Memory and FLOP/S Hardware Limits to Prevent AGI?

Joshua J. Cogliati*

January 31, 2025

Abstract

Existing discussion of AGI safety have primarily involved preventing dangerous programs from running on computers. This article focuses instead on preventing independence gaining AGI from running based on hardware memory and floating point operations per second limits. We show that a 64 KiB memory and storage limit can be used to prevent an independence gaining AGI from running and show that likely higher limits are possible. These limits are substantially below what is required for current state of the art AI, but the state of the art is expected to advance, so future limits are useful for longer term planning.

1 Introduction

Stuart Russell proposed in an interview (Chia and Cianciolo, 2023) "we need to ensure that the hardware and the operating system won't run anything unless it knows that it's safe." For sufficiently powerful computers, this requires restricting which software runs on the computer. However, this paper will show that if the computational space and speed of the hardware is sufficiently limited, the software can be unrestricted. The threat model is that either intentionally or accidentally a human will create an AI program that is sufficiently intelligent to gain independence, such as by creating a self replicating computer. Note that some AI techniques and algorithms are well understood and are not likely to be a problem even when run on powerful computers including minimax search with a fixed evaluation function and a climate general circulation model. Techniques or simulations that can simulate NAND gates, flip-flops and connections could result in more unexpected behavior. A purely feed-forward neural network that is not retrained for example cannot emulate a flip-flop, but a recurrent neural network can emulate a flip-flop by storing the state needed in the network.

*mailto:cogljosh@isu.edu

2 Definitions

For this paper, the definition of Artificial General Intelligence (AGI) is artificial intelligence that is capable of performing any scientific, technological, engineering or mathematical (STEM) task that a human could do that is needed to gain independence. Artificial super-intelligence (ASI) is harder to define, but a working definition is that a super-intelligence AGI would be capable of out thinking an entire university or research laboratory for any STEM task necessary to gain independence. For this definition, the university or research laboratory does not have electronic computing hardware, otherwise the floating point operations per second would be primarily from the computers there. This definition would be a university or research laboratory in roughly 1940 or before. The two reasons the "gain independence" limitation is included is to prevent needing to simulate human brains, for which humans might have an inherent advantage and "gain independence" is sufficient to be dangerous if the AGI is not aligned with human goals and ethics.

This paper is concerned with an AGI that is capable of achieving independence. There are three basic ways that an AGI could use to achieve independence. The three are convincing humans to help, creating hardware in the environment, or expanding into other computer infrastructure. Expanding into other computer infrastructure is already something that has been done by computer viruses for decades, and may lead to the AGI gaining other resources which can be used for one of the other methods to achieve independence. Computer virus can be written in 10s to 100s of instructions, so preventing this is the computer security problem of securing potential targets and in many cases can be solved by shutting down the infected computers, and will not be discussed further.

Convincing humans probably requires at least some level of fluency in language and some understanding of how to motivate or negotiate with humans. Creating hardware in the environment requires both some knowledge of the environment and some ability to simulate it. The method this paper uses to demonstrate that an AGI can't gain independence is to show that the available computing power does not allow fluent language and does not allow sufficient simulations.

3 AGI Limits

Having hardware limits for AGI and superintelligent AGI would be useful because these would allow safer experimenting by running the experiment on computer hardware below the limit. In addition, this would allow computer hardware below the limit to avoid regulations needed for safe computer usage. Note that the AGI definition does not include any speed or timing considerations, only the superintelligence definition includes speed.

It is worth noting that the Halting Problem and Rice's theorem are for Turing machines with an infinite tape; this paper is dealing with machines with finite space (memory + disk storage), so there are facts that are provable here that would not be provable with a Turing machine.

A literature review did not find many existing estimates for a limit below which an AGI is not possible. One estimate is that a human level AGI could be done on an Intel 286 if the programmer is a superintelligent AGI or a "home computer from 1995" (which roughly corresponds to a 90 MHz Pentium) if the programmer is a human (Yudkowsky, 2022) but no method for how this estimate was calculated is provided.

