologies for artificial learners, [and] even the development of ethical values."[20]

Lead institution: Reykjavik University

Partner institutions: Icelandic Institute for Intelligent Machines
- Thorisson publishes AGI research with researchers at various academic institutions, though none are listed as official institutional partnerships.[21]

Type of institution: academic

Open-source: no

Military connection: no
- Thórisson criticizes military AI in the Icelandic Institute for Intelligent Machines (IIIM) ethics policy.[22]

Lead country: Iceland

Partner countries: none (2017 survey: Switzerland)

Stated goals: humanitarian, intellectualist
- Thórisson's website links to the ethics policy of IIIM, which aims to "to advance scientific understanding of the world, and to enable the application of this knowledge for the benefit and betterment of humankind." It further aims to be "for the betterment of society, human livelihood and life on Earth." This mention of life on Earth suggests ecocentrism, but all other text is humanitarian or intellectualist.[23]

Engagement on safety: active
- The AERA group has written about enhancing safety during AGI self-improvement, arguing that certain design principles would make it easy to achieve safety.[24]

Size: small-medium

[18] Nivel et al. (2013)
[19] http://www.ru.is/faculty/thorisson
[20] http://www.ru.is/faculty/thorisson
[21] http://alumni.media.mit.edu/~kris/select_publ.html
[22] http://www.iiim.is/ethics-policy/
[23] http://www.iiim.is/ethics-policy/3
[24] Steunebrink et al. (2016)

AGI Brain*

Main website: https://www.researchgate.net/project/Artificial-General-Intelligence-Brain-AGI-Brain

AGI Brain was founded by Mohammadreza Alidoust in 2009, though it only shows activity in 2018-2019.[25] The project's stated goal is "to build an Artificial General Intelligence (AGI) Brain, i.e., a thinking machine that can perform any intellectual task that a human being can."[26] Alidoust is affiliated with AGT Co. in Mashhad, Iran.[27]

Lead institution: AGI Brain

Partner institutions: none

Type of institution: none

Open-source: no

Military connection: unspecified

Lead country: Iran

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[25] https://www.researchgate.net/project/Artificial-General-Intelligence-Brain-AGI-Brain. An earlier paper (Alidoust and Rouhani 2015) presents work that appears to be a precursor to the 2018-2019, though this earlier paper is not explicitly identified as part of the AGI Brain project.
[26] https://www.researchgate.net/project/Artificial-General-Intelligence-Brain-AGI-Brain
[27] https://ir.linkedin.com/in/mohammadreza-alidoust

AGI Laboratory*

Main website: https://agilaboratory.com

AGI Laboratory is a private corporation that leads several AGI-related activities, including the R&D project Mediated Artificial Superintelligence (MASI). It is an active research group with publications in a variety of notable conferences and journals.[28]

Lead institution: AGI Laboratory

Partner institutions: Microsoft, Institute for Education, Research, & Scholarships (IFERS), and Unanimous AI[29]

Type of institution: private corporation
- Artificial General Intelligence, Inc. (AGI Laboratory) is a privately held 'C' Corp.[30]

Open-source: no

Military connection: unspecified

Lead country: USA[31]

Partner countries: none

Stated goals: humanitarian, transhumanism
- A website for the MASI project lists its goal as "uplifting humanity through the digital transformation, collective superintelligence, and governance."[32]
- An AGI Lab paper explores the possibility of "identical neural substrates" for "humans and non-biological human counterparts."[33]
- A separate paper studies the ethical treatment of sentient AGI systems.[34]
- Chief Research Scientist David Kelley is also the founder of IAmTranshuman.[35]

Engagement on safety: active
- AGI Laboratory has an ethics webpage with extensive attention to AGI safety.[36]

Size: small

[28] https://agilaboratory.com/research/
[29] https://agilaboratory.com/support/
[30] https://agilaboratory.com/support/; See also author locations: Kelley and Atreides (2020)
[31] https://agilaboratory.com/support/
[32] https://uplift.bio/
[33] Dambrot (2020)
[34] Kelley and Atreides (2020)
[35] https://iamtranshuman.org/
[36] https://agilaboratory.com/ethics/

AGi3*

Main website: https://agiinnovations.com

AGi3, also known as AGI Innovations Inc., is a private corporation that was founded in 2013 by Peter Voss with a "long-term goal of developing and commercializing" AGI. Its R&D is based on Adaptive A.I. Inc (a2i2), founded in 2001.[37] AGi3 has an active commercial product called Aigo, which aims to provide "highly intelligent, hyper-personalized conversational assistants for everyone."[38]

Lead institution: AGi3

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified

Lead country: USA
- AGi3 lists its location as Los Angeles.[39]

Partner countries: none

Stated goals: humanitarian, profit
- AGi3's website states a "goal of developing and commercializing" AGI.[40]
- In a 2015 interview about AGi3, Voss predicts that AGI will solve many of humanity's problems.[41]

Engagement on safety: moderate
- In the same 2015 interview, Voss says there are "risks and dangers" involved with AGI.[42]

Size: small-medium

[37] https://agiinnovations.com/history
[38] https://www.aigo.ai/about-us
[39] http://www.agi-3.com/contact.html
[40] https://agiinnovations.com/history
[41] https://www.youtube.com/watch?v=brLeWnrq7N8 (4:40)
[42] https://www.youtube.com/watch?v=brLeWnrq7N8 (5:10)

AIBrain*

Main website: https://aibrain.com

AIBrain is an AI company that seeks to build a "human-like AI" that automates "the entire process of cognitive reasoning."[43]AIBrain was founded by Dr. Richard H. Shinn in Seoul, South Korea in 1997. It expanded to Palo Alto, California in 2012 and to Shenzhen, China, in 2017.[44] AIBrain focuses on three pieces of technology: AIBrain, AI Core, and Memory Graph.[45]

Lead institution: AIBrain

Partner institutions: none

Type of institution: private corporation

Open-source: yes[46]

Military connection: unspecified

Lead country: USA[47]

Partner countries: China, Germany, South Korea[48]

Stated goals: intellectualist
- AI Brain states that it has the "grand aim to augment human intelligence with AI."[49]

Engagement on safety: unspecified
- AI Brain Chief Scientist Steve Omohundro previously worked on AI safety,[50] though this work was not done for AIBrain.

Size: small-medium

[43] https://aibrain.com/about/
[44] https://aibrain.com
[45] https://aibrain.com/solutions/
[46] https://github.com/aibraininc
[47] See the AIBrain website footer.
[48] See the AIBrain website footer.
[49] https://aibrain.com/about/
[50] https://aibrain.com/about/who-we-are/ and Omohundro (2016)

AIXI

Main website: http://www.hutter1.net/ai/aixigentle.htm

AIXI is led by Marcus Hutter of Australian National University. AIXI is based on a "meta-algorithm" that searches the space of algorithms to find the best one for AGI.[51] Hutter proves that AIXI will find the most intelligent AI if given infinite computing power. While this is purely a theoretical result, it has led to approximate versions implemented in computer code.[52] Hutter is currently on leave from the Australian National University and is working at Google DeepMind in London, possibly pausing his work on AIXI.[53]

Lead institution: Australian National University

Partner institutions: none

Type of institution: academic

Open-source: yes[54]

Military connection: unspecified

Lead country: Australia

Partner countries: none

Stated goals: humanitarian (2017 survey: unspecified)

- In his book *Universal Artificial Intelligence*, Hutter states that human-level AI "would have enormous implications on our society" and is therefore worth pursuing "to reap the benefits."[55]

Engagement on safety: unspecified

Size: small

[51] Goertzel (2014) p.25
[52] Veness et al. (2011)
[53] http://www.hutter1.net/index.htm
[54] https://github.com/aslanides/aixijs and http://www.hutter1.net/aixijs/index.html
[55] http://www.hutter1.net/ai/uaibook.htm#oneline

Animats

Main website: https://github.com/ni1s/animats

Animats is a small project developed for the First International Workshop on Architectures for Generality & Autonomy[56] and the 2017 AGI conference,[57] which seeks to build AGI based on animal intelligence. [58] The project was originally a collaboration between researchers at universities in Sweden and the United States, and the Swiss company NNAISENSE,[59] but it seems most researchers from the 2017 project have since moved on to other projects. Animats currently seems to be led by Claes Strannegård.[60]

Lead institution: Chalmers University of Technology

Partner institutions: University of Gothenburg, Harvard University, Deakin University, NNAISENSE

Type of institution: academic (2017 survey: also private corporation)

Open-source: yes[61]

Military connection: unspecified (2017 survey: no)

Lead country: Sweden

Partner countries: Australia,[62] Switzerland, USA[63] (2017 survey: Switzerland and USA only)

Stated goals: unspecified

Engagement on safety: moderate[64]

Size: small

[56] http://cadia.ru.is/workshops/aga2017
[57] Strannegård et al. (2017a)
[58] Strannegård et al. (2017b)
[59] Strannegård et al. (2017b)
[60] https://www.chalmers.se/en/staff/Pages/claes-strannegard.aspx
[61] https://github.com/ni1s/animats
[62] Strannegård et al. (2020).
[63] Strannegård et al. (2017b)
[64] Strannegård et al. (2018), p.78.

ANSNA*

Main website: https://github.com/patham9/ANSNA

The ANSNA (Adaptive Neuro-Symbolic Network Agent) project was created by Patrick Hammer, a research assistant at Temple University in Dr. Pei Wang's research team.[65] ANSNA is situated within the "field of AGI," [66] and is a variation of the NARS and OpenNARS projects.[67] Hammer states that the goal of the ANSNA project is to try "to combine [his] most valuable insights about Dr. Pei Wang's NARS (Non-Axiomatic Reasoning System), Jeffrey Hawkins HTM (Hierarchical Temporal Memory), Tony Lofthouse's ALANN (Adaptive Logic and Neural Network), Rod Rinkus's Sparsey, Pentti Kanerva's SDM (Sparse distributed memory), and [his] own projects of the last decade, for creating an autonomous sensorimotor agent that starts with zero knowledge and organizes its experience into conceptual units in such an efficient way that it can directly be applied for autonomous systems with rich sensory data."[68]

Lead institutions: Temple University

Partner institutions: none

Type of institution: academic

Open-source: yes[69]

Military connection: unspecified

Lead country: USA

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

- Hammer is a research assistant of Dr. Pei Wang, who has discussed the ethics of AGI, and thereby Hammer may share similar views.[70]

Size: small

[65] https://www.linkedin.com/in/patrick-hammer-27a248b5/ and Wang et. al (2016)
[66] https://github.com/patham9/ANSNA/blob/master/README.md
[67] https://www.linkedin.com/in/patrick-hammer-27a248b5/
[68] https://github.com/patham9/ANSNA#readme
[69] https://github.com/patham9/ANSNA
[70] See NARS in Appendix 1

Apollo Program for AGI*

Main website: https://montrealartificialintelligence.com

The Apollo Program for AGI is a project of Montréal Artificial Intelligence. The Apollo Program was launched in 2019; Montréal AI was established in 2003.[71] The Apollo program seeks to "pioneer the general-purpose technology, the *holy grail of AI*, that will define the future" (emphasis original).[72] Montréal AI's overarching goal is to create the largest AI community possible that would intersect with every area of society, including academia, government, and business.[73]

Lead institution: Montréal AI

Partner institutions: none

Type of institution: private corporation[74]

Open-source: no
- GitHub pages associated with the project do not include AGI code.[75]

Military connection: unspecified

Lead country: Canada

Partner countries: none

Stated goals: humanitarian, intellectualist
- The Montréal AI website states the goals are to "orchestrate AGI-First systemic breakthroughs to solve planetary-scale challenges and to pioneer a new era of superhuman scientific discoveries."[76]
- Boucher describes Montréal AI as "a research Company at the forefront of the AI field developing and commercializing the most significant technology ever created to leverage enterprises and governments."[77]

Engagement on safety: moderate
- Montréal AI is "preparing a Global Network of Education Centers to pioneer an impactful understanding of AI and to foster a vector for safe humanitarian artificial general intelligence (AGI)."[78]

Size: small-medium

[71] https://www.linkedin.com/pulse/montrealai-research-vincent-boucher/
[72] https://montrealartificialintelligence.com
[73] http://www.montreal.ai/unicorn.jpg
[74] https://www.linkedin.com/company/montreal.ai/about/
[75] https://github.com/MontrealAI/MontrealAI.github.io and https://github.com/MontrealAI/MontrealAIMaster
[76] https://montrealartificialintelligence.com
[77] https://www.linkedin.com/in/montrealai/
[78] https://montrealartificialintelligence.com/

Astrum*

Main website: https://www.astrum.ai

Astrum was founded in 2019 by Srikanth Srinivas.[79] Srinivas has since taken a position at Scale AI, potentially leaving Astrum with no full-time employees. Astrum's mission is "dedicated to making AGI a reality within the next decade."[80]

Lead institution: Astrum

Partner institutions: none

Type of institution: private corporation[81]

Open-source: no

Military connection: unspecified

Lead country: Canada[82]

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[79] https://www.linkedin.com/in/srikanthsrnvs
[80] https://www.astrum.ai/about
[81] https://www.linkedin.com/company/astrum-ai/about/
[82] https://www.crunchbase.com/organization/astrum-ai

Baidu Research

Main website: http://research.baidu.com

Baidu Research is an AI research group within Baidu which claimed to be working on AGI in 2018, though this does not currently appear to be a significant theme for the group.[83] It has offices in Beijing, Shenzhen, and Sunnyvale, California.[84] Baidu Research has achieved success in "zero-shot learning" in language processing, in which the AI "is able to understand unseen sentences."[85] Some observers rate this as a significant breakthrough.[86]

Lead institution: Baidu

Partner institutions: none

Type of institution: public corporation

Open-source: no
- Baidu releases some work open-source,[87] but it is not AGI.

Military connection: unspecified
- Baidu receives AI funding from the Chinese government for "computer vision, biometric identification, intellectual property rights, and human-computer interaction."[88]

Lead country: China

Partner countries: USA

Stated goals: unspecified

Engagement on safety: unspecified
- Former Baidu Research Chief Scientist Andrew Ng says that Baidu takes safety and ethics seriously. He expresses his personal views that AI will help humanity, but says that "fears about AI and killer robots are overblown."[89] This was not in the context of AGI. No direct discussion of safety by Baidu Research was found.

Size: small-medium

[83] http://research.baidu.com/Blog/index-view?id=108
[84] http://bdl.baidu.com/contact_b.html
[85] http://research.baidu.com/Blog/index-view?id=99
[86] Griffin (2017)
[87] https://github.com/baidu
[88] Gershgorn (2017)
[89] Maddox (2016)

Big Mother*

Main website: https://bigmother.ai/

Big Mother was founded by Aaron Turner in 2018.[90] Its main goal is "to build a maximally-safe, maximally-benevolent, maximally-trustworthy, autonomous, goal-directed, super-intelligent machine (Artificial General Intelligence) called Big Mother, that is publicly owned by all mankind."[91]

Lead institution: Big Mother

Partner institutions: none

Type of institution: nonprofit

Open-source: restricted
- Big Mother members are allowed access to additional resources.[92]

Military connection: unspecified

Lead country: UK[93]

Partner countries: none
- Big Mother lists volunteers in Canada, the Netherlands, Poland, Russia, the UAE, the US, and the UK.[94]

Stated goals: humanitarian
- The Big Mother website states, "if we equate 'human happiness' with 'the extent to which true human preferences are satisfied' then, simply stated, the net effect of Big Mother's dominant goal (top-level objective) will be to (broadly) maximize human happiness and (broadly) minimize 'happiness inequality'."[95]

Engagement on safety: active
- Big Mother's mission statement emphasizes "maximally-safe" AGI.[96]
- Though the Big Mother website states this project could be active for the next 50-100 years, it also says that the project may take longer "if safety requires."[97]

Size: small-medium[98]

[90] Turner (2019)
[91] https://bigmother.ai
[92] https://bigmother.ai/page-1075267
[93] Turner is based in Cambridge, England: https://www.meetup.com/CambridgeAGI1/
[94] https://bigmother.ai
[95] https://bigmother.ai
[96] https://bigmother.ai
[97] https://bigmother.ai
[98] https://bigmother.ai/Sys/Login?ReturnUrl=%2fpage-1075213

Binary Neurons Network*

Main website: http://8iter.ru/ai.html[99]

Binary Neurons Network is a project by Ilya Shishkin. A version of the website was available in 2016.[100] It is "an attempt to create artificial intelligence in the original meaning: AGI, Strong AI, HLAI, True AI."[101] Shiskin also has a GitHub project called Device, which is the "practical implementation of a physical device with the binary neurons."[102]

Lead institution: Binary Neurons Network

Partner institutions: none

Type of institution: none

Open-source: yes[103]

Military connection: unspecified

Lead country: Russia

Partner countries: none

Stated goals: unspecified

Engagement on safety: moderate

- Shiskin writes, "AI is certainly unsafe, it makes decisions by itself, the more developed it is, the more unpredictable these decisions are. Attempts to instill something in him, he will see through in no time and then there will be no mercy. Our only hope is that AI will look at us orphaned and poor and, within the framework of humanitarian aid, will solve all our problems for us."[104]

Size: small

[99] English translation: http://translate.google.com/translate?js=n&sl=ru&tl=en&u=http://8iter.ru/ai.html
[100] https://web.archive.org/web/20161123081444/http://8iter.ru/ai.html
[101] https://github.com/cortl0/binary_neurons_network
[102] https://github.com/cortl0/device
[103] https://github.com/cortl0/binary_neurons_network
[104] http://translate.google.com/translate?js=n&sl=ru&tl=en&u=http://8iter.ru/ai.html

Blue Brain

Main website: http://bluebrain.epfl.ch

Blue Brain is a research project led by Henry Markram that has been active since 2005. Its website states that its goal is "to build biologically detailed digital reconstructions and simulations of the mouse brain."[105] A documentary on Blue Brain that has been ten years in the making will release in 2020.[106] Markram also founded the Human Brain Project, which shares a research strategy with Blue Brain.

Lead institution: École Polytechnique Fédérale de Lausanne

Partner institutions: none

- The 2017 survey listed 8 partner institutions across 4 countries on a webpage that is now inactive. Blue Brain may retain these or other partnerships, though insufficient evidence was identified. Blue Brain researchers have published papers co-authored by researchers at various institutions.

Type of institution: academic

- Blue Brain lists the public corporation IBM as a contributor, making "its researchers available to help install the BlueGene supercomputer and set up circuits that would be adequate for simulating a neuronal network." This contribution was judged to be too small to code Blue Brain as part public corporation.

Open-source: yes[107]

Military connection: unspecified

Lead country: Switzerland

Partner countries: none (2017 survey: Israel, Spain, UK, USA)

Stated goals: humanitarian, intellectualist

- The Blue Brain website states that "understanding the brain is vital, not just to understand the biological mechanisms which give us our thoughts and emotions and which make us human, but for practical reasons,"[108] the latter including computing, robotics, and medicine. These applications are broadly humanitarian, though a more muted humanitarianism than what is found for other projects.

Engagement on safety: unspecified

Size: large (2017 survey: medium-large)

[105] https://www.epfl.ch/research/domains/bluebrain/
[106] Hutton (2016)
[107] https://github.com/BlueBrain
[108] https://www.epfl.ch/research/domains/bluebrain/blue-brain/about/why-is-this-important/

Brain Simulator II*

Main website: https://futureai.guru/brainsim.aspx

Brain Simulator II is a "Neural Simulator for AGI research and development" created by Charles Simon.[109]

Lead institution: Future AI

Partner institutions: none

Type of institution: private corporation[110]

Open-source: yes[111]

Military connection: unspecified

Lead country: USA[112]

Partner countries: none

Stated goals: unspecified

Engagement on safety: moderate

- Simon wrote a book called *Will Computers Revolt?* and dedicates an entire chapter to different near-term and long-term threats posed by AGI. He concludes that humanity has a reason to proceed cautiously.[113]

Size: small

[109] https://futureai.guru/BrainSim.aspx
[110] https://futureai.guru/PR/Startup.aspx
[111] https://github.com/FutureAIGuru/BrainSimII
[112] https://futureai.guru/brainsim.aspx
[113] Simon (2018)

Brain2Bot*

Main website: https://web.archive.org/web/20180513230636/http://brain2bot.tech/

Brain2Bot was founded and run by Gunnar Newquist from between 2015 and 2019.[114] Their LinkedIn page states, "We are going beyond AI to develop an Artificial General Intelligence system with the intuition and emotion of a real living being."[115]

Lead institution: Brain2Bot

Partner institutions: none

Type of institution: private corporation[116]

Open-source: no

Military connection: unspecified

Lead country: United States
- Brain2Bot is listed in Crunchbase as being based in Reno, Nevada.[117]

Partner countries: none

Stated goals: transhumanist
- Brain2Bot seeks to build AGI with "intuition and emotion of a real living being," that will have "heart & soul," and "can become your companions, not just tools that you use."[118]

Engagement on safety: unspecified

Size: small

[114]https://www.linkedin.com/in/gunnar-newquist
[115]https://www.linkedin.com/company/brain2bot
[116]https://www.crunchbase.com/organization/brain2bot-inc
[117]https://www.crunchbase.com/organization/brain2bot-inc
[118]https://www.linkedin.com/company/brain2bot/about

Cerenaut Research*

Main website: https://cerenaut.ai

Cerenaut Research was founded as ProjectAGI by Gideon Kowadlo and David Rawlinson and dates to at least 2016.[119] It is "explicitly working towards Artificial General Intelligence (AGI) using the best tricks from Machine Learning, Computational Neuroscience and Artificial Intelligence."[120]

Lead institution: Incubator 491

Partner institutions: WBAI and Luria Neuroscience Institute[121]

Type of institution: private corporation

Open-source: yes
- Some AGI code is available on GitHub,[122] the Cerenaut Research website lists multiple joint projects which have additional content available upon request.[123]

Military connection: unspecified

Lead country: Australia[124]

Partner countries: Japan, USA[125]

Stated goals: humanitarian, intellectualist
- Cerenaut Research's vision is "a world enriched and secured by ubiquitous artificial general intelligence, contributing to advances in science, society, industry, and our understanding of ourselves."[126]

Engagement on safety: dismissive
- A Cerenaut blog post states, "We're not worried about runaway 'paperclip maximizers' or 'skynet'-style machine coups… Despite good intentions, AI risk-awareness groups such as MIRI may cause more harm than good by focusing public debate on the more distant *existential* risks of AI, while distracting from the immediate risks and harm being perpetrated **right now** using AI & ML" (emphasis original).[127]

Size: small-medium

[119] https://www.crunchbase.com/organization/project-agi
[120] https://research.cerenaut.ai/mission/
[121] https://research.cerenaut.ai/team/
[122] https://github.com/Cerenaut
[123] https://cerenaut.ai/requests-for-research/
[124] See website footer: https://research.cerenaut.ai
[125] https://research.cerenaut.ai/team/
[126] https://research.cerenaut.ai/mission/
[127] https://research.cerenaut.ai/2019/03/19/ai-is-already-harming-us-but-not-the-way-you-think/

China Brain Project

Main website: none found

China Brain Project is a research project of the Chinese Academy of Sciences focused on basic and clinical neuroscience and brain-inspired computing. As of July 2016, the Chinese Academy of Sciences had announced the project and said it would launch soon,[128] but in September 2020 no project website was found. Project lead Mu-Ming Poo describes the project as "learning from information processing mechanisms of the brain is clearly a promising way forward in building stronger and more general machine intelligence" and that "the China Brain Project will focus its efforts on developing cognitive robotics as a platform for integrating brain-inspired computational models and devices."[129]

Lead institution: Chinese Academy of Sciences

Partner institutions: Peking University, Tsinghua University, and Academy of Military Medical Sciences[130]

Type of institution: government
- The Chinese Academy of Sciences is a public institution under the Chinese government.[131]
- Mu-Ming Poo lists the Chinese Natural Science Foundation and Ministry of Science and Technology as guiding organizations for the China Brain Project.[132]

Open-source: no

Military connection: yes (2017 survey: unspecified)
- The China Brain Project has partnered with the Academy of Military Medical Sciences.[133]

Lead country: China

Partner countries: none

Stated goals: humanitarian, intellectualist
- Mu-Ming Poo describes the project's goals as "understanding the neural basis of human cognition" and "reducing the societal burden of major brain disorders."[134]

Engagement on safety: unspecified

Size: small

[128] Chen (2016)
[129] Poo et al. (2016)
[130] Cyranoski (2018)
[131] http://english.cas.cn/about_us/introduction/201501/t20150114_135284.shtml
[132] Poo et al. (2016)
[133] Cyranoski (2018)
[134] Poo et al. (2016)

CLARION

Main website: https://sites.google.com/site/drronsun/clarion/clarion-project

CLARION is a cognitive architecture project led by Ron Sun of Rensselaer Polytechnic Institute, aiming to investigate the human mind's fundamental structures by synthesizing many intellectual ideas into a unified, coherent model of cognition.[135]

Lead institution: Rensselaer Polytechnic Institute[136]

Partner institutions: none

Type of institution: academic

Open-source: yes[137]

Military connection: yes[138]

Lead country: USA

Partner countries: none

Stated goals: intellectualist

- The CLARION website states that it "aims to investigate the fundamental structures of the human mind," with "the ultimate goal of providing unified explanations for a wide range of cognitive phenomenon."[139]

Engagement on safety: unspecified

Size: small-medium

[135] http://www.clarioncognitivearchitecture.com
[136] http://www.clarioncognitivearchitecture.com/team
[137] http://www.clarioncognitivearchitecture.com/downloads
[138] The homepage states "The CLARION cognitive architecture project is headed by Professor Ron Sun and has been supported by such agencies as ONR, ARI, and others." Additionally, the 2017 Survey listed funding by the US Office of Naval Research and Army Research Institute.
[139] http://www.clarioncognitivearchitecture.com/home

Cognitive Science & Solutions*

Main website: https://cogscisol.com

Cognitive Science & Solutions (CogSciSol) was founded in 2015 by David Sherwood and Terry Higbee.[140] Their AGI is called "elastic representations" in order to emphasize that their code is not brittle.[141] The CogSciSol website appears to have been made public in 2018,[142] suggesting a possible stealth mode state from 2015 to 2017. Their mission is to "design and develop disruptive new advances in AI that allow information processing systems to demonstrate human-like understanding and reasoning."[143]

Lead institution: Cognitive Science & Solutions

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified
- Co-founders Sherwood and Higbee both have military service backgrounds and recently published an article on AGI possibly being used as military technology.[144]

Lead country: USA[145]

Partner countries: none

Stated goals: unspecified
- The CogSciSol website includes a page dedicated to business opportunities and says, "there is no shortage of opportunities for CogSciSol to monetize our technology."[146]

Engagement on safety: active
- They describe their AI as "Trustworthy AI - Explainable - Reliable – Repeatable."[147]

Size: small

[140] https://cogscisol.com/about-us-2/
[141] https://cogscisol.com/our-technology/
[142] https://web.archive.org/web/20180806205351/http://cogscisol.com/
[143] https://cogscisol.com/about-us-2/
[144] Sherwood and Higbee (2020)
[145] https://cogscisol.com/contact-us/
[146] https://cogscisol.com/business-opportunities/
[147] https://cogscisol.com/

CommAI

Main website: https://research.fb.com/projects/commai

CommAI is a project of Facebook AI Research (FAIR). FAIR is primarily led by Marco Baroni and Tomas Mikolov. The FAIR website states that CommAI is "a project aiming at developing general-purpose artificial agents that are useful for humans in their daily endeavours."[148] CommAI is used in the General AI Challenge sponsored by GoodAI.[149]

Lead institution: Facebook

Partner institutions: none

Type of institution: public corporation

Open-source: yes[150]

Military connection: unspecified
- Facebook does not appear to have any US defense contracts.[151]

Lead country: USA

Partner countries: Canada, France, Israel, UK[152] (2017 survey: France only)

Stated goals: humanitarian
- The CommAI website states that it aims "at developing general-purpose artificial agents that are *useful* for humans in their daily endeavours" (emphasis original).[153]
- Facebook is also a founding partner of the Partnership on AI to Benefit People & Society, which has humanitarian goals.[154]

Engagement on safety: moderate
- Facebook is a founding partner of the Partnership on AI to Benefit People & Society, which expresses concern about AI safety,[155] but CommAI shows no direct involvement on safety.

