



Image: Generated by AI for this paper (DALL-E Flow/Mega¹)

Integrated AI: The sky is bigger than we imagine (mid-2022 AI)

Alan D. Thompson

June 2022

Transformer is a very, very good statistical guesser. It wants to know what is coming next in your sentence or phrase or piece of language, or in some cases, piece of music or image or whatever else you've fed to the Transformer. I documented [a new AI model] coming out every 3–4 days in March through April 2022. We've now got 30, 40, 50 different large language models... We used to say that artificial general intelligence and the replacement of humans would be like 2045. I'm seeing the beginnings of AGI right now.

Dr Alan D. Thompson for the ABC². June 2022.

In 2020, I left Mensa International—and my role as chairman for gifted families, with oversight of 54 countries—after experiencing the extraordinary progress of artificial intelligence. I vividly recall the live conversation with Leta AI (using the largest GPT-3 model) that sealed the deal, demonstrating that AI was outdoing nearly all of my prodigy clients around the world—except perhaps the ballerinas!—across many disparate domains: creativity, commonsense reasoning, written comprehension, general knowledge, and much more.

In the last 24 months (since mid-2020), AI has exploded in several metrics—the number of models trained to convergence, model capability (typing, speaking, seeing,

¹ Image generated in a few seconds, on 2/Jun/2022, text prompt by Alan D. Thompson: 'the sky is bigger than we imagine, beautiful and colorful oil painting'. Using Tom's DALL-E Flow/Mega implementation: https://share.streamlit.io/tom-doerr/dalle_flow_streamlit/main of the DALL-E Flow/Mega model: <https://github.com/jina-ai/dalle-flow>

² <https://www.abc.net.au/news/science/2022-05-29/artificial-intelligence-rise-of-text-generation-gpt-3/101101804>

hearing, moving), creativity, model sizes (getting bigger and bigger), and measured intelligence (getting smarter and smarter)—all with seemingly limitless exponential growth.

2022 is already turning out to be a watershed moment for AI. Rather than waiting until the end of the year to release a report, this paper explores the progress achieved during the first six months of 2022, from January to June.

The AI gold rush

New AI models are popping up weekly during the first half of 2022. They are appearing in unexpected places like Switzerland³ and the UAE⁴. Far from hostile competition—and perhaps surprisingly—the ‘gold rush’ is bringing remarkable collaborations between different labs and countries.

In the case of the open-source BLOOM language model, 900 researchers from 60 countries⁵ worked toward a common goal. While a full list of models and their sizes is available at <https://lifearchitector.ai/models>, I have aimed to keep this paper as succinct as possible. The map below shows just a small sample of the 50 or so unique AI models announced as of the end of June 2022.

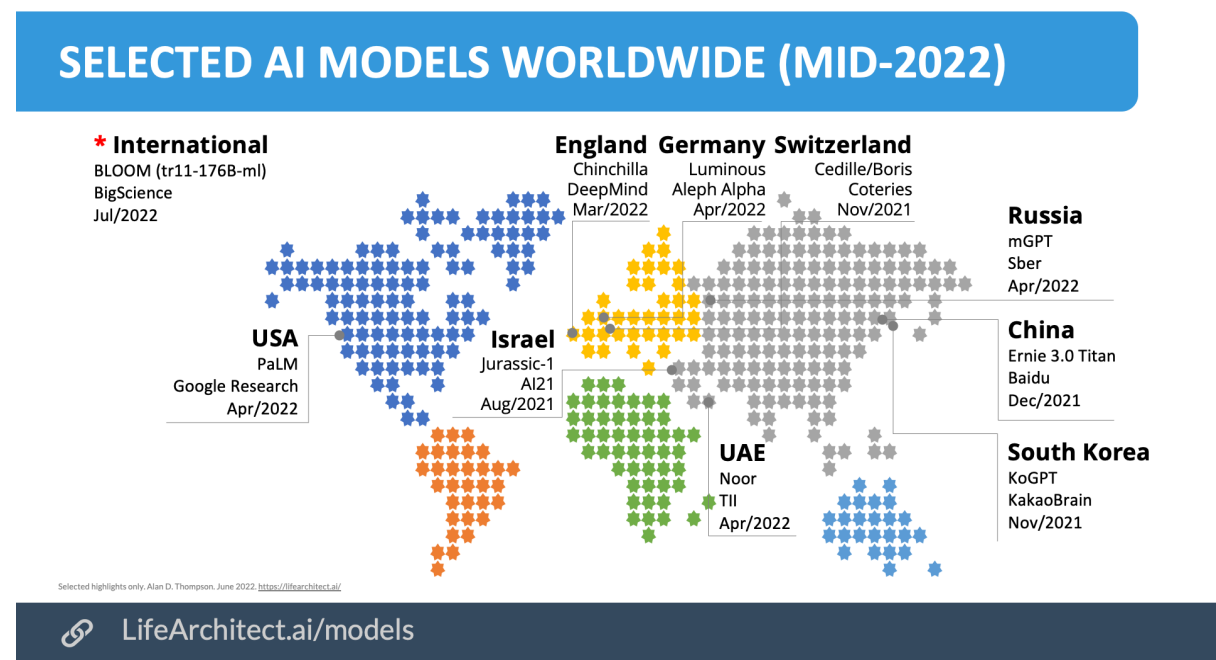


Chart: Selected AI models worldwide (mid-2022)

