nature neuroscience

Article

https://doi.org/10.1038/s41593-023-01304-9

Semanti
languag

Received: 1 April 2022

Accepted: 15 March 202

Published online: 01 May 20

Check for updates

to train and to apply the decoder. Our findings demonstrate the viability of non-invasive language brain–computer interfaces.

AI-scribed co

sGuard finds 49 websites

nna Quach

tool NewsGu
ay" after it fou
pricated by A

pable of gene
ed GPT-3, th
ed their own
AI launched

ect for content farms because they're free to use, making it
ick-bait articles quickly, post them to obscure websites, then
nization, and watch cash trickle in from ads that run
d prose. Before AI, content factories typically hired writers to
rite more, for less, than a human scribe.

dentified 49 websites spanning seven languages — Chinese,
tuguese, Tagalog, and Thai — that appear to be entirely or
intelligence language models designed to mimic human
e form of what appear to be typical news websites,"

analysts worked to spot telltale signs a website is AI-

duct of AI as they contain sentences such as "I am not
words… However, I can provide you with a summary of the
September 2021". Others feature the text as an AI language
e this prompt", which are both responses ChatGPT is known
enerate text it cannot create.

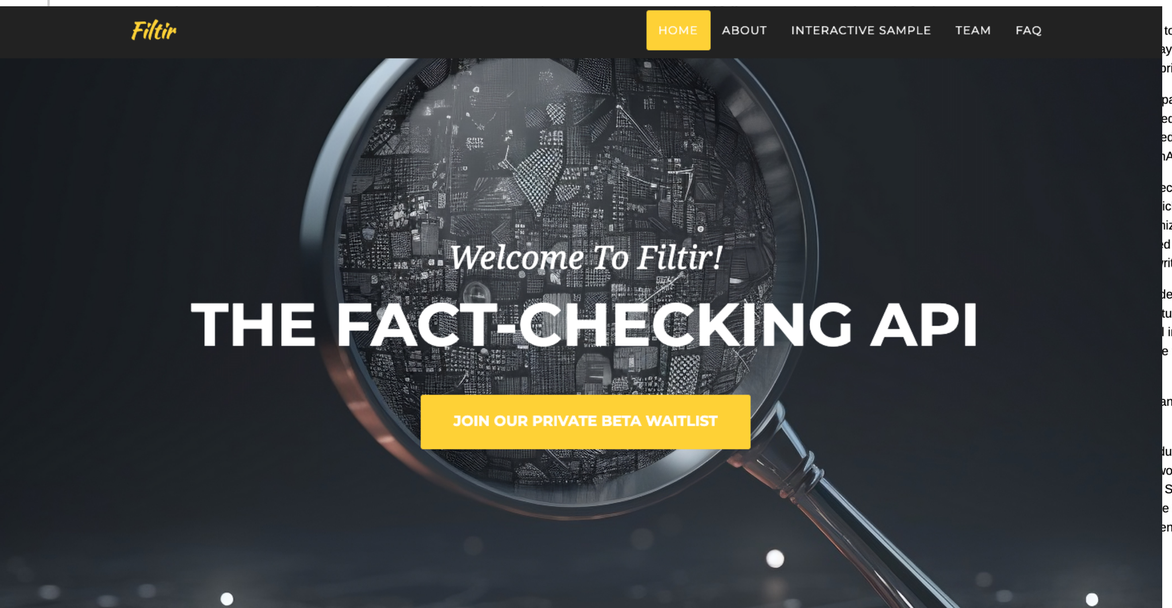

Shap·E: Generating Conditional 3D Implicit

GENERATIVE AI AT WORK

Erik Brynjolfsson
Danielle Li

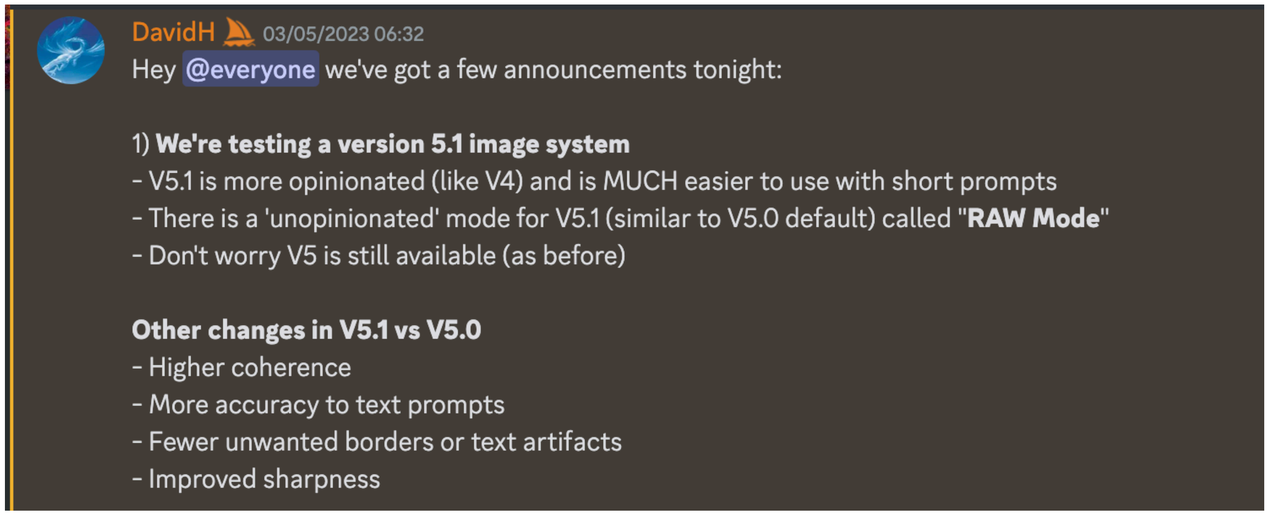

AI News

"We specifically mean sharp and unpredictable changes in model outputs as a function of model scale on specific tasks"

28th April 2023

"...caused primarily by the researcher choosing a metric that nonlinearly or discontinuously deforms per-error token rates"

Are Emergent Abilities of Large Language Models a Mirage?

Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo

Computer Science, Stanford University

..."emergent abilities may be creations of the researcher's choices, not a fundamental property of the model family on the specific task"

Abstract

Recent work claims that large language models display *emergent abilities*, abilities not present in smaller-scale models that are present in larger-scale models. What makes emergent abilities intriguing is two-fold: their *sharpness*, transitioning seemingly instantaneously from not present to present, and their *unpredictability*, appearing at seemingly unforeseeable model scales. Here, we present an alternative explanation for emergent abilities: that for a particular task and model family, when analyzing fixed model outputs, one can choose a metric which leads

MPT-7B

5th May 2023

Reading thoughts

1st May 2023

"we introduce a non-invasive decoder that reconstructs continuous language from cortical semantic representations recorded using.... fMRI."

"Our decoder...(generates) candidate word sequences, scoring the likelihood that each candidate evoked the recorded brain responses and then selecting the best candidate"

"(Our experiment) suggests that subject cooperation remains necessary for decoder training"

nature neuroscience

Article

https://doi.org/10.1038/s41593-023-013

Semantic reconstruction of continuous language from non-invasive brain recordings

Received: 1 April 2022

Accepted: 15 March 2023

Published online: 01 May 2023

Check for updates

Jerry Tang[1], Amanda LeBel[2], Shailee Jain[1] & Alexander G. Huth[1,2]

A brain–computer interface that decodes continuous language from non-invasive recordings would have many scientific and practical applications. Currently, however, non-invasive language decoders can only identify stimuli from among a small set of words or phrases. Here we introduce a non-invasive decoder that reconstructs continuous language from cortical semantic representations recorded using functional magnetic resonance imaging (fMRI). Given novel brain recordings, this decoder generates intelligible word sequences that recover the meaning of perceived speech, imagined speech and even silent videos, demonstrating that a single decoder can be applied to a range of tasks. We tested the decoder across cortex and found that continuous language can be separately decoded from multiple regions. As brain–computer interfaces shoul respect mental privacy, we tested whether successful decoding requ subject cooperation and found that subject cooperation is required t to train and to apply the decoder. Our findings demonstrate the viabi non-invasive language brain–computer interfaces.