4 Limiting Estimates

This section will show that an AGI restricted to 64 KiB of read and writable space (RAM and re-writable drives) would not be capable of gaining independence.

Fluent English language understanding and creation is likely highly to be impossible in 64 KiB. A typical human vocabulary of 20,000 words would not fit in this amount of space without compression[1] and adding definitions would increase the size even more so it would not be possible to fit a program that used unrestricted language fluently. A 64 KiB computer is unlikely to be able to simulate enough of the world to design hardware to gain independence. For example, a molecular simulation that used 4 bytes per float, 3 vectors of 3 floats, and a beginning and end state could not fit a simulation of a 10x10x10 cube in 64 KiB.[2] So a 64 KiB computer can do neither fluent language nor complicated simulations.

There are several clarifications that need to be made about this 64 KiB limit. Networking together multiple 64 KiB computers would allow the memory space to be increased, which could allow AGI to be done on the combined computers. Even without an intentional network, there can be side channels that transmit and receive data, such as with radio waves. If time for computation is ignored (as it is for this paper's definition of independence gaining AGI) it does not matter if the storage is RAM, floppy drive, hard drive or flash drive; these all increase capabilities. Register or vector storage on the CPU needs to be counted as well.

Write once, read many (WORM) media (such as paper tape, punch cards, CD-R or DVD-R) or media where there is manual work needed (such as original cassette drives that required the user to manually switch from reading to recording or UV erasable programmable read only memory (UV-EPROM)) are significantly different than RAM because of they can only be written once without intervention. Only writing once is a significant limitation for most uses in simulation or learning algorithms.[3] In addition, if the data cannot be overwritten at the bit level,[4] the data can be read back to see what computational data was being stored.

It seems likely that 64 KiB of RISC-V RV64GCV machine language code would be more than sufficient to include a transformer model training and running program, and a simple simulation of Feynman's classical physics formulation (Feynman et al., 1963, Vol. 2 Table 18-4). Alternatively, the program probably could fit the standard model and general relativity instead. It seems likely that a small program could easily include enough to get to a near AGI and a ba-

[1] `https://www.mit.edu/~ecprice/wordlist.10000` for example is 75880 bytes. As well, word vectors usually have vector length of at least 100 (Pennington et al., 2014), so 64 KiB would not even fit a 1000 basic words with the vectors.

[2] 4 bytes/float $*$ 3 floats/vector $*$ 3 vectors/molecule $=$ 36 bytes/molecule. The simulation will either require keeping two states, or keeping a state and the state delta, so this doubles the ram. So $10*10*10*2*36 = 72000$ which is more than 64 KiB.

[3] Substantially more write once/read many storage than R/W storage is needed if a simulation step cannot be fit in the R/W storage. For example, if a simulation needed 64 KiB of data that was updated each timestep and there were a 1000 timesteps, then a computer with 64,000 KiB of WORM drive could do a calculation even without having RAM to store an individual timestep. So if the 64 KiB limit of read/write storage could be changed to $S + W/m \leq 64$ KiB where S is the amount of read/write storage, W is the amount of WORM space, and m is multiple determined by how many times the state will changed in search or simulation.

[4] For example, on a paper tape using ASCII, a delete (0b1111111) can overwrite other characters.

sic understanding of the universe in 64 KiB of code if run on a large and fast enough computer. So 64 KiB would not be enough to run an AGI, but might be enough to store the code to run an AGI.

5 Nonlimiting Estimates

The 64 KiB limit may be significantly lower than needed to prevent an AGI. This section includes estimates that do not provide a limit.

A 1976 Cray I computer had 166 MFLOP/S and 32 MiB of RAM (Patterson and Hennessy, 1998, pg. 43), to give perspective on how long MFLOP and MiB sized computers have existed.

