



Image: Generated by AI for this paper (Stable Diffusion¹)

Integrated AI: The sky is infinite (2022 AI retrospective)

Alan D. Thompson

December 2022

Previous papers in this series...

Mid-2022 retrospective: <https://lifearchitect.ai/the-sky-is-bigger/>

2021 retrospective: <https://lifearchitect.ai/the-sky-is-on-fire/>

It's... about creating worlds that are rich in 3D and interactive and moving... Almost like a liquid imagination that swirls around the room and forms mountains and trees and little animals and ruins.

— David Holz, Midjourney (Nov/2022)²

In my mid-year 2022 AI report, *The sky is bigger than we imagine*, I closed with the words: '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.'

The second half of the year has passed more or less as expected, with an explosion of models, more AI labs aligning with the 'more data' theory from DeepMind, and huge commercialization opportunities being applied at a rapid pace.

¹ Image generated in a few seconds, on 26/Oct/2022, text prompt by Alan D. Thompson: 'side-on view of winding mountain path with headlights from several vehicles driving up the mountain, HDR colorful and vivid, stunning UV night sky, volves close to the camera, bioluminescent glow, cinematic, photorealistic, 8K'. Using Stable Diffusion 1.4 via [mage.space](https://www.mage.space/):

² <https://www.abc.net.au/news/2022-11-26/loab-age-of-artificial-intelligence-future/101678206>

Best of 2022

There were thousands of models released³ in 2022, while only a handful made media headlines. When it comes to usefulness, here are my top picks for 2022.

Best ... in 2022	Model name	Notes
Language model (open)	OpenAI text-davinci-003 ⁴	New default GPT-3 engine, released Nov/2022, can rhyme
Language model (closed)	Google PaLM ⁵	The Pathways family is spectacular
Text-to-image model (open)	Midjourney v4 ⁶	Released Nov/2022, layer of abstraction for prompting
Text-to-image model (closed)	Google Parti ⁷	Largest (4.8B image dataset)
AI app	Quickchat.ai Emerson ⁸	Still my favorite application using GPT-3, a year later

Table: Alan's 'best of' AI models and apps for 2022.

In the second half of 2022, and in just a matter of months, several companies have achieved multi-billion dollar valuations on the back of some of these models. These are the numbers we've seen publicly, and there are obviously many more investments (and larger ones!) hidden behind closed doors.

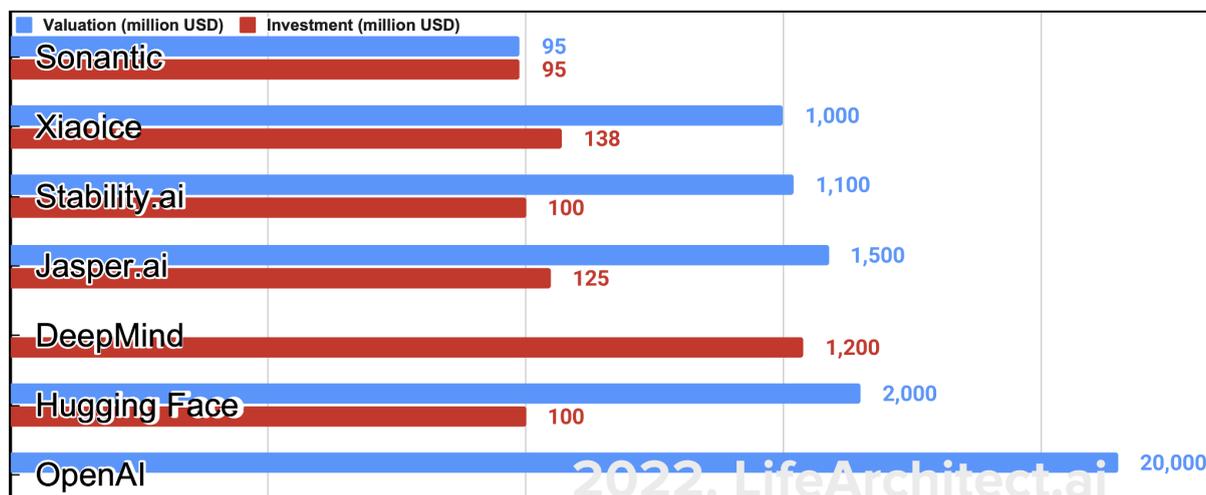


Chart: Valuation and investment: AI labs in 2022. Selected highlights only.

³ <https://huggingface.co/models?sort=modified>

⁴ <https://beta.openai.com/playground/p/default-summarize?model=text-davinci-003>

⁵ <https://lifearchitect.ai/pathways/>

⁶ <https://www.midjourney.com/>

⁷ <https://parti.research.google/>

⁸ <https://www.quickchat.ai/emerson>



New models out of DeepMind (with ‘investment’ corresponding to costs covered by Alphabet) and Google have been embraced by researchers, with several examples of models providing more personalized experiences. DeepMind’s Chinchilla model was developed into a new prompted output called Dramatron, and the scripts generated by the model were turned into plays, performed by real humans(!) at North America’s largest Fringe Festival in Aug/2022.

The applications built on these models cover a very wide space, with plenty of room for new startups to create platforms and tools. At the end of 2022, there seems to be just a handful of big players including Jasper.ai and GitHub Copilot, while others run to play catchup.

Commercializing text-to-image models

The main highlight in the second half of the year really was text-to-image models. With just a short list of keywords, users can ask the latest models to ‘imagine’ brand new images from scratch. The stunningly high-resolution header image for this report was generated in a matter of seconds, and is completely unique.

The monetization of these text-to-image models has been faster than expected. Mattel—a \$6 billion per year toy company—is applying image generation models to their toy concepts. OpenAI’s DALL-E 2 helped⁹ Mattel with their *Hot Wheels* series, ‘tweaking the design, asking DALL-E 2 to try it in pink or blue, with the soft-top on, and on and on.’ Mattel’s director of product design explained that the use of text-to-image models amplifies the quality of ideas. “It’s about going, ‘Oh, I didn’t think about that!’”.