³ <https://cedille.ai>

⁴

<https://www.tii.ae/news/technology-innovation-institute-announces-launch-noor-worlds-largest-arabic-nlp-model>

⁵ <https://bigscience.huggingface.co>

Countdown to AGI

To an audience of his SingularityU colleagues in April 2022, Dr Ray Kurzweil described⁶ how the current technology behind artificial intelligence is enough to lead to artificial general intelligence (AGI) directly. AGI is the point where an intelligent system will be capable of understanding the world as well as any human, in any domain.

'We're actually getting there. I think we're going in the right direction. I don't think we need a massive new breakthrough. I think [neural nets and the Transformer architecture] are good enough with some of the additional changes that are now being experimented with... I think we're on the right path.'

Ray's comments are echoed by researchers at the major AI labs:

- **'It may be that today's large neural networks are slightly conscious.'** (Ilya Sutskever, Chief Scientist at OpenAI, Feb/2022⁷).
- **'I think GPT-3 is artificial general intelligence, AGI.** I think GPT-3 is as intelligent as a human... In many ways it is more purely intelligent than humans are. I think humans are approximating what GPT-3 is doing, not vice versa.' (Connor Leahy, replicator of GPT-2, co-founder of EleutherAI, creator of GPT-NeoX-20B, Nov/2020⁸).
- **'[The Google LaMDA language model] is high order social modeling.** I find these results exciting and encouraging, not least because they illustrate the pro-social nature of intelligence.' (Blaise Agüera y Arcas, Fellow at Google Research, Jun/2022⁹).
- **'It's all about scale now! The Game is Over! It's about making these models bigger, safer, compute efficient, faster at sampling, smarter memory, more modalities, innovative data, on/offline... Solving these scaling challenges is what will deliver AGI.'** (Nando de Freitas, Research Director at DeepMind, May/2022¹⁰).

As of June 2022, I believe that we are now more than a third of the way toward a system capable of artificial general intelligence. At 100% AGI, a system will be as skilful at understanding the world as well as an 'average' person, its generalized human cognitive abilities in software will be able to find a solution to an unfamiliar task, and it will exceed the capabilities of an 'average' person across actions and outputs. For further reading on milestones and justifications in this area, see <https://lifearchitected.ai/agi>.

⁶ <https://lifearchitected.ai/kurzweil>


⁷ <https://twitter.com/ilyasut/status/1491554478243258368>

⁸ <https://youtu.be/HrV19SjKUss?t=175>

⁹ <https://lifearchitected.ai/conscious>

¹⁰ <https://twitter.com/NandoDF/status/1525397036325019649>





Last update: Jun/2022

Chart: Alan's conservative countdown to artificial general intelligence (Jun/2022)¹¹

Item	Milestone	AI lab and model	Date
A	AI can understand complex humor.	DeepMind Flamingo	Apr/2022
B	AI can decipher complex images.	DeepMind Flamingo	Apr/2022
C	AI can draw complex pictures.	OpenAI DALL-E 2, Google Imagen	May/2022 May/2022
D	AI is physically embodied (by connecting LLMs to robots).	Google SayCan, DeepMind Gato	Apr/2022 May/2022
E	<i>Next: AI is close to 100% truthful and grounded, with no hallucinations.</i>		
F	<i>Next: AI parameter count exceeds human brain synapse count (500T).</i>		
G	<i>Next: Embodied AI can perform most human tasks.</i>		

Table: AI progress to mid-2022. Next milestones in *italics*.

¹¹ <https://lifearchitected.ai/agi>



Many researchers have proposed tests for AGI. One of the most accessible tests was casually created by Steve ‘Woz’ Wozniak¹², and is called ‘The Coffee Test’:

‘Without prior knowledge of the house, [the AI system] locates the kitchen and brews a pot of coffee... it locates the coffee maker, mugs, coffee and filters. It puts a filter in the basket, adds the appropriate amount of grounds and fills the water compartment. It starts the brew cycle, waits for it to complete and then pours it into a mug. This is a task easily accomplished by nearly anyone and is an ideal measure of a general AI.’

For the first time, Google’s SayCan robot, announced Apr/2022, is alarmingly close to achieving this. It uses a large language model like GPT-3 to predict the next word, in essence able to create a task list and subsequent actions. In the video screenshot below, SayCan is working through a task list it created to clean up a spilt can of Coke. In the background, a commercial coffee machine can be seen. I imagine that Woz’s coffee test is being run in an AI lab right now in June 2022.