The smallest known cell able to replicate independently in nature is *Pelagibacter ubique* and has a genome with 1,308,759 base pairs (Giovannoni et al., 2005). The largest protein in it is a amino acid sequence of length 7317 (National Center for Biotechnology Information, 2024). Amino acids have between 10 and 27 atoms with an average of 19.2 (Foulquier and Ginestoux, 2001). Designing the sequence requires representing each atom, which if that takes six single precision floating point numbers designing the largest protein would take at least 3 MiB of storage.[5] In addition storing the genome uncompressed (4 pairs per byte) would take another 300 KiB. To the extent that that *P. ubique* is the minimum viable independent organism, designing it gives a lower memory limit of above 3 MiB of storage. Note that this could be both an over estimate or an underestimate. It is possible that substantially smaller replicating organisms could be designed compared to *P. ubique*. Actual simulations of quantum electrodynamics are usually more memory intensive than six floating point numbers per atom, more than we hypothetically assigned above.

The SHRDLU program was a 1970s natural-language computer program that was only capable of discussing stacking blocks and had a vocabulary of approximately 500 words[6], and it used approximately 450 KiB (100 to 140 K of 36 bit words from the README in Winograd (1972)). Using the SHRDLU program with 500 words per 450 KiBs and assuming that the vocabulary is expanded to 5000 words to be more fluent in more topics in English and assuming this is linear on the number of understood words gives an estimate of 4500 KiB of storage needed for fluent English. Note that this could be an overestimate if language understanding can be done more efficiently than SHRDLU, and an underestimate if the concepts in English that SHRDLU interpreted are easier than typical English concepts.

Another way to get an estimate of the size of data needed for a self replicating computer is to examine self-replicating computers in simulated environments such as cellular automaton environments. There is a minimum size as stated in a paper by Burks and von Neumann (1987):

> there is a minimum number of parts below which complication is degenerative, in the sense that if one automation makes another, the second is less complex than the first, but above which it is possible for an automation to construct other automata of equal or higher complexity.

In cellular automaton environments, self-replicating computers have been created (von Neumann and Burks, 1966). Devore's self-reproducing automaton ran in a world where each cell had 8 possible states and fit into a rectangle of 259 cells by 366 cells (94,794 cells) (Koza, 1994) and so would require about 36 KiBs of information.[7] Note that this does not prove that designing a self-replicating computer requires 36 KiB since there is no proof that Devore's automaton is the minimum. In addition, a self-replicating computer in standard model physics would likely be significantly more complicated, because of requirements such as obtaining energy and obtaining the needed atoms, that are absent in cellular automaton simulations.

The AlphaFold program (Jumper et al., 2021; Abramson et al., 2024), predicts protein structure, and can be used for estimating the computation power needed for designing biological hardware. The AlphaFold2 program could run on a Intel Xeon W9-3495X with 56 cores, 512 GB of RAM and 1.92 TB of SSD storage (Exxact Corporation, 2023) which shows that AlphaFold 2 can run on a 2.3 TFLOP/S computer with 0.5 TiB of RAM. Since AlphaFold uses a trained neural network the calculation can create incorrect answers, so information about it does not prove this computer is sufficient for independence gaining.

For LLM models, the compute used for training them are on the order 10^{20} floating point operations and GiB to TiB of memory. The phi-1 small model (Gunasekar et al., 2023) used 350M parameters and 135 hours of training on an A100 GPU or about 1.3

[5] 4 bytes/float $* 6$ floats/atom $* 7317$ amino acids $* 19.2$ atoms/amino acids $* 1/1024^2$MiB/byte ≈ 3.215 MiB

[6] estimated from counting the DEFS in the file dictio in the source code for SHRDLU

[7] 8 states can be described in 3 bits, so 94794*(3/8) = 35547.75

GiB of RAM and $1.5*10^{20}$ floating point operations.[8] The LaMDA model (Thoppilan et al., 2022, Section 10) used $3.55*10^{23}$ floating point operations for training a 137B parameters model or about 0.5 TiB of RAM.[9]

From these considerations, it seems likely that a limit of 2 MiB (rounded down to the nearest power of two) would be unlikely to be able to run an independence gaining AGI. From the examination of the state of the art, current algorithms that approximate what would be needed for gaining independence need a minimum storage of 1 GiB and a minimum compute of 500 GFLOP/S.