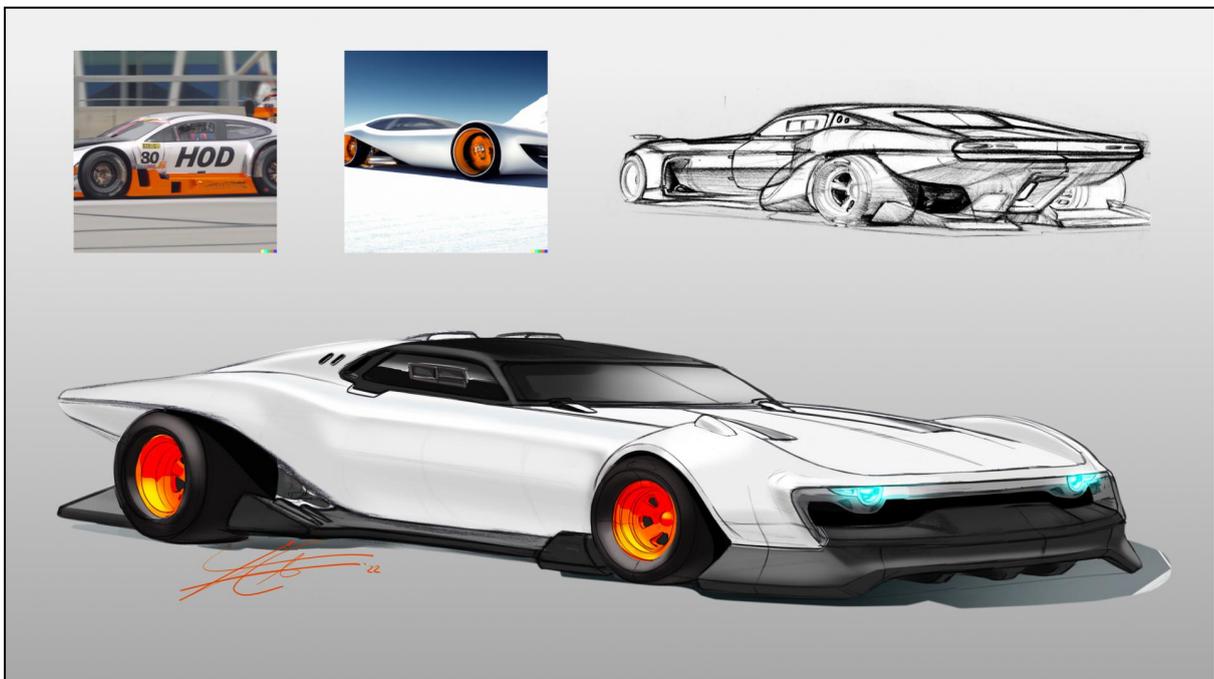


Image: Mattel (founded 1945) using DALL-E 2 (2022) for their *Hot Wheels* series.

⁹

<https://blogs.microsoft.com/ai/from-hot-wheels-to-handling-content-how-brands-are-using-microsoft-ai-to-be-more-productive-and-imaginative/>



In November 2022, German brewer Brauquadrat¹⁰ used the very recent release of Midjourney v4 to generate artwork for their line of sour beer. The results are astonishing, and the prompt was as simple as ‘commercial shot of raspberry, teal background, splashes, juicy --v 4’.



Image: Brauquadrat used Midjourney v4 to generate artwork for their product.

In the same month, Stability.ai announced¹¹ that the largest open-source text-to-image model, Stable Diffusion, had been licensed by over 200,000 software developers.

While DALL-E 2 and Stable Diffusion are two of the biggest players, there are now several options to choose from when it comes to these kinds of models. And all of them are being ‘baked in’ to the biggest software apps, including Canva¹², DeviantArt¹³, and even older players like AutoCAD, and Photoshop.

¹⁰ <https://www.brauquadrat.de/>

¹¹

<https://www.prnewswire.com/news-releases/stability-ai-announces-101-million-in-funding-for-open-source-artificial-intelligence-301650932.html>

¹² <https://www.canva.com/newsroom/news/text-to-image-ai-image-generator/>

¹³ <https://www.deviantart.com/team/journal/Create-AI-Generated-Art-Fairly-with-DreamUp-933537821>