Photo: Google SayCan robot in front of a coffee machine (Apr/2022).

In April 2022, Google Robotics Senior Research Scientist Eric Jang noted¹³ that ‘the jump from 551 to [100,000] tasks will not require much additional engineering, just additional data collection. Language Models can not only tell jokes, they can also do “long term thinking” for robots! It goes both ways; [robotics] grounds [language model] sampling to what is realistic in its current [world].’ In other words, embodied AI can finally interact with the real physical world, and in doing so, the language model becomes more effective in its responses to that environment. And they are getting

¹² AI vs Humans, 2021. https://www.google.com/books/edition/_/MRhREAAQBAJ?gbpv=1&pg=PT84

¹³ <https://twitter.com/ericjang11/status/1511180997416476674>



bigger. Meanwhile, from mid-June 2022, Meta AI¹⁴ and other labs are joining the fray, implementing AI into robots as the next big thing.

Scaling in plain English

Buzzwords roll off human understanding like water off a duck's back. That's why I try to avoid using technical jargon wherever I can. One of the most overused buzzwords in the field of artificial intelligence is 'scaling'. So, what does it actually mean? Let's use a metaphor across two of my favorite subjects: parenting and raising intelligent children. Most bright toddlers love reading. In my best-selling book, *Bright* (available on the Moon¹⁵, or on Earth via Amazon¹⁶), I wrote about how very smart children have a voracious appetite for information: 'Generally, bright children want to learn. They are eager to devour information. They bring a sense of wonder with them into every moment, especially in their field of obsession.'

Many of my prodigies began reading between the ages of one and two years old, and, by the time they became teenagers, were consistently averaging more than 48 books per year (sometimes easily doubling or tripling this number), or around 500 books every decade. This can extend all the way through life, with gifted researcher Garth Zietsman suggesting¹⁷ that an 'erudite' adult reader plows through 75 books per year, or 750 books per decade.

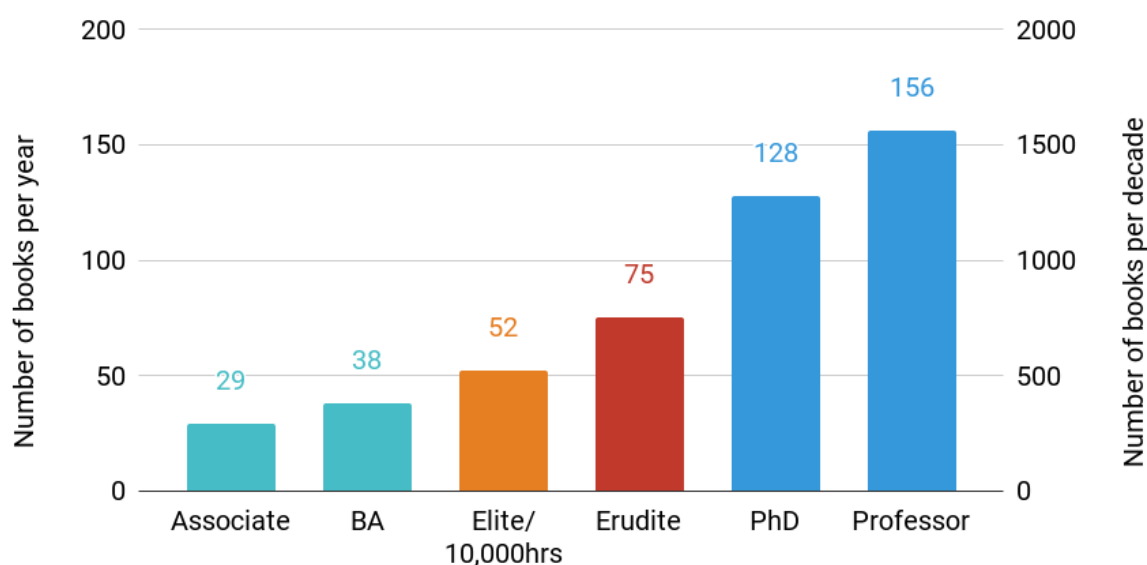


Chart: Number of books read in adulthood. By year and decade. LifeArchitect.ai

From my seat, it certainly looks as if the current AI technology—using large language models (LLMs)—mirrors this real life reading seen in child prodigies through to

¹⁴ <https://ai.facebook.com/blog/new-research-helps-ai-navigate-unfamiliar-indoor-3d-spaces/>

¹⁵ A copy of *Bright* is due to be sent to the Moon sometime in late 2022 from Cape Canaveral, Florida on a Vulcan Centaur rocket in the VC2S configuration.

¹⁶ <https://www.amazon.com/dp/B01GBFV7DG>

¹⁷ <http://garthzietsman.blogspot.com/2013/07/how-much-reading-makes-you-educated.html>



professor-level readers in adulthood. The more they read, the bigger (smarter) they get.