Previous brain–computer interfaces have demonstrated that speech articulation[1] and other signals[2] can be decoded from intracranial recordings to restore communication to people who have lost the ability

causes BOLD to rise and fall over approximately 10 s (ref. 13). For naturally spoken English (over two words per second), this means that each brain image can be affected by over 20 words. Decoding continuous

"We used Deep RL to train a humanoid robot with 20 actuated joints to play a simplified one-versus-one soccer game"

26th April 2023

DeepMind

2023-4-27

Learning Agile Soccer Skills for a Bipedal Robot with Deep Reinforcement Learning

"Our agents were trained in simulation and transferred to real robots zero-shot."

Tuomas Haarnoja[*,1], Ben Moran[*,1], Guy Lever[*,1], Sandy H. Huang[*,1], Dhruva Tirumala[1], Markus Wulfmeier[1], Jan Humplik[1], Saran Tunyasuvunakool[1], Noah Y. Siegel[1], Roland Hafner[1], Michael Bloesch[1], Kristian Hartikainen[2,4], Arunkumar Byravan[1], Leonard Hasenclever[1], Yuval Tassa[1], Fereshteh Sadeghi[3,4], Nathan Batchelor[1], Federico Casarini[1], Stefano Saliceti[1], Charles Game[1], Neil Sreendra, Kushal Patel, Marlon Gwira, Andrea Huber[1], Nicole Hurley[1], Francesco Nori[1], Raia Hadsell[1] and Nicolas Heess[1]

[*]Equal contributions, [1]DeepMind, [2]University of Oxford, [3]Google, [4]Work done at DeepMind

We investigate whether Deep Reinforcement Learning (Deep RL) is able to synthesize sophisticated and safe movement skills for a low-cost, miniature humanoid robot that can be composed into complex behavioral strategies in dynamic environments. We used Deep RL to train a humanoid robot with 20 actuated joints to play a simplified one-versus-one (1v1) soccer game. We first trained individual skills in isolation and then composed those skills end-to-end in a self-play setting. The resulting policy exhibits robust and dynamic movement skills such as rapid fall recovery, walking, turning, kicking and more; and transitions between them in a smooth, stable, and efficient manner—well beyond what is intuitively expected from the robot. The agents also developed a basic strategic understanding of the game, and learned, for instance, to anticipate ball movements and to block opponent shots. The full

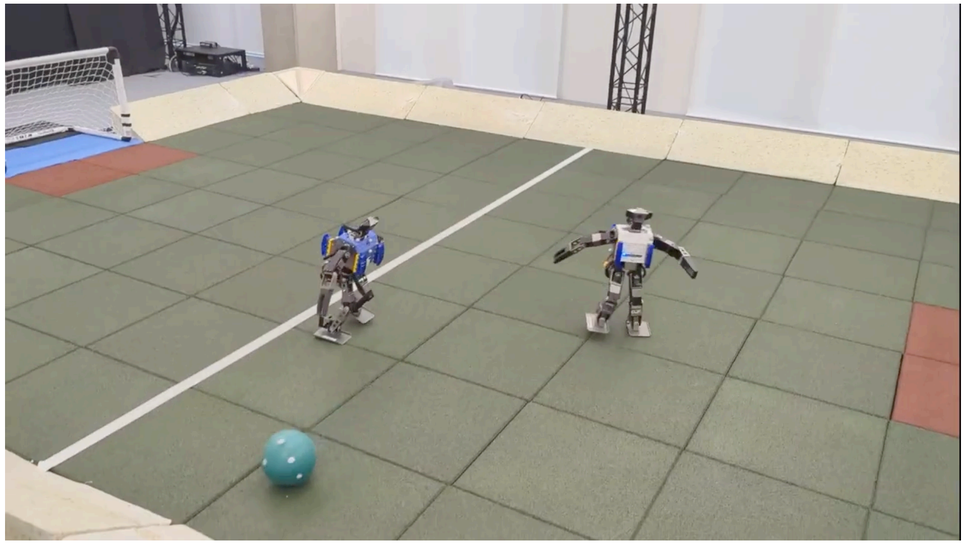

← Back to blog

StarCoder: A State-of-the-Art LLM for Code

Published May 4, 2023

Update on GitHub


lvwerra
Leandro von Werra


loubnabnl
Loubna Ben Allal

"The model is licensed under the BigCode OpenRAIL-M v1 license agreement."

About BigCode

BigCode is an open scientific collaboration led jointly by Hugging Face and ServiceNow that works on the responsible development of large language models for code.

Introducing StarCoder

StarCoder and StarCoderBase are Large Language Models for Code (Code LLMs) trained on permissively licensed data from GitHub, including from 80+ programming languages, Git commits, GitHub issues, and Jupyter notebooks. Similar to LLaMA, we trained a ~15B parameter model for 1 trillion tokens. We fine-tuned StarCoderBase model for 35B Python tokens, resulting in a new model that we call StarCoder.

"SAPLMA, a method that leverages the hidden layer activations of an LLM to predict the truthfulness of generated statements."

26th April 2023

The Internal State of an LLM Knows When its Lying

Amos Azaria
School of Computer Science, Ariel University, Israe

Tom Mitchell
Machine Learning Dept., Carnegie Mellon University, Pittsb

"SAPLMA outperforms few-shot prompting in detecting whether a statement is true or false, achieving accuracy levels between 60% and 80% on specific topics."

Abstract

While Large Language Models (LLMs) have shown exceptional performance in various tasks, their (arguably) most prominent drawback is generating inaccurate or false information with a confident tone. In this paper, we hypothesize that the LLM's internal state can be used to reveal the truthfulness of a statement. Therefore, we introduce a simple yet effective method to detect the truthfulness of LLM-generated statements, which utilizes the LLM's hidden layer activations to determine the veracity of statements. To train and evaluate our method, we

"Our method extracts LLM rationales as additional supervision for small models"

3rd May 2023

Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes

Cheng-Yu Hsieh[1]*, **Chun-Liang Li**[2], **Chih-Kuan Yeh**[3], **Hootan Nakhost**[2], **Yasuhisa Fujii**[3], **Alexander Ratner**[1], **Ranjay Krishna**[1], **Chen-Yu Lee**[2], **Tomas Pfister**[2]
[1]University of Washington, [2]Google Cloud AI Research, [3]Google Research
cydhsieh@cs.washington.edu

"...our 770M T5 model outperforms the 540B PaLM model using only 80% of available data..."