Size: medium

[148] https://research.fb.com/downloads/commai/
[149] https://research.fb.com/downloads/commai/
[150] https://github.com/facebookresearch/CommAI-env
[151] https://www.fpds.gov/ezsearch/fpdsportal?q=facebook+DEPARTMENT_FULL_NAME%3A"DEPT+OF+DEFENSE"
[152] https://research.fb.com/category/facebook-ai-research-/ and https://research.fb.com/commai-fellowships-and-visiting-researcher-programs/
[153] https://research.fb.com/downloads/commai/
[154] https://www.partnershiponai.org/tenets
[155] https://www.partnershiponai.org/tenets

Core AI*

Main website: https://www.akin.com

Core AI is the project of Akin, a company building personalized AI for "space habitat management, long-duration companionship, and complex tasks."[156] Their three projects are Core AI, Akin Space, and Akin Home.[157] Liesl Yearsly founded Akin in 2017 when IBM Watson acquired Yearsly's AI chatbot company Cognea.[158] The Akin website appears to have launched in 2018.[159]

Lead institution: Akin

Partner institutions: none

Type of institution: private corporation[160]

Open-source: no

Military connection: unspecified

Lead country: Australia[161]

Partner countries: USA[162]

Stated goals: humanitarian
- Akin's homepage says they are building AI "to help people organize their lives, achieve goals, and free up time for things they love."[163]
- Akin's website also states that "the next generation of AI will improve wellbeing across the matrix of life."[164]

Engagement on safety: moderate
- Though Yearsly acknowledges risks with future AGI, she seems more concerned with near-term risks.[165]

Size: small-medium

[156] https://www.akin.com
[157] https://www.akin.com
[158] https://www.linkedin.com/in/liesl/
[159] https://web.archive.org/web/20180809091039/https://www.akin.com/
[160] https://www.linkedin.com/company/personal-ai/about/
[161] https://www.crunchbase.com/organization/akin-fff7
[162] https://twitter.com/a_kin_ai
[163] https://www.akin.com
[164] https://www.akin.com/applied
[165] https://www.akin.com/blog/we-need-to-talk-about-the-power-of-ai-to-manipulate-humans

Curious AI*

Main website: https://thecuriousaicompany.com

The Curious AI Company (referred to as Curious AI) aims "to generalise (AGI) our machine learning technology components for a software solution capable of human-level knowledge work: a digital co-worker."[166] Curious AI was founded by Harri Valpola in 2015.[167] System 2 is "is the more reasoned approach to problem solving, based on internal models and simulations, drawing on imagination, planning and analysis by simulating potential root causes."[168]

Lead institution: The Curious AI Company

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified

Lead country: Finland[169]

Partner countries: none

Stated goals: profit

- The Curious AI Company homepage states, "Our AI solutions offer immediate savings over existing IT systems and processes. And the new intelligence unlocks completely new business opportunities for your organisation."[170]

Engagement on safety: unspecified

Size: small-medium

[166] https://thecuriousaicompany.com
[167] https://www.linkedin.com/in/harri-valpola-1805772/?originalSubdomain=fi
[168] https://thecuriousaicompany.com/terms/
[169] https://www.linkedin.com/company/the-curious-ai-company/about/
[170] https://thecuriousaicompany.com

Cyc

Main website: http://www.cyc.com

Cyc is led by Doug Lenat, who founded Cyc in 1984. Lead institution Cycorp, based in Austin, Texas, uses Cyc for services including corporate and government consulting.[171] The Cycorp website describes Cyc as "a leading provider of semantic technologies that bring a new level of intelligence and common sense reasoning to a wide variety of software applications. The Cyc software combines an unparalleled common sense ontology and knowledge base with a powerful reasoning engine and natural language interface to enable the development of novel knowledge-intensive applications."[172]

Lead institution: Cycorp

Partner institutions: Lucid AI[173]
- Cyc partners with Lucid AI in an attempt to commercialize their research, particularly applying Cyc's general knowledge software to a virtual assistance.[174]

Type of institution: private corporation

Open-source: restricted
- Cyc was available open-source from 2001 to 2017 and now requires a license. [175]

Military connection: yes
- Cycorp received a $25 million contract to analyze terrorism for the US military.[176]

Lead country: USA

Partner countries: none

Stated goals: humanitarian (2017 survey: unspecified)
- The Cyc website says, "AI makes us humans *better*. With AI, we can avoid falling prey to cognitive biases, and we can stretch our human intelligence to do amazing things" (emphasis original).[177]

Engagement on safety: active (2017 survey: unspecified)
- The Lucid website states, "Lucid.AI provides detailed, logical justifications for (and against!) each conclusion it reaches. Human experts can easily review Lucid.AI's reasoning chains. The result is greater trust."[178]

Size: medium

---

[171] http://www.cyc.com/enterprise-solutions
[172] https://www.cyc.com/about-us
[173] https://lucid.ai
[174] Knight (2016)
[175] https://github.com/asanchez75/opencyc/blob/master/README.md
[176] Deaton et al. (2005); https://www.cyc.com/about-us/frequently-asked-questions/what-differentiates-the-cyc-ontology
[177] https://www.cyc.com/about-us
[178] https://lucid.ai/

DeepBrainz*

Main website: https://deepbrainz.com/

DeepBrainz is a company founded by Arunkumar Venkataramanan in 2015.[179] Its website appears to have been made public in 2019.[180] The DeepBrainz website states that it is "building the human-level responsible AI for enterprises."[181] It lists several projects, the most AGI-oriented of which is called Universal AI.

Lead institution: DeepBrainz

Partner institution: Google Cloud[182]

Type of institution: private corporation[183]

Open-source: no
- Universal AI has a GitHub profile, but it does not provide code.[184]

Military connection: unspecified

Lead country: India[185]

Partner countries: USA[186]

Stated goals: humanitarian
- DeepBrainz's main goal is to "find smart ways of using technology that will help build a better tomorrow for everyone, everywhere."[187]

Engagement on safety: moderate
- The DeepBrainz website says that "AI systems should be robust and reliably operate in accordance with our intended purpose throughout our lifecycle."[188]

Size: small

[179] https://www.linkedin.com/in/arunkumarramanan/
[180] https://web.archive.org/web/20191128154944/https://deepbrainz.com/
[181] https://deepbrainz.com/
[182] https://deepbrainz.com/assistant
[183] https://profile.thecapitalnet.com/#!/public/deepbrainz
[184] https://github.com/Deep-Brainz/Universal-AI
[185] https://www.linkedin.com/company/deepbrainz/about/
[186] https://deepbrainz.com/assistant
[187] https://deepbrainz.com/why-deepbrainz/
[188] https://deepbrainz.com/why-deepbrainz/

DeepMind

Main website: http://deepmind.com

DeepMind is an AI corporation led by Demis Hassabis and Shane Legg.[189] It was founded in 2010 and acquired by Google in 2014 for $650 million.[190] It seeks to develop "systems that can learn to solve any complex problem without needing to be taught how," and it works "from the premise that AI needs to be general."[191]

Lead institution: Google

Partner institutions: OpenAI[192]

Type of institution: public corporation

Open-source: yes[193]

Military connection: unspecified
- Google has extensive defense contracts in the US,[194] but these appear to be unrelated to DeepMind.
- In 2018, DeepMind researchers pledged not to work on lethal autonomous weapons.[195]

Lead country: UK

Partner countries: Canada, France, USA[196] (2017 survey: Canada and USA only)

Stated goals: humanitarian
- DeepMind writes, "We use our technologies for widespread public benefit and scientific discovery, and collaborate with others on critical challenges, ensuring safety and ethics are the highest priority."[197] AI will be "helping humanity tackle some of its greatest challenges, from climate change to delivering advanced healthcare."[198]

Engagement on safety: active
- DeepMind collaborates with OpenAI on long-term AI safety projects.[199]

Size: large[200]

---

[189] https://deepmind.com/about#our_global_community
[190] Gibbs (2014)
[191] Beattie et al. (2016)
[192] Amodei et al. (2017)
[193] https://github.com/deepmind and https://deepmind.com/research?filters=%7B%22collection%22:%5B%22OpenSource%22%5D%7D
[194] https://www.fpds.gov/ezsearch/fpdsportal?q=google+DEPARTMENT_FULL_NAME%3A"DEPT+OF+DEFENSE"
[195] Tucker (2018)
[196] https://deepmind.com/careers
[197] https://deepmind.com/about
[198] Leike et al. (2017)
[199] https://blog.openai.com/deep-reinforcement-learning-from-human-preferences
[200] Shead (2020)

Drayker*

Main website: https://www.drayker.com

Drayker, led by Hyadhuad Lucer, is "a unified intelligence, organization, and computing system" that seeks to build an AGI called DK.[201] The Drayker roadmap states the first main projects will be completed between 2027 and 2030.[202]

Lead institution: Drayker

Partner institutions: none

Type of institution: none
- Drayker refers to itself as a federation and embassy.[203]

Open-source: restricted[204]
The Drayker GitHub repositories do not include code, but additional resources may be available to Drayker account holders.[205]

Military connection: unspecified

Lead country: Brazil[206]

Partner countries: none

Stated goals: humanitarian[207]
The Drayker website states, "Humanity has immense potential that we plan to unlock and cultivate."[208]

Engagement on safety: unspecified

Size: small
- Lucer seems to be the only public team member.[209]

[201] https://www.drayker.com
[202] https://dknowledge.drayker.org/roadmap/global-main-roadmap.html
[203] https://www.drayker.com/organization/daf
[204] https://github.com/draykerdk/DAF
[205] https://www.drayker.com/community
[206] https://www.linkedin.com/in/hyadhuad-lucer-586b37185
[207] https://medium.com/drayker/about
[208] https://www.drayker.com/home
[209] https://landr.me/hyadhuad.id.blockstack

Fairy Tale Artificial General Intelligence Solutions (FTAGIS)*

Main website: https://github.com/fairy-tale-agi-solutions

FTAGIS is an AGI project created by Răzvan Flavius Panda. The project states that it is "Helping humanity create safe artificial general intelligence through large scale remote collaboration."[210] Its GitHub page includes a section "Attempts at creating Human-Level Artificial Intelligence."[211] A private corporation for the project was set up in Dublin in September 2017 and was dissolved in March 2020,[212] though their Facebook and GitHub have continued to remain active.[213] According to Panda's LinkedIn, he is still actively self-employed at Fairy Tale.[214] On another GitHub profile, he says, "we are creating free or commercial open-source software. With the final goal to help humanity create Safe Artificial General Intelligence."[215]

Lead institution: Fairy Tale - Artificial General Intelligence Solutions Limited

Partner institutions: none

Type of institution: private corporation

Open-source: yes[216]

Military connection: unspecified

Lead country: Ireland

Partner countries: none

Stated goals: animal welfare, ecocentric, humanitarian, intellectualist

- The FTAGIS GitHub page states goals of "accelerating the pace of science done by humanity," a "good future for sentient beings," the "pursuit of truth using science for the betterment of mankind and other life forms," a focus "on the most important problems humanity is facing," and "helping reduce resource consumption." The emphasis on sentient beings is additionally suggestive of a goal of transhumanism.

Engagement on safety: active

- The FTAGIS GitHub page contains frequent emphasis on its pursuit of safe AGI.

Size: small

[210]https://github.com/fairy-tale-agi-solutions/
[211]https://github.com/fairy-tale-agi-solutions/Human-Level-Artificial-Intelligence
[212]https://www.solocheck.ie/Irish-Company/Fairy-Tale-Artificial-General-Intelligence-Solutions-Limited-612462
[213]https://www.facebook.com/fairy.tale.artificial.general.intelligence/
[214] https://www.linkedin.com/in/razvan-flavius-panda/?originalSubdomain=ie
[215] https://github.com/razvan-flavius-panda/blog
[216] https://github.com/fairy-tale-agi-solutions

FLOWERS (FLOWing Epigenetic Robots and Systems)

Main website: https://flowers.inria.fr

FLOWERS is led by Pierre-Yves Oudeyer of Inria (Institut National de Recherche en Informatique et en Automatique, or French National Institute for Research in Computer Science and Automation) and David Filliat of ENSTA ParisTech. The project studies "mechanisms that can allow robots and humans to acquire autonomously and cumulatively repertoires of novel skills over extended periods of time."[217]

Lead institutions: Inria and ENSTA ParisTech

Partner institutions: none

Type of institution: academic, government
- Inria is a government research institute; ENSTA ParisTech is a public college

Open-source: yes[218]

Military connection: unspecified[219]

Lead country: France

Partner countries: none

Stated goals: intellectualist
- The FLOWERS website focuses exclusively on intellectual aspects of its AI research and also cognitive science alongside AI as one of its two research strands.[220]

Engagement on safety: active
- FLOWERS has explored safety in the context of human-robot interactions.[221]

Size: small-medium[222]

[217] https://flowers.inria.fr
[218] https://flowers.inria.fr/software
[219] Funding reported in recent publications comes mainly from government science foundations.
[220] https://flowers.inria.fr
[221] Oudeyer et al. (2011)
[222] https://flowers.inria.fr/team/

GoodAI

Main website: https://www.goodai.com

GoodAI is a private corporation led by computer game entrepreneur Marek Rosa. Rosa states that "GoodAI is building towards my lifelong dream to create general artificial intelligence. I've been focused on this goal since I was 15 years old."[223] The website states GoodAI's goal is to "Develop safe general artificial intelligence – as fast as possible – to help humanity and understand the universe."[224]

Lead institution: GoodAI

Partner institutions: none

Type of institution: private corporation

Open-source: yes[225]

Military connection: unspecified

Lead country: Czech Republic

Partner countries: none

Stated goals: humanitarian,[226] intellectualist
The GoodAI website states that its "mission is to develop general artificial intelligence - as fast as possible - to help humanity and understand the universe." It aims "to build general artificial intelligence that can find cures for diseases, invent things for people that would take much longer to invent without the cooperation of AI, and teach us much more than we currently know about the universe."[227] It emphasizes that building AGI "is not a race. It's not about competition, and not about making money.[228]

Engagement on safety: active[229]
- GoodAI reports having a dedicated AI safety team led by Marek Havrda.[230]
- The GoodAI FAQ addresses goals of AI safety.[231]

Size: medium

[223] https://www.goodai.com/about
[224] https://www.goodai.com
[225] https://github.com/GoodAI
[226] https://www.goodai.com/9-ways-agi-could-shape-the-world-for-the-better/
[227] https://www.goodai.com/about
[228] https://www.goodai.com/about
[229] https://www.goodai.com/roadmapping-the-ai-race-to-help-ensure-safe-development-of-agi/
[230] https://www.goodai.com/about
[231] https://www.goodai.com/research/

Graphen*

Main website: https://www.graphen.ai/

Graphen was founded in 2017 by Dr. Ching-Yung Lin, former Chief Scientist for Graph Computing at IBM.[232] It claims to be "building next-generation AI platforms based on graphs to produce full brain functions and generate novel industry solutions."[233]

Lead institution: Graphen

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified

Lead country: USA

Partner countries: China, Singapore, Taiwan[234]

Stated goals: humanitarian
- The Graphen website states, "We strive to advance AI technologies and use them to make an impact in real life!"[235]
- Graphen emphasizes their "enterprise-ready" solutions, which they also call "products."[236] This is suggestive of a profit goal though was not clear enough to be coded as such.

Engagement on safety: unspecified

Size: medium

[232] https://www.graphen.ai/#home
[233] https://www.graphen.ai/
[234] https://www.graphen.ai/#contact
[235] https://www.graphen.ai/#about
[236] https://www.graphen.ai/products/index.html

He4o*

Main website: http://www.jiaxiaogang.cn and https://github.com/jiaxiaogang/he4o

He4o is a GitHub project created by Jia Xiaogang. It is described as "an AGI system, which is a practical project of spiral theory on the information entropy reduction machine model, so he4o is essentially an information entropy reduction machine."[237]

Lead institutions: He4o

Partner institutions: none

Type of institution: none

Open-source: yes[238]

Military connection: unspecified

Lead country: China[239]

Partner countries: none

Stated goals: humanitarian

- Xiaogang has stated in reference to he4o, "I hope to promote the great wheel of human civilization and improve everyone's life."[240]

Engagement on safety: unspecified

Size: small

[237] Translated by Google Translate; https://github.com/jiaxiaogang/he4o
[238] https://github.com/jiaxiaogang/he4o
[239] http://jiaxiaogang.cn/html/MyResume.html
[240] http://jiaxiaogang.cn/html/MyResume.html

HTM (Hierarchical Temporal Memory)

Main website: https://numenta.com

HTM is led by Jeffrey Hawkins, who previously founded Palm Computing. HTM was developed by the Numenta corporation and an open-source community that the corporation hosts. HTM is based on a model of the human neocortex, cortical theory, and Numenta's Thousand Brains Theory released in 2018.[241] Their website states, "We believe that a neuroscience-based approach is the fastest path to creating general intelligence."[242]

Lead institution: Numenta

Partner institutions: none
- Numenta lists Cortical.io, Grok, and Intelletic Trading Systems (ITS) as partners, but none work specifically on AGI research and development, so they are omitted here.[243]

Type of institution: private corporation

Open-source: yes[244]

Military connection: unspecified
- HTM was used in a 2008 Air Force Institute of Technology student thesis.[245]

Lead country: USA

Partner countries: none

Stated goals: humanitarian, intellectualist
- Numenta's mission states that it pursues AI "to create machines that can understand the world and add great value to humanity."[246]
- Hawkins writes that "the future success and even survival of humanity may depend on" humanity "building truly intelligent machines," citing applications in energy, medicine, and space travel.[247]

Engagement on safety: dismissive
- Hawkins has dismissed concerns about AGI as a catastrophic risk, stating, "I don't see machine intelligence posing any threat to humanity."[248]

Size: small-medium (2017 survey: medium)

---

[241] https://numenta.com
[242] https://numenta.com
[243] https://numenta.com
[244] https://numenta.com/machine-intelligence-technology/licensing-and-partners/
[245] Bonhoff (2008)
[246] https://numenta.com
[247] Hawkins (2017). The article is discussed on the Numenta blog: https://numenta.com/blog/2017/06/IEEE-special-edition-article-by-Jeff
[248] Hawkins (2015) and https://lukemuehlhauser.com/reply-to-jeff-hawkins-on-ai-risk/

Human Brain Project (HBP)

Main website: http://www.humanbrainproject.eu

HBP is a project for neuroscience research and brain simulation. It is sponsored by the European Commission, with $1 billion committed over ten years which began in 2013.[249] Initially led by Henry Makram, it was reorganized following extended criticism.[250] It hosts brain simulation platforms, neuromorphic computing, understanding cognition, medical informatics, neurorobotics, and massive computing. Markram also founded Blue Brain, which shares a research strategy with HBP.[251]

Lead institution: École Polytechnique Fédérale de Lausanne

Partner institutions: There are 140 institutions listed as active participants on the HBP website.[252]

Type of institution: academic

Open-source: restricted
- Additional resources may be available to those who request a HBP account.[253]

Military connection: no
- As an EU Horizons 2020 project, the project is constrained to civil purposes.[254]
- An HBP report discusses restrictions to military applications of openly published science.[255]

Lead country: Switzerland

Partner countries: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Portugal, Slovenia, Spain, Sweden, Turkey, and UK[256]

Stated goals: animal welfare, humanitarian, intellectualist
- HBP pursues brain simulation to "reduce the need for animal experiments" and "study diseases."[257] It also lists "understanding cognition" as a core theme.[258]

Engagement on safety: unspecified

Size: large

[249] http://www.humanbrainproject.eu/en/science/overview/
[250] Theil (2015)
[251] http://bluebrain.epfl.ch/page-52741-en.html
[252] http://www.humanbrainproject.eu/en/open-ethical-engaged/contributors/partners
[253] https://www.humanbrainproject.eu/en/hbp-platforms/getting-access/
[254] European Commission (2020)
[255] https://www.humanbrainproject.eu/en/social-ethical-reflective/dual-use/
[256] https://www.humanbrainproject.eu/en/open-ethical-engaged/contributors/partners/
[257] http://www.humanbrainproject.eu/en/brain-simulation
[258] http://www.humanbrainproject.eu/en/understanding-cognition

Intelligent Artifacts*

Main website: https://www.intelligent-artifacts.com

Intelligent Artifacts is a company founded by Sevak Avakians in 2010.[259] It was previously known as COGNITUUM.[260] It claims to have built "a true AGI that eliminates current machine intelligence shortcomings" which is referred to as Genie.[261]

Lead institution: Intelligent Artifacts

Partner institutions: none
- Intelligent Artifacts has partnered with Bombora, Sikorsky, and Trimedx, but none seem to focus on AGI.[262]

Type of institution: private corporation

Open-source: yes
- Some code is available on GitHub[263] and BitBucket,[264] but their AGI platform, requires an account for access.[265]

Military connection: unspecified

Lead country: USA[266]

Partner countries: none

Stated goals: humanitarian
- Intelligent Artifacts seeks to "evolve your business, humanity, and the world tomorrow. Plus, building a true AGI is just plain fun."[267]

Engagement on safety: moderate
- Avakians writes, "Some of the risks involved with AI/AGI are real. Others are pure fantasy."[268]

Size: small-medium

[259] https://www.linkedin.com/in/avakians/
[260] https://www.linkedin.com/in/avakians/
[261] https://www.intelligent-artifacts.com/
[262] https://www.intelligent-artifacts.com
[263] https://github.com/intelligent-artifacts
[264] https://bitbucket.org/intelligent-artifacts/geniesdk-python/src/master/
[265] https://bitbucket.org/intelligent-artifacts/geniesdk-python/src/master/
[266] https://www.linkedin.com/company/cognituum/about/
[267] https://www.intelligent-artifacts.com/about
[268] https://www.intelligent-artifacts.com/post/why-do-people-keep-trying-to-make-ai-better-despite-the-several-risks-it-may-have

Leabra

Main website: https://ccnlab.org

Leabra, which stands for Local, Error-driven, and Associative, Biologically Realistic Algorithm, is led by Randall O'Reilly of the University of Colorado Boulder. Leabra is a cognitive architecture project emphasizing modeling neural activity. Emergent was the original implementation of Leabra.[269] A 2016 paper states that Leabra is "a long-term effort to produce an internally consistent theory of the neural basis of human cognition." It also states, "More than perhaps any other proposed cognitive architecture, Leabra is based directly on the underlying biology of the brain, with a set of biologically realistic neural processing mechanisms at its core."[270]

Lead institution: University of California Davis
- Relocated from University of Colorado Boulder to the Center for Neuroscience at UC Davis. [271]

Partner institutions: none

Type of institution: academic

Open-source: yes[272]

Military connection: yes
- The Leabra group reports funding from the US Office of Naval Research and Army Research Lab. [273]

Lead country: USA

Partner countries: none

Stated goals: intellectualist
- The Leabra website states their research is in effort to "understand a wide range of different cognitive and behavioral phenomena."[274]

Engagement on safety: active (2017 survey: unspecified)
- A 2017 paper describes "some initial ideas to make neuromorphic AGI safer."[275]

Size: small-medium[276]

[269] https://en.wikipedia.org/wiki/Leabra
[270] O'Reilly et al. (2016)
[271] https://ccnlab.org
[272] https://github.com/emer/emergent
[273] https://ccnlab.org/funding/ and http://www.e-cortex.com/projects.html
[274] https://ccnlab.org/research/
[275] Jilk et al. (2017)
[276] https://ccnlab.org/people/

LIDA (Learning Intelligent Distribution Agent)

Main website: http://ccrg.cs.memphis.edu/projects.html

LIDA, which stands for Learning Intelligent Decision Agent (the project recently changed the acronym letter 'D' from Distribution to Decision),[277] is led by Stan Franklin of the Cognitive Computing Research Group (CCRG) at the University of Memphis. It is based on Bernard Baars's Global Workspace Theory and integrates "various forms of memory and intelligent processing in a single processing loop."[278]

Lead institution: University of Memphis

Partner institutions: none

Type of institution: academic

Open-source: restricted
- Registration is required to download code; commercial use requires a commercial license. [279]

Military connection: yes
- A past project called IDA was used by the Navy for "job distribution" and was funded by Office of Naval Research and Naval Personnel Research, Studies, & Technology. [280]

Lead country: USA

Partner countries: none

Stated goals: humanitarian, intellectualist
- The LIDA website says that it seeks "a full cognitive model of how minds work" and focuses predominantly on academic research aims.[281]
- A LIDA paper states that robots need ethics "to constrain them to actions beneficial to humans."[282]

Engagement on safety: active
- A LIDA paper addresses AGI safety challenges like the subtlety of defining human ethics with the precision needed for programming.[283]
- Franklin has collaborated with AI ethicists on getting AGIs to make correct moral decisions.[284]

Size: medium

[277] Kugele and Franklin (2020)
[278] Goertzel (2014), p.24
[279] http://ccrg.cs.memphis.edu/framework.html
[280] http://ccrg.cs.memphis.edu/projects.html
[281] http://ccrg.cs.memphis.edu
[282] Madl and Franklin (2015)
[283] Madl and Franklin (2015)
[284] Wallach et al. (2010)

M3-CLIC*

Main website: http://m3-ip.com

M3-CLIC (Cognitive Linguistic Intelligent Catalyst) is an R&D project founded by Vidur (Sonny) Nanda in 1995.[285] It attempts to "connect thoughts to information required by the user, based on natural (i.e., human) intelligence of language."[286] M3-CLIC claims to have a human-centric AGI prototype.