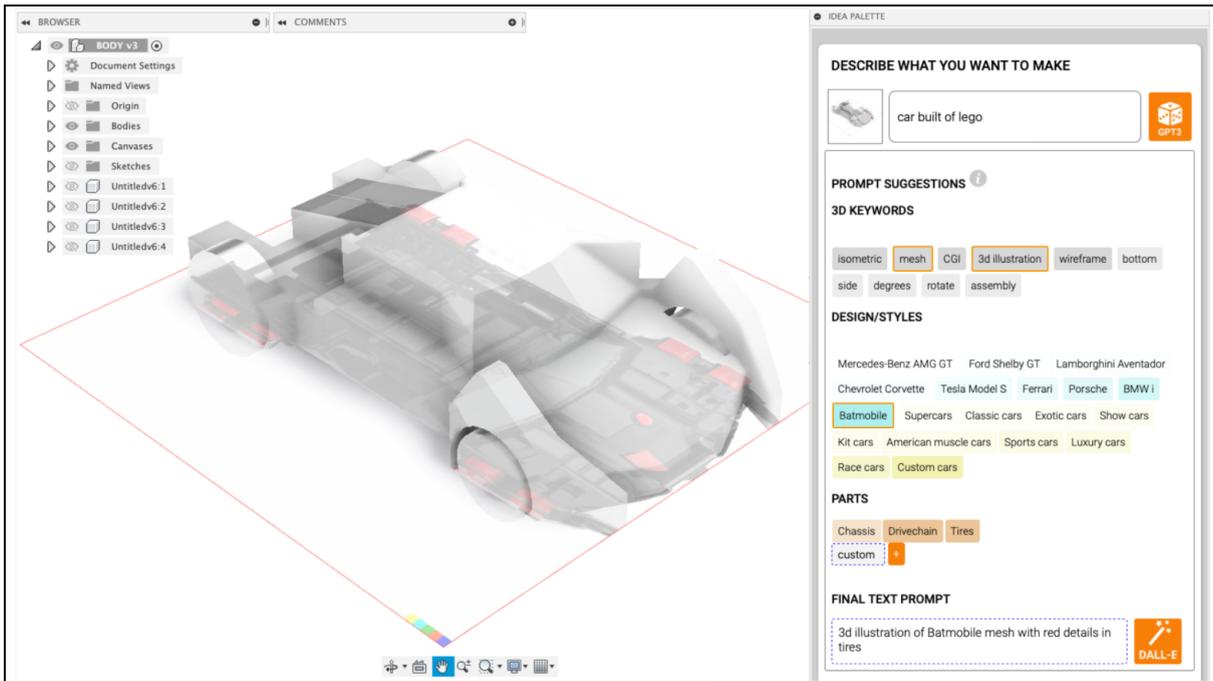


Image: AutoCAD by Autodesk with DALL-E 2 and GPT-3 (for prompting)¹⁴.

Anticompetitive and monopolistic¹⁵ software company Adobe is building text-to-image models inside its suite of applications, including Photoshop, perhaps for release in 2023. The results will be as astounding as the current crop of models in 2022.

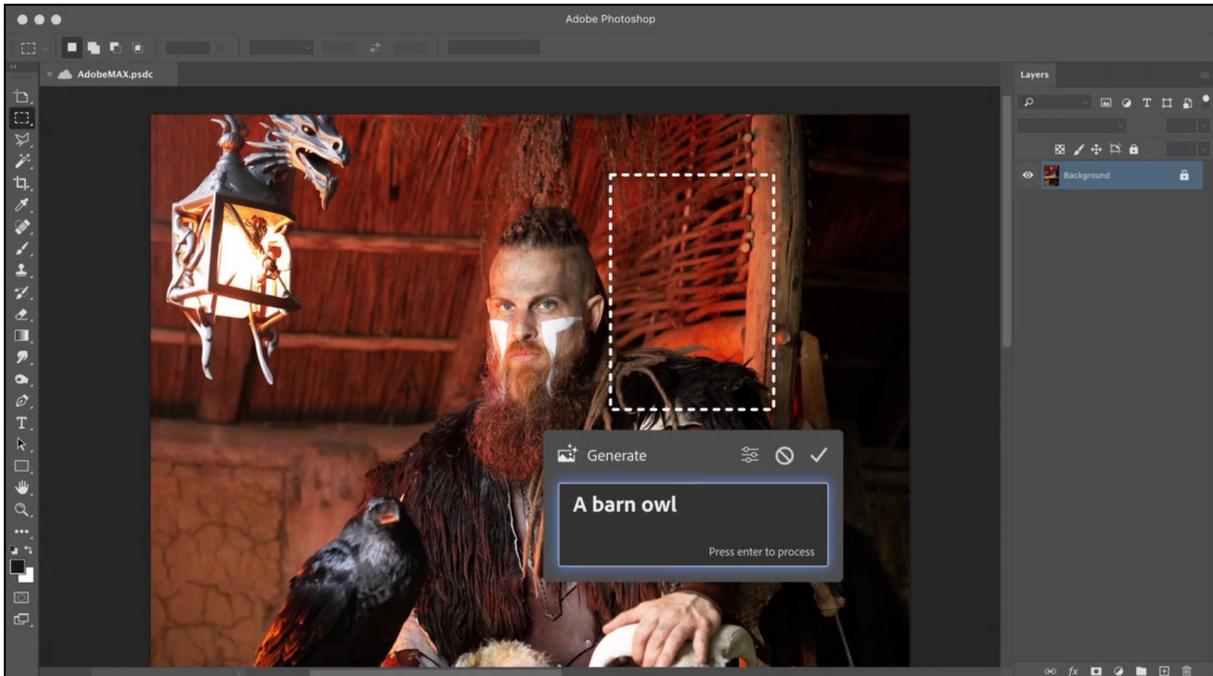


Image: Photoshop with text-to-image model integration concept¹⁶.

¹⁴ <https://arxiv.org/abs/2210.11603>

¹⁵ <https://www.vice.com/en/article/3kgw83/is-adobes-creative-cloud-too-powerful-for-its-own-good>

¹⁶ <https://blog.adobe.com/en/publish/2022/10/18/bringing-next-wave-ai-creative-cloud>



A picture is worth a thousand words

As large language models captured the imaginations of thousands of computer programmers over the last 24 months (Paul Graham observed¹⁷ that ‘Hackers are fascinated by GPT-3. To everyone else it seems a toy. Pattern seem familiar to anyone?’), it has taken visual illustrations to really allow people to see the power of current AI.

And why not? The phrase ‘A picture is worth a thousand words’ appears in many cultures. As Leonardo da Vinci reminded¹⁸ us, a writer would be ‘overcome by sleep and hunger before being able to describe with words what a painter is able to depict in an instant.’

The current avalanche of text-to-image models are a testament to that, producing masterpieces in a matter of seconds. Here are just some of the standalone (not fine-tuned) text-to-image models that made their way into the public zeitgeist this year. Each of these models represents millions or billions of text-image pairs, the equivalent of hundreds of years of training, and billions of dollars of potential revenue.

AI lab	Model	Type	Release	Notes
OpenAI	DALL-E 1	Autoregressive	Jan/2021	
Tsinghua	Cogview 1	Autoregressive	May/2021	4B params
Midjourney	Midjourney v1	Diffusion	Apr/2022	
OpenAI	DALL-E 2	Diffusion	Apr/2022	1M users in 3 months
Google	Imagen	Diffusion	May/2022	
Google	Parti	Autoregressive	Jun/2022	20B params
Tsinghua	CogView2	Autoregressive	Jun/2022	24B params, Pathways
Microsoft	NUWA-Infinity	Autoregressive	Jul/2022	‘Infinite’ canvas
Stability.ai	Stable Diffusion	Diffusion	Aug/2022	1M users, 50 countries
Craiyon	Craiyon (new)	GAN	Aug/2022	2.5B, ex-DALL-E Mini
Baidu	ERNIE-ViLG 2.0	Diffusion	Oct/2022	24B params
NVIDIA	eDiff-I	Diffusion	Nov/2022	
Midjourney	Midjourney v4	Diffusion	Nov/2022	

Table: Text-to-image models to end-2022. Selected highlights only.