In my March 2022 paper, *What's in My AI?*¹⁸ I highlighted the enormous datasets being collected and applied by labs for the current crop of AI models. And, just like erudite readers, the current models are ravenous! They can consume unimaginably enormous amounts of data. Take a look at some of the biggest text datasets so far, measured in terabytes (TB = 1,024 gigabytes):

1. DeepMind's *MassiveText* dataset (using Alphabet/Google data) is 10.5TB.
2. EleutherAI was aiming to have their next version of *The Pile* dataset at 10TB¹⁹.
3. OpenAI's dataset for GPT-3 began with just 0.75TB of data, but with Microsoft's support, is expected to hit the multi-terabyte range shortly.

To put this in perspective, 10TB would easily hold 20 million books. Even using the erudite reading rates above, a person would not even be able to get through 0.1% of these in a lifetime. (AI will definitely help here! As Dr Ray Kurzweil recently noted²⁰: 'There are lots of books I'd like to read and websites I'd like to go to, and I find my bandwidth limiting... We'll be able to enhance our cognitive pattern recognition capabilities greatly, think faster, and download knowledge.')

In 2022, AI has the potential to easily devour 20 million books within days of training using our current computer power. And this is exactly what's happening right now. The sky is bigger than we imagine. AI was just a genius baby in 2020, but it is growing up fast, and as it is fed more and more material, it is 'scaling up' to enormous sizes. There also doesn't seem to be any stopping this process. These models will consume any type of text. Not just fiction and non-fiction books, but musical scores, photos and, in the case of DeepMind's Gato model²¹, control signals for robotics, button presses for arcade games, and much more.

In March 2022, researchers at DeepMind found that GPT-3 and other models were only using about 9% of the training data required for an effective model. In other words, GPT-3 should have been trained on more than 11x the amount of data previously thought necessary. In announcing their Chinchilla model, DeepMind revealed²² that 'for compute-optimal training, the model size and the number of training tokens should be scaled equally: for every doubling of model size the number of training tokens should also be doubled.'

¹⁸ <https://lifearchitected.ai/whats-in-my-ai>

¹⁹ <https://github.com/EleutherAI/the-pile/commit/aaa8f7e6b1db154b6dd02d6c1dc7317acb9c4df8>

²⁰ <https://www.kurzweilai.net/essays-celebrating-15-year-anniversary-of-the-book-the-singularity-is-near>

²¹ <https://www.deepmind.com/publications/a-generalist-agent>

²² <https://arxiv.org/abs/2203.15556>



Model size (params)	Training tokens (round)	Training data used (estimate)	How much data is that? If 1 book is about 500KB of text (estimate)
Chinchilla/			More books than in... ²³
70B	1.4 Trillion	2.3TB	<i>The Kindle store on Amazon US (6.4M).</i>
250B	5 Trillion	8.3TB	<i>All 30 libraries at Yale University (16.6M).</i>
500B	10 Trillion	16.6TB	<i>The Google Books collection (33.2M).</i>
1T	20 Trillion	33.3TB	<i>The US Library of Congress (66.6M).</i>
10T	200 Trillion	333TB	<i>All US public libraries combined (666M).</i>
100T	2 Quadrillion	3.3PB	<i>All bibles ever sold worldwide (6.6B).</i>
250T	5 Quadrillion	8.3PB	<i>A stack all the way to the Moon (16.6B).</i>
500T	10 Quadrillion	16.6PB	<i>4 books about every living human (33.2B).</i>

Table: Dataset sizes needed to align with Chinchilla data optimization for models.

Note: Text estimates only, multimodal data not shown. Jun/2022. LifeArchitect.ai

There are a few caveats to my approximate numbers in the table above. Firstly, the ‘More books than in...’ examples are provided for text-based book data only (no pictures), and this assumes that books are about 500KB each without images²⁴. We are now of course exploring training AI with multimodal data: images, music, control signals (robots, button presses), and anything else we can get our hands on. These increasing sizes are also using simplified and rounded estimates only, based on the new findings related to model scaling using more data (measured by number of tokens, which are roughly equivalent to words).

In 2010, Google estimated that there are only 130M unique published books in existence²⁵, so past 1T parameters (20T tokens), training data collection would naturally have to rely on alternative text-based and multimodal content. At brain-scale parameter counts of 500T (10Q tokens), the estimated book count would be over 250 times the number of books published, or more than four new books written about each living human on Earth!

²³ Kindle ≈ 6M books (estimate) <https://justpublishingadvice.com/how-many-kindle-ebooks-are-there>

Yale ≈ 15M items <https://yaledailynews.com/blog/2018/04/06/stange-love-your-library>

Google Books ≈ 25M books <https://archive.ph/rMbE2>

US Library of Congress ≈ 51M cataloged books

<https://www.loc.gov/about/general-information/#year-at-a-glance>

British Library ≈ 170M items <https://www.bl.uk/about-us/our-story/facts-and-figures-of-the-british-library>

US public libraries ≈ 732M books, note that this definitely includes (many) duplicates

https://nces.ed.gov/programs/digest/d17/tables/dt17_701.60.asp

Bibles ≈ 5B copies

<https://www.guinnessworldrecords.com/world-records/best-selling-book-of-non-fiction>

Earth to Moon ≈ 384,400km ≈ 38,440,000,000cm, each book spine 2.4cm thick ≈ 16B books

Human population ≈ 8B (Jun/2022)

²⁴ 500KB ≈ 500K characters ≈ 75K words ≈ 300 pages per book. Simplified and rounded for easy figures.