Lead institution: M3-IP Ltd.

Partner institutions: none

Type of institution: private corporation[287]

Open-source: no

Military connection: unspecified

Lead country: UK[288]

Partner countries: India[289]

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[285] https://uk.linkedin.com/in/vidur-sonny-nanda-4289651
[286] https://www.linkedin.com/company/m3-ip-ltd/about/
[287] https://www.linkedin.com/company/m3-ip-ltd/about/
[288] http://m3-ip.com
[289] https://uk.linkedin.com/in/vidur-sonny-nanda-4289651

MARAGI*

Main website: https://maragi.io

MARAGI, which stands for Microservices Architecture for Robotics and Artificial General Intelligence,[290] was created by Dave Shapiro.[291] MARAGI claims to have pluggable architecture and aims to facilitate collaboration in order to make AGI accessible to everyone.[292] The ultimate goal is to "democratize access to AGI" like "a giant App Store for AGI microservices."[293]

Lead institution: MARAGI

Partner institutions: none

Type of institution: none

Open-source: yes[294]

Military connection: unspecified

Lead country: unspecified
- Neither the website nor Shapiro's GitHub profile list a location.

Partner countries: none

Stated goals: intellectualist
- The MARAGI website states that it is designed "so that we can work together to develop a solid theory of intelligence."[295]

Engagement on safety: unspecified

Size: small

[290] https://maragi.io
[291] https://github.com/daveshap
[292] https://maragi.io
[293] See "Mission" section of https://maragi.io
[294] https://github.com/daveshap/maragi
[295] See "Facilitate Experimentation" section of https://maragi.io

Mauhn*

Main website: https://mauhn.com

Mauhn is a "human-level AI company with a focus on ethics and safety,"[296] founded by Berg Severens in 2018.[297]

Lead institution: Mauhn

Partner institutions: none

Type of institution: private corporation
- Mauhn is a capped profit company.[298]

Open-source: no

Military connection: unspecified

Lead country: Belgium[299]

Partner countries: none

Stated goals: unspecified
- Severens writes that he is inspired by the Effective Altruism community.[300]

Engagement on safety: moderate
- The Mauhn website motto states a focus on "ethics and safety."[301]

Size: small

[296] https://www.linkedin.com/company/maughn/about/
[297] https://www.linkedin.com/in/berg-severens-843a2884/
[298] https://mauhn.com
[299] See website footer
[300] https://www.linkedin.com/in/berg-severens-843a2884/
[301] https://mauhn.com

Microsoft Research AI (MSR AI)

Main website: https://www.microsoft.com/en-us/research/lab/microsoft-research-ai

MSR AI is an AI "research and incubation hub" at Microsoft announced in July 2017[302] which is led by Eric Horvitz.[303] The project seeks "to probe the foundational principles of intelligence, including efforts to unravel the mysteries of the human intellect, and use this knowledge to develop a more general, flexible artificial intelligence."[304] The project pulls together more than 100 researchers from different branches of AI at Microsoft's Redmond headquarters,[305] though much of the work is not focused specifically on AGI.

Lead institution: Microsoft

Partner institutions: none

- MSR AI regularly collaborates on projects with various institutions, but it is unclear which projects are specifically related to AGI.[306]

Type of institution: public corporation

Open-source: no

Military connection: unspecified

- Microsoft has a Military Affairs program,[307] but its link to MSR AI is unclear.

Lead country: USA

Partner countries: none

- Microsoft Research also has personnel in China and India, but it is unclear whether these researchers are working on projects related to AGI.

Stated goals: intellectualist (2017 survey: also humanitarian)

- The MSR AI website states that it aims "to solve some of the toughest challenges in AI" and "probe the foundational principles of intelligence."[308]

Engagement on safety: unspecified

- The MSR AI group on Aerial Informatics and Robotics has extensive attention to safety, but this is for the narrow context of aircraft, not for AGI.[309]

Size: medium-large

[302] https://blogs.microsoft.com/blog/2017/07/12/microsofts-role-intersection-ai-people-society
[303] https://www.microsoft.com/en-us/research/lab/microsoft-research-ai/people/?#
[304] https://www.microsoft.com/en-us/research/lab/microsoft-research-ai
[305] Etherington (2017)
[306] https://www.microsoft.com/en-us/research/lab/microsoft-research-redmond/groups/?#
[307] https://military.microsoft.com/about
[308] https://www.microsoft.com/en-us/research/lab/microsoft-research-ai
[309] https://www.microsoft.com/en-us/research/group/air

Mind Simulation*

Main website: https://mind-simulation.com/en/

Mind Simulation is "a research laboratory whose mission is to achieve Artificial General Intelligence."[310] It was founded in 2014 by Leonid Derikyants, Sergey Pankratov, and Vasily Mazin,[311] and its website was publicly launched in 2020.[312] Mind Simulation uses the Intellectual Core, a software package using a hierarchical multi-agent system.[313] They also have a technology called CyberMind which endows non-playable video game characters (NPC's) with more depth through the use of AI and eventually through AGI.[314]

Lead institution: Mind Simulation

Partner institutions: none

Type of institution: private corporation[315]

Open-source: no

Military connection: unspecified

Lead country: Russia
- The Mind Simulation website states, "Russia AI lab brings Geralt of Rivia to life in a way that could change video games forever."[316]

Partner countries: none

Stated goals: unspecified
- The Mind Simulation website states, "Our main goal is to achieve General Artificial Intelligence and maximize its application in various fields."

Engagement on safety: unspecified
- In a Mind Simulation presentation, they specify one of their main principles is "friendliness" and "anthropocentricity."[317]

Size: small-medium

[310] https://mind-simulation.com/en/
[311] https://mind-simulation.com/en; https://www.linkedin.com/in/leonid-derikyants-785414164
[312] https://web.archive.org/web/20200428013004/https://mind-simulation.com/en/
[313] https://mind-simulation.com/en/technology.html
[314] https://mind-simulation.com/en/gaming.html; https://mind-simulation.com/en/blog/gaming/agi-technology-and-making-videogame-characters-alive.html; https://www.vg247.com/2020/05/19/witcher-3-interview-russian-ai-geralt/
[315] https://www.crunchbase.com/organization/mind-simulation
[316] https://mind-simulation.com/en/
[317] https://mind-simulation.com/en/technology.html; see Slide 5 of the "Mind Simulation" slide deck

Mindtrace*

Main website: https://www.mindtrace.ai

Mindtrace, founded in 2017 by Hoon Chung,[318] builds "brain-inspired intelligent systems to maximally exploit the functionality and performance of edge devices in markets such as smart devices and inspection." Its goal is "to empower machines in such a way that they are capable of human levels of intelligence."[319] Its website was publicly launched in 2018.[320]

Lead institution: Mindtrace

Partner institutions: none

Type of institution: private corporation[321]

Open-source: no

Military connection: unspecified

Lead country: UK[322]

Partner countries: none

Stated goals: humanitarian

- The Mindtrace website states, "Our vision is a world in which machines possess levels of intelligence which allow them to efficiently perceive, understand & act on the world, in ways which make it a healthier, safer, more creative and better informed environment from which everyone can benefit."[323]

Engagement on safety: unspecified

Size: small-medium

[318] https://www.vbprofiles.com/companies/mindtrace-ai-5c701e44105eb57dc7c64662
[319] https://mindtrace.ai/technology.html
[320] https://web.archive.org/web/20180323013706/http://mindtrace.ai/
[321] https://www.vbprofiles.com/companies/mindtrace-ai-5c701e44105eb57dc7c64662
[322] https://www.mindtrace.ai/contact.html
[323] https://www.mindtrace.ai/index.html

Monad*

Main website: https://monad.ai

Monad is a company that was founded by Jovan Williams in 2013.[324] It seeks to build "the world's first Artificial General Intelligence (AGI) Computational System." [325] Monad's AGI project is referred to as "Consciousness Centric AI" on Reddit[326] and YCombinator.[327]

Lead institution: Monad

Partner institutions: none

Type of institution: private corporation[328]

Open-source: no

Military connection: unspecified

Lead country: USA[329]

Partner countries: none

Stated goals: unspecified

Engagement on safety: moderate

- Williams dismisses AGI safety as something that need not be thought about if AGI is properly created.[330] This implies that for an AGI to be properly created, it must be designed safely.

Size: small

[324] https://monad.ai/about/
[325] https://monad.ai
[326] https://www.reddit.com/r/singularity/comments/50oug5/monadai_consciousness_centric_framework_agi/
[327] https://news.ycombinator.com/item?id=12408493
[328] https://www.linkedin.com/company/monad.ai/about/
[329] https://www.welcome.ai/monad-ai
[330] https://www.quora.com/Can-a-sentient-AI-have-constraints-put-upon-it/answer/Jovan-Williams-2, https://www.quora.com/Who-should-be-in-charge-of-AI-ethics/answer/Jovan-Williams-2, https://www.quora.com/How-could-we-create-safe-artificial-general-intelligence/answer/Jovan-Williams-2, https://www.quora.com/What-can-we-do-to-regulate-artificial-intelligence-A-I/answer/Jovan-Williams-2, and https://www.quora.com/Is-Sam-Harris-right-to-be-worried-about-AI/answer/Jovan-Williams-2

NARS

Main website: https://cis.temple.edu/~pwang/NARS-Intro.html

NARS is an AGI research project led by Dr. Pei Wang of Temple University. NARS is an acronym for Non-Axiomatic Reasoning System, in reference to the AI being based on tentative experience and not axiomatic logic, consistent with its "assumption of insufficient knowledge and resources."[331] In a 2011 interview, Wang suggests that NARS may achieve human-level AI by 2021.[332] NARS has an open-source version called OpenNARS.[333]

Lead institution: Temple University

Partner institutions: none

Type of institution: academic

Open-source: yes[334]

Military connection: unspecified

Lead country: USA

Partner countries: none

Stated goals: intellectualist (2017 survey: also humanitarian)

- The NARS website states that it is a "project aimed at the building of a general-purpose intelligent system, i.e., a 'thinking machine' (also known as 'AGI') that follows the same principles as the human mind and can solve problems in various domains."[335]
- The OpenNARS website states, "The ultimate goal of this research is to build a thinking machine."[336]

Engagement on safety: active

- Wang has written on NARS's safety issues, such as "motivation management," a factor in NARS's ability to pursue its goals and not be out of control reliably.[337]

Size: small-medium[338] (2017 survey: medium)

[331] https://cis.temple.edu/~pwang/NARS-Intro.html
[332] Goertzel (2011)
[333] https://github.com/opennars
[334] https://github.com/opennars/opennars
[335] https://cis.temple.edu/~pwang/NARS-Intro.html
[336] http://opennars.org
[337] Wang (2012); See also Bieger et al. (2015)
[338] https://phillyagiteam.github.io/Website/

NDEYSS*

Main website: https://github.com/KarjamP/NDEYSS

NDEYSS (pronounced "Indies"), which stands for Network Designed to Eventually Yield Sentience and Sapience,[339] was created by KaramP.[340] It is "a series of projects designed to explore implementing vessels for conscious mind entirely within code."[341] One NDEYSS project works toward "the true means of mind uploading into a machine."[342]

Lead institution: NDEYSS

Partner institutions: none

Type of institution: none

Open-source: yes[343]

Military connection: unspecified

Lead country: unspecified

Partner countries: none

Stated goals: transhumanist

- The NDEYSS project description states, "Eventually, the project would give way to an entire race of beings either existing within a virtual space, or existing within the mechanical bodies of real life robots." It further states, "The project also extends to providing technology relating to providing brain-computer interfacing between the human mind and the computer, in addition to copying the human mind and, eventually, allowing for the true means of mind uploading into a machine, and allowing for one to achieve digital immortality."[344]

Engagement on safety: unspecified

Size: small

[339] https://github.com/KarjamP/NDEYSS
[340] https://github.com/KarjamP
[341] https://github.com/KarjamP/NDEYSS#readme
[342]https://github.com/KarjamP/NDEYSS
[343] https://github.com/KarjamP/NDEYSS
[344] https://github.com/KarjamP/NDEYSS#readme

New Sapience*

Main website: https://www.newsapience.com

New Sapience, founded by Bryant Cruse in either 2005[345] or 2014,[346] aims to turn "computers into thinking machines" by scaling a chatbot service.[347]

Lead institution: New Sapience

Partner institutions: none

Type of institution: private corporation[348]

Open-source: no

Military connection: unspecified

Lead country: USA[349]

Partner countries: none

Stated goals: humanitarian
- New Sapience's vision is for "a society in which the basic necessities of life are virtually or literally free."[350]
- The New Sapience Welcome AI listing also states an aim to develop AGI for use by businesses and consumers and also licensed to others. Although this seems to be profit oriented, it was found to be insufficient for New Sapience to be coded as such.[351]

Engagement on safety: active
- The New Sapience vision statement states that they "have thought deeply" about safety risks and they have selected a "fully deterministic" paradigm that is safer than other AGI approaches.[352]
- A 2018 New Sapience whitepaper expresses concern about "stochastic AI" and contrasts its approach as "reliable servants whose world view has been received directly from the hands of their creators."[353]

Size: small-medium

[345] https://www.linkedin.com/in/bryantcruse/
[346] https://www.newsapience.com/our-team/
[347] https://www.newsapience.com/what-we-do/
[348] https://www.newsapience.com/welcome/
[349] https://www.linkedin.com/company/new-sapience/about/h
[350] https://www.newsapience.com/our-vision/
[351] https://www.welcome.ai/new-sapience
[352] https://www.newsapience.com/our-vision/
[353] https://www.newsapience.com/wp-content/uploads/2020/06/New-Sapience-101.pdf

Nigel

Main website: http://kimera.ai

Nigel is the AGI project of Kimera, an AI corporation based in Oregon which was founded by Mounir Shita and Nicholas Gilman. Kimera's website states, "Nigel AGI is designed to bring Artificial GENERAL Intelligence (AGI) to the global networks and all the devices connected to it" (emphasis original). Furthermore, Kimera claims that Nigel will one day transform "the global economy and what it means to be human on planet earth."[354] In 2016, Kimera unveiled Nigel, claiming it is "the first commercially deployable artificial general intelligence technology."[355] Nigel has been described as a personal assistant bot similar to Apple's SIRI and Amazon's Alexa.[356] Additionally, Kimera has a cryptocurrency called NigelCoins, which are earned based on how much sensor data one creates.[357]

Lead institution: Kimera

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified

Lead country: USA

Partner countries: Australia, Austria, Germany, Italy, Norway, Russia[358] (2017 survey: none)

Stated goals: humanitarian

- Kimera presents a humanitarian vision for AGI, writing that "Artificial General Intelligence has the power to solve some - or all - of humanity's biggest problems, such as curing cancer or eliminating global poverty."[359]

Engagement on safety: unspecified

Size: small-medium

[354] https://kimera.ai
[355] http://kimera.ai
[356] Boyle (2016)
[357] https://kimera.ai/early-access-program/
[358] https://kimera.ai/about-us/
[359] http://kimera.ai/company

NiHa*

Main website: https://sites.google.com/view/imrcuilhr

Nature-inspired Humanoid (NiHa) is a project led by Dr. Wajahat Mahmood Qazi at the COMSATS Institute of Information Technology.[360] It seeks to build a "self-aware artificial general intelligence."[361]

Lead institution: COMSATS Institute of Information Technology

Partner institutions: none
- NiHa collaborates with a few institutions, but it is unclear which collaborations are related to AGI.[362]

Type of institution: academic

Open-source: no

Military connection: unspecified

Lead country: Pakistan

Partner countries: none

Stated goals: humanitarian, intellectualist
- The NiHA homepage lists multiple objectives, including "development of cognitive and conscious artifacts to improve quality of life," "establishment of science of consciousness and its role in the universe," and "redefining humanity."[363]
- Qazi says his goal with NiHa "is on applying 'applied artificial intelligence' to solve problems of Industry 4.0 and bioinformatics."[364]

Engagement on safety: unspecified

Size: small-medium

[360] https://sites.google.com/view/imrcuilhr/members?authuser=0
[361] https://www.thinkmind.org/index.php?view=article&articleid=cognitive_2018_5_30_40064
[362] https://sites.google.com/view/imrcuilhr/collaborations?authuser=0
[363] https://sites.google.com/view/imrcuilhr/home?authuser=0
[364] https://lahore.comsats.edu.pk/Employees/985

NNAISENSE

Main website: https://nnaisense.com

NNAISENSE is a private corporation based in Lugano, Switzerland. Several of its team members have ties to the Dalle Molle Institute for Artificial Intelligence (IDSIA, a Swiss nonprofit research institute), including co-founder and Chief Scientist Jürgen Schmidhuber. Its website states that it delivers "advanced neural network solutions that improve how products are made and how they work."[365] Co-founder Bas Steunebrink is listed as Director of Artificial General Intelligence,[366] and previously worked on the Animats project.[367]

Lead institution: NNAISENSE

Partner institutions: none

Type of institution: private corporation[368]

Open-source: yes[369] (2017 survey: no)

Military connection: unspecified

Lead country: Switzerland

Partner countries: USA[370] (2017 survey: none)

Stated goals: intellectualist (2017 survey: also profit)

- Schmidhuber is described as a "consummate academic" who founded the company to prevent other companies from poaching his research talent; NNAISENSE reportedly "chooses projects based on whether they'll benefit the machine's knowledge, not which will bring in the highest fees."[371]
- The NNAISENSE mission states that NNAISENSE, "was formed in 2014 in order to build large-scale neural network solutions for industrial process inspection, modeling, and control."[372] Additionally, the website footer says, "Let's discuss how we can automate your business."[373] These quotes insinuate heavy ties to profit-based goals, but insufficient evidence was identified.

Engagement on safety: moderate (2017 survey: unspecified)

- Stuenebrink seems moderately concerned with AGI safety and proposed a project to the Future of Life Institute on the safety of experience-based AI (EXPAI) in 2016.[374]

Size: medium (2017 survey: small-medium)

---

[365] https://nnaisense.com
[366] https://nnaisense.com/company/#team
[367] Strannegård et al. (2016)
[368] https://www.realwire.com/releases/NNAISENSE-Concludes-Successful-Series-B-Investment-Round
[369] https://github.com/nnaisense
[370] Co-founder Faustino Gomez lives in Austin, Texas: https://nnaisense.com/company/#team
[371] Webb (2017)
[372] https://nnaisense.com/company/#mission
[373] https://nnaisense.com
[374] https://futureoflife.org/ai-researcher-bas-steunebrink/?cn-reloaded=1

Olbrain*

Main website: https://olbrain.com

Olbrain is a company seeking to build AGI for robotics, including 3D visual representation, grasping, few-shot learning, real-time motion planning, and object affordances.[375] Alok Gautam, Nishant Singh, and Mayank Kumar founded Olbrain in 2017,[376] and its website was publicly launched in 2018.[377] They also have AGI-SR and AGI-IR (Space Robots and Industrial Robots, respectively),[378] and ROVIS (robotics vision system) to streamline robot training.[379] Olbrain states that it won the 2019 Animal-AI Olympics in the Advanced Preferences Category.[380]

Lead institution: Olbrain

Partner institutions: none

Type of institution: private corporation
- Olbrain was founded in early 2017 and incorporated in early 2019 as Delaware C Corp.[381]

Open-source: no

Military connection: unspecified

Lead country: India[382]

Partner countries: none

Stated goals: humanitarian
- The Olbrain LinkedIn page states, "The Future of Humanity will be shaped by Robots - our new Marco Polos" and that "The wave of Robotization will first make the lands of wider space fertile so that the Human Civilization of Tomorrow can grow."[383]

Engagement on safety: unspecified
- The Olbrain website states, "The core objective function of humans has been to ensure the survival of the species as a whole, for that we have developed many capabilities that are needed just to maintain a relationship with fellow humans. AGI is supposed to understand these relationships in order to perform well in the Human World."[384]

Size: small

[375] https://olbrain.com
[376] https://olbrain.com
[377] https://web.archive.org/web/20180107122355/http://www.olbrain.com/
[378] http://olbrain.com/agi-sr/
[379] https://olbrain.com
[380] https://olbrain.com
[381] https://olbrain.com
[382] https://tracxn.com/d/companies/olbrain.com
[383] https://www.linkedin.com/company/0lbrain/about/
[384] http://olbrain.com/agi-sr/

Omega*

Main website: https://github.com/celestial-intellect

Omega is an "open-ended, modular, self-improving Omega AI unification architecture" that "embodies several crucial principles of general intelligence."[385] It is the primary project of Celestial Intellect Cybernetics, a company founded by Eray Özkural in 2010.[386] Celestial Intellect Cybernetics focuses on the "long-term memory of a general-purpose AI system."[387]

Lead institution: Celestial Intellect Cybernetics

Partner institution: none

Type of institution: private corporation[388]

Open-source: yes[389]

Military connection: unspecified

Lead country: Turkey

Partner countries: none

Stated goals: unspecified
- Özkural is the co-founder of the "Scientific Transhumanism" Facebook group.[390]

Engagement on safety: dismissive
- In a recent podcast on AGI, Özkural is described in the following way: "Eray is extremely critical of Max Tegmark, Nick Bostrom and MIRI founder Elizier Yodokovsky and their views on AI safety. Eray thinks that these views represent a form of neoludditism and they are capturing valuable research budgets with doomsday fear-mongering and effectively want to prevent AI from being developed by those they don't agree with."[391]
- In a 2017 blog post on his personal website, Özkural criticizes AGI risk as "comical" and a "schizophrenic delusion."[392]

Size: small

[385] Özkural (2018)
[386] https://tr.linkedin.com/in/erayozkural
[387] https://tr.linkedin.com/in/erayozkural
[388] https://tr.linkedin.com/company/celestial-intellect
[389] https://github.com/celestial-intellect
[390] https://www.facebook.com/groups/scientific.transhumanism/about
[391] https://www.youtube.com/watch?v=pZsHZDA9TJU
[392] https://examachine.net/blog/simulation-argument-and-existential-ai-risk

OpenAI

Main website: https://openai.com

OpenAI is a nonprofit AI research organization founded by several prominent technology investors who have pledged $1 billion to the project. The website tagline says, "Discovering and enacting the path to safe artificial general intelligence."[393] It is part of the Partnership on AI to Benefit People & Society.[394] In 2019, Microsoft invested $1 billion in OpenAI explicitly to develop AGI by eventually scaling Microsoft Azure capabilities.[395]

Lead institution: OpenAI

Partner institution: DeepMind,[396] Microsoft[397]

Type of institution: nonprofit, private corporation (2017 survey: nonprofit only)
- In 2019, OpenAI established a capped-profit company called OpenAI LP to secure additional funding.[398] OpenAI LP is owned by OpenAI Nonprofit.

Open-source: yes[399]

Military connection: unspecified

Lead country: USA

Partner countries: none

Stated goals: humanitarian
- OpenAI seeks that to build AGI that "leads to a good outcome for humans,"[400] and that "AGI's benefits are as widely and evenly distributed as possible."[401]

Engagement on safety: active
- Safe AGI is part of OpenAI's mission. While it releases much of its work openly, its website states, "We will attempt to directly build safe and beneficial AGI."[402]
- OpenAI also collaborates with DeepMind on long-term AI safety.[403]

Size: large

[393] https://openai.com
[394] https://www.partnershiponai.org/partners
[395] https://openai.com/blog/microsoft/
[396] https://blog.openai.com/deep-reinforcement-learning-from-human-preferences
[397] https://openai.com/blog/microsoft/
[398] https://openai.com/blog/openai-lp/
[399] https://github.com/openai
[400] https://openai.com/jobs
[401] https://openai.com/about
[402] https://openai.com/about/
[403] https://blog.openai.com/deep-reinforcement-learning-from-human-preferences

OpenCog

Main website: http://wiki.opencog.org/w/CogPrime_Overview

OpenCog is an open-source project led by Ben Goertzel. The vision of The Open Cognition Project is "to create an open-source framework for Artificial General Intelligence, intended to one day express general intelligence at the human level and beyond."[404] A revised version of OpenCog called OpenCog Hyperon is in early stages of development.[405] OpenCog was listed as CogPrime in the 2017 survey, but the name has since deprecated.