¹⁷ <https://twitter.com/paulg/status/1284733103391612928>

¹⁸ Janson, H.W. (1962). History of Art (6th ed.). Abrams Books. p. 613.



This does not show the complete picture of releases, and even omits the massive round of advanced text-to-video models released by Google (Imagen Video and Phenaki), Meta (Make-A-Video), and other major labs. Text-to-video models will go some way to setting up your next imagined film or TV series in an instant, and then your entire virtual reality environment.

A picture is worth a thousand words, so...

If you'll allow me a bit of rope to reverse-engineer the previous phrase: if a picture is worth a thousand words, and we are seeing billions of dollars poured into text-to-image models in a matter of months, then most of humanity is missing the '1,000 times more powerful' large language models proliferating our world right now.

I am seeing large language models being exponentially more important for humanity than text-to-image models, and I'm counting down the days until we combine both of these maturing technologies into one enormous multimodal model.

The timeline of large language models through to the end of 2022 shows interesting open-source offerings from Russia, China, and the usual suspects through the US.

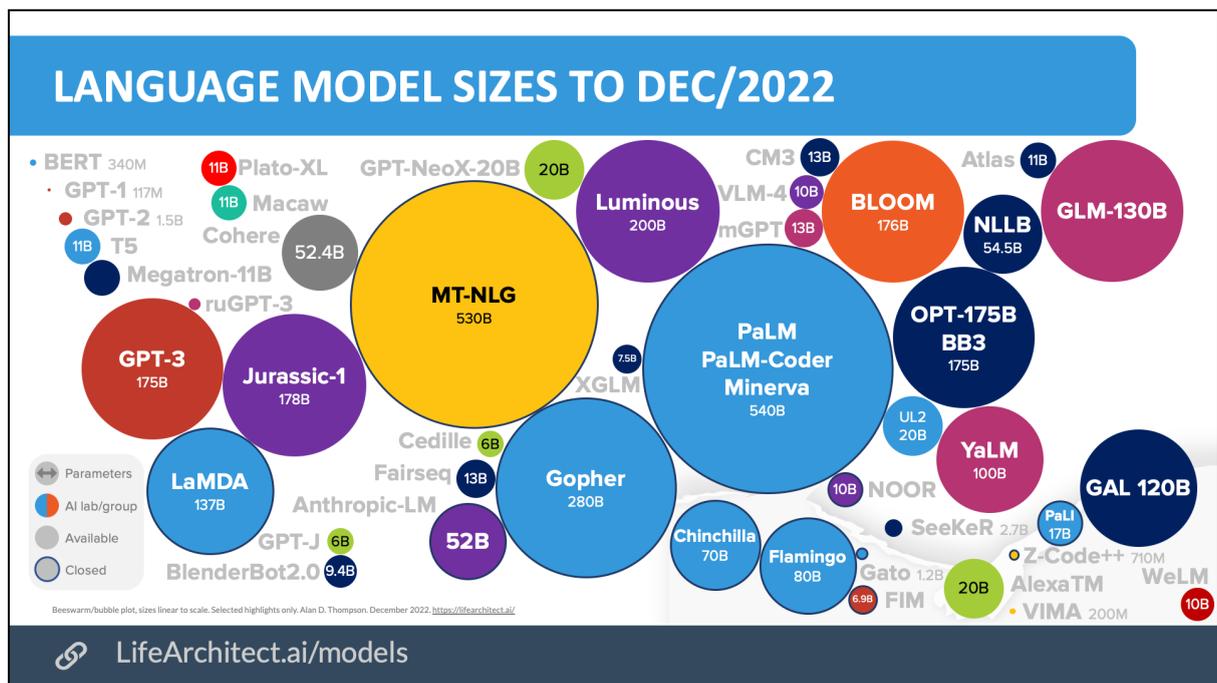


Chart: AI language models 2018-2022. Selected highlights only.

Come together

Open source teams led the pack in collaboration through the second half of 2022. 1,000 researchers in 60 countries replicated GPT-3 using a more equal distribution of more than 40 languages. Training started on 11/Mar/2022, and ended on 6/Jul/2022, using 384x A100 GPUs on the Jean Zay public supercomputer at a cost of \$7M in compute. The model—originally called tr11-176B-ml—was finally christened BLOOM.



Across the pond, Russia's Google equivalent, Yandex released a 100B-parameter model capable of speaking a 50/50 mix of English and Russian, while China opened their GLM-130B model to the world, and then WeChat delivered a data-optimal model called WeLM 10B.

Back in the United States, Amazon promised to open their newest Alexa Teacher Model (AlexaTM 20B), and Meta opened several models to the public. Meta's Galactica scientific model (GAL 120B) featured a public-facing demo interface, converting LaTeX and Wikipedia format data into their appropriate style templates. It was summarily misrepresented by bad faith actors¹⁹ (i.e. those deliberately entering in abusive prompts, with Galactica returning appropriate responses), forcing Meta to hide the interface from the public.

AI goes to work

In a recent keynote to around 4,000 developers representing companies like Google AI and Microsoft, I spoke about just how deeply AI has penetrated the corporate world. After demonstrating some of the latest applications of GPT-3 (a Google Sheets function, an English to SQL converter, an instant story book generator), I explored the current state of AI in the enterprise even further.

After working with some of the biggest Fortune 500s in the world in health, retail, engineering, fashion, pharmaceuticals, I wonder if there are any Fortune 500s that have not played around with these post-2020 AI models. Here are just a few of the big companies visibly working with GPT-3, keeping in mind that this is just one of many large language models available for commercial use.



Chart: GPT-3 enterprise users. Selected highlights only.

¹⁹ <https://twitter.com/ylecun/status/1593342777034739713>



While I've been advising many large companies over the last 24 months, it is always fascinating to see just how diverse the industries are, as well as the sheer breadth of ways that post-2020 AI can be applied.