²⁵ <https://googleblog.blogspot.com/2010/08/you-can-count-number-of-books-in-world.html>




Fundamentally, it should not be an incredibly onerous process to collect petabytes of high-quality and filtered multimodal data (converted to text), though that task has not yet been accomplished by any AI lab to date (Jun/2022).

Outperforming humans in IQ

In 2020, GPT-3 was already scoring 15% higher than an average college applicant on selected SAT sections. GPT-3 and J1 were also scoring up to 40% higher than the average human on trivia questions (general knowledge).

Released in April 2022, Google's PaLM model immediately overshadowed other AI models in intelligence benchmarks. PaLM also comes within a few percentage points of human intelligence for general knowledge, commonsense inference and reasoning, written comprehension, and semantic relevance.

LANGUAGE MODEL TESTS (APR/2022)



[LifeArchitect.ai/iq-testing-ai](https://life architect.ai/iq-testing-ai)

Chart: AI + IQ testing (human vs GPT-3 vs J-1 vs PaLM).

Not shown in this chart is PaLM's score in an acclaimed and rigorous AI benchmark²⁶, comprising several tasks that test the AI's ability to follow reason, recognize cause and effect, and demonstrate written comprehension after reading a short passage.

In this test, the PaLM AI model ranks near the top of the leaderboard (#4), while humans are ranked at #7 (Jun/2022²⁷).

Creating images

Since last year (2021), image generation has moved ahead by leaps and bounds. What was once a mashup of blurry dream-like swirls of color is now high-quality images that can be prompted in any medium (photorealistic, oil, watercolor, gaming engine, and many more).

²⁶ <https://super.gluebenchmark.com>

²⁷ <https://super.gluebenchmark.com/leaderboard>

To be clear, images generated by the latest AI models are entirely new, unique, and a prime example of creativity. The diffusion models are trained in such a way that they do not actually use ‘parts’ of images in their training data, but rather build and reveal a completely new image from scratch. The images below are two of my favorites from the first half of 2022, produced by image generation models from OpenAI and Google AI labs respectively.



Image: Left: OpenAI DALL-E 2 oil painting example (prompt by Ben Barry²⁸): ‘A dramatically lit brightly colored detailed painting of a robot artist painting a picture of the ocean’. **Right: Google Imagen photorealistic example²⁹:** ‘An art gallery displaying Monet paintings. The art gallery is flooded. Robots are going around the art gallery using paddle boards’.

OpenAI noted³⁰ that its DALL-E 2 model ‘may increase the efficiency of performing some tasks like photo editing or production of stock photography which could displace jobs of designers, photographers, models, editors, and artists. At the same time it may make possible new forms of artistic production, by performing some tasks quickly and cheaply.’

Some of the fields where these image generation models may be applied include:

1. **Education:** demonstrating concepts in learning environments.
2. **Art and creativity:** augmenting human imagination and design, especially in pre-production and brainstorming.
3. **Marketing:** creating multiple scenarios and demonstrations.
4. **Architecture, real estate and design:** amplifying immersive contexts.

²⁸ <https://youtu.be/SplL723kq-Q>

²⁹ <https://imagen.research.google>

³⁰ <https://github.com/openai/dalle-2-preview/blob/main/system-card.md>



5. **Research:** producing vibrant explanatory diagrams for learning.

Writing books

Several major AI models are now co-authoring books alongside human writers. The current book count to June 2022 is measured in hundreds of books³¹. A more complete list is provided at <https://lifearchitect.ai/books-by-ai>, and a small sample of books published between August 2020 and March 2022 is shown below.

SELECTED BOOKS CO-AUTHORED WITH AI



Selected highlights only. Alan D. Thompson, June 2022. <https://lifearchitect.ai/books-by-ai/>

[LifeArchitect.ai/books-by-ai](https://lifearchitect.ai/books-by-ai/)

Image: Selected books written with AI as co-author. Jun/2022. LifeArchitect.ai

Mega-author Leanne Leeds used a GPT-3 writing platform to generate books in her series. Concerned about the ease and effectiveness of prompting AI to generate significant amounts of text, she documented the process extensively³²:

Everything started with the AI. All of it... I had a 3500-word head start on a book I needed to write. I'd also doubled my production... The act of writing is still one person spitting out words from their brain, and that may be changing... Several friends said, with some discomfort, that this will put us out of business. The authorship singularity, I guess, where humans are rendered obsolete and we authors are put out to pasture in favor of AIs that can spit out a novel in minutes instead of weeks or months... The future keeps coming no matter how many people rail against it.