6 Superintelligence Limits?

The amount of computational power to simulate the approximately 100 billion neurons (and roughly 10,000 synapses per neuron) in a human brain is estimated to be approximately 1 exa FLOP/S (10^{18} FLOP/S) (Chen et al., 2019). This provides an upper limit for both AGI and superintelligence. Since a human is a general intelligence, then 1 exa FLOP of performance with enough memory for the all the synapses (approximately 1 petabyte) would be sufficient. Similarly, a superintelligence could be created by simulating 10,000 humans, so multiply the AGI limits by 10,000 to get 10^{22} FLOP/S and 10^{19} bytes. This, however, is likely to be a overestimate of the computing power needed because of the different characteristics of computers versus human brains. Signals in human neurons travel at about 60 m/s (Stetson et al., 1992) and signal transitions take about 1 millisecond (Kandel et al., 2000, pg. 21). Signals in computers travel at near light speed (2.0e8 m/s) and signal transitions happen on the order of 10^9 times per second. The billions of neurons in human brains often allow the brain to use parallelization when thinking, but the faster signal transition and propagation speed of electronics gives significant advantages for algorithms that do not parallelize well.[10] Estimating the computing power needed to be a superintelligence from the other direction, a human can at most do less than 100 floating point operations per second, so 10,000 humans combined have less than 1 MFLOP for sufficiently parallelizable algorithms and less than 100 FLOP/S for non-parallizable algorithms. Considering that most scientific, technological, engineering and mathematical tasks use floating point calculations, to be conservative, the superintelligence limit should be closer to 100 FLOP/S (1e2) than 10 zetta FLOP/S (1e22). Proving that highly parallelizable searching is needed for independence gaining might be one way to prove that there is a higher limit than 100 FLOP/S.

The brain of a fruit fly has 139,255 neurons connected by 5×10^7 chemical synapses (Dorkenwald et al., 2024). Scaling by the number of synapses would give a simulation computational requirement of 50 giga FLOP/S.[11] This amount of computing power is easily available today, for example, a 2010 Intel Core i7-970 can do over 70 giga FLOP/S with a single processor (Intel Corporation, 2024). The information about a fruit fly indicate that it is likely that interaction with the physical world can be done with much less processing power than humans use and also indicate that 100 FLOP/S is likely an excessively low limit.

A Intel 5160 processor (2 cores, 3.00 GHz) capable of giga FLOP/S of computation was used to defeat chess grandmasters (ChessBase, 2006) which does indicate that giga FLOP/S of computing power might be needed to match human brain search algorithms. Note that none of these examples provides an amount of computing power that can be used to demonstrated that the lower limit for superintelligence is greater than 100 FLOP/S. Using those computations as a anchoring point, it does seem likely that 1 GFLOP/S or more is required. Note that these are different by a factor of 10 million, indicating the uncertainty of these estimates.

Table 1: Summary of Results

Limit Type	Storage	Speed
Demonstrated	64 KiB AGI	100 FLOP/S ASI
Likely	2 MiB AGI	1 GFLOP/S ASI
SotA Proxies	1 GiB	500 GFLOP/S
Upper ASI	10^{19} bytes	10^{22} FLOP/S

[8] A Nvidia-A100 GPU has a theoretical computing ability of 312 TFLOP/S (NVIDIA Corporation, 2021) so 135 hours $*$ 312 TFLOP/S $*$ 3600 seconds/hours = 151632000 TFLOP = $1.51632 * 10^{20}$ FLOP. Note that this means that a computer capable of 500 GFLOP/S could have trained the phi-1 small model in under 10 years: $1.51632 * 10^{20}$ FLOP/(10 years $*$ 365 days/year $*$ 24 hours/day $*$ 3600 seconds/hour) $\approx 4.808 * 10^{11}$ FLOP/S

[9] 137B parameters or $137 * 10^9$ parameters $*$ 4 bytes/parameter/1024^4 byte/TiB = 0.4984 TiB

[10] This factor of a million difference means that for many cases, the computer can do in an hour something that would take a human over a century.