Lead institution: OpenCog Foundation

Partner institutions: none

Type of institution: nonprofit[406]

Open-source: yes[407]

Military connection: unspecified
- OpenCog has expressed concerns about military AGI.[408]
- Recently, Ben Goertzel tweeted his concern about military AGI.[409]

Lead country: USA
- While OpenCog is highly international, its website refers to an application for 501(c)(3) nonprofit status,[410] implying an administrative home in the US.

Partner countries: China, Ethiopia[411]

Stated goals: unspecified
- Elsewhere, Goertzel has advocated cosmism.[412]

Engagement on safety: active
- OpenCog states that AGI "should be able to reliably achieve a much higher level of commonsensically ethical behavior than any human being," adding that their "explorations in the detailed design of OpenCog's goal system have done nothing to degrade this belief."[413]

Size: medium (2017 survey: medium-large)

[404] https://wiki.opencog.org/w/The_Open_Cognition_Project
[405] https://wiki.opencog.org/w/Hyperon
[406] https://opencog.org/about/
[407] https://github.com/opencog
[408] Vepstas (2014)
[409] https://twitter.com/bengoertzel/status/1229709919512756225
[410] http://opencog.org/about
[411] https://wiki.opencog.org/w/Vision
[412] Goertzel (2010)
[413] http://wiki.opencog.org/w/CogPrime_Overview

Optimizing Mind*

Main website: https://www.optimizingmind.com

Optimizing Mind is a project led by Tsvi Achler[414] and was founded in 2017.[415] Optimizing Mind aims to bridge the connection between neuroscience and computer science to create AI that "thinks like a human."[416]

Lead institution: Optimizing Mind

Partner institution: none

Type of institution: private corporation[417]

Open-source: yes[418]

Military connection: unspecified

Lead country: USA[419]

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[414] https://www.linkedin.com/in/tsvi-achler-8419664/
[415] https://www.linkedin.com/company/optimizingmind/about/
[416] https://www.optimizingmind.com
[417] https://www.linkedin.com/company/optimizingmind/about/
[418] https://github.com/Optimizing-Mind and https://github.com/OptimizingMind
[419] https://www.linkedin.com/company/optimizingmind/

ORBAI*

Main website: https://www.orbai.ai

ORBAI was founded by Brent Oster in 2018, [420] and is "developing Artificial General Intelligence that is more like what we expect AI to be – more like us."[421] They have started by creating "approachable" narrow AI in various fields,[422] but intend to reach true AGI by 2025.[423] ORBAI filed a patent in 2019 for NeuroCAD, the tool they intend to use to create AGI.[424]

Lead institution: ORBAI

Partner institutions: none

Type of institution: private corporation[425]

Open-source: no

Military connection: unspecified

Lead country: USA[426]

Partner countries: none

Stated goals: humanitarian

- ORBAI's mission webpage states, "An AGI can use its vast knowledge and memory to make fair and unbiased decisions, to help us, to guide us, to bring us justice, fairness, prosperity, and hope. This is the only way, after all these millennia of us failing to do this in our civilization, otherwise this will be the last millennia of our civilization."[427]

Engagement on safety: unspecified

Size: small

[420] https://www.linkedin.com/in/brentosterorbai/
[421] https://www.crunchbase.com/organization/orbai
[422] Cole (2020); https://www.linkedin.com/in/justine-falcon-75a163196/?fbclid=IwAR2Mih9jW8JFIyYrfzdEkp_gIUP3e6IAJHxl8MFOw5F5iw073pxX7gqjqx8
[423] https://www.orbai.ai/mission.htm
[424] https://www.linkedin.com/pulse/orbai-files-patent-strong-agi-brent-oster/
[425] https://www.linkedin.com/company/orbai/about/
[426] https://www.linkedin.com/company/orbai/about/
[427] https://www.orbai.ai/mission.htm

Prome*

Main website: http://www.prome.ai

Prome is a self-funded company founded in 2011 by Sean Everett which aims to develop AGI software.[428] They describe their AGI as "Biologic Intelligence," which "emulates an animal's brain and nervous system into software and robotics." [429] Prome also offers consulting for companies involved in any sort of AI development.[430]

Lead institution: Prome

Partner institutions: none

Type of institution: private corporation[431]

Open-source: no

Military connection: unspecified

Lead country: USA[432]

Partner countries: none

Stated goals: humanitarian
- In a blog post about Prome, Everett writes, "we want to push humanity forward."[433]

Engagement on safety: unspecified

Size: small

[428] http://www.prome.ai/about.html and https://www.linkedin.com/in/seanmeverett/
[429] http://www.prome.ai
[430] http://www.prome.ai/consulting.html
[431] https://www.crunchbase.com/organization/prome
[432] https://twitter.com/interintel; https://www.miaeverett.com/contact
[433] Everett (2017)

Research Center for Brain-Inspired Intelligence (RCBII)

Main website: http://bii.ia.ac.cn

RCBII is a "long term strategic scientific program proposed by the Institute of Automation, Chinese Academy of Sciences."[434] The group is based in Beijing.[435] RCBII does research in cognitive brain modeling, brain-inspired information processing, and neuro-robotics.[436] It states that "Brain-inspired Intelligence is the grand challenge for achieving Human-level Artificial Intelligence."[437]

Lead institution: Chinese Academy of Sciences

Partner institutions: none

- RCBII often publishes research papers with other institutions, but no official partnerships could be identified.[438]

Type of institution: government

- The Chinese Academy of Sciences is a public institution under the Chinese government[439]

Open-source: no

Military connection: yes (2017 survey: unspecified)

- Researchers from RCBII recently co-wrote a paper with the Academy of Military Medical Sciences.[440]

Lead country: China

Partner countries: none

Stated goals: intellectualist

- The RCBII website only lists intellectual motivations, stating, "The long-term goal of CASIA Brain Simulation effort is to decode the mechanisms and principles of human intelligence and develop Brain-inspired intelligent systems."[441]

Engagement on safety: unspecified

Size: medium (2017 survey: small-medium)

[434] http://bii.ia.ac.cn/about.htm
[435] http://english.ia.cas.cn/au/fu/
[436] http://bii.ia.ac.cn/about.htm
[437] http://bii.ia.ac.cn/about.htm
[438] http://bii.ia.ac.cn/publication.htm
[439] http://english.cas.cn/about_us/introduction/201501/t20150114_135284.shtml
[440] Kong et al. (2018)
[441] http://bii.ia.ac.cn

Robot Brain Project

Main website: https://github.com/brohrer/robot-brain-project and https://e2eml.school/robot_brain_project.html

Fomerly known as Becca, the Robot Brain Project is led by Brandon Rohrer, who is currently at iRobot.[442] According to its GitHub, the Robot Brain Project "is a general learning program for use in any robot or embodied system"; it "aspires to be a brain for any robot, doing anything."[443] Though the GitHub has not been recently updated, the Robot Brain Project seems to be more active on Twitter.[444]

Lead institutions: Robot Brain Project

Partner institutions: none

Type of institution: none

Open-source: yes[445]

Military connection: unspecified

- The project began while Rohrer was at Sandia National Laboratories,[446] but this connection appears to be inactive.

Lead country: USA[447]

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[442] https://www.linkedin.com/in/brohrer
[443] https://github.com/brohrer/robot-brain-project
[444] https://twitter.com/_brohrer_becca
[445] https://github.com/brohrer/robot-brain-project
[446] https://www.linkedin.com/in/brohrer
[447] https://github.com/brohrer/

Sanctuary AI*

Main website: https://www.sanctuary.ai

Sanctuary AI was founded in 2018 by Suzanne Gildert and Geordie Rose. Its goal is "to build synthetic humans – "synths" – that are indistinguishable from us physically, cognitively and emotionally."[448] Its website states that it is "building and scaling embodied artificial general intelligence."[449]

Lead institutions: Sanctuary AI

Partner institutions: none

Type of institution: private corporation[450]

Open-source: no

Military connection: unspecified

Lead country: Canada[451]

Partner countries: none

Stated goals: humanitarian, intellectualist

- The Sanctuary AI LinkedIn page states that its goal is "to create general purpose robots that can autonomously perform most economically valuable work."[452]
- In an interview, Gildert said, "the Sanctuary AI approach will offer a platform of sorts to advance research in areas such as sequence learning, computer vision, cognitive architectures, and sensor fusion."

Engagement on safety: dismissive

- In a video interview, Gildert says, "some people are worried about AIs becoming superintelligent, but I actually believe it's better to have more intelligence in the world"; she continues, "I don't see [AGI] as a threat to humanity, but a mirror that's held up to our civilization that allows us to ask deep questions about ourselves."[453]

Size: medium

[448] https://www.crunchbase.com/organization/sanctuary-ai and https://www.linkedin.com/in/geordie-rose-a53abb/
[449] https://www.sanctuary.ai
[450] https://www.crunchbase.com/organization/sanctuary-ai
[451] https://www.crunchbase.com/organization/sanctuary-ai
[452] https://www.linkedin.com/company/sanctuaryai/
[453] https://www.youtube.com/watch?v=l4iAJORJGdc (~3:40)

Sigma

Main website: http://cogarch.ict.usc.edu

Sigma, led by Paul Rosenbloom, focuses on "embodying, and integrating together, the range of capabilities necessary for human(-like) intelligence."[454] It has a publication record dating to 2009 and won awards at the 2011 and 2012 AGI conferences.[455] Rosenbloom was previously a co-PI of Soar.[456]

Lead institution: University of Southern California

Partner institution: California State Northridge University[457]

Type of institution: academic

Open-source: yes[458]

Military connection: yes
- Funding is reported from the Us Army,[459] the US Air Force Office of Scientific Research, and the US Office of Naval Research.[460]

Lead country: USA

Partner countries: none

Stated goals: intellectualist, transhumanist (2017 survey: intellectualist only)
- The Sigma aims "to develop a sufficiently efficient, functionally elegant, generically cognitive, grand unified, cognitive architecture in support of virtual humans (and hopefully intelligent agents/robots – and even a new form of unified theory of human cognition – as well)."[461]
- Rosenbloom also hints at transhumanist views in a 2013 interview, stating, "I see no real long-term choice but to define, and take, the ethical high ground, even if it opens up the possibility that we are eventually superseded – or blended out of pure existence – in some essential manner."[462]

Engagement on safety: unspecified
- In a 2013 interview, Rosenbloom hints at being dismissive, questioning "whether superhuman general intelligence is even possible," but also explores some consequences if it is possible, all while noting his lack of "any particular expertise" on the matter.[463]

Size: small-medium (2017 survey: medium)

---

[454] https://cogarch.ict.usc.edu/research/
[455] https://sites.usc.edu/rosenbloom/recent-publications/
[456] https://sites.usc.edu/rosenbloom/bio-past-research/
[457] Ustun et al. (2018)
[458] https://bitbucket.org/sigma-development/sigma-release/wiki/Home
[459] Rosenbloom and Ustun (2019)
[460] Rosenbloom (2013)
[461] http://cogarch.ict.usc.edu
[462] https://intelligence.org/2013/09/25/paul-rosenbloom-interview
[463] https://intelligence.org/2013/09/25/paul-rosenbloom-interview

SingularityNET

Main website: https://singularitynet.io

SingularityNET is an AGI project led by Ben Goertzel which was publicly launched in 2017.[464] It aims to bring AI and blockchain together to create a decentralized open market for AIs and improve their interoperability[465] and "ultimately generate coordinated artificial general intelligence."[466]

Lead institution: SingularityNET Foundation

Partner institutions: OpenCog Foundation, Hanson Robotics, Novamente, Vulpem[467]

Type of institution: nonprofit,[468] private corporation (2017 survey: nonprofit only)
- SingularityNET has a private corporation branch called Singularity Studio.[469]

Open-source: yes[470]

Military connection: unspecified
- Though decentralized, this open marketplace could end up being a dual-use issue and is open to use by anyone with a SingularityNET account, including military personnel.

Lead country: Netherlands[471] (2017 survey: China; 2020 recoding of 2017 data: unspecified)[472]

Partner countries: Brazil, China, India, Russia, South Korea, USA[473] (2017 survey: also Australia, Canada, Germany, and Portugal)

Stated goals: animal welfare, ecocentric, humanitarian, transhumanist
- SingularityNET is described as "for the People (and the Robots!)" and "the happiness of sentient beings," with "benefits for all people, and for all life."[474]
- SingularityNET describes profit as a means to other goals, seeking "to direct the profit thus generated to apply AI for global good."[475]

Engagement on safety: unspecified

Size: medium (2017 survey: small-medium)

[464] Blog posts at https://blog.singularitynet.io date to October 2017.
[465] https://singularitynet.io/aboutus/
[466] Goertzel et al. (2017); See also https://public.singularitynet.io/whitepaper.pdf
[467] Goertzel et al. (2017)
[468] Goertzel et al. (2017) p.8
[469] https://blog.singularitynet.io/singularitynet-announces-a-for-profit-spin-off-singularity-studio-f650a82d7455
[470] https://github.com/singnet
[471] See the website footer; See also https://www.linkedin.com/company/singularitynet/about/
[472] The information available in 2017 is coded as China per the 2017 methodology and unspecified per the 2020 methodology. Additional information has appeared since 2017 indicating Netherlands as the lead country.
[473] SingularityNET personnel include Cassio Pennachin, Brazil, http://www.pennachin.com; Raam Baranidharan, India, https://www.linkedin.com/in/raam-baranidharan-ba78215; Alexey Potapov, Russia, Potapov et al. (2016); Youngsook Park, Korea, https://www.linkedin.com/in/youngsook-park-14237144; Matt Ikle, USA, https://www.linkedin.com/in/matthewikle
[474] Goertzel (2017b); Goertzel et al. (2017).
[475] Goertzel (2017b)

Soar

Main website: http://soar.eecs.umich.edu and https://soartech.com

Soar was founded in 1981 and is led by John Laird of the University of Michigan. [476] Soar is an acronym for State, Operator Apply Result, and it is "a general cognitive architecture for developing systems that exhibit intelligent behavior."[477] A spinoff corporation, called SoarTech, is also based in Michigan.

Lead institution: University of Michigan, SoarTech

Partner institutions: Bar Ilan University, Cogniteam (a private corporation), Pace University, Pennsylvania State University, University of Portland, University of Portsmouth, University of Southern California[478]

Type of institution: academic, private corporation

Open-source: yes[479]

Military connection: yes

- SoarTech lists customers including research laboratories of the US Air Force, Army, Navy, DARPA, and the US Department of Transportation.[480]

Lead country: USA

Partner countries: Israel, UK

Stated goals: intellectualist

- The Soar website describes it as an investigation into "an approximation of complete rationality" aimed at having "all of the primitive capabilities necessary to realize the complete suite of cognitive capabilities used by humans."[481]

Engagement on safety: unspecified

Size: medium (2017 survey: medium-large)

[476] http://ai.eecs.umich.edu/people/laird
[477] https://soar.eecs.umich.edu
[478] https://soar.eecs.umich.edu/groups
[479] https://github.com/SoarGroup, https://soar.eecs.umich.edu/Downloads
[480] http://soartech.com/about
[481] http://soar.eecs.umich.edu

Susaro

Main website: http://www.susaro.com

Susaro, short for Surfing Samuri Robots Inc., is an AI corporation based in New York[482] led by Richard Loosemore.[483] Its website states that its goal is to "build *artificial general intelligence* systems" (emphasis original) using an approach that "is a radical departure from conventional AI."[484]

Lead institution: Susaro

Partner institutions: none

Type of institution: private corporation

Open-source: no
- Susaro has a GitHub page, but it does not contain any code.[485]

Military connection: unspecified

Lead country: USA (2017 survey: UK)

Partner countries: UK[486] (2017 survey: none)

Stated goals: ecocentric, humanitarian
- The Susaro website states that it aims to advance "human and planetary welfare… without making humans redundant."[487]

Engagement on safety: active
- The Susaro website states that "the systems we build will have an unprecedented degree of safety built into them… making it virtually impossible for them to become unfriendly."[488]
- Loosemore has also written, "This entire class of [AGI] doomsday scenarios is found to be logically incoherent at such a fundamental level that they can be dismissed as extremely implausible."[489]

Size: small[490]

[482] https://opengovus.com/sam-entity/831207811
[483] https://www.seamless.ai/c/richard-loosemore-Ab7JOtY5bKEn4 and https://www.linkedin.com/in/rloosemore/
[484] http://www.susaro.com
[485] https://github.com/susaroltd
[486] https://github.com/susaroltd
[487] http://www.susaro.com
[488] http://www.susaro.com
[489] Loosemore (2014)
[490] https://www.linkedin.com/search/results/all/?keywords=Susaro%2C%20Ltd

Tencent AI Lab (TAIL)

Main website: http://ai.tencent.com/ailab

TAIL is the AI group of Tencent, the Shenzhen-based Chinese technology company. Its website lists several research areas, including games, social networking, and platform-based tools AI, though it does not explicitly state research interest in AGI.[491] In 2019, Tencent's AI agent *Juewu* beat the top multiplayer online battle arena players in the game *Wangzhe Rongyao (*roughly translated as "Honor of Kings" or "King of Glory"), and Tencent reported that this AI agent would ultimately advance them toward AGI development.[492]

Lead institution: Tencent

Partner institutions: none
- TAIL lists many additional partner institutions, but none explicitly working on AGI.[493]

Type of institution: public corporation

Open-source: no
- Tencent releases some work open-source,[494] but not its AGI.
- Its website states that "Tencent will open-source its AI solutions in the areas of image, voice, security to its partners through Tencent Cloud," but it does not state that its AGI research is open-source.[495]

Military connection: unspecified

Lead country: China

Partner countries: USA
- TAIL maintains an office in Seattle.[496]

Stated goals: unspecified

Engagement on safety: unspecified

Size: small-medium

[491] https://ai.tencent.com/ailab/en/about
[492] Yuan (2019)
[493] https://ai.tencent.com/ailab/en/together/detial2?num=0 and https://ai.tencent.com/ailab/en/together/detial2?num=1
[494] https://github.com/Tencent
[495] http://ai.tencent.com/ailab
[496] Mannes (2017)

True Brain Computing*

Main website: http://truebraincomputing.com/en/truebraincomputing/

True Brain Computing is a brain-inspired AGI project founded by Alexey Redozubov in 2018.[497] They have created a new model explaining human brain functioning called context-semantic.[498] Their goal is "to develop a *strong* artificial intelligence using custom approach to neuromorphic computing that is *not* based on conventional neural networks or deep learning" (emphasis original).[499]

Lead institutions: True Brain Computing

Partner institutions: Institute of the Human Brain of the Russian Academy of Science[500]

Type of institution: private corporation[501]
Team member Svyatoslav V. Medvedev is a member of the Russian Academy of Sciences.[502]

Open-source: no

Military connection: unspecified

Lead country: Russia

Partner countries: none

Stated goals: humanitarian, intellectualist

- True Brain Computing states goals including "new methods of treatment of brain diseases" and to "explain the phenomenon of consciousness."[503]

Engagement on safety: unspecified

Size: small

[497] https://www.linkedin.com/company/truebraincomputing/about/ and http://truebraincomputing.com/en/team/
[498] http://truebraincomputing.com/en/truebraincomputing/
[499] http://truebraincomputing.com/en/truebraincomputing/
[500] https://www.linkedin.com/company/truebraincomputing/
[501] https://www.linkedin.com/company/truebraincomputing/about/
[502] https://neurologycongress.com/scientific-committee/member/svyatoslav-medvedev and http://truebraincomputing.com/en/team/
[503] https://www.linkedin.com/company/truebraincomputing/about/

Uber AI Labs

Main website: https://www.uber.com/info/ailabs

Uber AI Labs is the AI research division of Uber. Uber AI Labs began in 2016 with the acquisition of Geometric Intelligence, [504] a private company founded in 2014 by Gary Marcus, Kenneth Stanley, and Zoubin Ghahramani in 2014 with incubation support from NYU.[505] Geometric Intelligence was based on Marcus's ideas for AGI, especially how to "learn with less training data" than deep learning.[506] Uber AI Labs was reportedly part of Uber's attempt to expand beyond the private taxi market, similar to how Amazon expanded beyond books.[507] Due to the COVID-19 pandemic, Uber recently fired 3,000 employees and shut down Uber AI Labs.[508] Ghahramani has since joined Google Brain.[509]

Lead institution: Uber

Partner institutions: none

Type of institution: public corporation[510] (2017 survey: private corporation)

Open-source: no
- Uber AI Labs has a GitHub profile, but this does not appear to include its AGI.[511]

Military connection: unspecified
- Uber AI Labs does not seem to have any military contracts, but Uber has partnered with the US Army Research Lab to work on a flying taxi service named UberAIR.[512]

Lead country: USA

Partner countries: Canada (2017 survey: UK)
- Uber AI Labs had a residency program available in Canada.[513]
- Uber has offices around the globe, but Uber AI Labs's presence in them could not be determined.

Stated goals: unspecified (2017 survey: humanitarian)

Engagement on safety: unspecified

Size: medium[514]

---

[504] Temperton (2016)
[505] https://www.nyu.edu/about/news-publications/news/2016/december/nyu-incubated-start-up-geometric-intelligence-acquired-by-uber.html
[506] Chen (2017)
[507] Metz (2016)
[508] BBC News (2020) and Sagar (2020)
[509] Yuan (2020)
[510] Hawkins (2019)
[511] https://github.com/uber
[512] Miller (2018)
[513] https://www.uber.com/info/ailabs
[514] https://www.uber.com/blog/ai-meet-the-team-jingchen-liu/

Vicarious

Main website: https://www.vicarious.com

Vicarious is a privately held AI corporation founded in 2010 by Scott Phoenix and Dileep George and based in San Francisco. It has raised tens of millions of dollars in investments from several prominent investors.[515] It states that it "aims to bring about a robotic golden age by using AI to automate more and more general tasks until we reach artificial general intelligence."[516] In an interview, Phoenix says that Vicarious is working towards AGI, with "plenty of value created in the interim,"[517] and that AGI would be "virtually the last invention humankind will ever make."[518]

Lead institution: Vicarious

Partner institutions: none

Type of institution: private corporation

Open-source: yes[519]

Military connection: unspecified

Lead country: USA

Partner countries: none

Stated goals: humanitarian

- Vicarious is a Flexible Purpose Corporation, reportedly so that it can "pursue the maximization of social benefit as opposed to profit."[520] Scott Phoenix says that Vicarious aims to build AI "to help humanity thrive."[521]
- Its principles state, "developing these technologies can help us solve many of the world's largest problems."[522]

Engagement on safety: active (2017 survey: moderate)

- The Vicarious website discusses safety at length and states that it will "maintain high standards of industrial safety" and will "address many safety issues as we develop our products." [523]

Size: medium-large

[515] Cutler (2014); High (2016); https://www.vicarious.com/company/
[516] https://www.vicarious.com/posts/principles/
[517] High (2016)
[518] TWiStartups (2016)
[519] https://github.com/vicariousinc
[520] High (2016)
[521] High (2016)
[522] https://www.vicarious.com/posts/principles/
[523] https://www.vicarious.com/posts/principles/

Whole Brain Architecture Initiative (WBAI)

Main website: https://wba-initiative.org

WBAI is a nonprofit organization based in Tokyo and led by Hiroshi Yamakawa. Yamakawa is the Director of AI at Dwango and also affiliated with Tamagawa University and the Japanese Society for Artificial Intelligence. WBAI's mission "aims to support and promote research and development activities to realize artificial general intelligence (AGI) with human-like intellectual capabilities while learning from the architecture of the entire brain." They "aim for the construction of artificial general intelligence (AGI) to surpass the human brain capability around the year 2030."[524] WBAI receives support from, among others, Panasonic, Toshiba, and Toyota.[525]

Lead institution: Whole Brain Architecture Initiative

Partner institutions: Cerenaut Research[526]

Type of institution: nonprofit

Open-source: yes[527]

Military connection: unspecified

Lead country: Japan

Partner countries: Australia (2017 survey: none)

Stated goals: humanitarian

- WBAI promotes AI development that is "best for the human society."[528] Additionally, in a slideshow about WBAI, they quote Yamakawa as stating that "the grace and wealth that EcSIA [ecosystem of shared intelligent agents] affords needs to be properly distributed to everyone."[529]

Engagement on safety: active

- Safety is a significant theme for WBAI. For example, their website states, "we could say it is a relatively safe choice to build the first AGI in a form similar to us."[530] The implication here is that WBAI seeks to build brain-like AGI in part because that would be safer.