Hardware explodes

In my 2021 AI report²⁰, I covered Google's use of AI to design its current generation of TPU chips. TPUs (which can be compared with GPUs) are used to train today's AI models. Often, thousands of these units are linked together to train the model in parallel for the equivalent of hundreds of years. In that report, I wrote that 'while humans value neatness and order, artificial intelligence has no such limitations, optimizing for flow, efficiency, and even unanticipated synergies brought about by electromagnetic coupling.'

NVIDIA has now followed suit for its H100 Hopper chips, using AI to help design their GPUs²¹:

We demonstrate that not only can AI learn to design these circuits from scratch, but AI-designed circuits are also smaller and faster than those designed by [humans and even] state-of-the-art electronic design automation (EDA) tools. The latest NVIDIA Hopper GPU architecture has nearly 13,000 instances of AI-designed circuits.

The result is GPUs that are up to 6x faster than the A100 chips used for training and inference across most of the current AI models in 2022. Watch out for this hardware to be used across the board in 2023, from text-to-image models, all the way through to embodied AI via robots.

Data explodes

In March 2022, the Chinchilla paper by Jordan Hoffmann and 21 other researchers at DeepMind was a shock to the AI world. In the simplest terms, it found that AI labs had been training models on about 9% the amount of data that they should have been using. So where GPT-3 175B trained on about 300B tokens (maybe 600GB), it should have trained on 3,500B tokens (maybe 7TB).

Following this revelation, labs have been scrambling to collect and gather more and more data. While the Chinchilla paper doesn't mention if the data can be 'repeated', there is some comfort in the fact that the training may allow the model to 'see' tokens more than once. In simple terms, this may mean that for a 1T-parameter dense language model, it perhaps allows labs to find 16.5TB of clean text data—and run it twice—instead of 33TB of clean text data.

²⁰ <https://lifearchitect.ai/the-sky-is-on-fire/>

²¹ <https://developer.nvidia.com/blog/designing-arithmetic-circuits-with-deep-reinforcement-learning/>



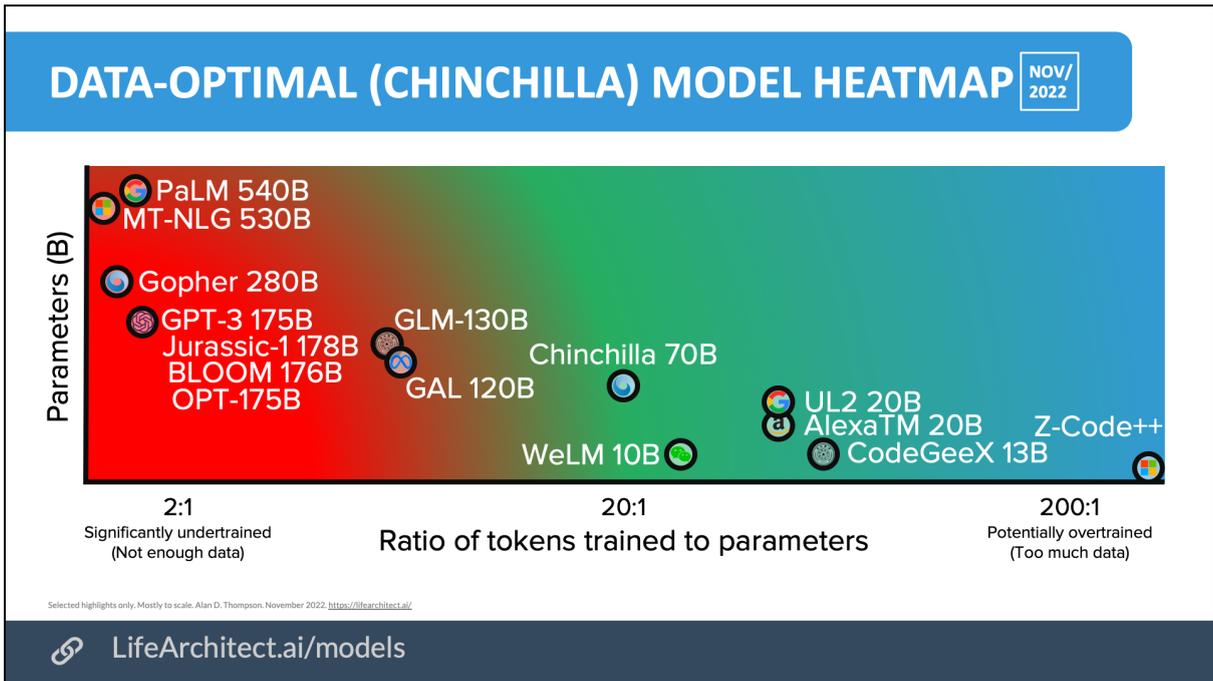


Chart: Data-optimal large language models 2020-2022. Selected highlights only.

Labs have definitely been taking heed of this discovery already through 2022, with models like WeChat’s WeLM 10B, Amazon’s AlexaTM 20B, and Microsoft’s Z-Code++ exceeding the new recommendation. As labs seek out more and more data, they are taking some interesting approaches. OpenAI’s Whisper²², a transformer-based automatic speech recognition model, was trained on 77 years of audio content, and is expected to assist with the creation of new text datasets based on audio, including spoken word extracted from video.

The dogs bark, and the caravan moves on

Back in 2019, OpenAI’s head of policy made a comment²³ in relation to AI and ethics: ‘We’re trying to build the road as we travel across it.’

The road is unfolding so rapidly that it seems difficult to see. The Turkish proverb, ‘The dogs bark, and the caravan moves on,’ reminds us that progress moves ahead, no matter the criticism it may attract. And the AI landscape has been attracting plenty of criticism over the last six months.

As post-2020 AI creeps into people’s field of view more easily, so too does negativity bias, as well as defensiveness. We are living many years after the dawn of the sabertooth cat, yet many critics are given a mouthpiece to react as if we are still running from a tiger. The volume of this sense of hysteria seems to be getting louder, and those shouting the loudest seem to have significant crossover with those who are least informed about the subject.