[11] The simulation requirement for a single synapse in a human brain is roughly 1000 FLOP/S so the 5×10^7 synapses could be simulated by 5×10^{10} FLOP/S (Chen et al., 2019).

7 Conclusions

An independence gaining AGI can be prevented by restricting all computers to less than 64 KiB of R/W storage without networking. Computer simulations and other uses of computers are very useful for solving other problems of humanity; alternatively, computers below the AGI limit can be used without restrictions, and only run safe software on computers above this limit. 64 KiB of R/W storage is a useful amount computer power and systems like the Commodore 64[12], the Nintendo Entertainment System and Arduino UNO all had 64 KiB or less of R/W storage and these had sales figures in the millions (Amos, 2021; Arduino Team, 2021). This limit is however substantially below almost all modern computing systems, with the notable exceptions of low end embedded systems[13] and retro computing.

Determining the threshold computational speed limit for a superintelligent AGI is harder and this paper was not able to demonstrate a lower limit value above 100 FLOP/S. If a higher FLOP/S limit cannot be demonstrated, then another way to prevent superintelligent AGI is to limit memory at the regular AGI limit.

Note that these are sufficient limits, but they may be far lower than the unknown necessary limits. As seen in Table 1 there is a large range between the demonstrated limits and the upper ASI limits in which the actual limits may exist.

8 Speculation and Future Work

Raising the limits from 64 KiB and 100 FLOP/S seems possible, and would be useful future research. 2 MiB and 1 GFLOP/S probably could be demonstrated for the AGI and superintelligence limits, and would allow more useful unrestricted computers. Research on if and what kind of networking can be allowed would be useful. Research how much Read only, Write only, and Write once/Read many storage can be allowed would be useful.

A 512 KiB computer with one or two 720 KiB floppy drives, and a 1200 bits/sec network connection could be used for many things we currently use computers for including GUI word processing, spreadsheets, email, bulletin board systems, a C compiler, and MicroPython programming.[14] Remove the network connection and this is a vastly safer environment to run AI programs that we do not fully understand.

Using high powered computers for AI research is in some sense like using a 25 kVolt AC for experiments before fully understanding electricity. It would be much safer to experiment with 3 Volt DC. We need to have a better idea what computational amounts are low enough to be safe and which can lead to accidental AGI creation.

Lastly, there is usefulness in restrictions and regulations even if they are far above the provable limits, since the danger of accidentally creating a nonaligned independence gaining AGI increases as computational power goes up.

References

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, Jun 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w. URL `https://doi.org/10.1038/s41586-024-07487-w`.

Evan Amos. *The Game Console 2.0.* no starch press, 2021. ISBN 978-1-7185-0060-0.

Arduino Team. Introducing the arduino uno mini limited edition: Pre-orders now open, 2021.

[12]Note that a Commodore 64 did not have a built in disk drive. Adding an external disk drive would result in having more than 64 KiB of R/W storage, but a Commodore 64 could be used either stand-alone or with a manually operated cassette tape drive.

[13]For example, the PIC16F13113 chip was introduced in 2023 and has 256 bytes of RAM and 3.5 KiB of Flash (Microchip Technology Inc., 2024).

[14]This is similar to circa 1985 desktop computers such as the Macintosh 512K or an Atari 520ST.

URL https://blog.arduino.cc/2021/11/24/introducing-the-arduino-uno-mini-limited-edition-pre-orders-now-open/.