Size: small-medium

[524] https://wba-initiative.org/en/wba
[525] https://wba-initiative.org/en/supporting-members
[526] https://cerenaut.ai/requests-for-research/
[527] https://github.com/wbap
[528] https://wba-initiative.org/en/about/vision
[529] https://www.slideshare.net/HiroshiYamakawa/2017-0512gatsby-wsv10b-75941913 (slide 7)
[530] https://wba-initiative.org/en/2071

WILLIAM*

Main website: https://occam.com.ua

WILLIAM is the project of the Odessa Competence Center for Artificial Intelligence and Machine Learning (OCCAM). OCCAM was founded by Dr. Arthur Franz and Michael Löffler in 2017.[531] It states that its "fundamental research is focused on so-called Artificial General Intelligence."[532] Its mission is "building thinking machines" by "trying to make universal induction tractable by building efficient data compression algorithms."[533] Franz and Löffler won the Kurzweil Prize at the 2019 AGI Conference.[534]

Lead institution: OCCAM

Partner institution: Odessa National Mechnikov University[535]

Type of institution: nonprofit[536]

Open-source: no

Military connection: unspecified

Lead country: Ukraine[537]

Partner countries: none

Stated goals: humanitarian, intellectualist

- The WILLIAM website states that its mission is to "advance both fundamental research and practical application in the fields of artificial intelligence and machine learning."[538]
- The WILLIAM website states that it pursues AGI "Because we really need help in order to end suffering, finally and forever."[539]
- The WILLIAM website states, "we believe that profit aspirations in this particular challenge are harmful to our goal," because pursuit of short-term profits impedes long-term AGI development.[540]

Engagement on safety: unspecified

Size: small

[531] https://lifeboat.com/ex/bios.arthur.franz and https://occam.com.ua
[532]https://occam.com.ua/
[533] https://occam.com.ua
[534] https://www.facebook.com/Occam.com.ua/posts/1471022709726971
[535] https://occam.com.ua/app/uploads/2018/08/ic_proceedings.pdf
[536] https://occam.com.ua
[537] https://www.facebook.com/Occam.com.ua/
[538] https://occam.com.ua
[539] https://occam.com.ua/why-we-do-it/
[540] https://occam.com.ua/why-we-do-it/

Xephor Solutions*

Main website: http://xephor-solutions.com/en/

Xephor Solutions was founded by Isabell Kunst in 2012[541] with the mission "to create the most sophisticated Artificial General Intelligence in the world and to help our customers to become innovation leaders in their field."[542] The actual AGI seems to be called "Xephor Solution,"[543] with several different products for different applications - Xephor Finance, Xephor Marketing, Xephor Healthcare, Xephor Security, and Xephor Disposition.[544] Xephor Solutions is partially funded by the European Union.[545]

Lead institution: Xephor Solutions

Partner institutions: none

Type of institution: private corporation

Open-source: no

Military connection: unspecified

Lead country: Austria[546]

Partner countries: none

Stated goals: profit

- Xephor Solutions's website states, "Our mission is to create the most sophisticated Artificial General Intelligence in the world and to help our customers to become innovation leaders in their field." [547]
- Additionally, Xephor Solutions's LinkedIn states, "Our mission is to make artificial general intelligence technology accessible to businesses worldwide."[548]

Engagement on safety: unspecified

Size: small

[541] https://www.crunchbase.com/organization/xephor-solutions and https://www.linkedin.com/company/xephor-solutions-agi/
[542] http://xephor-solutions.com/en/company/
[543] http://xephor-solutions.com/en/product/
[544] https://www.crunchbase.com/organization/xephor-solutions
[545] See website footer
[546] http://xephor-solutions.com/en/contact/
[547] http://xephor-solutions.com/en/
[548] https://www.linkedin.com/company/xephor-solutions-agi/

Appendix 2. Inactive AGI R&D Projects

The following pages document five AGI R&D projects that were active in 2017, inactive in 2020, and for which there is new information that affects how the project is coded.

The entries below follow the same format as Appendix 1, listing the project website, a brief summary of the project, and details the project's attributes.

As with Appendix 1, projects with an asterisk (*) next to their name indicate that the project is new to the 2020 survey. Only one such project is listed below (BasicAI).

The other four projects listed below (AIDEUS, Alice in Wonderland, Icarus, and Maluuba) appear in the 2017 survey and are recoded below due to changes in coding methodology between the 2017 and 2020 surveys. As with Appendix 1, red font indicates a data point that has been changed due to a change in coding methodology between the 2017 and 2020 surveys. For example, the 2020 survey uses a more restrictive standard for coding partner countries. As a result, AIDEUS is coded in the 2020 survey as having no partner countries, whereas in the 2017 survey it was coded as having a partner country of France. Full methodology changes are detailed in Section 4.

Appendix 3 presents summary information about AGI R&D projects that were identified in the 2017 survey and for which there is no new coding information. For full coding of these projects, see the 2017 survey.

Referenced websites were active when project data was collected and coded during June to September 2020. Some websites may have since become inactive. Many of these websites can be viewed via the Internet Archive (https://archive.org).

AIDEUS

Main website: http://aideus.com

AIDEUS was led by Alexey Potapov of ITMO University in Saint Petersburg and Sergey Rodionov of Aix Marseille Université. Potapov is an advisor to SingularityNET.[549] The AIDEUS website states their goal as the "creation of a strong artificial intelligence."[550] Its approach is to "proceed from universal prediction models on the basis of algorithmic probability used for choosing optimal actions."[551]

Lead institutions: AIDEUS

Partner institutions: none

Type of institution: none

Open-source: no

Military connection: unspecified
- Funding sources include "Government of Russian Federation, Grant 074-U01," which does not appear to be military, but this could not be confirmed.

Lead country: Russia

Partner countries: none (2017 survey: France)

Stated goals: humanitarian, intellectualist
- The project aims to build superintelligence in order to "help us better understand our own thinking and to solve difficult scientific, technical, social and economic problems."[552]

Engagement on safety: active
- AIDEUS has published AGI safety research, e.g. Potapov and Rodionov (2014).

Size: small

[549] https://singularitynet.io
[550] http://aideus.com
[551] http://aideus.com/research/research.html
[552] http://aideus.com/community/community.html

Alice In Wonderland (AIW)

Main website: https://github.com/arnizamani/aiw and https://www.gu.se/om-universitetet/hitta-person/claesstrannegard

Alice in Wonderland (AIW) was led by Claes Strannegård of Chalmers University of Technology in Sweden. A paper about AIW in the *Journal of Artificial General Intelligence* describes it as being similar to NARS.[553] A separate paper describes it as a prototype for implementing new ideas about "bridging the gap between symbolic and sub-symbolic reasoning."[554] The most recent GitHub commit was in 2015.[555] Strannegård's website notes that his research focuses on Animats, an animal-based AGI project (found in Appendix 1).[556]

Lead institutions: Chalmers University of Technology

Partner institutions: none

Type of institution: academic

Open-source: yes[557]

Military connection: unspecified (2017 survey: no)

Lead country: Sweden

Partner countries: none

Stated goals: unspecified

Engagement on safety: unspecified

Size: small

[553] Strannegård et al. (2016b)
[554] Strannegård et al. (2016a)
[555] https://github.com/arnizamani/aiw
[556] https://www.chalmers.se/en/staff/Pages/claes-strannegard.aspx
[557] https://github.com/arnizamani/aiw

BasicAI*

Main Website: https://web.archive.org/web/20170327193351/http://www.basicai.org/index.html

BasicAI was the project of Sean Markan, a former researcher at MIT CSAIL.[558] BasicAI emphasizes the need for a long-term strategy in order to successfully develop human-level AI, and rejects more near-term strategies of other projects.[559]

Lead institution: BasicAI

Partner institutions: none

Type of institution: none

Open-source: no

Military connection: unspecified

Lead country: unspecified

Partner countries: none

Stated goals: humanitarian
- The BasicAI homepage stated, "our long-term mission is to develop HLAI and ensure that it is beneficial for everyone."[560]

Engagement on safety: moderate
- In a blog post Markan writes, "an advanced AI could (in a theoretical sense) do a lot of damage. But the 'extreme danger' narrative has gone too far…it may be physically possible to build a very dangerous AI. But nobody wants to do that, and—in my view—it looks quite avoidable. (Especially if we build a transparent, understandable system.)."[561]

Size: small

[558] https://web.archive.org/web/20170401063907/http://www.basicai.org/about.html
[559] https://web.archive.org/web/20170327193351/http://www.basicai.org/index.html and http://markan.net/hlai.html
[560] https://web.archive.org/web/20170327193351/http://www.basicai.org/index.html
[561] https://web.archive.org/web/20170327193254/http://www.basicai.org/blog/effective-altruism-ai.html

Icarus

Main website: https://web.archive.org/web/20190127212722/http://cll.stanford.edu/research/ongoing/icarus

Icarus was led by Pat Langley of Stanford University. Icarus was a cognitive architecture project similar to ACT-R and Soar, emphasizing perception and action in physical environments.[562] Their research was funded by DARPA IPTO, the Office of Naval Research, and the National Science Foundation. Support for earlier work came from the Air Force Office of Scientific Research, NASA Ames Research Center, and DaimlerChrysler Research and Technology.[563]

Lead institution: Stanford University

Partner institutions: University of Kansas

Type of institution: academic

Open-source: no

Military connection: yes
- Funding was reported from the US Office of Naval Research, the US Navy Research Lab, and DARPA.[564]

Lead country: USA

Partner countries: none

Stated goals: intellectualist
- Choi and Langley (2017) write that "our main goal" is "achieving broad coverage of cognition functions" in "the construction of intelligent agents."

Engagement on safety: unspecified

Size: small-medium (2017 survey: small)

[562] Goertzel (2014); Choi and Langley (20178)
[563] https://web.archive.org/web/20181028051348/http://cll.stanford.edu:80/research/ongoing/icarus/
[564] Choi and Langley (2017)

Maluuba

Main website: https://web.archive.org/web/20170201195341/http://www.maluuba.com/

Maluuba was a Montréal-based AI company acquired by Microsoft in 2017. The Maluuba website stated, "our vision has been to solve artificial general intelligence by creating literate machines that could think, reason and communicate like humans."[565] They have since dissolved and become part of MSR AI's Montréal AI team.[566]

Lead institution: Microsoft

Partner institutions: none

Type of institution: public corporation

Open-source: yes[567]

Military connection: unspecified

Lead country: Canada

Partner countries: USA

Stated goals: intellectualist (2017 survey: intellectualist and profit)

- Maluuba writes that they aim "to solve fundamental problems in language understanding, with the vision of creating a truly literate machine."[568]
- Microsoft is a founding partner of the Partnership on AI to Benefit People & Society, which has humanitarian goals,[569] but this does not appear to have transferred to Maluuba's goals.

Engagement on safety: moderate

- Maluuba researcher Harm van Seijen writes that "I think such discussions [about AI safety] are good, although we should be cautious of fear mongering."[570] Microsoft is also a founding partner of the Partnership on AI to Benefit People & Society, which expresses concern about AI safety.[571] No direct safety activity by Maluuba was identified.

Size: medium

[565] https://web.archive.org/web/20170126001505/http://www.maluuba.com/blog/2017/1/13/maluuba-microsoft
[566] https://www.microsoft.com/en-us/research/lab/microsoft-research-montreal/
[567] https://github.com/Maluuba
[568] https://web.archive.org/web/20170126001505/http://www.maluuba.com/blog/2017/1/13/maluuba-microsoft
[569] https://www.partnershiponai.org/tenets
[570] Townsend (2016)
[571] https://www.partnershiponai.org/tenets

Appendix 3. Other Notable Projects

The following projects include inactive AGI R&D projects or other similar projects that are nonetheless notable for their relation to AGI R&D.

The marking * denotes projects new to the 2020 survey.

The marking † denotes projects that were identified as active AGI R&D projects in the 2017 survey, are inactive in 2020, and for which there is no new coding information. For full coding of these projects, see the 2017 survey. Projects with new coding information are presented in Appendix 2.

4CAPS

Main website: http://www.ccbi.cmu.edu/projects_4caps.html

4CAPS is led by psychologist Marcel Just of Carnegie Mellon University's Center for Brain Imaging (CCBI). The Center has been applying machine learning to fMRI brain imaging data, "making it possible for the first time to relate patterns of brain activity to specific thoughts."[572] 4CAPS is "a hybrid of a computational neuroscience model and a symbolic AI system."[573] It "can account for both traditional behavioral data and, more interestingly, the results of neuroimaging studies."[574] The project reports funding by both the Office of Naval Research and the Multidisciplinary Research Program of the University Research Initiative.[575]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Not R&D; not AGI

4D/RCS (Real-time Control Systems Architecture)

Main website: https://www.nist.gov/el/intelligent-systems-division-73500/rcs-real-time-control-systems-architecture

4D/RCS was led by James Albus at the US National Institute of Standards and Technology. It consists of "hard-wired architecture and algorithms… augmented by learning."[576] It provided a theoretical foundation for designing, engineering, integrating, and testing intelligent systems software for crewless vehicle systems.

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive; not explicitly AGI

[572] http://www.ccbi.cmu.edu/index.html
[573] (Goertzel 2014, p.19)
[574] http://www.ccbi.cmu.edu/projects_4caps.html
[575] http://www.ccbi.cmu.edu/projects_4caps.html
[576] (Goertzel 2014, p.24)

Achler

Main website: None

This unnamed project by Tsvi Achler of Los Alamos National Labs used neural networks in "a novel approach to bridging the symbolic-subsymbolic gap."[577]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

AGINAO

Main website: http://aginao.com

AGINAO was a project of Wojciech Skaba in Poland. It was a privately sponsored project designed to apply machine learning to a Nao robot and deploy in a pre-school environment to develop AGI.[578] It was active during between 2011 and 2013,[579] and has showed no recent activity.

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

Alibaba

Main website: https://damo.alibaba.com/labs/ai

Alibaba has a series of AI Labs named the DAMO (discovery, adventure, momentum, and outlook) Academy.[580] Within the AI Labs, five distinct labs are working on narrow AI Projects: City Brain Lab,[581] Speech Lab,[582] Vision Lab,[583] Language Technology Lab,[584] and Decision Intelligence Lab.[585]

Reason for consideration: Alibaba is a major computing technology company

Reason for exclusion: No indications of AGI projects were found

[577] Goertzel (2014, p.17)
[578] http://aginao.com/
[579] For example, Skaba (2012; 2012b); http://aginao.com/page2.php
[580] https://damo.alibaba.com/labs/ai
[581] https://damo.alibaba.com/labs/city-brain/
[582] https://damo.alibaba.com/labs/speech/
[583] https://damo.alibaba.com/labs/vision/
[584] https://damo.alibaba.com/labs/language-technology/
[585] https://damo.alibaba.com/labs/decision-intelligence/

Amazon

Main website: https://aws.amazon.com/amazon-ai

Amazon has an AI group within its Amazon Web Services (AWS) division, but it does not appear to work on AGI. Amazon has donated AWS resources to OpenAI.[586]

Reason for consideration: Amazon is a major computing technology company

Reason for exclusion: No indications of AGI projects were found

Apple

Main website: https://www.apple.com/jobs/uk/teams/machine-learning-and-ai.html

Apple has an AI group that does not appear to work on AGI. However, Apple has a reputation for secrecy and has a minimal website.[587] Apple is said to have less capable AI than companies like Google and Microsoft because Apple has stricter privacy rules, denying itself the data used to train AI.[588] Likewise, at least some AI research may be oriented towards learning from limited data or synthetic data.[589] Its recent AI company acquisitions are for narrow AI.[590] While it may be possible that Apple is working on AGI, no indications of this were found. Their machine learning journal does not indicate AGI.[591]

Reason for consideration: Apple is a major computing technology company

Reason for exclusion: No indications of AGI projects were found

Araya*

Main website: https://www.araya.org/en/

Araya was founded by Ryota Kanai in 2013.[592] According to LinkedIn, their primary goal is "in the development of artificial consciousness, strong AI technologies grounded on computational theories of consciousness, with combinations of neuroscience and information science."[593] The website states, "we are conducting research and development in anticipation of an era when robots have Artificial Consciousness and begin high-level interactions with humans."[594]

Reason for consideration: R&D related to AGI

[586] https://blog.openai.com/infrastructure-for-deep-learning
[587] https://www.geekwire.com/2020/exclusive-apple-acquires-xnor-ai-edge-ai-spin-paul-allens-ai2-price-200m-range/
[588] Vanian (2017a)
[589] Vanian (2017b)
[590] Tamturk (2017) and Sun (2020_
[591] https://machinelearning.apple.com
[592] https://www.araya.org/en/features/
[593] https://www.linkedin.com/company/araya-ai/about/
[594] https://www.araya.org/en/features/

Reason for exclusion: Not enough information to establish a focus on AGI

Artificial Brain Laboratory

Main website: None

The Artificial Brain Laboratory (ABL) was led by Hugo de Garis at Xiamen University. The project appears to have ended upon de Garis's firing around 2010.[595] Xiamen University now has a Brain-like Intelligent Robotic Systems Group,[596] but this is not necessarily related to the ABL.

Reason for consideration: Included in 2017 Survey

Reason for exclusion: Apparently inactive

Automatski*

Main website: http://automatski.com/index.html

A "Fundamental research" company with the mission of solving the world's most challenging problems, one of which is AGI.[597] Their website's structure has outlined scientific problems they aim to tackle within the decade, century, and millennium. AGI falls under the millennium category. No indications were found of Automaski having an active AGI project.[598]

Reason for consideration: Explicit interest in future AGI development

Reason for exclusion: Inactive

Automorph*

Main website: https://www.automorph.com

Automorph is an organization founded by Andreas H.[599] that is developing "(AGI): Computer programs that are potentially capable of understanding, learning, and doing anything."[600] They claim that their programs will be open-source and trustworthy in order to make life more enjoyable. According to the website copyright, the website has not been updated since 2017, and its Twitter has not been active since 2018.[601] According to Andreas H.'s LinkedIn, Automorph was active briefly from April to September 2017.[602]

[595] https://profhugodegaris.wordpress.com/punishing-the-mob-money-over-brains-psychology-of-females/
[596] http://information.xmu.edu.cn/en/?mod=departments&id=31
[597] https://www.crunchbase.com/organization/automatski-fundamental-research
[598] http://automatski.com/artificial-general-intelligence.html
[599] https://www.linkedin.com/in/andreas-h-2a7802142/
[600] https://www.automorph.com
[601] https://twitter.com/AutomorphHQ
[602] https://www.linkedin.com/in/andreas-h-2a7802142/

Reason for consideration: Explicitly AGI focused

Reason for exclusion: Apparently inactive

Brain Imaging and Modeling Section (BIMS)

Main website: https://www.nidcd.nih.gov/research/labs/brain-imaging-and-modeling-section

BIMS is a research project led by Barry Horwitz of the US National Institutes of Health. The project combines brain imaging with computer modeling to advance basic neuroscience and treatment of brain disorders. There have been no updates since Horowitz retired in 2017.[603]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Inactive; focused on basic neuroscience, not the development of an AGI

BRAIN Initiative

Main website: https://www.braininitiative.nih.gov

BRAIN Initiative is a research project aimed at understanding the human brain. BRAIN is an acronym for Brain Research through Advancing Innovative Neurotechnologies. The project is based at the US National Institutes of Health and partners with several other US government agencies and private organizations.[604] Its website states that "by accelerating the development and application of innovative technologies, researchers will be able to produce a revolutionary new dynamic picture of the brain that, for the first time, shows how individual cells and complex neural circuits interact in both time and space."[605]

Reason for consideration: A large-scale brain research project similar to Blue Brain

Reason for exclusion: Focused on basic neuroscience, not the development of an AGI

Brain/MINDS

Main website: http://brainminds.jp/en

Brain/MINDS is a neuroscience research project. Brain/MINDS is an acronym for Brain Mapping by Integrated Neurotechnologies for Disease Studies. The project is sponsored by the Japanese Ministry of Education, Culture, Sports, Science, and Technology (MEXT). It focuses on studying non-human primate brains, neural networks of brain disorders, and improving cooperation between basic and clinical neuroscience.[606]

[603] https://www.nidcd.nih.gov/news/2017/retirement-symposium-barry-horwitz-september-8
[604] https://www.braininitiative.nih.gov/about/index.htm
[605] https://www.braininitiative.nih.gov
[606] http://brainminds.jp/en/overview/greeting

Reason for consideration: A large-scale brain research project similar to Blue Brain

Reason for exclusion: Focused on basic neuroscience, not the development of an AGI

C-BRIC*

Main website: https://engineering.purdue.edu/C-BRIC

The Center for Brain-inspired Computing Enabling Autonomous Intelligence is a five-year project supported by $27 million in funding from the Semiconductor Research Corp. (SRC). The mission of C-BRIC is to "deliver key advances in cognitive computing, to enable a new generation of autonomous intelligent systems such as self-flying drones and interactive personal robots."[607] It includes 10 US universities and funding by DARPA.

Reason for consideration: Largely funded university research group focusing on AGI related topics

Reason for exclusion: Not explicitly AGI

Carboncopies

Main website: https://www.carboncopies.org

Carboncopies is a nonprofit based in San Francisco that "provides support to scientists in fields related to Whole brain emulation."[608]

Reason for consideration: A research project focused on related technical details of AGI

Reason for exclusion: Not explicitly AGI; not R&D

CERA-CRANIUM

Main website: https://www.conscious-robots.com

CERA-CRANIUM was a cognitive architecture project sometimes discussed in the context of AGI.[609] It was led by Raúl Arrabales of the University of Madrid. It was used for computer games, winning a competition in 2010.[610] Though Arrabales is still actively publishing, the most recent work on CERA-CRANIUM is from 2013.[611]

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

[607] https://engineering.purdue.edu/C-BRIC
[608] https://www.carboncopies.org/mission
[609] For example, Ng et al. (2017)
[610] Arrabales and Muñoz (2010)
[611] Arrabales et al. (2013)

CHAI (Center for Human-Compatible AI)

Main website: http://humancompatible.ai

CHAI is a research group based at the University of California, Berkeley. Its website states that its goal is "to develop the conceptual and technical wherewithal to reorient the general thrust of AI research towards provably beneficial systems." This is primarily in the context of "machines that are more capable than humans across a wide range of objectives and environments," which it sees as likely to exist eventually.[612]

Reason for consideration: A research project focused on technical details of AGI

Reason for exclusion: Focused on safety aspects of AGI, not on the development of an AGI

CHREST

Main website: http://chrest.info

CHREST is led by Fernand Gobet of the University of Liverpool. Gobet started CHREST in 1992 and traces it to the 1959 EPAM system.[613] CHREST is an acronym for Chunk Hierarchy and REtrieval STructures. It is "a cognitive architecture that models human perception, learning, memory, and problem solving."[614] A paper on CHREST describes its strengths in categorization and understanding as complementary to other projects (e.g., ACT-R, Soar) in problem-solving.[615] Its website lists no publications since 2017.[616]

Reason for consideration: An AGI project

Reason for exclusion: Apparently inactive

Cognitive Computing Project

Main website: http://research.ibm.com/cognitive-computing and
https://www.research.ibm.com/cognitive-computing/neurosynaptic-chips.shtml

CCP is part of a suite of IBM AI projects, which also includes the famed Watson system. Goertzel (2014) discusses a project led by Dharmendra Modha to build computer hardware and software systems modeled after the human brain. The project has produced a new programming language and a new computer chip called TrueNorth, which Modha postulates as a "turning point in the history of computing."[617] The chip was introduced in a 2014 article in *Science*.[618] The chip development was supported by the DARPA SyNAPSE program aimed at making "low-power electronic neuromorphic

[612] http://humancompatible.ai/about
[613] http://www.chrest.info/history.html
[614] http://www.chrest.info
[615] Lane and Gobet (2012)
[616] http://chrest.info/publications.html
[617] http://www.research.ibm.com/articles/brain-chip.shtml
[618] Merolla et al. (2014)

computers that scale to biological levels."[619] Part of IBM's recent CCP work includes extensive research on neuromorphic computing[620] and neurosymbolic AI[621] with the hope of giving machines the ability to perform "more general tasks."[622]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Project focused on hardware development related to AGI, not AGI itself

Cognitive Systems Toolkit (CST)

Main website: http://cst.fee.unicamp.br

CST is a project led by Ricardo Gudwin of the University of Campinas in Brazil. It is "a Java-based toolkit to allow the construction of Cognitive Architectures."[623]

Reason for consideration: A project related to technical aspects of AGI

Reason for exclusion: Focused on tools that could be used to develop AGI, not on building an AGI

Comirit

Main website: http://www.comirit.com

Comirit (Commonsense Intelligence and Reasoning through Integrative Technologies) was a project of Benjamin Johnston used for his Ph.D. at the University of Technology Sydney.[624] It aimed to build "robotic and software systems with commonsense intelligence" with a short-term focus of "weeks or months, rather than decades." However, it is "inspired by the long term goals of creating systems that have deep human-like understanding of the real world," and thus pursued designs that "can be gradually evolved into more capable systems."[625] It was active mainly between 2010 and 2011.

Reason for consideration: An R&D project with AGI aspirations

Reason for exclusion: Apparently inactive

Covariant*

Main website: https://covariant.ai

[619] http://www.darpa.mil/program/systems-of-neuromorphic-adaptive-plastic-scalable-electronics
[620] https://www.ibm.com/blogs/research/2018/07/synaptic-architecture/
[621] https://mitibmwatsonailab.mit.edu/category/neuro-symbolic-ai/
[622] Moskovitch (2020)
[623] http://cst.fee.unicamp.br
[624] https://www.comirit.com/papers/dissertation.pdf
[625] http://www.comirit.com

Covariant was founded by Rocky Duan, Tianhao Zhang, Pieter Abbeel, and Peter Chen in 2017.[626] Its goal is "universal AI that allows robots to see, reason, and act on the world around them."[627] They have reported funding from Baidu.[628] In an interview, Chen said, "the company is on a quest to solve the hot research challenge of how do you build general AI for robotics."[629]

Reason for consideration: R&D related to AGI

Reason for exclusion: Not enough information to determine AGI research

Dav & SAIL

Main website: http://www.cse.msu.edu/~weng/research/LM.html

Dav & SAIL were projects of Juyang Weng of Michigan State University between 1998 and 2010. The two robots were designed to learn as human children do. The project aimed at achieving the "machine's human-level performance through autonomous development."[630] Since 2010, this project was built upon by the Where What Network (WWN). However, the latest iteration of this is in 2015, and other publications, as recent as 2020, do not pertain to Dav & SAIL.[631] The work received funding from the National Science Foundation and DARPA.[632]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

Decoupled Multimodal Learning*

Main website: https://github.com/Jakobovski/decoupled-multimodal-learning

Decoupled Multimodal Learning is "a decoupled, generative, unsupervised, multimodal neural architecture" created by Zohar Jackson.[633] Most of the updates on its GitHub page are from 2017, with at least one from 2018.[634]

Reason for consideration: GitHub page tagged as AGI

Reason for exclusion: Not clearly focused on AGI

[626] https://www.linkedin.com/in/rocky-duan-1a82662a/
[627] https://covariant.ai/about-us
[628] https://covariant.ai/about-us
[629] Cai (2020)
[630] http://www.cse.msu.edu/~weng/research/LM.html
[631] http://www.cse.msu.edu/~weng/research/LM.html
[632] Weng et al. (1999)
[633] https://www.linkedin.com/in/zohar-jackson-61362368
[634] https://github.com/Jakobovski/decoupled-multimodal-learning

DeSTIN (Deep SpatioTemporal Inference Network) †

Main website: http://wiki.opencog.org/w/DeSTIN
DeSTIN was initially developed by Itamar Arel and colleagues at the University of Tennessee. It is also being developed by the OpenCog open-source AI project, formerly known as CogPrime. DeSTIN used deep learning for pattern recognition. The OpenCog website states that OpenCog "has adopted this academic project to prepare it for open-source release."[635] Goertzel (2014, p.17) notes that DeSTIN "has been integrated into the CogPrime architecture… but is primarily being developed to serve as the center of its own AGI design."