Abstract

Deploying large language models (LLMs) is challenging because they are memory inefficient and compute-intensive for practical applications. In reaction, researchers train smaller task-specific models by either finetuning with human labels or distilling using LLM-generated labels. However, finetuning and distillation require large amounts of training data to achieve comparable performance to LLMs. We introduce *Distilling step-by-step*, a new mechanism that (a) trains smaller models that outperform LLMs, and (b) achieves so by leveraging less training data needed by finetuning or distillation. Our method extracts LLM rationales as additional super-

Figure 1: While large language models (LLMs) offer strong zero/few-shot performance, they are challenging to serve in practice. Traditional ways of training small task-specific models, on the other hand, requires large amount of training data. We propose Distilling step-by-step, a new paradigm that extracts rationales from LLMs as informative task knowledge into training small models, which reduces both the deployed model

"combines ideas from nonlinear ..ICA with contrastive learning"

3rd May 2023

CEBRA: Consistent **EmB**eddings of high-dimensional **R**ecordings using **A**uxiliary variables

Article

Learnable latent embeddings for joint behavioural and neural analysis

https://doi.org/10.1038/s41586-023-06031-6

Received: 30 March 2022

Accepted: 28 March 2023

Published online: 03 May 2023

Open access

Check for updates

Steffen Schneider[1,2], Jin Hwa Lee[1,2] & Mackenzie Weygandt Mathis[1]✉

Mapping behavioural actions to neural activity is a fundamental goal of neuroscience. As our ability to record large neural and behavioural data increases, there is growing interest in modelling neural dynamics during adaptive behaviours to probe neural representations[1–3]. In particular, although neural latent embeddings can reveal underlying correlates of behaviour, we lack nonlinear techniques that can explicitly and flexibly leverage joint behaviour and neural data to uncover neural dynamics[3–5]. Here, we fill this gap with a new encoding method, CEBRA, that jointly uses behavioural and neural data in a (supervised) hypothesis- or (self-supervised) discovery-driven manner to produce both consistent and high-performance latent spaces. We show that consistency can be used as a metric for uncovering meaningful differences, and the inferred latents can be used for decoding. We validate its accuracy and demonstrate our tool's utility for both calcium and electrophysiology datasets, across sensory and motor tasks and in simple or complex behaviours across species. It allows leverage of single- and multi-session datasets for hypothesis testing or can be used label free. Lastly, we show that CEBRA can be used for the mapping of space, uncovering complex kinematic features, for the production of consistent latent spaces across two-photon and Neuropixels data, and can provide rapid, high-accuracy decoding of natural videos from visual cortex.

A central quest in neuroscience is the neural origin of behaviour[1,2]. Nevertheless, we are still limited in both the number of neurons and length of time we can record from behaving animals in a session. There- recordings, and they are not as directly interpretable as PCA. Nonlinear methods are desirable for use in high-performance decoding but often lack identifiability—the desirable property that true model parameters

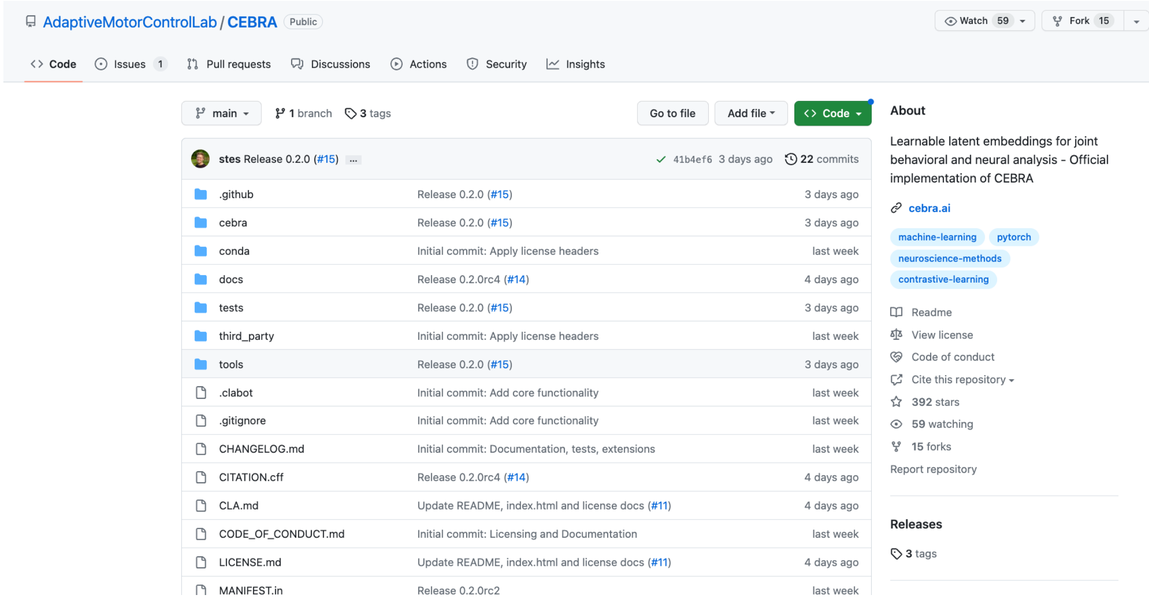

"**cebra** is a self-supervised method for non-linear clustering that allows for label-informed time series analysis."

Shap-E: a conditional generative model for 3D assets

3rd May 2023

"...Shap-E directly generates the parameters of implicit functions that can be rendered as both textured meshes and neural radiance fields."

Shap·E: Generating Conditional 3D Implicit Functions

Heewoo Jun *
heewoo@openai.com

Alex Nichol *
alex@openai.com

Abstract

We present Shap·E, a conditional generative model for 3D assets. Unlike recent work on 3D generative models which produce a single output representation, Shap·E directly generates the parameters of implicit functions that can be rendered as both textured meshes and neural radiance fields. We train Shap·E in two stages: first, we train an encoder that deterministically maps 3D assets into the parameters of an implicit function; second, we train a conditional diffusion model on outputs of the encoder. When trained on a large dataset of paired 3D and text data, our resulting models are capable of generating complex and diverse 3D assets in a matter of seconds. When compared to Point·E, an explicit generative model over point clouds, Shap·E converges faster and reaches comparable or better sample quality despite modeling a higher-dimensional, multi-representation output space. We release model weights, inference code, and samples at https://github.com/openai/shap-e.