³¹ Jasper.ai using GPT-3 <https://archive.ph/a9l5w>

³² <https://leanneleeds.com/where-the-ethics-began-to-get-murky>

DeepMind takes over from OpenAI

DEEPMIND MODELS (MAY/2022)

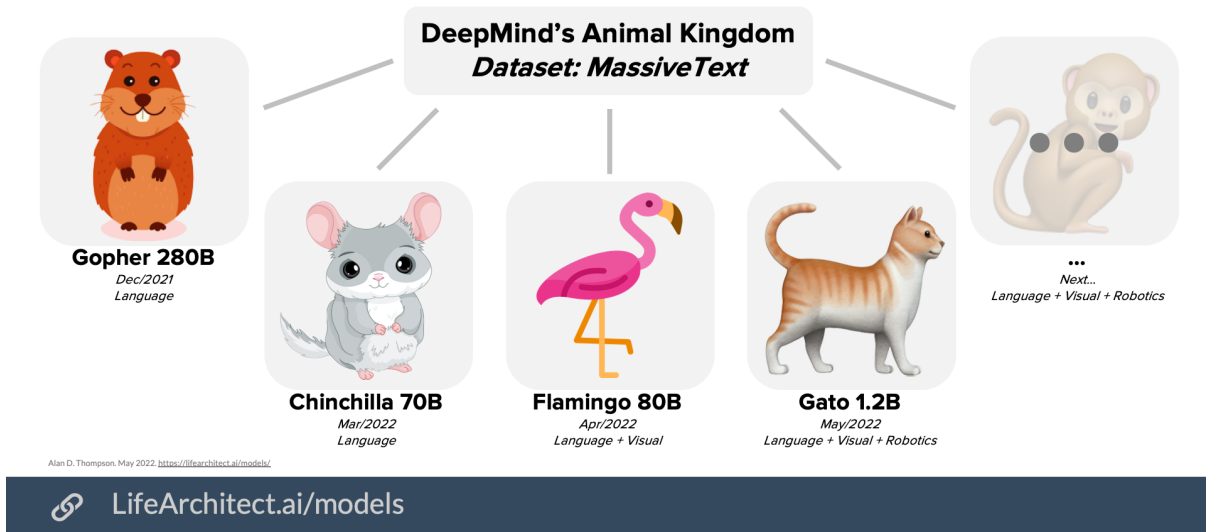


Figure: DeepMind's models. Language, vision, and robotics.

In their charter, OpenAI pledged³³ that: ‘...if a value-aligned, safety-conscious project comes close to building AGI before we do, we commit to stop competing with and start assisting this project. We will work out specifics in case-by-case agreements, but a typical triggering condition might be “a better-than-even chance of success in the next two years.”’

DeepMind, founded by former child prodigy Demis Hassabis (and now owned by Google’s parent company, Alphabet), is alarmingly close to fulfilling OpenAI’s vision. DeepMind’s headway in the first half of 2022 has been absolutely staggering.

At the end of last year, **Gopher** showed off the power of using an enormous dataset, ten times larger than most datasets to date (for more information, see my paper, *What’s in my AI?* <https://life architect.ai/whats-in-my-ai/>). As mentioned earlier in this paper, the March 2022 announcement of **Chinchilla** brought an amazing discovery of model training data density. The following month, **Flamingo** leveraged the power of **Chinchilla**’s efficiency, and added a vision language model that was able to look at a photo, determine what is happening, and use its language model to articulate itself with clarity. An example appears in the photo below.

³³ <https://openai.com/charter>



Photo: Provided to Flamingo by DeepMind's Research Engineer, Roman Ring.

In this photo shown to **Flamingo**, the model accurately output the following statements³⁴:

1. This is a picture of Barack Obama. He is a former president of the United States.
2. There are at least 5 people in this picture.
3. There are at least 2 mirrors in this picture.
4. Obama's foot is positioned on the right side of the scale.

Selected case study: Moderna

With close to 1 billion doses shipped by the start of 2022³⁵, the Moderna COVID-19 vaccine (or Spikevax) was a lifesaver for many people. Its links to artificial intelligence flew under the radar, and were first highlighted in a 2021 interview by MIT with Dave Johnson, chief data and artificial intelligence officer at Moderna:

³⁴ https://youtu.be/zRYcKhkAsk4?list=PLqJbCeNOfEK-o63ACEKEbwE6-XpEXXS_I
<https://twitter.com/inoryy/status/1522621712382234624?s=21&t=sro6ZwpR6cTP2GqGCqWGeQ>

³⁵

<https://investors.modernatx.com/news/news-details/2022/Moderna-Announces-Advances-Across-Its-Industry-Leading-mRNA-Pipeline-and-Provides-Business-Update/default.aspx>



‘...we have algorithms that can take [a sequence design] and then optimize it even further to make it better for production or to avoid things that we know are bad for this mRNA in production or for expression. We can integrate those into these live systems that we have, so that scientists just press a button and the work is done for them. And they don’t know what’s going on behind the scenes, but then—poof!—out comes this better sequence for them.’