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

DSO-CA †

Main website: none found

DSO-CA was a project of Gee Wah Ng when he was at DSO National Laboratories, Singapore's primary national defense research agency. It is "a top-level cognitive architecture that models the information processing in the human brain," with similarities to LIDA, OpenCog, and other AGI cognitive architectures.[636] Gee Wah Ng is now at Home Team Science and Technology Agency (HTX) and is still actively publishing on AGI.[637] This is Singapore's statutory board formed under the Ministry of Home Affairs to develop science and technology capabilities for Home Team operations, though there is no information that HTX is developing AGI.

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

DUAL

Main website: http://alexpetrov.com/proj/dual

DUAL was led by Boicho Kokinov at New Bulgarian University. It was active from around 1999 to 2005. It was based on Marvin Minsky's *Society of Mind*, in which minds are made from interacting sub-mind "agents." DUAL integrates symbolic and emergentist approaches and integrates declarative learning (learning about information one can readily talk about) and procedural learning (learning that is more habitual and harder to talk about).[638]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

[635] http://wiki.opencog.org/w/DeSTIN
[636] Ng et al. (2017)
[637] Ng and Leung (2020)
[638] http://alexpetrov.com/proj/dual

Reason for exclusion: Apparently inactive

Entropica

Main website: http://entropica.com

Entropica is an AGI project and private company led by Alexander Wissner-Gross, and the main website now links to Wissner-Gross's website.[639] It was included in a 2017 list of AGI projects.[640] However, no activity has been identified since 2013. In August 2017, Wissner-Gross did not include Entropica on a list of companies on his website[641] or his CV.[642] Entropica is based on ideas Wissner-Gross published in a research paper in the same year.[643] A video describing the project states that it is "broadly applicable to a variety of domains" and shows it functioning in several seemingly different domains.[644] Media coverage described it as a breakthrough to AGI and superintelligence,[645] but other AI researchers have been critical[646] and some observers suspect it to be a hoax.[647] Forbes describes his new company Gemedy as being focused on AGI,[648] but the Gemedy website does not indicate AGI.[649]

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

Einstein

Main website: https://www.salesforce.com/products/einstein

Einstein is a project of SalesForce that applies AI to its customer service business. Einstein grew out of the private company MetaMind, which SalesForce acquired in 2016.[650] Chief Scientist Richard Socher reportedly aspires to build AGI[651] and has vocalized interest in AGI.[652] However, no indications were found that Einstein is working on AGI.

Reason for consideration: An AI R&D project led by someone who seeks to build AGI

Reason for exclusion: No AGI R&D found

[639] https://www.alexwg.org
[640] http://2ai.org/landscape
[641] http://www.alexwg.org/companies
[642] http://www.alexwg.org/AWG-CV.pdf
[643] Wissner-Gross and Freer (2013)
[644] https://www.youtube.com/watch?v=cT8ZqChv8P0
[645] Dvorsky (2013)
[646] Marcus and Davis (2013)
[647] https://www.quora.com/How-can-we-prove-that-Entropica-is-a-hoax
[648] Schmelzer (2020)
[649] https://www.gemedy.com/
[650] Novet (2016)
[651] The Economist (2016)
[652] https://twitter.com/richardsocher/status/1084859265121083392?lang=en-gb

EPIC

Main website: None
Executive-Process Interactive Control (EPIC) was a cognitive architecture led by computer scientist David Kieras and David Meyer at the University of Michigan. Goertzel (2014, p.16) writes that "it has been connected to SOAR for problem solving, planning and learning."

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

EvoGrid

Main website: http://www.evogrid.org

EvoGrid was an open-source artificial life project initiated by Bruce Damer and was submitted as his Ph.D. thesis.[653] It sought to overcome the computing challenge of artificial life by accessing a distributed network of computer hardware similar to that used by projects like SETI@home. The project website shows no updates since around 2010, though Damer's work on artificial life continues.[654]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

GLAIR

Main website: none found

Grounded Layered Architecture with Integrated Reasoning (GLAIR) was a project of Stuart Shapiro of State University of New York at Buffalo. It was active between 1993 and 2013.[655] GLAIR aimed for "computational understanding and implementation of human-level intelligent behavior without necessarily being bound by the actual implementation of the human mind."[656]

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

GMU BICA

Main website: https://krasnow.gmu.edu/asl/bica-project/

[653] http://www.evogrid.org/index.php/Main_Page
[654] http://www.damer.com
[655] https://www.cse.buffalo.edu/~shapiro/Papers
[656] Shapiro and Bona (2010, p. 307)

GMU Bica was a project of Alexei Samsonovich of George Mason University, and it was active between 2006 and 2007. The project has moved to the Adaptive Systems Laboratory and is listed as a past project by the name of "Integrated Self-Aware Cognitive Architecture (DARPA BICA)."[657] Samsonovich seems to still be regularly publishing around the concept of an eBICA (emotional Biologically Inspired Cognitive Architecture) and seems to work in the field of "Artificial Emotional Intelligence."[658]

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

Goedel Machine

Main website: http://people.idsia.ch/~juergen/goedelmachine.html

Goedel Machine was a project of Jürgen Schmidhuber of the Dalle Molle Institute for Artificial Intelligence Research in Switzerland. The Goedel Machine proceeds by taking the action it proves to be best at each step in its activity, which requires infinite computing power.[659] Schmidhuber writes on his website that "since age 15 or so, the main goal of professor Jürgen Schmidhuber has been to build a self-improving Artificial Intelligence (AI) smarter than himself, then retire."[660]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

Google Brain

Main website: https://research.google.com/teams/brain

Google Brain is an AI research group at Google. Its researchers have collaborated with DeepMind on AGI research,[661] but its work is mainly focused on machine learning.[662]

Reason for consideration: A research group with links to AGI

Reason for exclusion: Not sufficiently focused on AGI

Hanson Robotics*

Main website: https://www.hansonrobotics.com

[657] https://krasnow.gmu.edu/asl/research/
[658] https://scholar.google.com/citations?hl=en&user=AIvKvXcAAAAJ&view_op=list_works&sortby=pubdate
[659] Goertzel (2014) p.25
[660] http://people.idsia.ch/~juergen
[661] Fernando et al. (2017) and Wiggers (2020)
[662] https://research.google.com/teams/brain/about.html

Sophia is the humanoid robot of Hanson Robotics, founded by David Hanson in 2013.[663] Hanson Robotics describes itself as "an AI and robotics company dedicated to creating socially intelligent machines that enrich the quality of our lives."[664] Hanson Robotics is a partner of Ben Goertzel and SingularityNET.[665] Sophia has her own Twitter.[666]

Reason for consideration: Connections to AGI research

Reason for exclusion: Not explicitly AGI focused

HUMANOBS

Main website: none

HUMANOBS (Humanoids that Learn Socio-communicative Skills by Observation) was a project for developing robots that "can learn social interaction," which is "a big step towards the ultimate goal of creating intelligence that is both self-sufficient and adaptable in a wide variety of environments." The EU Grant agreement for the project ran from 2009 to mid- 2012, but HUMANOBS continued to be active until 2014.[667]

Reason for consideration: A former AGI R&D project

Reason for exclusion: Apparently inactive

IM-CLEVER

Main website: none

IM-CLEVER (Intrinsically Motivated Cumulative Learning Versatile Robots) was a project to design robots that could learn independently and apply their knowledge across contexts. It has been inactive since around 2013, and as of 2020, their original website no longer works.[668] It was funded by an EU grant, which ran from 2009 to early-2013.[669] It was led by Gianluca Baldassarre of Istituto di Scienze e Tecnologie della Cognizione (Institute of Cognitive Sciences and Technologies) in Italy. It included collaborators from around Europe and one group in the United States.[670] It worked on a robotics platform "iCub" for robots that could learn new skills and apply them to diverse tasks.[671]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

[663] https://www.linkedin.com/company/hanson-robotics/about/
[664] https://www.hansonrobotics.com/about/
[665] https://www.hansonrobotics.com/research/
[666] https://twitter.com/RealSophiaRobot
[667] https://cordis.europa.eu/project/id/231453
[668] http://www.im-clever.eu
[669] https://cordis.europa.eu/project/id/231722
[670] https://ec.europa.eu/digital-single-market/en/blog/im-clever-towards-intelligent-humanoids
[671] https://icub.iit.it/about-us/icub-history

Intel

Main website: https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html

Intel does not currently indicate any research explicitly on AGI but is active in AI research.[672] It has acquired several AI companies.[673] In a 2017 white paper, they write, "while there isn't a commonly accepted definition for AI, Intel views it as a computerized system that performs tasks we normally associate with people. But in spite of the remarkable advances of computing power and sophisticated algorithms, there is still a long way to go before what is called General AI becomes a reality."[674] However, their Loihi Chip is a piece of neuromorphic (SNN) hardware which "supports dramatically accelerated learning in unstructured environments for systems that require autonomous operation and continuous learning." It has been compared to Tsinghua University's Tianjic Chip, which aims to advance AGI.

Reason for consideration: Hardware possibly used for AGI development

Reason for exclusion: Not explicitly AGI

IPSEL*

Main website: https://www.researchgate.net/project/Human-level-cognitive-architecture

Information Processing Systems with Emergent Logic (IPSEL) is a recent cognitive architecture project created by Bryan Fruchart and Beniot Leblanc in 2019.[675] Although it can be considered AGI, its purpose is to analyze the relationship between humans and future AI systems.[676]

Reason for consideration: Exploring AGI-related cognitive architectures

Reason for exclusion: Not R&D

Israel Brain Technologies (IBT)

Main website: http://israelbrain.org

IBT is a neuroscience research project. It is an Israeli nonprofit based in Ramat HaSharon. It receives funding from the Israeli government, philanthropists, and corporations.[677] Its mission is "to accelerate the development of innovative treatments and cures for brain disease."[678] The website's latest blog post

[672] https://www.intel.com/content/www/us/en/analytics/artificial-intelligence/overview.html and https://www.intel.ie/content/www/ie/en/analytics/artificial-intelligence/overview.html and https://newsroom.intel.com/press-kits/artificial-intelligence/#gs.hjl520
[673] Tamturk (2017)
[674] https://blogs.intel.com/policy/files/2017/10/Intel-Artificial-Intelligence-Public-Policy-White-Paper-2017.pdf
[675] Fruchart and Le Blanc (2019)
[676] Fruchart and Le Blanc (2020)
[677] http://israelbrain.org/donate
[678] http://israelbrain.org/about-us/mission

was in 2018, and its rolling header is advertising the BrainTech2019 conference (March 4-5 2019).[679] The latest tweet was on December 2019[680] and the latest Facebook posts were in June 2020,[681] though the posts are not on the work of IBT.

Reason for consideration: A large-scale brain research project similar to Blue Brain

Reason for exclusion: Focused on basic neuroscience, not the development of AGI

Kuipers Group

Main website: https://www.cs.utexas.edu/users/qr/robotics/bootstrap-learning.html

Benjamin Kuipers of the University of Texas led a group that developed robots that would learn a wide range of information about the world on their own from their own experiences. They were actively publishing between 1997 and 2007. Benjamin Kuipers now supervises the Intelligent Robotics Lab who are currently publishing papers including on the topics of "Bootstrap Learning of Foundational Representations" and "Ethics of AI and Robotics."[682]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

Large-Scale Model of Mammalian Thalamocortical Systems (LSMMTS)

Main website: https://www.izhikevich.org/publications/large-scale_model_of_human_brain.htm

LSMMTS was a project of Eugene Izhikevich and Gerald Edelman of The Neurosciences Institute. LSMMTS is notable for being a brain simulation "on a scale similar to that of the full human brain itself."[683] Eugene Izhikevich left the Neurosciences Institute in 2009 and is now CEO of Brain Corp, which is an AI company focusing on autonomous robots.[684]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

Machine Intelligence Research Institute (MIRI)

Main website: https://intelligence.org

[679] http://israelbrain.org/
[680] https://twitter.com/IsraelBrainTech
[681] https://www.facebook.com/israelbraintech?ref=hl
[682] https://web.eecs.umich.edu/~kuipers/research/
[683] Goertzel (2014) p.19
[684] http://www.scholarpedia.org/article/User:Eugene_M._Izhikevich

MIRI is an independent nonprofit research group focused on "foundational mathematical research to ensure smarter-than-human artificial intelligence has a positive impact."[685] It states that its mission is "to develop formal tools for the clean design and analysis of general-purpose AI systems, with the intent of making such systems safer and more reliable when they are developed."[686]

Reason for consideration: A research project focused on technical details of AGI

Reason for exclusion: Focused safety aspects of AGI, not on the development of an AGI

MicroPsi†

Main website: http://cognitive-ai.com

MicroPsi was led by Joscha Bach of the Harvard Program for Evolutionary Dynamics until 2019.[687] Bach's mission was reportedly "to build a model of the mind is the bedrock research in the creation of Strong AI, i.e., cognition on par with that of a human being."[688] Bach now works as the Vice President of Research at the AI Foundation.[689]

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

MLECOG†

Main website: none

MLECOG is a cognitive architecture project led by Janusz Starzyk of Ohio University. A paper on MLECOG describes it as similar to NARS and Soar.[690] MLECOG is an acronym for Motivated Learning Embodied Cognitive Architecture.

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive; not R&D

Neuralink*

Main website: https://www.neuralink.com

[685] https://intelligence.org
[686] https://intelligence.org/about
[687] https://www.linkedin.com/in/joschabach/
[688] http://bigthink.com/experts/joscha-bach
[689] https://www.aifoundation.com/about/
[690] Starzyk and Graham (2015)

Neuralink is an Elon Musk project "developing ultra-high bandwidth brain-machine interfaces to connect humans and computers."[691] The project aims to create neural implants in order to be connected to mobile devices at all times. There does not seem to be any connection to AGI.

Reason for consideration: Popular technology company

Reason for exclusion: Not AGI

Neurogrid

Main website: https://web.stanford.edu/group/brainsinsilicon/neurogrid.html

Neurogrid is computer hardware designed for running low-cost brain simulations. It is part of Stanford University's Brains in Silicone group, which is led by Kwabena Boahen.[692] A 2014 version of Neurogrid claims to be "9,000 times faster and using significantly less power than a typical PC," but still much less energy-efficient than the human brain.[693]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Focused on hardware development related to AGI, not AGI itself

NOMAD (Neurally Organized Mobile Adaptive Device)

Main website: none

NOMAD was a project of the Neurosciences Institute, a nonprofit research institute in California led by the late Nobel Laureate Gerald Edelman. The website listed in the 2017 survey redirects to The Neurosciences Institute with a new focus of "understanding how consciousness arises from the activity of the nervous system" and is likely inactive as well.[694]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

OSCAR

Main website: http://johnpollock.us/ftp/OSCAR-web-page/oscar.html

OSCAR was a project of the late John Pollock of the University of Arizona, active from around 1995 to 2005.[695] Pollock writes, "the 'grand problem' of AI has always been to build artificial agents of

[691] https://neuralink.com
[692] https://web.stanford.edu/group/brainsinsilicon/index.html
[693] http://news.stanford.edu/pr/2014/pr-neurogrid-boahen-engineering-042814.html, a press release discussing Benjamin et al. (2014)
[694] http://www.nsi.edu/~nomad
[695] http://johnpollock.us

human-like intelligence… capable of operating in environments of real-world complexity… OSCAR is a cognitive architecture for GIAs [generally intelligent agents], implemented in LISP."[696] Pollock described Oscar's main features as the ability to reason defeasibly about perception, change, and persistence, causation, probabilities, plan construction and evaluation, and decision.[697]

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

PAGI World

Main website: http://rair.cogsci.rpi.edu/projects/pagi-world

Psychometric Artificial General Intelligence (PAGI) World is a project led by John Licato of the University of South Florida and based at Rensselaer Polytechnic Institute, where Licato was a Ph.D. student. PAGI world is "a simulation environment written in Unity 2D, which allows AI and AGI researchers to test out their ideas."[698] The website hosted by the Rensselaer AI and Reasoning Lab (RAIR)[699] states that PAGI World is no longer maintained by the RAIR Lab and points to a site hosted by the Advancing Machine and Human Reasoning Lab (AMHR - though that Lab also states that it does not maintain PAGI).[700]

Reason for consideration: A project on an aspect of AGI R&D

Reason for exclusion: Inactive; focused on tools for evaluating AGI, not on developing an AGI

PolyScheme

Main website: https://dspace.mit.edu/handle/1721.1/8325

PolyScheme was developed for the Ph.D. thesis of Nicholas Cassimatis at the Massachusetts Institute of Technology. It "integrates multiple methods of representation, reasoning, and inference schemes for general problem solving."[701] The project is no longer active. Cassimatis now works at Dry.io.[702]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

[696] Pollock (2008)
[697] https://johnpollock.us/ftp/OSCAR-web-page/oscar.html
[698] http://rair.cogsci.rpi.edu/projects/pagi-world
[699] https://rair.cogsci.rpi.edu/past-projects/pagi-world/
[700] https://sites.google.com/view/amhr/research/pagi-world
[701] (Goertzel 2014, p.24)
[702] https://www.linkedin.com/in/nickcassimatis/

project-origin*

Main website: https://github.com/kourgeorge/project-origin

project-origin is an artificial life simulator for investigating noogenesis created by George Kour, a member of the Machine Learning Technologies Group at the IBM research lab and a Ph.D. candidate in the Sagol Department of Neurobiology at the Haifa University.[703] Though it does not explicitly state a goal or mission of AGI, it can be compared to Unscripted (Appendix 3).

Reason for consideration: AGI related

Reason for exclusion: Apparently inactive; not R&D

Quantum Artificial Intelligence Lab (QAIL)

Main website: https://research.googleblog.com/2013/05/launching-quantum-artificial.html

QAIL is a project between Google[704] and NASA[705] seeking to use quantum computing to advance AI. It is a "hub for assessing the potential of quantum computers to impact computational challenges faced by the agency in the decades to come." QAIL is listed as an AGI project by 2AI,[706] and some commentators propose that quantum computing could be important for developing AGI.[707] However, QAIL gives no indications of aiming for AGI.

Reason for consideration: Listed as an AGI project by 2AI

Reason for exclusion: No apparent AGI focus

Real AI†

Main website: http://realai.org

Real AI was a private company in Hong Kong led by Jonathan Yan, also on the EthicsNet team.[708] Its mission was "to ensure that humanity has a bright future with safe AGI."[709] It worked on strategy for safe AGI and technical research in deep learning, the latter on the premise that deep learning can scale up to AGI.[710]

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

[703] https://github.com/kourgeorge
[704] Neven (2013)
[705] https://ti.arc.nasa.gov/tech/dash/groups/quail/
[706] http://2ai.org/landscape
[707] For example, Wang (2014), DeAngelis (2014)
[708] https://www.ethicsnet.org/about
[709] https://web.archive.org/web/20170806002223/http://realai.org/about
[710] https://web.archive.org/web/20170806011538/http://realai.org/prosaic

Robust AI*

Main website: https://www.robust.ai

A project led by Gary Marcus, "building the world's first industrial-grade cognitive engine." Their mission is to make robots "smart, collaborative, robust, safe, flexible, and genuinely autonomous."

Reason for consideration: Referenced in Marcus and Davis (2019)

Reason for exclusion: Not explicitly AGI

SAL (Synthesis of ACT-R and Leabra)

Main website: None

SAL cognitive architecture was developed by David Jilk, Christian Lebiere, and colleagues at the eCortex Corporation of Boulder, Colorado, the University of Colorado, and Carnegie Mellon University.[711] The project produced a brief publishing record[712] and appears inactive since 2008. As its name implies, SAL is based on ACT-R (see dedicated ACT-R entry) and Leabra, a neural simulation.

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

Scene Based Reasoning (SBR)

Main website: http://agi-conf.org/2015/wp-content/uploads/2015/07/agi15_bergmann.pdf

SBR is an AGI R&D project by Frank Bergmann and Brian Fenton presented at the 2015 AGI conference, but apparently inactive since. It is described as "a cognitive architecture based on the notions of scene and plan. Scenes represent real-world 3D scenes as well as planner states."

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

Shruti

Main website: http://www1.icsi.berkeley.edu/~shastri/shruti

Shruti was a project of Lokendra Shastri of the University of California, Berkeley, and was active between 1996 and 2007. The project developed computational tools based on fast, reflex-like human

[711] Jilk et al. (2008)
[712] E.g., Jilk et al. (2008)

inference. It has been funded by, among others, the US National Science Foundation, Office of Naval Research, and Army Research Institute.[713]

Reason for consideration: Listed in the AGI review paper Goertzel (2014)

Reason for exclusion: Apparently inactive

SiMA†

Main website: http://sima.ict.tuwien.ac.at/description

SiMA was a project led by Dietmar Dietrich of Vienna University of Technology. SiMA is an acronym for the Simulation of the Mental Apparatus & Applications. The project aims "to develop a broad human-like intelligent system that is able to cope with complex and dynamic problems rather than with narrowly and well-defined domains."[714] It includes extensive attention to psychoanalysis, especially Freud and other German-language scholars. Dietrich started SiMA in 1999.[715]

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

SNePS (Semantic Network Processing System) †

Main website: http://www.cse.buffalo.edu/sneps

SNePS was led by Stuart Shapiro at State University of New York at Buffalo, with a publication record dating to 1969.[716] According to the SNePS website, its long-term goal is "to understand the nature of intelligent cognitive processes by developing and experimenting with computational cognitive agents that are able to use and understand natural language, reason, act, and solve problems in a wide variety of domains."[717]

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

SonyAI*

Main website: https://ai.sony

[713] http://www1.icsi.berkeley.edu/~shastri/shruti
[714] http://sima.ict.tuwien.ac.at/description
[715] Brandstätter et al. (2015, p.V)
[716] http://www.cse.buffalo.edu/sneps/Bibliography
[717] http://www.cse.buffalo.edu/sneps

SonyAI is the research portion of Sony. SonyAI hopes to use AI to "help unleash human imagination and creativity." Sony recently acquired an AI company called Cogitai.[718] No indications of AGI could be identified.[719]

Reason for consideration: Large technology company

Reason for exclusion: Not AGI; not R&D

Tianjic Chip*

Main website: https://www.cbicr.tsinghua.edu.cn/english/

A neuromorphic chip created by the Center for Brain-Inspired Computing Research at Tsinghua University. The chip can support both SNN and ANN, including convolutional neural networks (CNNs), multilayer perceptrons (MLPs), and recurrent neural networks (RNNs). The chip's goal is "to stimulate AGI development by paving the way to more generalized hardware platforms."[720] It has been compared with Intel's Loihi Chip.

Reason for consideration: Goal of AGI development

Reason for exclusion: Hardware focused

Twitter

Main website: https://cortex.twitter.com

Twitter has an AI group called Cortex, but no indications of AGI work were found. It also uses IBM Watson for its natural language processing to remove abusive messages.[721]

Reason for consideration: Twitter is a major computing technology company

Reason for exclusion: No indications of AGI projects were found

Unscripted*

Main website: https://github.com/gnitr/unscripted

Unscripted is a "naive proof-of-concept for a collaborative online virtual world to experiment with unscripted bots. Primary Objective: experiment with adaptive, artificial human-like agents (0 prior knowledge about world, no training data, general purpose, 100% algorithmically controlled); a sort of

[718] https://www.sony.net/SonyInfo/News/Press/201605/16-052E/
[719] https://ai.sony
[720] Pei et al. 2019
[721] https://www.bernardmarr.com/default.asp?contentID=1373

artificial general intelligence." Updates to its GitHub page are mostly dated 2017, with at least one dated 2018.[722]

Reason for consideration: AGI oriented

Reason for exclusion: A platform for studying AGI-like entities, not for building AGI

VICTOR†

Main website: http://2ai.org/victor

VICTOR was the main project of 2AI, a subsidiary of the private company Cifer Inc. 2AI is led by Timothy Barber and Mark Changizi.[723] The VICTOR website stated, "the VICTOR project will bring 'other-awareness' to AI" and refers to the project as the 'emotion chip'."[724]

Reason for consideration: Included in the 2017 AGI survey

Reason for exclusion: Apparently inactive

Xapagy

Main website: http://www.xapagy.com

Xapagy is a cognitive architecture designed "to perform narrative reasoning, that is to model/mimic the mental processes humans perform with respect to stories."[725] A 2014 paper on Xapagy won the "Kurzweil Best AGI Idea Prize 2014."[726]

Reason for consideration: A cognitive architecture presented at an AGI conference[727]

Reason for exclusion: Inactive; focus on narrative makes it ultimately narrow AI

Ymir

Main website: http://alumni.media.mit.edu/~kris/ymir.html

Ymir (an implementation of Gandalf)[728] was a project of Kristinn Thórisson and was active between 1999 and 2009.[729] It "was created with the goal of endowing artificial agents with human-like communicative and manipulation capabilities in the form of embodied multimodal task-oriented dialog

[722] https://github.com/gnitr/unscripted
[723] https://web.archive.org/web/20201026190747/http://2ai.org/legal
[724] https://web.archive.org/web/20181113082134/http://victor.ai
[725] http://www.xapagy.com/?page_id=26
[726] Bölöni (2014)
[727] Bölöni (2014)
[728] http://alumni.media.mit.edu/~kris/gandalf.html
[729] http://alumni.media.mit.edu/~kris/ymir.html

skills."[730] The most recently published paper was in 2009,[731] and the Ymir website was last updated in 2010.