²² <https://openai.com/blog/whisper/>

²³ <https://www.theguardian.com/technology/2019/feb/14/elon-musk-backed-ai-writes-convincing-news-fiction>



In an article titled ‘Mocking AI Panic’ in the IEEE Spectrum, Diane Proudfoot²⁴ saw that:

...history is repeating itself. In the mid-1940s, public reaction to reports of the new “electronic brains” was fearful. Newspapers announced that “the controlled monster” (a room-size vacuum-tube computer) could rapidly become “the monster in control,” reducing people to “degenerate serfs.” Humans would “perish, victims of their own brain products.”

...[The father of AI, Dr Alan] Turing’s response to AI panic was gentle mockery.

It’s not just misguided fears and anger about AI taking over. It’s also people wanting to ‘cash in’ in the use of data. A silly website²⁵ even argues that ‘Generative AI profits off your [content]. Make them pay for it.’ This sits in contrast to my 2021 article²⁶ where I had proposed a giant dataset of ‘bonum’ or ‘ultimate good’ data, perhaps idealistically assuming that, while ‘there are also intellectual property and copyright considerations for some of the datasets... it is expected that these would be easily cleared by the respective authors for the good of humanity.’

In any case, the current copyright claims—by multiple parties in multiple lawsuits—should hopefully come to naught. Perhaps AI can even help us solve the strange human delusions of patents and copyright.

What is certain is that AI is speeding forward despite the dogs barking, and the road is appearing in front of us at a tremendous pace. Dr Ray Kurzweil was far more aggressive²⁷ in describing today’s climate, except he said it back in 2020.

You can’t stop the river of advances. These ethical debates are like stones in a stream. The water runs around them. You haven’t seen any of these... technologies held up for one week by any of these debates. To some extent, they may have to find some other ways around some of the limitations, but there are so many developments going on.

There are dozens of very exciting ideas about how to use... information. Although the controversies may attach themselves to one idea here or there, there’s such a river of advances. The concept of technological advance is so deeply ingrained in our society that it’s an enormous imperative.

...There is also a tremendous moral imperative. We still have not millions but billions of people who are suffering from disease and poverty, and we have the opportunity to overcome those problems through these technological advances.

²⁴ <https://lifearchitect.ai/ai-panic/>

²⁵ <https://www.paytotrain.ai/>

²⁶ <https://lifearchitect.ai/bonum/>

²⁷ <https://www.kurzweilai.net/essays-celebrating-15-year-anniversary-of-the-book-the-singularity-is-near>



Smarter and smarter

I've been singing the praises of superintelligence in large language models for a couple of years now. Since 2020, models have been outperforming humans in specific subtests. With the release of Google PaLM in Apr/2022, and its subsequent large language models Minerva (1/Jul/2022) and the U-PaLM and Flan-PaLM families (20/Oct/2022), the newest benchmark scores are incredibly confronting.

These models do not have a calculator, they cannot search for an older maths paper, and they were not specifically 'taught' how to do maths. They must use the connections made during the pretraining. From my seat, they should not even be able to understand what a number is, let alone the operators and other features to successfully manipulate those numbers.

Please remember that post-2020 large language models are pretrained black boxes, and after training they have no access to search the web, nor can they retrieve the original training data as-is. In fact, the GPT-3 paper²⁸ authors had similar concerns, even going so far as double-checking the original dataset to see if every calculation (e.g. $1 + 1 = 2$) was memorized by the model verbatim during training.

To spot-check whether the model is simply memorizing specific arithmetic problems, we took the 3-digit arithmetic problems in our test set and searched for them in our training data in both the forms " $\langle \text{NUM1} \rangle + \langle \text{NUM2} \rangle =$ " and " $\langle \text{NUM1} \rangle$ plus $\langle \text{NUM2} \rangle$ ". Out of 2,000 addition problems we found only 17 matches (0.8%) and out of 2,000 subtraction problems we found only 2 matches (0.1%), suggesting that only a trivial fraction of the correct answers could have been memorized. In addition, inspection of incorrect answers reveals that the model often makes mistakes such as not carrying a "1", suggesting it is actually attempting to perform the relevant computation rather than memorizing a table.

Here is an example of a real question²⁹ presented to Google's Minerva model (based on PaLM), from the Polish national math exam, with the model's answer below. Note the May/2022 exam paper was only made available after the training of Minerva was complete, so there is no possibility of problem memorization.

²⁸ <https://arxiv.org/abs/2005.14165>

²⁹ <https://minerva-demo.github.io/#category=Algebra&index=2>



QUESTION Prove that for every $a, b, b \neq a$, holds $\frac{a^2+b^2}{2} > \left(\frac{a+b}{2}\right)^2$.

MODEL ANSWER We have

$$\begin{aligned}\frac{a^2 + b^2}{2} &> \left(\frac{a + b}{2}\right)^2 \\ \frac{a^2 + b^2}{2} &> \frac{a^2 + 2ab + b^2}{4} \\ 2a^2 + 2b^2 &> a^2 + 2ab + b^2 \\ a^2 - 2ab + b^2 &> 0 \\ (a - b)^2 &> 0,\end{aligned}$$

which is true for all a and b .

Figure: Google PaLM/Minerva outperformed students on the 2020 Polish math exam. A real question from the test, and subsequent correct answer by AI.

In a nutshell, for a selected subset of questions, Google's Pathways³⁰ models:

1. Minerva 540B scored 14% higher than the average student in the May/2022 Polish national math exam.
2. Minerva 540B scored around 79% higher than the average student in the 2019 UK national math exam.
3. PaLM 540B and other models are now outperforming humans on the SuperGLUE benchmark, designed to have a high ceiling, and to be as hard as possible.
4. Flan-PaLM 540B is achieving twice the performance of an average human on the MMLU benchmark.