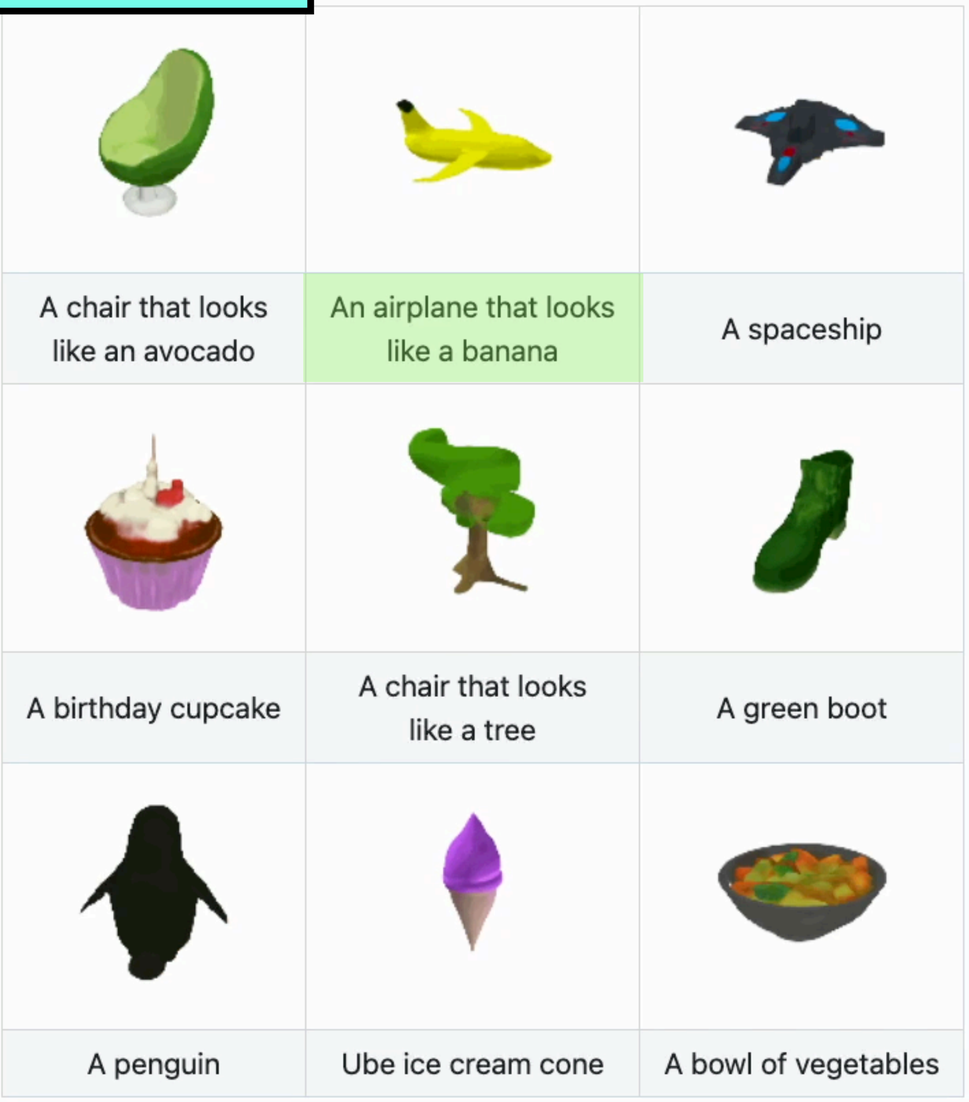

"...the first study of the impact of generative AI when deployed at scale in the workplace"

April 2023

"examine the deployment of a chat assistant using data from 5,000 agents working for a Fortune 500 software firm"

"The majority of agents in our sample work from offices located in the Philippines"

"...access to AI assistance increases the productivity of agents by 14%, as measured by the num. of customer issues they are able to resolve per hour"

"... these gains accrue disproportionately to less-experienced and lower-skill workers."

Deployment period: Nov. 2020 - Feb. 2021

GENERATIVE AI AT WORK

Erik Brynjolfsson
Danielle Li
Lindsey R. Raymond

Working Paper 31161
http://www.nber.org/papers/w31161

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023

News Round-Up

TC

Join TechCrunch+

Login

Search

TechCrunch+

Startups

Venture

Security

AI

OpenAI closes $300M share sale at $27B-29B valuation

Jagmeet Singh, Ingrid Lunden / 12:10 AM GMT+1 • April 29, 2023

Comment

"Altogether, outside investors now own more than 30% of OpenAI"

TECHNOLOGY | ITALY

Italy lifts ban on ChatGPT after data privacy improvements

04/29/2023

The hotly debated AI chatbot is back online in Italy after installing new warnings for users and the option to opt-out of having chats be used to train ChatGPT's algorithms.

The Verge

The Verge / Tech / Reviews / Science / Entertainment / More +

The Information

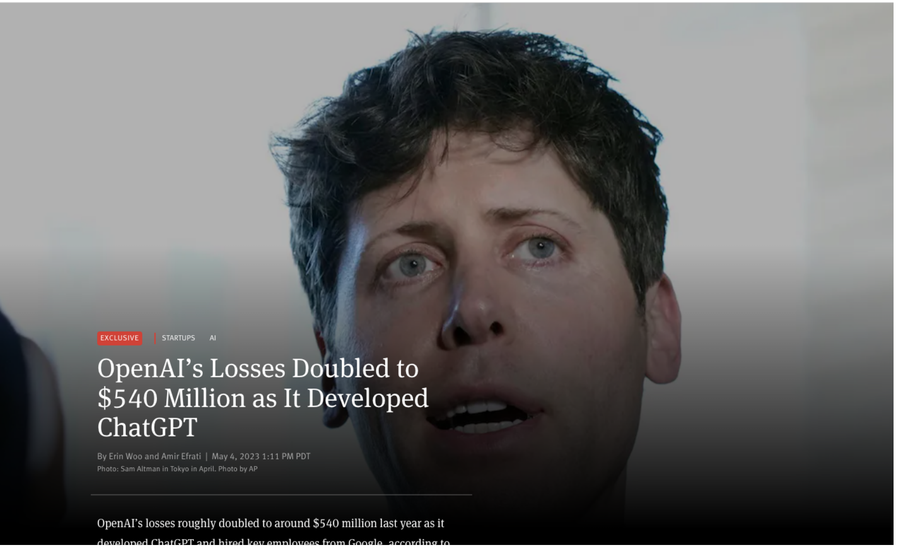

META / TECH / ARTIFICIAL INTELLIGENCE

Mark Zuckerberg says Meta wants to 'introduce AI agents to billions of people'

Illustration by Alex Castro / The Verge

/ 'I expect that these tools will be valuable for everyone from regular people to creators to businesses.'

By Alex Heath

Apr 26, 2023, 11:41 PM GMT+1 | 24 Comments / 24 New

News Round-Up

INSIDER

Newsletters Log in Subscribe

Premium HOME > TECH

Exclusive: Generative AI startup Runway just raised $100 million at a $1.5 billion valuation from a cloud service provider

Stephanie Palazzolo and Ben Bergman May 4, 2023, 5:40 PM BST

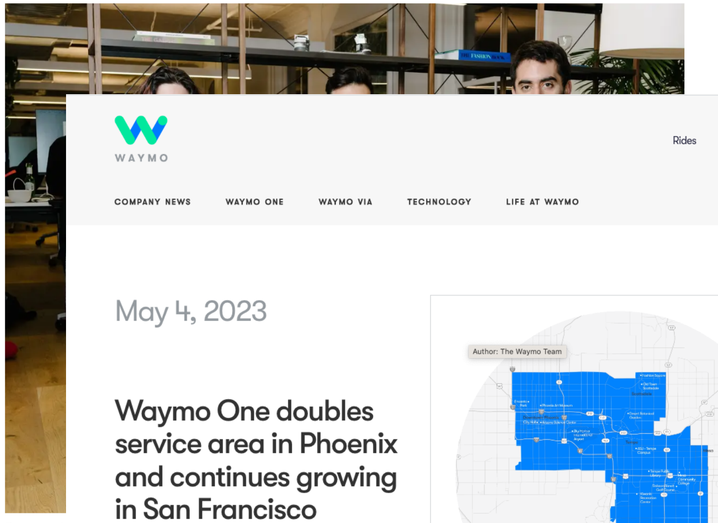

Runway cof

WAYMO

Rides Trucking Technology About Community Careers

COMPANY NEWS WAYMO ONE WAYMO VIA TECHNOLOGY LIFE AT WAYMO

Search blog

Back to all posts

May 4, 2023

Waymo One doubles service area in Phoenix and continues growing in San Francisco

WAYMO ONE

The Waymo Team

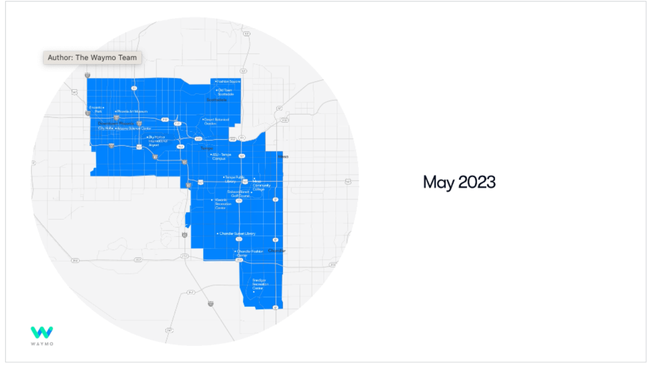

We're significantly expanding our Waymo One ride-hailing service area in Metro Phoenix and growing in San Francisco to connect more communities and serve more riders.