Reason for consideration: An AGI R&D project

Reason for exclusion: Apparently inactive

[730] Thórisson and Helgasson (2012), p.8
[731] http://alumni.media.mit.edu/~kris/ymir.html#papers

Appendix 4. Other Projects

The following projects have stated that they have a focus on AGI, but they otherwise demonstrate little identifiable activity on AGI.

- AGI Concept Map: https://github.com/tomzx/agi-concept-map
    - The AGI Concept Map is a GitHub project by Tom Rochette, which is "reconstructing all the internal knowledge I've acquired about artificial general intelligence over the years." It was last updated in May 2017.
- AGI Literature Bank: https://github.com/wasifferoze/agi-literature-bank
    - Artificial General Intelligence Literature Collection is a collection of AGI resources by Wasif Feroze.
- AI Foundation: https://aifoundation.com
    - AI Foundation is a project founded by Lars Butler that aims to give every person their own AI that shares their personal goals and values. They also hope to fight increasingly realistic deepfakes. They have recently received significant funding.[732]
- Artificial Intelligence: https://github.com/MrRobb/Artificial-Intelligence
    - A collection of AI projects by Roberto Ariosa Hernández. The project's last commit was in 2018. The three projects are on vehicles, Tetris AI, and genetic algorithms. None appear to be AGI, but the project is tagged as such.
- Awesome Artificial Intelligence: https://github.com/fairy-tale-agi-solutions/awesome-artificial-general-intelligence
    - A GitHub project created by Răzvan Flavius Panda with information and resources about the subject of AGI (including scientific articles, movies, anime, podcasts, etc.). Projects included in the 'Organizations' section have either been coded in 2017 or 2020.
- Awesome Continual Learning: https://github.com/szrlee/awesome-continual-learning
    - Awesome Continual Learning is a collection of resources about Continual Learning. It was last updated in 2018 and was created by Richard Li.
- Awesome Deep Reinforcement Learning: https://github.com/tigerneil/awesome-deep-rl
    - Awesome Deep Reinforcement Learning is an active GitHub project run by Xiaohu Zhu, whose GitHub description describes him as "Founder & Chief Scientist of University AI. A Watchful Guardian for AGI." The GitHub is a repository of papers, but not an AGI project. Zhu's LinkedIn lists him as the founder of multiple AGI related organizations, none of which have an internet presence.[733]
- Blog.TomRochette: https://github.com/tomzx/blog.tomrochette.com-content
    - Blog.tomrochette.com contains the content of the blog of Tom Rochette. Rochette works for ElementAI, an organization that does not appear to be working on any AGI projects.[734]
- Caesium: https://github.com/generic-github-user/Caesium
    - Caesium is a "General-purpose AI library with NEAT-style genetic algorithm." It is a collection of AI algorithms rather than an AGI project.
- Chatbot with Python: https://github.com/ganeshkavhar/Chatbot-with-Python-by-ganesh-kavhar
    - A chatbot written in Python by Ganesh Kavhar. No detection of AGI, but is tagged as such.
- Checkers AI: https://github.com/sirCamp/CheckersAI

[732] https://aifoundation.com/ai-foundation-q3-newsletter/
[733] https://www.linkedin.com/in/xiaohu-zhu-97686637/
[734] https://www.elementai.com/

    - Checkers AI is an artificial intelligence that can play checkers. No detection of AGI, but is tagged as such.
- Cognitive Biotechnology: https://www.crunchbase.com/organization/cognitive-biotechnology
    - According to Crunchbase, they "develop new treatments with Cognitive Computing using Artificial General Intelligence." However, their listed website does not work, nor is there additional information.
- Creating a Chatbot: https://github.com/spydaz/Creating-a-Chatbot
    - A chatbot project created by Leroy Dyer. This is not an AGI project, but it is tagged as such on GitHub. The new version of the chatbot is not tagged as AGI.[735] Dyer's LinkedIn also does not suggest he is working on any AGI projects.[736]
- Date and Time Library in Python: https://github.com/ganeshkavhar/Date-and-Time-Library-in-Python-by-ganesh-kavhar
    - DateTime Library is a project written in Python and created by Ganesh Kavhar. No AGI was detected.
- Distributed Deep Reinforcement Learning for Large Scale Robotic Simulations: https://github.com/AmrMKayid/KayDDRL
    - A project by Amr Kayid which appears to be Kayid's undergraduate thesis.
- Evolving ANN: https://github.com/AlexanderKoch-Koch/EvolvingANN
    - Spiking Artificial Neural Network is a project by Alexander Koch. It is an "attempt to build a simple spiking artificial neural network with a reward-driven Hebbian learning function."
- God: https://github.com/JordanMicahBennett/God
    - God is the GitHub project of Jordan Micah (God)[737] Bennett and is "a non-absolute, time-space complex optimal artificial brain (composed in matrix laboratory (MatLab)" and seeks to contribute to AGI research. His other projects are the Supermathematics and AGI project and Supersymmetric ANN (both found in Appendix 4), and he believes that the purpose of human life is to build AGI.[738]
- Helmholtz Machines: https://github.com/gautam1858/Helmholtz-Machines
    - Helmholtz Machines is a GitHub project created by Gautam Ramachandra with the goal of "implementation to understand learning." His LinkedIn describes his interest in Machine Learning but not AGI.[739]
- Human Artificial General Intelligence: https://github.com/OneCivilization/Human-Artificial_General_Intelligence
    - A Gitbook about AGI written in Chinese which was last updated in 2018.
- Incremental Machine Learning: https://github.com/florianwiech/incremental-machine-learning
    - Comparative Evaluation of Incremental Machine Learning is a project by Florian Wiech comparing machine learning techniques and was last updated in mid-2019.
- Jiva.ai: https://www.jiva.ai
    - Predictive analytics in healthcare through AI. No indication of AGI on their website, but their Crunchbase says, "The core technology puts fusion at the heart of its learning procedure, enabling multimodal AI. Fusion is the process by which one can iteratively build partial models independently and then join them together to create bigger and better representations. Model fusion is a fundamental prerequisite to artificial general intelligence.

[735] https://github.com/spydaz/Chatbot_2020_Tutorial
[736] https://www.linkedin.com/in/leroy-dyer-msc-data-science-744a8230/?originalSubdomain=uk
[737] Bennett (2020)
[738] Bennett (2015)
[739] https://www.linkedin.com/in/gautamrbharadwaj/?originalSubdomain=in

Using the platform as a basis, the Jiva team will build practical diagnostic solutions to human diseases."

- Kobe: https://github.com/minuteman1911/kobe
    - Kobe is "an open-source simulation tool written in Python which can model artificial spiking neural networks, along with its environment (using OpenAI Gym), with more focus on the computational or functional aspect rather than the biophysical one." It was created by Aishwarya Dabhade who describes it as an attempt to bring ANNs more in line with General Intelligence, to bridge the gap between "how a biological neural network works and how an artificial one works." The goal of Kobe "is not full brain simulation, such a thing is premature, inconceivable and might not even be possible at all. The goal is to help AI researchers, neuroscientists and enthusiasts in creating, testing and implementing novel algorithms, morphologies and methods which are closer to the biological brain while focusing on the computational rather than the biological aspect of the simulation."[740]
- Lisa: http://www.ai-machine.com/index.html
    - Lisa is the AI assistant from the organization AI Machine. They claim it is true AGI: "LISA is NOT a domain specific AI Bot that has been trained using Machine Learning models just to do one thing right. LISA is powered by a Cognitive Architecture with Deep Semantic understanding based on Artificial General Intelligence (AGI or Strong AI) and Natural Language Understanding (NLU) Techniques from the leading AI Research Labs."[741]
- Lya: https://github.com/lya-corp/lya
    - Lya is a GitHub repo by Lya Corp,[742] the developer of the Kalliope project, an "Alexa-esque" voice assistant.[743]
- Medsis: https://www.medsis.com
    - A medical tech company; The website does not reference AGI, but their description on Crunchbase does: "Our platform enables access, management, integration, consolidation, machine learning, sharing and distribution, analytics, and artificial general intelligence (AGI)."[744]
- MetaQuants: https://www.crunchbase.com/organization/metaquants
    - According to Crunchbase, they are a "Data-Driven Self-Adaptive Evolving Investment Platform using Cognitive Computing with Artificial General Intelligence." However, their listed website does not work, and there is no additional information.
- Neural GPU Algorithm Learner: https://github.com/prakashsellathurai/Neural_gpu_algorithm_learner
    - A "Neural Network that learns algorithms" created by Prakash Sellathurai. It is based on a 2015 preprint of a similar name.[745]
- PeakPerformance: https://www.crunchbase.com/organization/peakperformance-ai
    - According to Crunchbase, "PeakPerformance delivers a high-performance cognitive computing framework using Artificial General Intelligence (AGI) & Genetic Algorithms." However, their listed website does not work and there is no additional information.
- Pegasus: https://github.com/PegasusProject/Pegasus
    - Pegasus is a "multi-platform, intelligent and extremely modular daily life management program." It was created by the Pegasus Project GitHub user. It seems to be a sort of Alexa-

[740] https://minuteman1911.github.io/kobe/
[741] https://www.crunchbase.com/organization/ai-machines
[742] https://www.linkedin.com/company/lya-electronic-corp/
[743] https://kalliope-project.github.io/
[744] https://www.crunchbase.com/organization/medsis-medical-systems
[745] Kaiser and Sutskever (2015)

esque project, using "general" to mean a general-purpose AI voice assistant rather than AGI.

- Project Origin: https://github.com/kourgeorge/project-origin
    - A GitHub project which attempts "to create an artificial life environment that allows studying the emergence and evolution of intelligence." No detection of explicit AGI, but tagged as such.
- Reinforcement Learning Research Lab (RLRL): https://github.com/NaxAlpha/rl-rl
    - The RLRL toolkit is used to "build agents of any windows game you have using python." It was created by Nauman Mustafa and appears to be inactive.
- SORN: https://github.com/Saran-nns/sorn
    - Self-Organizing Recurrent Neural Networks created by Saranraj Nambusubramaniyam is a "class of neuro-inspired artificial network build based on plasticity mechanisms in biological brain and mimic neocortical circuits ability of learning and adaptation through neuroplasticity mechanisms." Use Cases and the API are stored in the separate repo PySORN_0.1.[746]
- Sigmind: https://sigmind.ai
    - A computer vision AI company. The website does not reference AGI, but the Crunchbase description does: "To make it more accelerated and public friendly, we are working on putting research hours to bring the holy grail of artificial general intelligence for practical real world engineering applications to ease human life."[747] Crunchbase also states they are collaborating with Deepmind, Microsoft, and Facebook.
- Spydaz Web AI Intelligent Agent: https://github.com/spydaz/SpydazWebAI_IntelligentAgent
    - Leroy Dyer created SpydazWebAI_IntelligentAgent as part of the SpydazWebAI project.[748] The project appears to be a chatbot and not an AGI.
- Supermathematics and Artificial General Intelligence: https://github.com/JordanMicahBennett/Supermathematics-and-Artificial-General-Intelligence
    - A GitHub project detailing Jordan Micah (God) Bennett's contribution to AGI called the "Supermanifold Hypothesis." Bennett has other projects that discuss this theme, such as Supersymmetric Artificial Neural Network and God (both Appendix 4).
- Supersymmetric ANN: https://github.com/JordanMicahBennett/Supersymmetric-artificial-neural-network
    - A "Lie Superalgebra aligned algorithmic learning model, based on evidence pertaining to Supersymmetry in the biological brain" created by Jordan Micah (God) Bennett.
- Thoughts on AI Trust: https://github.com/jonathan-smith-1/Thoughts-on-AI-Trust
    - A GitHub compilation of thoughts about AI Safety by Jonathan Smith.
- Tree Algebra: https://github.com/raviq/tree_algebra
    - A project that Manipulation of binary tree topologies by Hafi R. whose "Solving Tree Problems with Category Theory" was published in AGI 2018.
- Voice Activity Detection Final Project Work: https://github.com/samimoftheworld/Voice-Activity-Detection-FInal-Project-work
    - A project created for the BTech degree of Samim Ekram, a "Freelancing Machine Learning and Artificial Intelligence Developer and Enthusiast."

[746] https://github.com/Saran-nns/PySORN_0.1
[747] https://www.crunchbase.com/organization/sigmind
[748] http://www.spydazweb.co.uk

References

Alidoust M, Rouhani M, 2015. A computational behavior model for life-like intelligent agents. In: Romportl J, Zackova E, Kelemen J (eds) Beyond Artificial Intelligence. Cham, Switzerland: Springer, pp. 159-175.

Allen G, Kania EB, 2017. China is using America's own plan to dominate the future of artificial intelligence. Foreign Policy, http://foreignpolicy.com/2017/09/08/china-is-using-americas-own-plan-to-dominate-the-future-of-artificial-intelligence

Allison G, 2017. Destined for War: Can America and China Escape Thucydides's Trap? New York: First Mariner Books.

Amodei D, Christiano P, Ray A, 2017. Learning from human preferences. OpenAI, https://openai.com/blog/deep-reinforcement-learning-from-human-preferences/

Anderson JR, Bothell D, Byrne MD, Douglass S, Lebiere C, Qin Y, 2004. An integrated theory of the mind. Psychological Review, 111(4), 1036–1060.

Arel I, Rose D, Coop R, 2009. DeSTIN: A scalable deep learning architecture with application to high-dimensional robust pattern recognition. Proceedings of the AAAI Fall Symposium on Biologically Inspired Cognitive Architectures, pp. 11–15.

Armstrong S, Bostrom N, Shulman C, 2016. Racing to the precipice: A model of artificial intelligence development. AI & Society, 31(2), 201-206.

Arrabales R, Muñoz J, 2010. The awakening of conscious bots: Inside the mind of the 2K BotPrize 2010 winner.

Arrabales R, Muñoz J, Ledezma A, Gutierrez G, Sanchis A, 2013. A machine consciousness approach to the design of human-like bots. In Hingston P (Ed), Believable Bots. Berlin: Springer, pp. 171-191.

Auerbach C, Silverstein LB, 2003. Qualitative Data: An Introduction to Coding and Analysis. New York: NYU Press.

Bach J, 2009. Principles of synthetic intelligence PSI an architecture of motivated cognition. Oxford: Oxford University Press.

Barrett AM, Baum SD, 2017a. A model of pathways to artificial superintelligence catastrophe for risk and decision analysis. Journal of Experimental & Theoretical Artificial Intelligence, 29(2), 397-414.

Barrett AM, Baum SD, 2017b. Risk analysis and risk management for the artificial superintelligence research and development process. In Callaghan V, Miller J, Yampolskiy R, Armstrong S (Eds), The Technological Singularity. Berlin: Springer, pp. 127-140.

Baum SD, 2017a. A survey of artificial general intelligence projects for ethics, risk, and policy. Global Catastrophic Risk Institute Working Paper 17-1.

Baum SD, 2017b. On the promotion of safe and socially beneficial artificial intelligence. AI & Society, 32(4), 543-551.

Baum SD, 2018. Reconciliation between factions focused on near-term and long-term artificial intelligence. AI & Society, 33(4), 565-572.

Baum SD, 2020. Social choice ethics in artificial intelligence. AI & Society, 35 (1), 165-176.

Baum SD, de Neufville R, Barrett AM, Ackerman G, 2019. Lessons for artificial intelligence from other global risks. In Maurizio Tinnirello (Editor), The Global Politics of Artificial Intelligence. Boca Raton: CRC Press, in press.

Baum SD, Goertzel B, Goertzel TG, 2011. How long until human-level AI? Results from an expert assessment. Technological Forecasting and Social Change 78(1), 185-195.

BBC News, 2020. Coronavirus: Uber announces drastic cuts to secure its future. BBC Business, https://www.bbc.com/news/business-52711649

Beattie C, Leibo J, Petersen S, Legg S, 2016. Open-sourcing DeepMind lab. DeepMind, https://deepmind.com/blog/article/open-sourcing-deepmind-lab
Belfield H, 2020. Activism by the AI community: Analysing recent achievements and future prospects. In Proceedings of the 2020 AAAI/ACM Conference on AI, Ethics and Society, pp. 15-21.
Benjamin BV, Gao P, McQuinn E, Choudhary S, Chandrasekaran AR, Bussat JM, et al., 2014. Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations. Proceedings of the IEEE, 102(5), 699-716.
Bennett J, 2015. Why is the purpose of human life to create artificial general intelligence? https://www.researchgate.net/publication/319235750_Why_is_the_purpose_of_human_life_to_create_Artificial_General_Intelligence
Bennett J, 2020. Why I (an atheist) legally changed my name to "God." https://www.researchgate.net/publication/342328687_Why_I_an_atheist_legally_changed_my_name_to_God
Bieger J, Thórisson KR, Wang P, 2015. Safe baby AGI. In Bieger J, Goertzel B, Potapov A (Eds), Proceedings of AGI 2015, 8th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 46-49.
Bölöni L, 2014. Autobiography based prediction in a situated AGI agent. In Goertzel B, Orseau L, Snaider J (Eds), Proceedings of AGI 2014, 7th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 11-20.
Bonhoff GM, 2008. Using hierarchical temporal memory for detecting anomalous network activity. Masters Thesis, Department of Electrical and Computer Engineering, US Air Force Institute of Technology.
Bostrom N, 2014. Superintelligence: Paths, Dangers, Strategies. Oxford: Oxford University Press.
Bostrom N, 2017. Strategic implications of openness in AI development. Global Policy, 8(2), 135-148.
Boyle A, 2016. A bot is born: Kimera Systems adds 'Nigel' to the crowd of AI assistants. GeekWire, https://www.geekwire.com/2016/kimera-systems-nigel-ai-agi
Brandstätter C, Dietrich D, Doblhammer K, Fittner M, Fodor G, Gelbard F, et al., 2015. Natural Scientific, Psychoanalytical Model of the Psyche for Simulation and Emulation. Scientific Report III, Institute of Computer Technology, Vienna University of Technology, Version eV005.
Bringsjord S, 2012. Belief in the singularity is logically brittle. Journal of Consciousness Studies, 19(7), 14-20.
Brundage M, Avin S, Wang J, Belfield H, Krueger G, Hadfield G, et al., 2020. Toward trustworthy AI development: Mechanisms for supporting verifiable claims. https://arxiv.org/abs/2004.07213
Cai F, 2020. AI startup Covariant.ai building 'universal AI for robots'. Synced, https://syncedreview.com/2020/06/16/ai-startup-covariant-ai-building-universal-ai-for-robots/
Cassimatis NL, 2007. Adaptive algorithmic hybrids for human-level artificial intelligence. In Goertzel B, Wang P (Eds), Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms. Amsterdam: IOS Press, pp. 94-110.
In Goertzel B, Wang P (Eds). Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms, pp. 94–110.
Chen N, 2016. China Brain Project to launch soon, aiming to develop effective tools for early diagnosis of brain diseases. Chinese Academy of Sciences, http://english.cas.cn/newsroom/archive/news_archive/nu2016/201606/t20160617_164529.shtml
Chen X, 2017. Gary Marcus: The road to artificial general intelligence. Medium, https://medium.com/@Synced/gary-marcus-the-road-to-artificial-general-intelligence-ce5b1371aa02
Choi D, Langley P, 2018. Evolution of the Icarus cognitive architecture. Cognitive Systems Research, 48, 25-38.

Chong HQ, Tan AH, Ng GW, 2007. Integrated cognitive architectures: A survey. Artificial Intelligence Review, 28(2), 103-130.
Christiano P, 2016. Prosaic AI alignment. AI Alignment, https://ai-alignment.com/prosaic-ai-control-b959644d79c2
Cihon P, 2019. Standards for AI governance: International standards to enable global coordination in AI research & development. FHI Technical report.
Clark J, Hadfield GK, 2019. Regulatory markets for AI safety. https://arxiv.org/abs/2001.00078
Coffey A, Atkinson P, 1996. Making Sense of Qualitative Data: Complementary Research Strategies. Thousand Oaks, CA: Sage.
Cole S, 2020. This company made a 'cute' AI lawyer to deploy 'information warfare' for divorced men. Vice, https://www.vice.com/en/article/k7ew9z/this-company-made-a-cute-ai-lawyer-to-deploy-information-warfare-for-divorced-men
Conn A, 2016. The White House considers the future of AI. Future of Life Institute, https://futureoflife.org/2016/06/15/white-house-future-ai-1
Cutler KM, 2014. Vicarious grabs a huge, new $40M growth round to advance artificial intelligence. TechCrunch, https://techcrunch.com/2014/03/22/vicarious-grabs-a-huge-new-40m-growth-round-to-advance-artificial-intelligence
Cyranoski D, 2018. Beijing launches pioneering brain-science centre. Nature, 556 (7699), 157-158.
Dambrot SM, 2020. Theoretical and hypothetical pathways to real-time neuromorphic AGI/post-AGI ecosystems. Procedia Computer Science, 169, 110-122.
Danzig R, 2018. Technology roulette: Managing loss of control as many militaries pursue technological superiority. Center for a New American Security, https://www.cnas.org/publications/reports/technology-roulette
Dastin J, 2018. Amazon scraps secret AI recruiting tool that showed bias against women. Reuters, https://reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK086
de Garis H, 2005. The Artilect War. Pittsburgh: Etc Press.
de Garis H, Shuo C, Goertzel B, Ruiting L, 2010. A world survey of artificial brain projects, Part I: Large-scale brain simulations. Neurocomputing, 74(1), 3-29.
DeAngelis SF, 2014. Quantum computers will transform the world—we hope. Enterra Insights Blog, http://www.enterrasolutions.com/quantum-computers-will-transform-world-hope
Deaton C, Shepard B, Klein C, Mayans C, Summers B, Brusseau A, et al., 2005. The comprehensive terrorism knowledge base in Cyc. Proceedings of the 2005 International Conference on Intelligence Analysis.
Dewey D, 2015. Long-term strategies for ending existential risk from fast takeoff. In Müller VC (Ed), Risks of Artificial Intelligence. Boca Raton: CRC, pp. 243-266.
Dong D, Franklin S, 2014. The action execution process implemented in different cognitive architectures: A review. Journal of Artificial General Intelligence, 5(1), 49-68.
Duch W, Oentaryo RJ, Pasquier M, 2008. Cognitive architectures: Where do we go from here? In Wang P, Goertzel B, Franklin S (Eds), Frontiers in Artificial Intelligence and Applications, Vol. 171. Amsterdam: IOS Press, pp. 122-136.
Dvorsky G, 2013. How Skynet might emerge from simple physics. io9, http://io9.gizmodo.com/how-skynet-might-emerge-from-simple-physics-482402911
Etherington D, 2017. Microsoft creates an AI research lab to challenge Google and DeepMind. TechCrunch, https://techcrunch.com/2017/07/12/microsoft-creates-an-ai-research-lab-to-challenge-google-and-deepmind
Etzioni O, 2016. No, the experts don't think superintelligent AI is a threat to humanity. MIT Technology Review, https://www.technologyreview.com/s/602410/no-the-experts-dont-think-superintelligent-ai-is-a-threat-to-humanity

European Commission, 2020. EU Grants: Guidance note — Exclusive focus on civil applications: V1.1 — 07.01.2020.
Everett SM, 2017. Video explaining Biologic Intelligence. Medium, https://medium.com/the-mission/video-explaining-biologic-intelligence-76ce84476f54
Fernando C, Banarse D, Blundell C, Zwols Y, Ha D, Rusu AA, et al., 2017. PathNet: Evolution channels gradient descent in super neural networks. https://arxiv.org/abs/1701.08734
Franklin S, Strain S, Snaider J, McCall R, Faghihi U, 2012. Global workspace theory, its LIDA model and the underlying neuroscience. Biologically Inspired Cognitive Architectures, 1, 32-43.
Fruchart B, Le Blanc B, 2019. Architecture cognitive et comportements. In Sylvain H, Bernard A (Eds), L'activation du corps vivant. Émersions, hybridations, remédiations, Intellectica, 71, 139-156.
Fruchart B, Le Blanc B, 2020. Cognitive machinery and behaviors. https://www.researchgate.net/publication/342183742_Cognitive_Machinery_and_Behaviours
Gershgorn D, 2017. China is funding Baidu to take on the US in deep-learning research. Quartz, https://qz.com/916738/china-is-funding-baidu-to-take-on-the-united-states-in-deep-learning-research
Gibbs S, 2014. Google buys UK artificial intelligence startup Deepmind for £400m. The Guardian, https://www.theguardian.com/technology/2014/jan/27/google-acquires-uk-artificial-intelligence-startup-deepmind
Goertzel B, 2009. Will China build AGI first? The Multiverse According to Ben, http://multiverseaccordingtoben.blogspot.com/2009/07/will-china-build-agi-first.html
Goertzel B, 2010. A Cosmist Manifesto: Practical Philosophy for the Posthuman Age. Humanity+ Press.
Goertzel B, 2011. Pei Wang on the path to artificial general intelligence. h+ Magazine, http://hplusmagazine.com/2011/01/27/pei-wang-path-artificial-general-intelligence
Goertzel B, 2012a. Should humanity build a global AI nanny to delay the singularity until it's better understood? Journal of Consciousness Studies, 19(1-2), 96-111.
Goertzel B, 2014. Artificial general intelligence: Concept, state of the art, and future prospects. Journal of Artificial General Intelligence, 5(1), 1-48.
Goertzel B, 2015. Superintelligence: Fears, promises and potentials. Journal of Evolution and Technology, 25(2), 55-87.
Goertzel B, 2016. Infusing advanced AGIs with human-like value systems: Two theses. Journal of Evolution and Technology, 26(1), 50-72.
Goertzel B, 2017a. The corporatization of AI is a major threat to humanity. H+ Magazine, http://hplusmagazine.com/2017/07/21/corporatization-ai-major-threat-humanity
Goertzel B, 2017b. SingularityNET—AGI of the People, by the People and for the People (and the Robots!). Medium, 26 October, https://medium.com/ben-goertzel-on-singularitynet/singularitynet-agi-of-the-people-by-the-people-and-for-the-people-and-the-robots-7246a6886e90
Goertzel B, Pitt J, 2012. Nine ways to bias open-source AGI toward friendliness. Journal of Evolution & Technology, 22(1), 116-131.
Goertzel B, Giacomelli S, Hanson D, Pennachin C, Argentieri M, the SingularityNET team, 2017. SingularityNET: A decentralized, open market and inter-network for AIs. https://public.singularitynet.io/whitepaper.pdf
Goertzel B, Lian R, Arel I, De Garis H, Chen S, 2010. A world survey of artificial brain projects, Part II: Biologically inspired cognitive architectures. Neurocomputing, 74(1), 30-49.
Goertzel B, Pitt J, Wigmore J, Geisweiller N, Cai Z, Lian R, Huang D, Yu G, 2011. Cognitive synergy between procedural and declarative learning in the control of animated and robotic agents using the OpenCogPrime AGI architecture. Proceedings of AAAI-11, 2, pp. 1436-1441.