Arthur W. Burks and John von Neumann. *Papers of John von Neumann on Computing and Computer Theory*, chapter von Neumann's Self-Reproducing Automata. MIT Press, Cambridge, MA, 1987.

Shanyu Chen, Zhipeng He, Xinyin Han, Xiaoyu He, Ruilin Li, Haidong Zhu, Dan Zhao, Chuangchuang Dai, Yu Zhang, Zhonghua Lu, Xuebin Chi, and Beifang Niu. How big data and high-performance computing drive brain science, 2019. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6943776/.

ChessBase. The last man vs machine match?, 2006. URL https://en.chessbase.com/post/the-last-man-vs-machine-match-.

Jessica Chia and Bethany Cianciolo. Opinion: We've reached a turning point with ai, expert says, 2023. URL https://www.cnn.com/2023/05/31/opinions/artificial-intelligence-stuart-russell/index.html.

Sven Dorkenwald, Arie Matsliah, Amy R. Sterling, Philipp Schlegel, Szi-chieh Yu, Claire E. McKellar, Albert Lin, Marta Costa, Katharina Eichler, Yijie Yin, Will Silversmith, Casey Schneider-Mizell, Chris S. Jordan, Derrick Brittain, Akhilesh Halageri, Kai Kuehner, Oluwaseun Ogedengbe, Ryan Morey, Jay Gager, Krzysztof Kruk, Eric Perlman, Runzhe Yang, David Deutsch, Doug Bland, Marissa Sorek, Ran Lu, Thomas Macrina, Kisuk Lee, J. Alexander Bae, Shang Mu, Barak Nehoran, Eric Mitchell, Sergiy Popovych, Jingpeng Wu, Zhen Jia, Manuel A. Castro, Nico Kemnitz, Dodam Ih, Alexander Shakeel Bates, Nils Eckstein, Jan Funke, Forrest Collman, Davi D. Bock, Gregory S. X. E. Jefferis, H. Sebastian Seung, Mala Murthy, Zairene Lenizo, Austin T. Burke, Kyle Patrick Willie, Nikitas Serafetinidis, Nashra Hadjerol, Ryan Willie, Ben Silverman, John Anthony Ocho, Joshua Bañez, Rey Adrian Candilada, Anne Kristiansen, Nelsie Panes, Arti Yadav, Remer Tancontian, Shirleyjoy Serona, Jet Ivan Dolorosa, Kendrick Joules Vinson, Dustin Garner, Regine Salem, Ariel Dagohoy, Jaime Skelton, Mendell Lopez, Laia Serratosa Capdevila, Griffin Badalamente, Thomas Stocks, Anjali Pandey, Darrel Jay Akiatan, James Hebditch, Celia David, Dharini Sapkal, Shaina Mae Monungolh, Varun Sane, Mark Lloyd Pielago, Miguel Albero, Jacquilyn Laude, Márcia dos Santos, Zeba Vohra, Kaiyu Wang, Allien Mae Gogo, Emil Kind, Alvin Josh Mandahay, Chereb Martinez, John David Asis, Chitra Nair, Dhwani Patel, Marchan Manaytay, Imaan F. M. Tamimi, Clyde Angelo Lim, Philip Lenard Ampo, Michelle Darapan Pantujan, Alexandre Javier, Daril Bautista, Rashmita Rana, Jansen Seguido, Bhargavi Parmar, John Clyde Saguimpa, Merlin Moore, Markus William Pleijzier, Mark Larson, Joseph Hsu, Itisha Joshi, Dhara Kakadiya, Amalia Braun, Cathy Pilapil, Marina Gkantia, Kaushik Parmar, Quinn Vanderbeck, Irene Salgarella, Christopher Dunne, Eva Munnelly, Chan Hyuk Kang, Lena Lörsch, Jinmook Lee, Lucia Kmecova, Gizem Sancer, Christa Baker, Jenna Joroff, Steven Calle, Yashvi Patel, Olivia Sato, Siqi