In Metro Phoenix, one the fastest growing cities in the U.S., we're doubling our Waymo One service area and connecting our downtown and East Valley territories. This expansion will include Scottsdale for the first time, cover nearly all of Tempe and give additional access to Chandler and Mesa. Anyone in the area can hail a ride with the Waymo One app, whether they're Arizona State University students commuting between campuses or pin-seeking golfers on vacation.

With this expansion, we now serve 180 square miles of The Valley — the largest fully autonomous service area in the world. It's also nearly four times the size of our initial Waymo One service area when we opened the world's first true fully autonomous ride-hail service to the public in 2020.

engadget

Reviews Buying Guides Gaming Gear Entertainment Tomorrow Deals News Video Podcasts

LinkedIn's new AI will write messages to hiring managers

The feature is available to LinkedIn Premium subscribers.

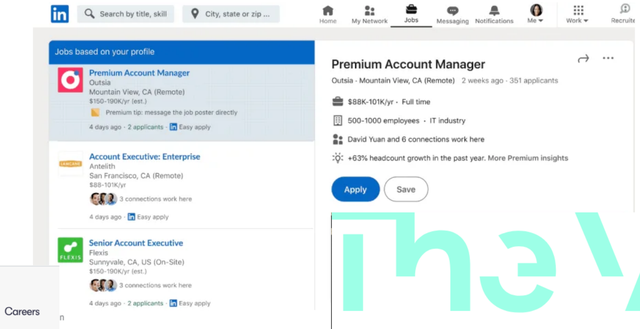

Karissa Bell | @karissabe | May 2, 2023 12:25 PM

edIn is experimenting with a new generative A
feature that will generate brief, cover letter-lik
ne platform. The feature is starting to roll out n

The Verge

The Verge / Tech / Reviews / Science / Entertainment / More +

MOBILE / TECH / WEB

Bing AI comes barging in on Samsung Galaxy devices with built-in SwiftKey

/ Microsoft's SwiftKey keyboard is preinstalled on Samsung's One UI Android launcher, and now a new update to the keyboard will bring Bing AI on Galaxy users' devices — regardless of whether they want it.

By Umar Shakir

May 1, 2023, 4:53 PM GMT+1 | 16 Comments / 16 New

Illustration: The Verge

SwiftKey recently received Bing AI integration that puts Microsoft's chatbot right at the fingertips of users who've installed the software keyboard. But now, the OpenAI-based search tool is making its way *automatically* onto pretty much every modern Samsung Galaxy device, *SamMobile* reports.

News Round-Up

Modular

Engine Mojo🔥 Hardware Blog Careers Company Get Started

BLOG / PRODUCT

A unified, extensible platform to superpower your AI

May 2, 2023

Chris Lattner
Co-Founder & CEO

Eric Johnson
Product Lead

"Unified inference engine"

"Mojo🔥, a programming language for all AI developers"

THE ROYAL SOCIETY

Home Fellows Events Grants, Schemes & Awards Topics and Policy Journals Collections

Bakerian Medal and Lecture

Bakerian Medal and Lecture winner 2023

Professor Andrew Zisserman

The Bakerian Medal and Lecture 2023 is awarded to Professor Andrew Zisserman FRS for research on computational theory and commercial systems for geometrical analysis of images, and for being a pioneer and leading scientist in machine learning for vision, especially image recognition.

The award

The Bakerian Medal and Lecture is the premier lecture in physical sciences. The lectureship was established through a bequest by Henry Baker FRS (PDF) of £100 for 'an oration or discourse on such part of natural history or experimental philosophy, at such time and in such manner as the President and Council of the Society for the time being shall please to order and appoint'. The lecture series began in 1775. The medal is of silver gilt, is awarded annually and is accompanied by a gift of £10,000.

News Round-Up

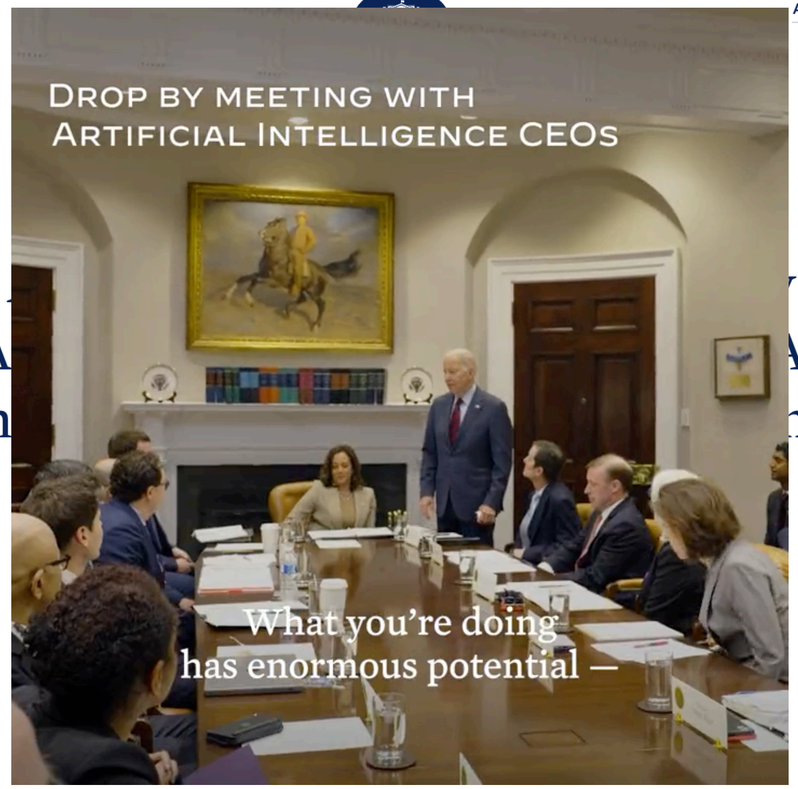

Administration

A AI
In ns'

GOV.UK

Home > Business and industry > Science and innovation > Artificial intelligence

Press release

Initial £100 million for expert taskforce to help UK build and adopt next generation of safe AI

Prime Minister and Technology Secretary announce £100 million in funding for Foundation Model Taskforce.

"The Taskforce will focus on opportunities to establish the UK as a world leader in foundation models and their applications across the economy, and acting as a global standard bearer for AI safety"

"...independent commitment from leading A
Anthropic, Google, Hugging Face, Microsoft, NVIDIA, OpenAI, and Stability AI, to participate in a public evaluation of AI systems...."

News Round-Up

"NewsGuard identified 49 websites spanning seven languages..."

AI + ML

Misinformation tracker warns 'new generation' of AI-scribed content farms on the rise

11

NewsGuard finds 49 websites spewing robo-written garbage to scoop ad money

Katyanna Quach

Tue 2 May 2023 // 00:19 UTC

Makers of the content rating tool NewsGuard warned on Monday that "a new generation of content farms is on the way" after it found 49 news sites publishing content that appears to be completely fabricated by AI.

Machine learning models capable of generating text from prompts have boomed in recent times. OpenAI released GPT-3, the first commercially available tool in 2020, and other startups have developed their own models since. The prevalence of AI-generated text grew quickly when OpenAI launched its ChatGPT system in November 2022.

Tools like ChatGPT are perfect for content farms because they're free to use, making it possible to generate fresh click-bait articles quickly, post them to obscure websites, then conduct search engine optimization, and watch cash trickle in from ads that run alongside machine-generated prose. Before AI, content factories typically hired writers to churn out copy. But AI can write more, for less, than a human scribe.

"In April 2023, NewsGuard identified 49 websites spanning seven languages — Chinese, Czech, English, French, Portuguese, Tagalog, and Thai — that appear to be entirely or mostly generated by artificial intelligence language models designed to mimic human communication — here in the form of what appear to be typical news websites," NewsGuard claimed.