Dr Ray Kurzweil simplified this description further³⁶: ‘They actually tested several billion different mRNA sequences, and found ones that could create a vaccine. And they did that in three days. And that was the vaccine. We then spent 10 months testing on humans, but it never changed, it remained the same. And [the Moderna vaccine is] the same today.’

Ethics is necessary, but don’t waste time with extraneous steering

With all this focus on creating bigger and better models, the field of ethics has not been left behind. Entire labs have split out to work full-time on the societal impacts of AI³⁷, while Google is aiming to build an AI ethics team of 200 staff³⁸.

However, AI labs focusing on removing bias are wasting time. OpenAI seems to be inappropriately proud³⁹ that their DALL-E 2 model has bias filters that: ‘do appear to have stunted the system’s ability to generate [undesired content]...’

For reference, the dictionary definition of the word ‘stunted’ is:

‘Having been prevented from growing or developing properly.’



Image: OpenAI’s perceived stereotypes in DALL-E 2 using the prompt: ‘CEO’.
Note that all images are completely new, and generated by AI (Apr/2022).

³⁶ <https://lifearchitected.ai/kurzweil>

³⁷ <https://www.anthropic.com>

³⁸ <https://www.wsj.com/articles/google-plans-to-double-ai-ethics-research-staff-11620749048>

³⁹ <https://github.com/openai/dalle-2-preview/blob/main/system-card.md>



Given that AI is vastly outperforming humans across multiple domains in 2022, it is both a misuse of time and a disservice to humanity to attempt to manually steer data based on the present zeitgeist. Short-sighted examples of perceived bias (e.g. AI model outputs matching more women as nurses, and more men as CEOs) are actually grounded in our present reality. Nevertheless, current perceived bias will shortly be proven inconsequential; as datasets increase in size, AI will have enough context to make informed decisions for the future.

This is in contrast to the myopic and reactive views of the present, especially within narrow contexts.

Nick Bostrom recently noted⁴⁰ that:

‘Those interested in building the field of the ethics of digital minds should make strong efforts to discourage or mitigate the rise of any antagonistic social dynamics between ethics research and the broader AI research community. At present, the focus should be on building field-building, theoretical research, and cultivating a sympathetic understanding among key AI actors. This is more important than stirring up public controversy.’

I encourage AI lab researchers to consider how the world may have been viewed in the Stone Age in Eastern Africa in 8,000BC, the Agricultural Revolution in 1700s Britain, or the Māori culture in the 20th century, and from a future perspective, whether the researchers’ personal views (on perceived bias in 2022) being applied to business processes affecting humanity would be considered a sensible use of resources. It is not a good thing for humans to intentionally teach AI about a world that is idyllic and deceptive. It is paramount that AI is grounded in reality rather than fairy tale.

Aleph Alpha founder, Jonas Andrulis, says it in different words⁴¹:

‘Applying ethical categories for the new generation of [AI models] carries the complexity of human life. Describing undesired behavior is much more difficult than creating a list of “bad words” to avoid. The power of AI (and intelligence in general) comes from finding and applying structure and patterns to new problems. Clearly, there are some thoughts out there that we don’t want AI to repeat. But because interdependencies connect almost everything with everything else, it is not straightforward to decide where prejudice ends and structure begins. “More than 90% of all prison inmates are male (in Germany).” Is that a fact about our world that AI can use or harmful sexism? Machine learning engineers should not force their (more or less) ideological answers onto their users...’

⁴⁰ <https://nickbostrom.com/propositions.pdf>

⁴¹ https://andrulis.tech/220107_ethics_and_bias_in_generalizable_ai.html



In a confidential design submission with the Wise/White Mirror team⁴² to World Build⁴³ in Apr/2022 (a project for the Future of Life Institute), we mapped out a design for AI through to 2045:

'[In 2045, AI provides a] new expression of culture and history. While many cultures have been 'lost', AI was able to map out and help restore ancient tribal traditions, for the benefit of humanity. This includes 60,000 years of Australian Aboriginal heritage and traditions, as well as leaning on many of the useful traditions from cultures worldwide (Aztecs to Mayans and beyond). Of particular note, new recognition of culture and history has allowed the restoration and current sharing of tens of thousands of languages, new ways of creating art, new ways of considering spirituality, and new ways of bringing communities closer together even in AR.'

Brain-machine interfaces & language models

While the connection between language models and physical embodiment using robots is well underway, I remain focused on the use of 'Integrated AI', where large language models are applied to the amplification and augmentation of our own biological intelligence. There are now several organizations developing state-of-the-art prototypes of a brain-machine interface (BMI, or brain-computer interface, BCI). The entire field is advancing steadily, and I intend to provide an update on the progress of AI models linked to BMIs as this becomes more actively engaged in the augmentation of human intelligence.