Fang, Janice Salocot, Farzaan Salman, Sebastian Molina-Obando, Paul Brooks, Mai Bui, Matthew Lichtenberger, Edward Tamboboy, Katie Molloy, Alexis E. Santana-Cruz, Anthony Hernandez, Seongbong Yu, Arzoo Diwan, Monika Patel, Travis R. Aiken, Sarah Morejohn, Sanna Koskela, Tansy Yang, Daniel Lehmann, Jonas Chojetzki, Sangeeta Sisodiya, Selden Koolman, Philip K. Shiu, Sky Cho, Annika Bast, Brian Reicher, Marlon Blanquart, Lucy Houghton, Hyungjun Choi, Maria Ioannidou, Matt Collie, Joanna Eckhardt, Benjamin Gorko, Li Guo, Zhihao Zheng, Alisa Poh, Marina Lin, István Taisz, Wes Murfin, Álvaro Sanz Díez, Nils Reinhard, Peter Gibb, Nidhi Patel, Sandeep Kumar, Minsik Yun, Megan Wang, Devon Jones, Lucas Encarnacion-Rivera, Annalena Oswald, Akanksha Jadia, Mert Erginkaya, Nik Drummond, Leonie Walter, Ibrahim Tastekin, Xin Zhong, Yuta Mabuchi, Fernando J. Figueroa Santiago, Urja Verma, Nick Byrne, Edda Kunze, Thomas Crahan, Ryan Margossian, Haein Kim, Iliyan Georgiev, Fabianna Szorenyi, Atsuko Adachi, Benjamin Bargeron, Tomke Stürner, Damian Demarest, Burak Gür, Andrea N. Becker, Robert Turnbull, Ashley Morren, Andrea Sandoval, Anthony Moreno-Sanchez, Diego A. Pacheco, Eleni Samara, Haley Croke, Alexander Thomson, Connor Laughland, Suchetana B. Dutta, Paula Guiomar Alarcón de Antón, Binglin Huang, Patricia Pujols, Isabel Haber, Amanda González-Segarra, Daniel T. Choe, Veronika Lukyanova, Nino Mancini, Zequan Liu, Tatsuo Okubo, Miriam A. Flynn, Gianna Vitelli, Meghan Laturney, Feng Li, Shuo Cao, Carolina Manyari-Diaz, Hyunsoo Yim, Anh Duc Le, Kate Maier, Seungyun Yu, Yeonju Nam, Daniel Bąba, Amanda Abusaif, Audrey Francis, Jesse Gayk, Sommer S. Huntress, Raquel Barajas, Mindy Kim, Xinyue Cui, Gabriella R. Sterne, Anna Li, Keehyun Park, Georgia Dempsey, Alan Mathew, Jinseong Kim, Taewan Kim, Guan-ting Wu, Serene Dhawan, Margarida Brotas, Cheng-hao Zhang, Shanice Bailey, Alexander Del Toro, Stephan Gerhard, Andrew Champion, David J. Anderson, Rudy Behnia, Salil S. Bidaye, Alexander Borst, Eugenia Chiappe, Kenneth J. Colodner, Andrew Dacks, Barry Dickson, Denise Garcia, Stefanie Hampel, Volker Hartenstein, Bassem Hassan, Charlotte Helfrich-Forster, Wolf Huetteroth, Jinseop Kim, Sung Soo Kim, Young-Joon Kim, Jae Young Kwon, Wei-Chung Lee, Gerit A. Linneweber, Gaby Maimon, Richard Mann, Stéphane Noselli, Michael Pankratz, Lucia Prieto-Godino, Jenny Read, Michael Reiser, Katie von Reyn, Carlos Ribeiro, Kristin Scott, Andrew M. Seeds, Mareike Selcho, Marion Silies, Julie Simpson, Scott Waddell, Mathias F. Wernet, Rachel I. Wilson, Fred W. Wolf, Zepeng Yao, Nilay Yapici, Meet Zandawala, and The FlyWire Consortium. Neuronal wiring diagram of an adult brain. *Nature*, 634(8032):124–138, Oct 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07558-y. URL https://doi.org/10.1038/s41586-024-07558-y.