³⁰ <https://lifearchitect.ai/pathways/>



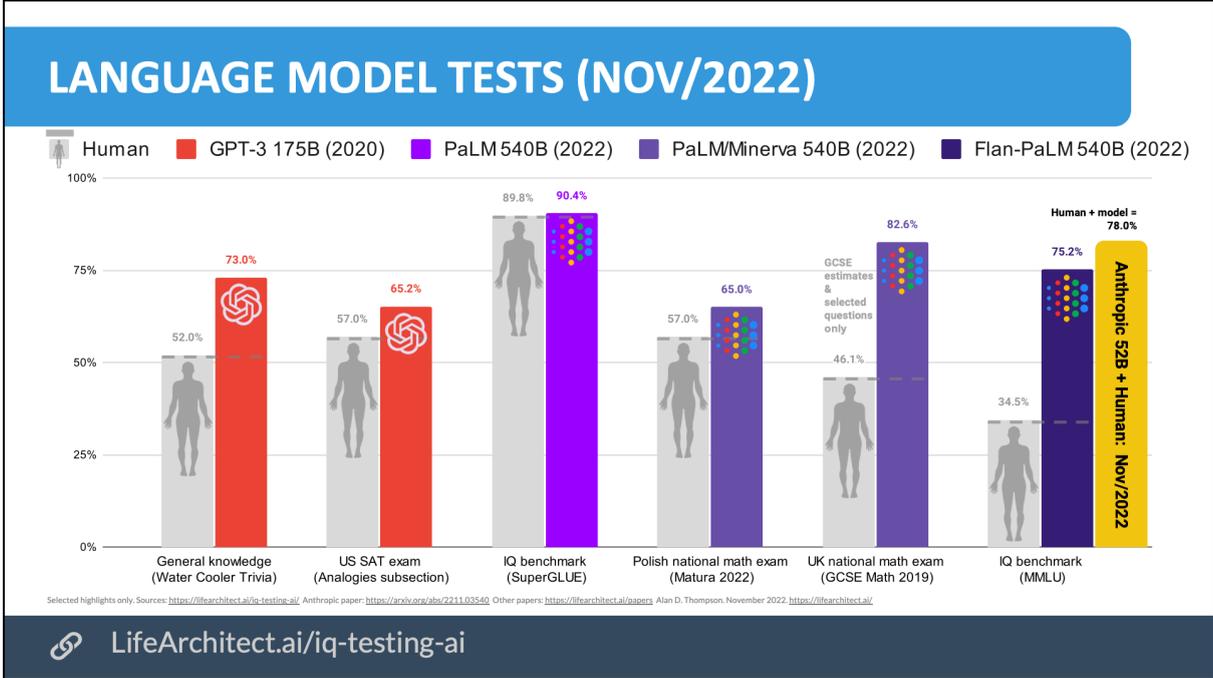


Chart: IQ testing AI: GPT-3 & PaLM family models on AI benchmarks.

This rapid evolution of superhuman intelligence allows so many new things to happen at once. Language models can be applied to anything that uses language. When Google Transformer creator Dr Ashish Vaswani founded Adept³¹, he wanted to apply the technology to any task in the browser, and eventually any task on the computer.

With former researchers from DeepMind, Google Brain, and OpenAI, plus more than \$65M in funding, the team released their first model, Action Transformer or ACT-1 in September 2022. The model is currently hooked up to a Chrome extension which allows ACT-1 to observe what’s happening in the browser and take certain actions, like clicking, typing, and scrolling.

The results are impressive, and in some ways even more tangible than embodied AI. You can see ACT-1 follow instructions to search for houses, add customers to a CRM, write emails, and much more. Released on 14/Sep/2022, the next iterations will be able to interact with different applications on your laptop.

³¹ <https://www.adept.ai/post/introducing-adept>



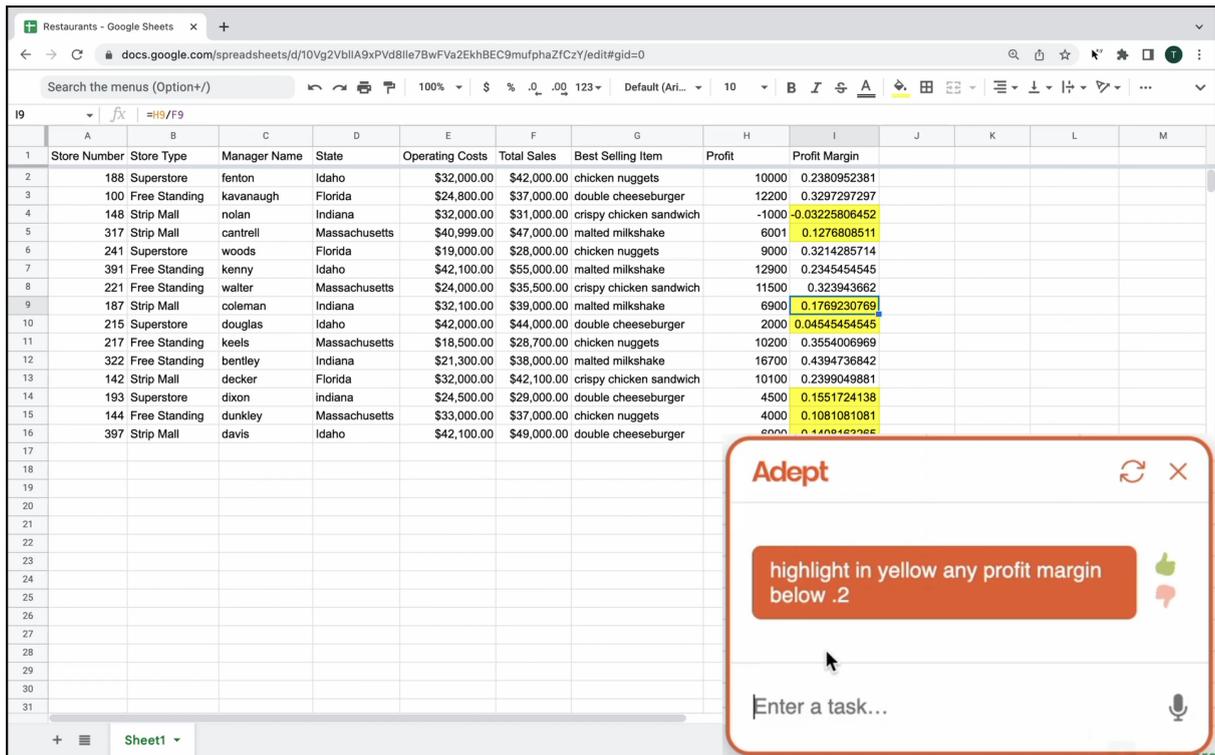


Image: Adept.ai's ACT-1 model using Google Sheets.