Good IJ, 1965. Speculations concerning the first ultraintelligent machine. Advances in Computers, 6, 31-88.
Griffin M, 2017. Baidu's AI just achieved zero shot learning. Fanatical Futurist, http://www.fanaticalfuturist.com/2017/04/baidus-ai-just-achieved-zero-shot-learning
Gupta A, Lanteigne C, Heath V, 2020. Report prepared by the Montreal AI Ethics Institute (MAIEI) for Publication Norms for Responsible AI by Partnership on AI. https://arxiv.org/abs/2009.07262
Hawkins J, 2007. On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines. New York: Times Books.
Hawkins J, 2015. The Terminator is not coming. The future will thank us. recode, https://www.recode.net/2015/3/2/11559576/the-terminator-is-not-coming-the-future-will-thank-us
Hawkins J, 2017. What intelligent machines need to learn from the neocortex. IEEE Spectrum, http://spectrum.ieee.org/computing/software/what-intelligent-machines-need-to-learn-from-the-neocortex
Hawkins AJ, 2019. Uber goes public: Everything you need to know about the biggest tech IPO in years. The Verge, https://www.theverge.com/2019/5/10/18564197/uber-ipo-stock-valuation-pricing-fares-drivers-public-market
Hibbard B, 2002. Super-Intelligent Machines. New York: Springer.
High P, 2016. Vicarious is the AI company that includes Zuckerberg, Bezos, Musk, and Thiel as investors. Forbes, https://www.forbes.com/sites/peterhigh/2016/04/11/vicarious-is-the-ai-company-that-includes-zuckerberg-bezos-musk-and-thiel-as-investors
Hughes JJ, 2007. Global technology regulation and potentially apocalyptic technological threats. In Allhoff F (Ed), Nanoethics: The Ethical and Social Implications of Nanotechnology. Hoboken, NJ: John Wiley, pp. 201-214.
Hutton N, 2016. Selections from Bluebrain: A 10 year neuroscience documentary. Medium, https://medium.com/labocine/selections-from-bluebrain-a-10-year-neuroscience-documentary-56d869599640
Jilk D, Herd S, Read S, O'Reilly R, 2017. Anthropomorphic reasoning about neuromorphic AGI safety. Journal of Experimental & Theoretical Artificial Intelligence, 29(6), 1337-1351.
Jilk DJ, Lebiere C, O'Reilly RC, Anderson JR, 2008. SAL: An explicitly pluralistic cognitive architecture. Journal of Experimental and Theoretical Artificial Intelligence, 20(3), 197-218.
Joy B, 2000. Why the future doesn't need us. Wired, 8(4), 238-263.
Kaczynski T, 1995. Industrial Society and Its Future.
Kaiser Ł, Sutskever I, 2015. Neural GPUs learn algorithms. https://arxiv.org/abs/1511.08228
Kaushal A, Altman R, Langlotz C, 2020. Health care AI systems are biased. Scientific American, https://www.scientificamerican.com/article/health-care-ai-systems-are-biased/
Kelley D, Atreides K, 2020. AGI protocol for the ethical treatment of artificial general intelligence systems. Procedia Computer Science, 169, 501-506.
Knight W, 2016. An AI with 30 years' worth of knowledge finally goes to work. MIT Technology Review, https://www.technologyreview.com/2016/03/14/108873/an-ai-with-30-years-worth-of-knowledge-finally-goes-to-work
Kong Q, Han J, Zeng Y, Xu B, 2018. Efficient coding matters in the organization of the early visual system. Neural Networks, 105, 218-226.
Kotseruba I, Gonzalez OJA, Tsotsos JK, 2016. A review of 40 years of cognitive architecture research: Focus on perception, attention, learning and applications. https://www.researchgate.net/profile/John_Tsotsos/publication/309483878_A_Review_of_40_Years_of_Cognitive_Architecture_Research_Focus_on_Perception_Attention_Learning_and_Applications/links/58148e1c08aedc7d8963b93b.pdf
Koza JR, 1994. Genetic programming as a means for programming computers by natural selection. Statistics and Computing, 4(2), 87-112.

Kugele S, Franklin S, 2020. Learning in LIDA. Cognitive Systems Research, in press, doi 10.1016/j.cogsys.2020.11.001.
Laird JE, 2012. The SOAR cognitive architecture. Cambridge, MA: MIT Press.
Lane PC, Gobet F, 2012. CHREST models of implicit learning and board game interpretation. In Bach J, Goertzel B, Ikle M (Eds), Proceedings of AGI 2012, 5th International Conference on Artificial General Intelligence. Berlin: Springer, pp. 148-157.
Langley P, 2005. An adaptive architecture for physical agents. Proceedings of IEEE/WIC/ACM International Conference on Intelligent Agent Technology, pp. 18-25.
Langley P, 2017. Progress and challenges in research on cognitive architectures. Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17), pp. 4870-4876.
Langley P, Laird J, Rogers S, 2008. Cognitive architectures: Research issues and challenges. Cognitive Systems Research 10, 141-160.
LeCun Y, Bengio Y, Hinton G, 2015. Deep learning. Nature, 521(7553), 436-444.
Lentzos F, 2011. Hard to prove: The verification quandary of the Biological Weapons Convention. Nonproliferation Review, 18(3), 571-582.
Legg S, Hutter M, 2008. Machine Super Intelligence. http://www.vetta.org/documents/Machine_Super_Intelligence.pdf
Leike J, Martic M, Legg S, 2017. Learning through human feedback. DeepMind, https://deepmind.com/blog/article/learning-through-human-feedback
Loosemore RPW, 2014. The maverick nanny with a dopamine drip: Debunking fallacies in the theory of AI motivation. Association for the Advancement of Artificial Intelligence Spring Symposium.
Love D, 2014. The most ambitious artificial intelligence project in the world has been operating in near secrecy for 30 years. Business Insider, http://www.businessinsider.com/cycorp-ai-2014-7
Maddox T, 2016. Should Baidu be your AI and machine learning platform? ZDNet, http://www.zdnet.com/article/should-baidu-be-your-ai-and-machine-learning-platform
Madl T, Franklin S, 2015. Constrained incrementalist moral decision making for a biologically inspired cognitive architecture. In Trappl R (Ed), A Construction Manual for Robots' Ethical Systems: Requirements, Methods, Implementations. Cham, Switzerland: Springer, pp. 137-153.
Mannes J, 2017. Tencent to open AI research center in Seattle. TechCrunch, https://techcrunch.com/2017/04/28/tencent-to-open-ai-research-center-in-seattle
Marcus G, Davis E, 2013. A grand unified theory of everything. New Yorker, http://www.newyorker.com/tech/elements/a-grand-unified-theory-of-everything
Marcus G, Davis E, 2019. Rebooting AI: Building Artificial Intelligence We Can Trust. New York: Pantheon Books.
Marquis C, Toffel, MW, Zhou Y, 2016. Scrutiny, norms, and selective disclosure: A global study of greenwashing. Organization Science, 27(2), 483-504.
McDermott D, 2012. Response to the singularity by David Chalmers. Journal of Consciousness Studies, 19(1-2), 167-172.
Miller S, 2018. Army Research Lab teams up with Uber. Defense Systems, https://defensesystems.com/articles/2018/08/15/uber-arl-ut-nasa.aspx
Merolla PA, Arthur JV, Alvarez-Icaza R, Cassidy AS, Sawada J, Akopyan F, et al., 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. Science, 345(6197), 668-673.
Metz C, 2016. Uber Buys a mysterious startup to make itself an AI company. Wired, https://www.wired.com/2016/12/uber-buys-mysterious-startup-make-ai-company
Moskovitch K, 2020. Neurosymbolic AI to give us machines with true common sense. Medium, https://medium.com/swlh/neurosymbolic-ai-to-give-us-machines-with-true-common-sense-9c133b78ab13

Muehlhauser L, Helm L, 2012. Intelligence explosion and machine ethics. In Eden A, Søraker J, Moor JH, Steinhart E (Eds), Singularity hypotheses: A Scientific and Philosophical Assessment. Berlin: Springer, pp.101-126.
Muggleton S, 1991. Inductive logic programming. New Genereration Computing, 8(4), 295-318.
Nestor A, Kokinov B, 2004. Towards active vision in the DUAL cognitive architecture. Information Theories & Applications 11(1), 9-15.
Neven H, 2013. Launching the quantum artificial intelligence lab. GoogleAI Blog, https://ai.googleblog.com/2013/05/launching-quantum-artificial.html
Ng KH, Du Z, Ng GW, 2017. DSO cognitive architecture: Unified reasoning with integrative memory using global workspace theory. In Everitt T, Goertzel B, Potapov A (Eds), Proceedings of AGI 2017, 10th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 44-53.
Ng GW, Leung WC, 2020. Strong artificial intelligence and consciousness. Journal of Artificial Intelligence and Consciousness, 7(1), pp. 63-72.
Nivel E, Thórisson KR, Steunebrink BR, Dindo H, Pezzulo G, Rodriguez M, et al., 2013. Bounded recursive self-improvement. Reykjavik University School of Computer Science Technical Report RUTR-SCS13006.
Novet J, 2016. Salesforce forms research group, launches Einstein A.I. platform that works with Sales Cloud, Marketing Cloud. Venture Beat, https://venturebeat.com/2016/09/18/salesforce-forms-research-group-launches-einstein-a-i-platform-that-works-with-sales-cloud-marketing-cloud
O'Reilly RC, Hazy TE, Herd SA, 2016. The Leabra cognitive architecture: How to play 20 principles with nature. In Chipman SEF (Ed), The Oxford Handbook of Cognitive Science. Oxford: Oxford University Press, pp. 91-116.
Omohundro S, 2016. AIBrain talk: AI and human safety. Self-Aware Systems, https://selfawaresystems.com/2016/08/17/aibrain-talk-ai-and-human-safety
Oreskes N, Conway EM, 2010. Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming. New York: Bloomsbury.
Oudeyer PY, Ly O, Rouanet P, 2011. Exploring robust, intuitive and emergent physical human-robot interaction with the humanoid Acroban. Proceedings of the IEEE-RAS International Conference on Humanoid Robots, Bled, Slovenia.
Özkural E, 2018. Omega: An architecture for AI Unification. arXiv:1805.12069.
Pei J, Deng L, Song S, Zhao M, Zhang Y, Wu S, et al., 2019. Towards artificial general intelligence with hybrid Tianjic chip
Pollock JL, 2008. OSCAR: An agent architecture based on defeasible reasoning. Proceedings of the 2008 AAAI Spring Symposium on Architectures for Intelligent Theory-Based Agents, pp. 55-60.
Poo MM, Du JL, Ip NY, Xiong ZQ, Xu B, Tan T, 2016. China Brain Project: Basic neuroscience, brain diseases, and brain-inspired computing. Neuron, 92(3), 591-596.
Potapov A, Rodionov S, 2014. Genetic algorithms with DNN-based trainable crossover as an example of partial specialization of general search.
Potapov A, Rodionov S, Potapova V, 2016. Real-time GA-based probabilistic programming in application to robot control. In Steunebrink B, Wang P, Goertzel B (Eds), Proceedings of AGI 2016, 9th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 95-105.
Ray R, 2020. 5 questions policymakers should ask about facial recognition, law enforcement, and algorithmic bias. The Brookings Institute, https://www.brookings.edu/research/5-questions-policymakers-should-ask-about-facial-recognition-law-enforcement-and-algorithmic-bias/
Richardson M, Domingos P, 2006. Markov logic networks. Machine Learning, 62(1-2), SPEC. ISS., pp. 107–-136.

Rosbe J, Chong RS, Kieras DE, 2001. Modeling with perceptual and memory constraints: An EPIC-Soar model of a simplified enroute air traffic control task. Air Force Research Laboratory Report AFRL-HE-WP-TR-2002-0231
Rosenbloom PS, 2013. The Sigma cognitive architecture and system. AISB Quarterly, 136, 4-13.
Rosenbloom PS, Ustun V, 2019. An architectural integration of Temporal Motivation Theory for decision making. Proceedings of the 17thAnnual Meeting of the International Conference on Cognitive Modeling.
RT, 2017. 'Whoever leads in AI will rule the world': Putin to Russian children on Knowledge Day. RT, https://www.rt.com/news/401731-ai-rule-world-putin
Sagar R, 2020. Uber winds down its AI labs: A look at some of their top work. Analytics India Mag, https://analyticsindiamag.com/uber-ai-labs-layoffs
Samsonovich AV, 2010. Toward a unified catalog of implemented cognitive architectures. Biologically Inspired Cognitive Architectures, 221, 195-244.
Scherer MU, 2016. Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies. Harvard Journal of Law & Technology, 29(2), 353-400
Schienke EW, Tuana N, Brown DA, Davis KJ, Keller K, Shortle JS, et al., 2009. The role of the NSF Broader Impacts Criterion in enhancing research ethics pedagogy. Social Epistemology, 23(3-4), 317-336.
Schmelzer R, 2020. Can't define AI? Try defining intelligence. Forbes, https://www.forbes.com/sites/cognitiveworld/2020/02/27/cant-define-ai-try-defining-intelligence
Schmidhuber J, 2007. Gödel machines: Fully self-referential optimal universal self-improvers. In Goertzel B, Pennachin C (Eds), Artificial General Intelligence. Berlin: Springer, pp. 199-226.
Shapiro SC, Bona JP, 2010. The GLAIR cognitive architecture. International Journal of Machine Consciousness, 2(2), 307-332.
Shapiro SC, Rapaport WJ, Kandefer MW, Johnson FL, Goldfain A, 2007. Metacognition in SNePS. AI Magazine, 28(1), 17–31.
Shastri L, Ajjanagadde V,1990. From simple associations to systemic reasoning: A connectionist representation of rules, variables and dynamic bindings. Behavioral and Brain Sciences, 16(3), 417-451.
Shead S, 2020. Why the buzz around DeepMind is dissipating as it transitions from games to science. CNBC, https://www.cnbc.com/2020/06/05/google-deepmind-alphago-buzz-dissipates.html
Sherwood D, Higbee T, 2020. When will artificial general intelligence be ready for smart weapon and sensor systems? Military Embedded Systems, https://militaryembedded.com/ai/deep-learning/when-will-artificial-general-intelligence-be-ready-for-smart-weapon-and-sensor-systems
Shevlane T, Dafoe A, 2020. The offense-defense balance of scientific knowledge: Does publishing AI research reduce misuse? Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, pp. 173-179.
Shulman C, 2009. Arms control and intelligence explosions. Proceedings of the 7th European Conference on Computing and Philosophy (ECAP), Bellaterra, Spain, July 2-4.
Simon CJ, 2018. Will computers revolt? Preparing for the future of artificial intelligence. Washington, DC: FutureAI Press.
Skaba W, 2012a. Binary space partitioning as intrinsic reward. In Bach J, Goertzel B, Ikle M (Eds), Proceedings of AGI 2012, 5th International Conference on Artificial General Intelligence. Berlin: Springer, pp. 282-291.
Skaba W, 2012b. The AGINAO self-programming engine. Journal of Artificial General Intelligence, 3(3), 74-100.
Sotala K, Yampolskiy RV, 2014. Responses to catastrophic AGI risk: a survey. Physica Scripta, 90(1), 018001.

Sowa JF, 1993. Building large knowledge-based systems: Representation and inference in the cyc project. D.B. Lenat and R.V. Guha. Artificial Intelligence, 61(1), 95-104.
Starzyk JA, Graham J, 2015. MLECOG: Motivated learning embodied cognitive architecture. IEEE Systems Journal, 11(3), 1272-1283.
Steunebrink BR, Thórisson KR, Schmidhuber J, 2016. Growing recursive self-improvers. In Steunebrink B, Wang P, Goertzel B (Eds), Proceedings of AGI 2016, 9th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 129-139.
Stilgoe J, Maynard A, 2017. It's time for some messy, democratic discussions about the future of AI. The Guardian, 1 February, https://www.theguardian.com/science/political-science/2017/feb/01/ai-artificial-intelligence-its-time-for-some-messy-democratic-discussions-about-the-future
Strannegård C, Nizamani AR, 2016a. Integrating symbolic and sub-symbolic reasoning. In Steunebrink B, Wang P, Goertzel B (Eds), Proceedings of AGI 2016, 9th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 171-180.
Strannegård C, Nizamani AR, Juel J, Persson U, 2016b. Learning and reasoning in unknown domains. Journal of Artificial General Intelligence, 7(1), 104-127.
Strannegård C, Svangård N, Bach J, Steunebrink B, 2017a. Generic animats. In Everitt T, Goertzel B, Potapov A (Eds), Proceedings of AGI 2017, 10th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 23-32.
Strannegård C, Svangård N, Lindström D, Bach J, Steunebrink B, 2017b. The animat path to artificial general intelligence. Proceedings of IJCAI-17 Workshop on Architectures for Generality & Autonomy.
Strannegård C, Svangård N, Lindström D, Bach J, Steunebrink B, 2018. Learning and decision-making in artificial animals. Journal of Artificial Intelligence, 9(1), 55-82.
Strannegård C, Xu W, Engsnr N, Endler JA, 2020. Combining evolution and learning in computational ecosystems. Journal of Artificial Intelligence, 11(1), 1-37.
Sun L, 2020. Apple is quietly expanding its artificial intelligence ecosystem. The Motley Fool, https://www.fool.com/investing/2020/01/23/apple-is-quietly-expanding-its-artificial-intellig.aspx
Sun R, Zhang X, 2004. Top-down versus bottom-up learning in cognitive skill acquisition. Cognitive Systems Research, 5(1), pp. 63-89.
Taatgen N, Anderson J, 2010. The past, present, and future of cognitive architectures. Topics in Cognitive Science 2(4), pp. 693-704.
Tamturk V, 2017. Google, Apple, Facebook, and Intel battle for AI supremacy. CMS Connected, http://www.cms-connected.com/News-Archive/April-2017/Google-Apple-Facebook-Intel-Microsoft-Salesforce-Twitter-Battle-AI-Supremacy
Temperton J, 2016. Uber acquires AI firm to boost self-driving car project. Wired, http://www.wired.co.uk/article/uber-artificial-intelligence-lab-self-driving-cars
The Economist, 2016. From not working to neural networking. The Economist, https://www.economist.com/news/special-report/21700756-artificial-intelligence-boom-based-old-idea-modern-twist-not
Theil S, 2015. Why the Human Brain Project went wrong—and how to fix it. Scientific American, https://www.scientificamerican.com/article/why-the-human-brain-project-went-wrong-and-how-to-fix-it
Thórisson K, Helgasson H, 2012. Cognitive architectures and autonomy: A comparative review. Journal of Artificial General Intelligence, 3(2), 1-30.
Totschnig W, 2019. The problem of superintelligence: Political, not technological. AI & Society, 34, 907-920.
Tucker P, 2018. Google DeepMind researchers join pledge not to work in lethal AI. Defense One, https://www.defenseone.com/technology/2018/07/google-deepmind-researchers-pledge-lethal-ai/149819

Turner A, 2019. Big Mother project history, 1983-2019. Big Mother AI, https://bigmother.ai/resources/2019-09-15%20SRL-MMM-BMAI%20history%20v04.pdf
TWiStartups, 2016. Vicarious co-founder Scott Phoenix on innovating in AI & the race to unlock the human brain to create artificial general intelligence, the last tech humans will invent. Medium, https://medium.com/@TWiStartups/vicarious-co-founder-scott-phoenix-on-innovating-in-ai-the-race-to-unlock-the-human-brain-to-9161264b0c09
Ustun V, Rosenbloom PS, Sajjadi S, Nuttall J, 2018. Controlling synthetic characters in simulations: A case for cognitive architectures and Sigma. Proceedings of I/ITSEC 2018.
Vanian J, 2017a. Apple's artificial intelligence guru talks about a sci-fi future. Fortune, http://fortune.com/2017/03/28/apple-artificial-intelligence
Vanian J, 2017b. Apple just got more public about its artificial intelligence plans. Fortune, http://fortune.com/2017/07/19/apple-artificial-intelligence-research-journal
Veness J, Ng KS, Hutter M, Uther W, Silver D, 2011. A Monte-Carlo AIXI approximation. Journal of Artificial Intelligence Research, 40(1), 95-142.
Vepstas L, 2014. What is consciousness? OpenCog, https://blog.opencog.org/2014/01/17/what-is-consciousness
Wallach W, Allen C, Franklin S, 2011. Consciousness and ethics: Artificially conscious moral agents. International Journal of Machine Consciousness, 3(1), 177-192.
Wallach W, Franklin S, Allen C, 2010. A conceptual and computational model of moral decision making in human and artificial agents. Topics in Cognitive Science, 2(3), 454-485.
Wang B, 2014. Quantum computing and new approaches to artificial intelligence could get the resources to achieve real breakthroughs in computing. Next Big Future, https://www.nextbigfuture.com/2014/03/quantum-computing-and-new-approaches-to.html
Wang P, 2012. Motivation management in AGI systems. In Bach J, Goertzel B, Ikle M (Eds.), Proceedings of AGI 2012, 5th International Conference on Artificial General Intelligence. Berlin: Springer, pp. 352-361.
Wang P, Li X, 2016. Different conceptions of learning: Function approximation vs. self-organization. In Steunebrink B, Wang P, Goertzel B (Eds.), Proceedings of AGI 2016, 9th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 140-149.
Wang P, Talanov M, Hammer P, 2016. The emotional mechanisms in NARS. In Steunebrink B, Wang P, Goertzel B (Eds.), Proceedings of AGI 2016, 9th International Conference on Artificial General Intelligence. Cham, Switzerland: Springer, pp. 150-159.
Webb A, 2017. AI pioneer wants to build the renaissance machine of the future. Bloomberg, https://www.bloomberg.com/news/articles/2017-01-16/ai-pioneer-wants-to-build-the-renaissance-machine-of-the-future
Webster G, Creemers R, Triolo P, Kania E, 2017. China's plan to 'lead' in AI: Purpose, prospects, and problems. New America, https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems
Weng J, Evans CH, Hwang WS, Lee Y-B, 1999. The developmental approach to artificial intelligence: Concepts, developmental algorithms and experimental results. Proceedings of NSF Design and Manufacturing Grantees Conference.
Wiggers K, 2020. Goole Brain and DeepMind researchers attack reinforcement learning efficiency. Venture Beat, https://venturebeat.com/2020/02/18/google-brain-deepmind-reinforcement-learning-efficiency
Wilson G, 2013. Minimizing global catastrophic and existential risks from emerging technologies through international law. Virginia Environmental Law Journal, 31, 307-364.
Wirzberger M, Borst JP, Krems JF, Rey GD, 2020. Memory-related cognitive load effects in an interrupted learning task: A model-based explanation. Trends in Neuroscience and Education, 20, 100139.

Wissner-Gross AD, Freer CE, 2013. Causal entropic forces. Physical Review Letters, 110(16), 168702.
Yampolskiy RV, 2013. Artificial intelligence safety engineering: Why machine ethics is a wrong approach. In Müller VC (Ed), Philosophy and Theory of Artificial Intelligence. Berlin: Springer, pp. 389-396.
Yampolskiy R, Fox J, 2013. Safety engineering for artificial general intelligence. Topoi, 32(2), 217-226.
Yang W, Liu W, Viña A, Tuanmu MN, He G, Dietz T, Liu J, 2013. Nonlinear effects of group size on collective action and resource outcomes. Proceedings of the National Academy of Sciences, 110(27), 10916-10921.
Yuan Y, 2019. Tencent AI 'Juewu' beats top MOBA gamers. Synced, https://syncedreview.com/2019/12/30/tencent-ai-juewu-beats-top-moba-gamers
Yuan Y, 2020. Ex-Uber AI chief scientist Zoubin Ghahramani joins Google Brain leadership team. Synced, https://syncedreview.com/2020/09/13/ex-uber-ai-chief-scientist-zoubin-ghahramani-joins-google-brain-leadership-team
Yudkowsky E, 2004. Coherent Extrapolated Volition. San Francisco: The Singularity Institute.
Zhang Q, Walsh MM, Anderson JR, 2016. The effects of probe similarity on retrieval and comparison processes in associative recognition. Journal of Cognitive Neuroscience, 29(2), 352-367.