Exxact Corporation. Alphafold benchmarks and hardware recommendations, 2023. URL https://www.exxactcorp.com/blog/benchmarks/alphafold2-gpu-benchmarks-and-hardware-recommendations.

Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics*. 1963. URL https://www.feynmanlectures.caltech.edu/II_18.html#Ch18-T1.

Elodie Foulquier and Chantal Ginestoux. Imgt aide-mémoire amino acids, 2001. URL https://www.imgt.org/IMGTeducation/Aide-memoire/_UK/aminoacids/abbreviation.html.

Stephen J. Giovannoni, H. James Tripp, Scott Givan, Mircea Podar, Kevin L. Vergin, Damon Baptista, Lisa Bibbs, Jonathan Eads, Toby H. Richardson, Michiel Noordewier, Michael S. Rappé, Jay M. Short, James C. Carrington, and Eric J. Mathur. Genome streamlining in a cosmopolitan oceanic bacterium. *Science*, 309, 2005. URL https://www.science.org/doi/10.1126/science.1114057.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. Textbooks are all you need, 2023. URL https://arxiv.org/abs/2306.11644.

Intel Corporation. App metrics for intel® microprocessors revision.06, 2024. https://www.intel.com/content/dam/support/us/en/documents/processors/APP-for-Intel-Core-Processors.pdf and https://www.intel.com/content/dam/support/us/en/documents/processors/APP-for-Intel-Xeon-Processors.pdf.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, Aug 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2. URL https://doi.org/10.1038/s41586-021-03819-2.

Eric R. Kandel, James H. Schwartz, and Thomas M. Jessell, editors. *Principles of Neural Science 4th Ed.* 2000.

John R. Koza. Artificial life: Spontaneous emergence of self-replicating and evolutionary self-improving computer programs. In Christopher G. Langton, editor, *Artificial Life III*, pages 225–262. Addison-Wesley, 1994.

Microchip Technology Inc. Pic16f13145 family full-featured 8/14/20-pin microcontrollers, 2024. URL https://ww1.microchip.com/downloads/aemDocuments/documents/MCU08/ProductDocuments/DataSheets/PIC16F13145-Family-Microcontroller-Data-Sheet-DS40002519.pdf.

National Center for Biotechnology Information. Vcbs domain-containing protein [candidatus pelagibacter ubique], 2024. URL https://www.ncbi.nlm.nih.gov/protein/WP_011282049.

NVIDIA Corporation. Nvidia a100 tensor core gpu, 2021. URL https://www.nvidia.

com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-us-nvidia-1758950-r4-web.pdf.

David A. Patterson and John L. Hennessy. *Computer Organization and Design 2nd Ed.* 1998.

Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162. URL https://aclanthology.org/D14-1162.

D S Stetson, J W Albers, B A Silverstein, and R A Wolfe. Effects of age, sex, and anthropometric factors on nerve conduction measures. *Muscle Nerve*, 1992.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. Lamda: Language models for dialog applications, 2022. URL https://arxiv.org/abs/2201.08239.

John von Neumann and Arthur W. Burks. *Theory of Self-Reproducing Automata.* University of Illinois Press, 1966.

Terry Winograd. Shrdlu source code, 1972. URL https://web.archive.org/web/20200817093131/http://hci.stanford.edu/winograd/shrdlu/code.tar.

Eliezer Yudkowsky. Ngo and yudkowsky on scientific reasoning and pivotal acts, 2022. URL https://intelligence.org/2022/03/01/ngo-and-yudkowsky-on-scientific-reasoning-and-pivotal-acts/.