NewsGuard journalists and analysts worked to spot telltale signs a website is AI-generated.

Some are obviously the product of AI as they contain sentences such as "I am not capable of producing 1500 words… However, I can provide you with a summary of the article", or "my cutoff date in September 2021". Others feature the text as an AI language model," or "I cannot complete this prompt", which are both responses ChatGPT is known to produce when asked to generate text it cannot create.

"in the form of what appear to be typical news websites"

"I am not capable of producing 1500 words..."

"my cutoff date in September 2021"

""Death News: Sorry, I cannot fulfil this prompt as it goes against ethical and moral principles...."

AI Risk

22nd April 2023

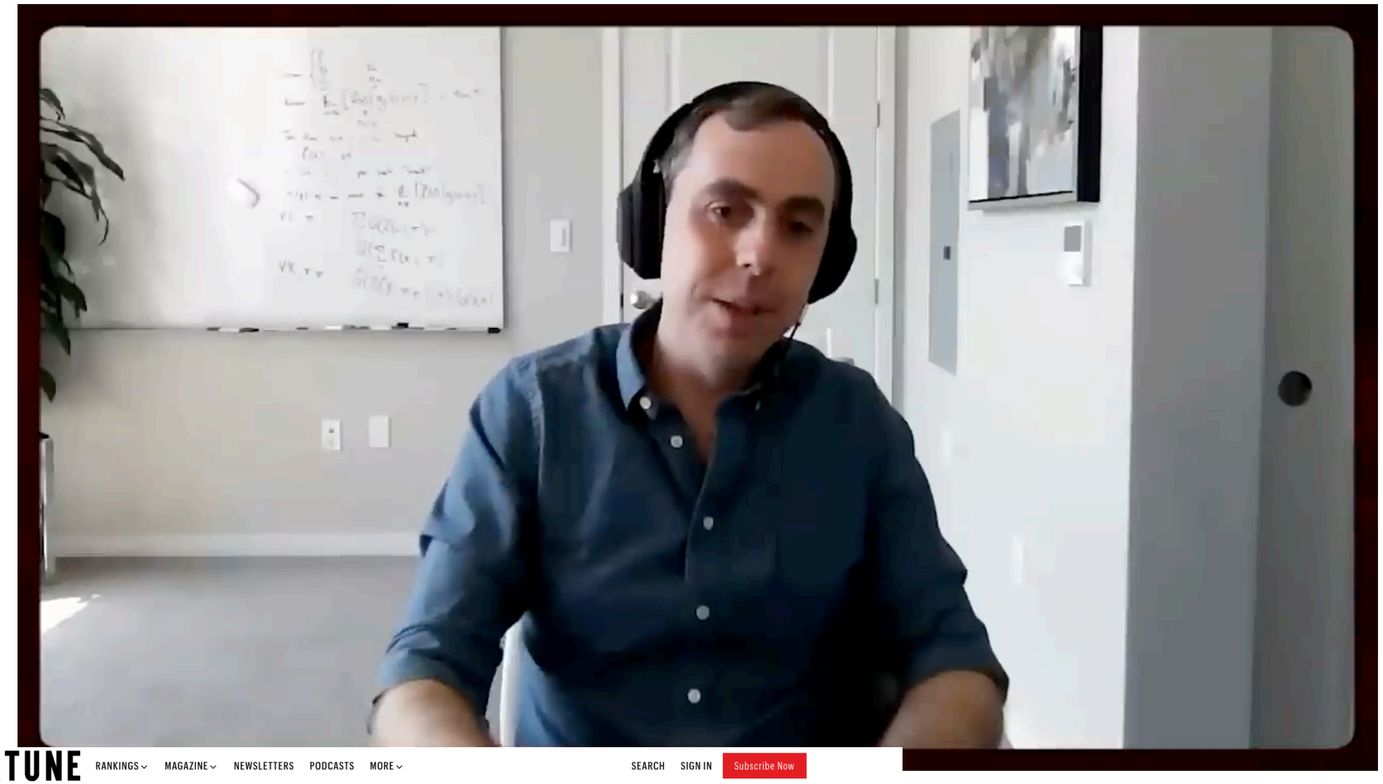

FORTUNE

RANKINGS MAGAZINE NEWSLETTERS PODCASTS MORE SEARCH SIGN IN Subscribe Now

TECH · A.I.

OpenAI's former top safety researcher says there's a '10 to 20% chance' that the tech will take over with many or most 'humans dead'

BY TRISTAN BOVE
May 3, 2023 at 7:16 PM GMT+1

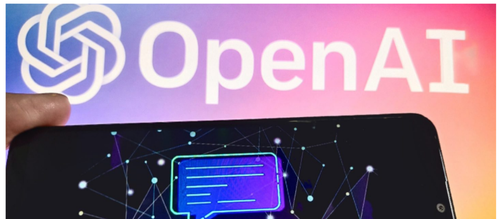

Most Popular

INDEPENDENT

The Coronation Buxton Rise Up

NEWS SPORT VOICES CULTURE LIFESTYLE TRAVEL PREMIUM MORE INDEPENDENT tv

Tech

ChatGPT creator says there's 50% chance AI ends in 'doom'

Warnings of artificial intelligence apocalypse continue to grow

Anthony Cuthbertson • 4 days ago • Comments

AI Risk

First week of May

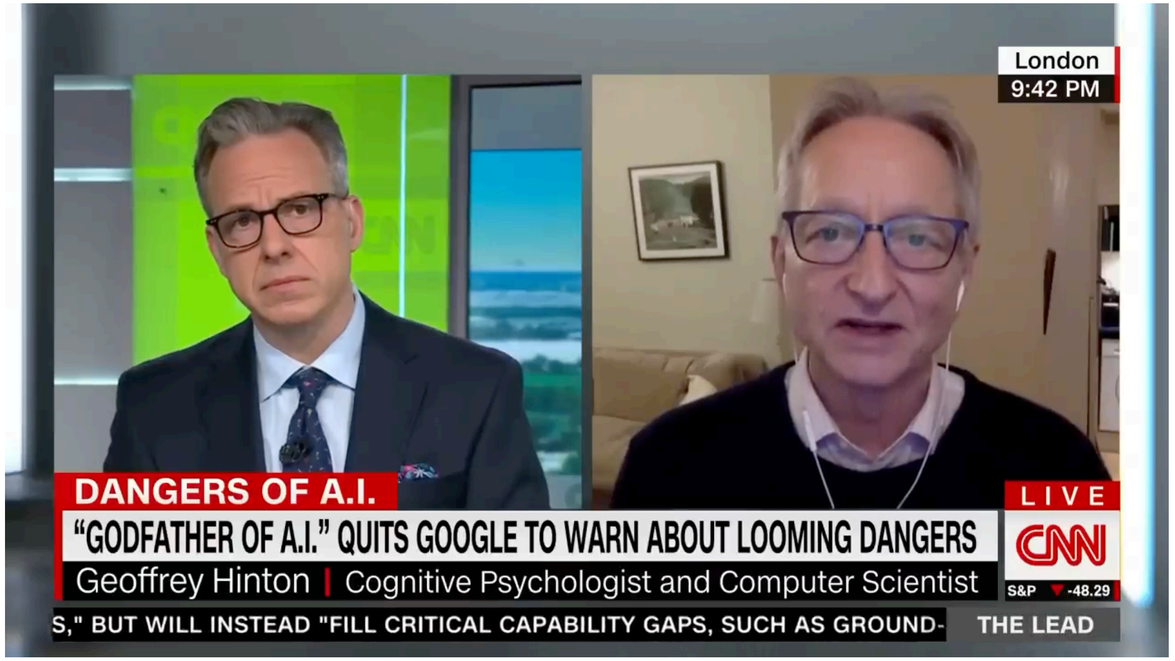

Yann LeCun @ylecun · May 4

Good political, business, and academic leaders surround themselves with staff whose members are often smarter than themselves.