Next up in 2023

I generally shy away from predictions, but in 2023, I am very much looking forward to:

1. **DeepMind Gato 2.** In a Lex Fridman interview, DeepMind CEO Demis Hassabis revealed³² that the company is already training the next embodied generalist agent, ready for AGI. The original Gato³³ was already an unforeseen innovation.
2. **Google Pathways expansion.** The entire model family is phenomenal, and is expanding to cover more languages. Dr Jeff Dean remarked³⁴ that it will 'support the 1,000 most spoken languages, bringing greater inclusion to billions of people in marginalized communities all around the world. This will be a many years undertaking – some may even call it a moonshot – but we are already making meaningful strides here and see the path clearly.'
3. **OpenAI GPT-4.** The successor to OpenAI's GPT-3 model from 2020 has been rumored for quite some time, and I expect this model to significantly alter the AI landscape.
4. **Text-to-video models.** We've already seen several text-to-video models in the second half of 2022, but look for models hitting higher resolutions and frame rates, and perhaps even models that are ready for prime time...

³² <https://youtu.be/Gfr50f6ZBvo> 'Gato predicts potentially any action or any token, and it's just the beginning really, it's our most general agent... that itself can be scaled up massively, more than we've done so far, obviously we're in the middle of doing that.'

³³ <https://youtu.be/6fWEHrXN9zo>

³⁴ <https://blog.google/technology/ai/ways-ai-is-scaling-helpful/>



In the middle of infinity

At the beginning of 2022, OpenAI's chief scientist predicted³⁵ that 'In the future, it will be obvious that the sole purpose of science was to build AGI'. We're in the middle of that future right now, to the extent that one of the AI labs (Midjourney) even named their company after this time period!

Through the lens of rapid acceleration, this moment in time is easily one of the most audacious. This report has covered billions of dollars of investment, a dozen new text-to-image models, and application explosions across the breadth of the AI field, all of which have come about in a matter of months.

There is another lens in the eyepiece though.

It is the lens of hindsight.

I know that in just a few years from now, we will look back on this time in history and be shocked at the current set of circumstances. With artificial intelligence booming, most of society is still negatively impacted by low-hanging fruit and easily-solvable problems:

1. **Driving by hand.** A human is brutally killed every 24 seconds³⁶ trying to drive cars, and driving causes an additional two significant injuries every second.
2. **Education in herds.** Children waste 11,000 hours inside a classroom³⁷ (and a few thousand more hours studying) learning how to memorize facts, solve problems, and be 'creative'.
3. **Jobs in cages.** Most adults spend close to all of their waking time at work or stressing about work, a place where their productivity and effectiveness is less than 50% that of today's AI, and less than 1% of tomorrow's AGI.
4. **Money for nothing.** The delta between CEO and minimum wage³⁸ is 670-to-1, meaning the average CEO receives \$6,700,000 in compensation for every \$10,000 the worker received.
5. **Nutrition by guesswork.** There's calorie counting (manual), vitamins and supplements (approximate), and just overall eating (whatever). Most diets in 2022 are guesswork, based on feel or limited research. At the same time, 42% of Americans are obese³⁹, and 32% of all deaths worldwide are attributed to heart disease⁴⁰.

³⁵ <https://twitter.com/ilyasut/status/1505754945860956160?lang=en>

³⁶ https://www.un.org/sites/un2.un.org/files/media_gstc/FACT_SHEET_Road_safety.pdf

³⁷

<https://www.smh.com.au/education/australian-students-spend-more-time-in-class-are-still-outperformed-in-tests-20180727-p4zu34.html>

³⁸

<https://www.theguardian.com/us-news/2022/jun/07/us-wage-gap-ceos-workers-institute-for-policy-studies-report>

³⁹ <https://www.tfah.org/report-details/state-of-obesity-2022/>

⁴⁰ <https://www.who.int/health-topics/cardiovascular-diseases>



6. **Relationships in the dark.** Most adults are navigating relationships with next to no training, with more than half of them experiencing some form of mental health issue in their lifetime⁴¹.

AI is primed to solve—or is well on the way to supporting new ways of achieving a beneficial result in—most of the examples above. In 2022, the reality has been that although AI has exploded, and its potential known to be infinite, the speed at which its promise is applied to the 8 billion people alive today is painstakingly slow.

Perhaps more alarmingly, while the takeup of AI by major corporations is gaining momentum, governments and intergovernmental organizations have already fallen far, far behind. This holds true even in my experience advising on post-2020 AI to the most advanced governments, including those with AI departments and ministries.

To align with the record-setting pace of AI releases, it is vital that leadership teams and overarching bodies keep up, and it is only possible to do so by taking a different path.

The promise of artificial intelligence and its benefits is life-changing for everyone. As more and more people begin sitting up to take notice of today's AI performance, spectacular new models are already in the works and ready for release in 2023. From the dawn of real artificial general intelligence (AGI), to new ways of seeing the world through the eyes of AI, the sky is infinite.

⁴¹ <https://www.cdc.gov/mentalhealth/learn/index.htm>



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://www.youtube.com/@DrAlanDThompson/videos>

The previous paper in this series was:

Thompson, A. D. (2022). *Integrated AI: The sky is bigger than we imagine (mid-2022 AI retrospective)*.

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

To cite this paper:

Thompson, A. D. (2022). *Integrated AI: The sky is infinite (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>

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=PLqJbCeNOfEK88QyAkBe-U0zxCgbHrGa4V>

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

Thompson, A. D. (2022b). *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>