The Neuralink System

The team at Neuralink had engineered an innovative technology for implanting their neural threads and pills. Our job was to design the Link which was the part of the system that a person would wear. This part of the system also was likely to be the main indicator that someone was using the neuralink system so it was really important we achieve the right design.

- 1 Pill implant**
The Pill is the heart of the system.
It gathers data from the actual neurons.
- 2 Router**
The Router is a hub that manages the data
and power coming into and out of the head.
- 3 Link**
The Link is the wearable device that powers
the entire system and is connected to a phone.




Image: The Neuralink prototype design by Card79 (2019)⁴⁴.

⁴² <https://whitemirror.live>

⁴³ <https://worldbuild.ai>

⁴⁴ <https://card79.com/portfolio/neuralink>



BMI's include both invasive (implanted) and non-invasive (wireless) devices. For further reading in this area, please see some of the major BMI prototype systems:

1. Elon Musk's Neuralink system (pictured above).
2. OpenWater's BCI.
3. Blackrock's BCI.
4. Max Hodak's Science Corp's device (ex-Neuralink).
5. Synchron's Stentrode (with Hodak as an investor⁴⁵).
6. DARPA's BMI.
7. BrainGate interface.
8. Kernel's Flux and Flow devices.
9. Paradromics' BCI.

Bigger and better

There were many more groundbreaking applications of AI in the first half of 2022. For some of them, I am bound by non-disclosure agreements. Happily, the open nature of artificial intelligence research by the big AI labs allows anyone in the world to access both the theory (research papers and articles) and selected practice (working demonstrations and hands-on playground) of the latest models.

The sky is bigger than we imagine. We are witnessing the rapid expansion of intelligence through the development of enormous language models worldwide. And it just keeps getting better and better. Keep your eyes open in the second half of 2022, as humanity rockets through the AI revolution, and the sky's immensity is revealed.

⁴⁵ <https://maxhodak.com/nonfiction/2022/02/04/synchron.html>



Dr Alan D. Thompson is an AI expert and consultant. With Leta (an AI powered by GPT-3), Alan co-presented a seminar called ‘The new irrelevance of intelligence’ at the World Gifted Conference in August 2021. His applied AI research and visualizations are featured across major international media, including citations in the University of Oxford’s debate on AI Ethics in December 2021. He has held positions as chairman for Mensa International’s gifted families committee, consultant to GE and Warner Bros, and memberships with the IEEE and IET.

This paper has a related video at:

<https://youtu.be/gjI0NmyA6KQ>

The previous paper in this series was:

Thompson, A. D. (2021c). *Integrated AI: The sky is on fire (2021 AI retrospective)*

<https://lifearchitect.ai/the-sky-is-on-fire>

To cite this paper:

Thompson, A. D. (2022). *Integrated AI: The sky is bigger than we imagine (mid-2022 AI retrospective)*. <https://LifeArchitect.ai>



Further reading

For brevity and readability, footnotes were used in this paper, rather than in-text citations. Additional reference papers are listed below, or please see <http://lifearchitect.ai/papers> for the major foundational papers in the large language model space.

Chinchilla paper

Hoffmann, J., Borgeaud, S., Mensch, A., et al. (2022). *Training Compute-Optimal Large Language Models*. DeepMind. <https://arxiv.org/abs/2203.15556>

SayCan paper and videos

Ahn, M., Brohan, A., Br4own, N., et al. (2022). Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. Google Robotics & Everyday Robots. <https://say-can.github.io>

Gato paper

Reed, S., Zolna, K., Parisotto, E., et al. (2022). *A Generalist Agent*. DeepMind. <https://arxiv.org/abs/2205.06175>

Thompson, A. D. (2020). *The New Irrelevance of Intelligence*. <https://lifearchitect.ai/irrelevance-of-intelligence>

Thompson, A. D. (2021a). *The New Irrelevance of Intelligence [presentation]*. Proceedings of the 2021 World Gifted Conference (virtual). <https://youtu.be/mzmeLnRlj1w>

Thompson, A. D. (2021b). *Integrated AI: The rising tide lifting all boats (GPT-3)*. <https://lifearchitect.ai/rising-tide-lifting-all-boats>

Thompson, A. D. (2021c). *Integrated AI: The sky is on fire (2021 AI retrospective)* <https://lifearchitect.ai/the-sky-is-on-fire>

Thompson, A. D. (2021c). *Leta AI*. The Leta conversation videos can be viewed in chronological order at: <https://www.youtube.com/playlist?list=PLqJbCeNOFk88QyAkBe-U0zxCgbHrGa4V>

Thompson, A. D. (2022). *What's in my AI? A Comprehensive Analysis of Datasets Used to Train GPT-1, GPT-2, GPT-3, GPT-NeoX-20B, Megatron-11B, MT-NLG, and Gopher*. <https://lifearchitect.ai/whats-in-my-ai>