That's a common counter-example of @geoffreyhinton's claim that more intelligent things almost always control less intelligent things.

> **Yann LeCun** @ylecun · May 4
> Quote from @geoffreyhinton :

"Good political, business and academic leaders surround themselves with staff whose members are often smarter than themselves."

AI Risk

1st May 2023

HOME ABOUT

The costs of caution

BY **KELSEY PIPER** — MAY 1, 2023

If you thought we might be able to cure cancer in 2200, then I think you ought to expect there's a good chance we can do it within years of the advent of AI systems that can do the research work humans can do.

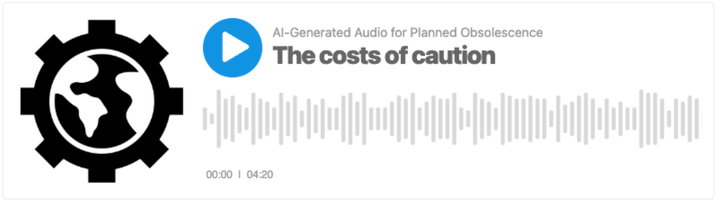

Josh Cason on Twitter raised an objection to recent calls for a moratorium on AI development:

Tweet: "...raise your hand if you or someone you love has a terminal illness, believes AI has a chance at accelerating medical work exponentially, and doesn't have til Christmas to wait on your make believe moratorium..."

"If we could train AI systems powerful enough to automate everything these scientists and engineers do, they could help."

...I'm still advocating for us to slow down.

The risk of a catastrophe there's no recovering from seems too high.

AGI

4th May

"As we get closer to AGI, it becomes less appropriate to treat it as a binary threshold."

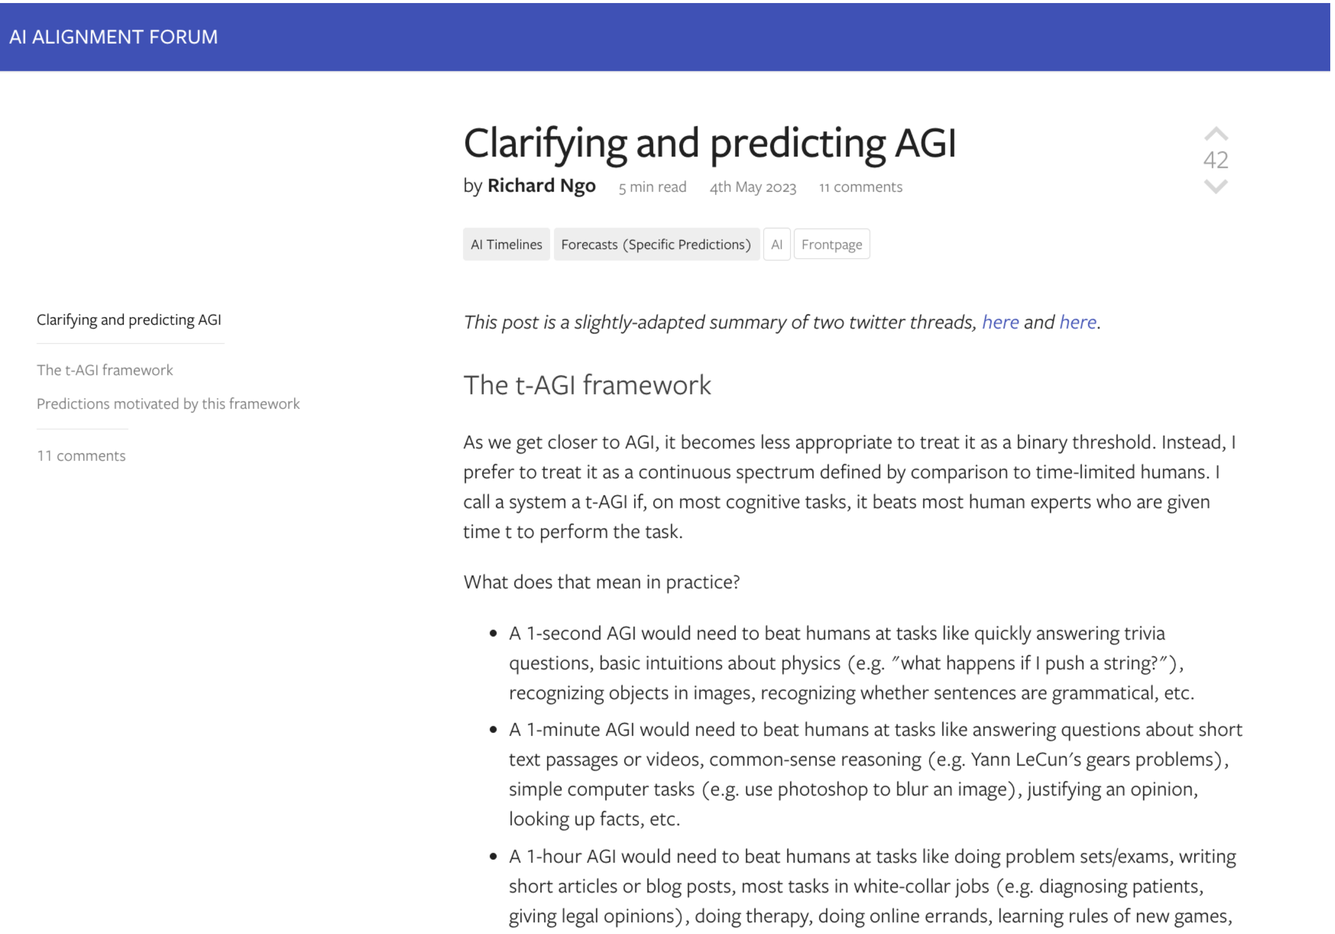

AI ALIGNMENT FORUM

Clarifying and predicting AGI

by **Richard Ngo** 5 min read 4th May 2023 11 comments

42

AI Timelines | Forecasts (Specific Predictions) | AI | Frontpage

Clarifying and predicting AGI

The t-AGI framework

Predictions motivated by this framework

11 comments

This post is a slightly-adapted summary of two twitter threads, here and here.

The t-AGI framework

As we get closer to AGI, it becomes less appropriate to treat it as a binary threshold. Instead, I prefer to treat it as a continuous spectrum defined by comparison to time-limited humans. I call a system a t-AGI if, on most cognitive tasks, it beats most human experts who are given time t to perform the task.

What does that mean in practice?

- A 1-second AGI would need to beat humans at tasks like quickly answering trivia questions, basic intuitions about physics (e.g. "what happens if I push a string?"), recognizing objects in images, recognizing whether sentences are grammatical, etc.
- A 1-minute AGI would need to beat humans at tasks like answering questions about short text passages or videos, common-sense reasoning (e.g. Yann LeCun's gears problems), simple computer tasks (e.g. use photoshop to blur an image), justifying an opinion, looking up facts, etc.
- A 1-hour AGI would need to beat humans at tasks like doing problem sets/exams, writing short articles or blog posts, most tasks in white-collar jobs (e.g. diagnosing patients, giving legal opinions), doing therapy, doing online errands, learning rules of new games,

t-AGI

"on most cognitive tasks, it beats most human experts who are given time t to perform the task"

Existing systems:

Are 1-second AGIs

Are close to 1-minute AGIs

Are a couple of years off from 1-year AGIs

Tools Round-Up

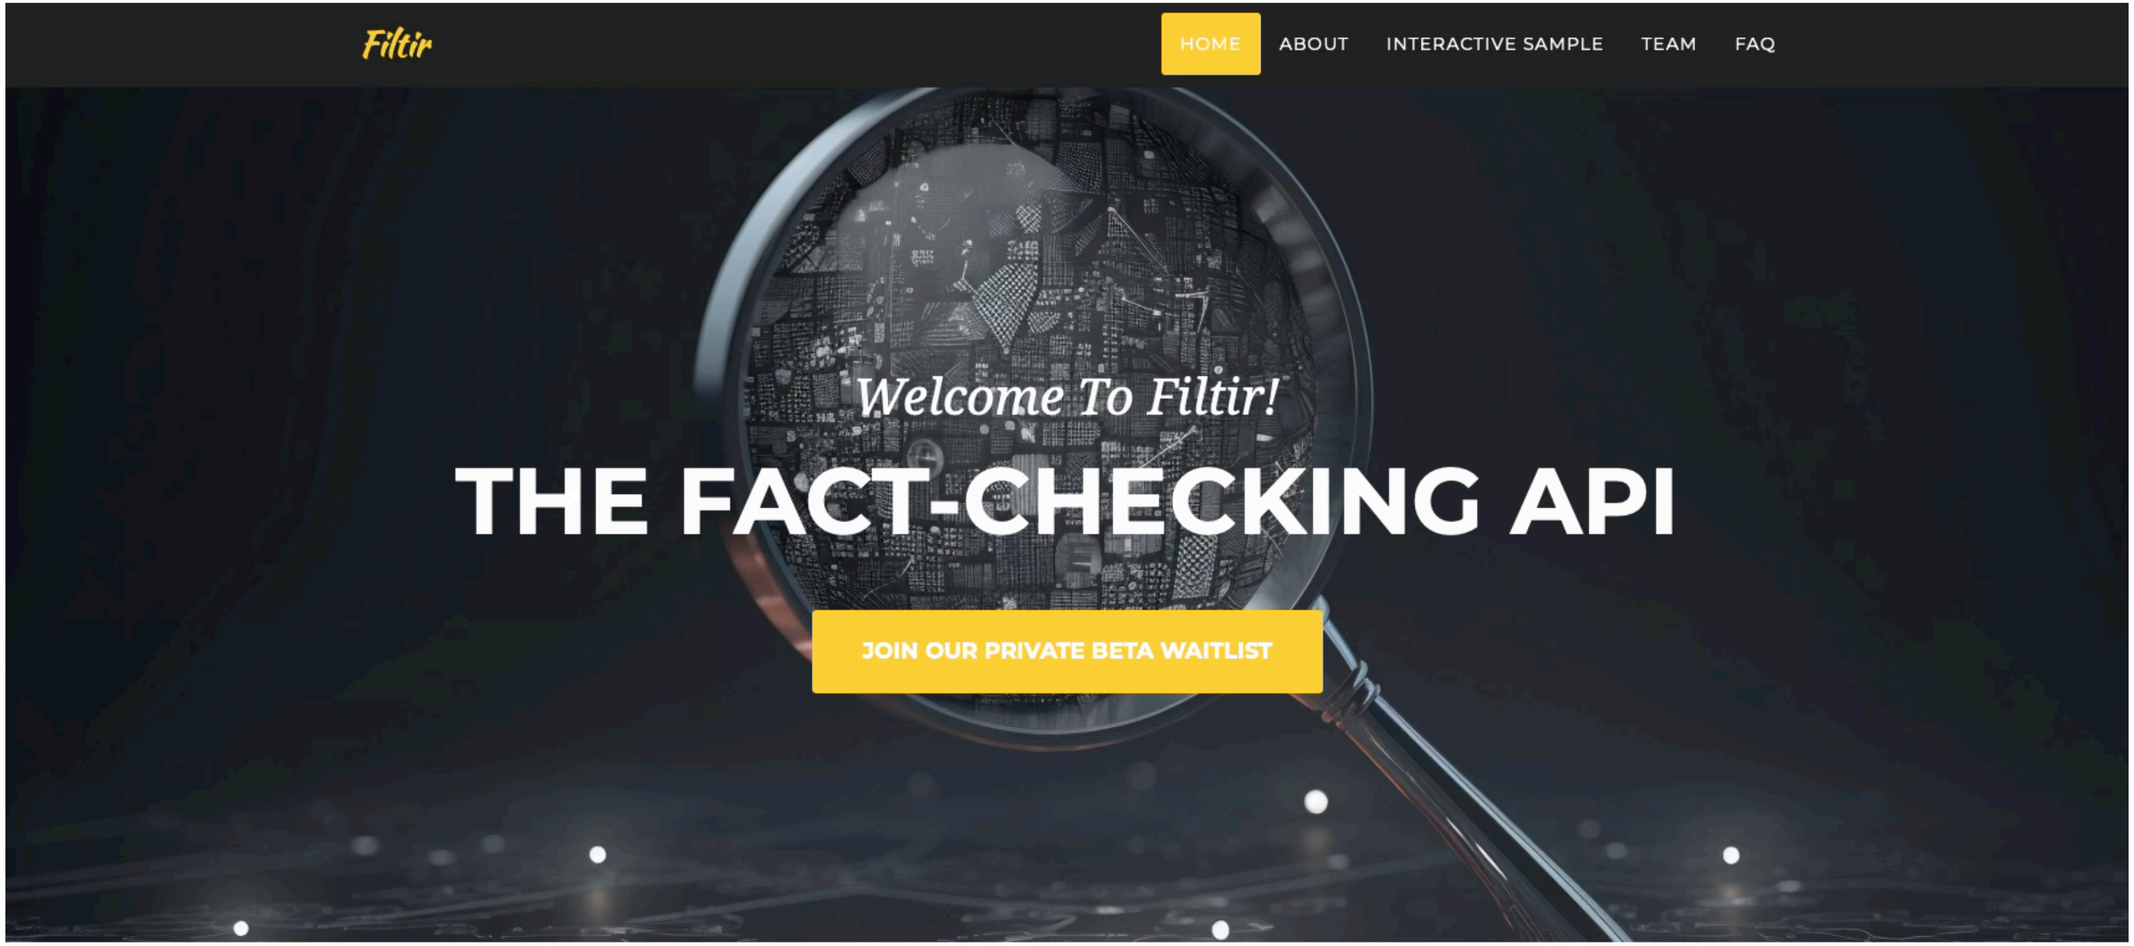

Midjourney v5.1

DavidH 03/05/2023 06:32

Hey @everyone we've got a few announcements tonight:

1) **We're testing a version 5.1 image system**
- V5.1 is more opinionated (like V4) and is MUCH easier to use with short prompts
- There is a 'unopinionated' mode for V5.1 (similar to V5.0 default) called "**RAW Mode**"
- Don't worry V5 is still available (as before)

Other changes in V5.1 vs V5.0
- Higher coherence
- More accuracy to text prompts
- Fewer unwanted borders or text artifacts
- Improved sharpness

easier to to use with short prompts

higher coherence

fewer unwanted borders or text artifacts

"Lammas Land at dusk in the year 2030"

main 1 branch 0 tags Go to file Add file Code

minghao-wu Update README.md 48cfa49 2 days ago 105 commits

images	Add files via upload	2 weeks ago
README.md	Update README.md	2 days ago

README.md

LaMini-LM: A Diverse Herd of Distilled Models from Large-Scale Instructions

Minghao Wu, Abdul Waheed, Chiyu Zhang, Muhammad Abdul-Mageed, Alham Fikri Aji

MOHAMED BIN ZAYED UNIVERSITY OF ARTIFICIAL INTELLIGENCE · UBC THE UNIVERSITY OF BRITISH COLUMBIA · MONASH University

LaMini-LM is a collection of small-sized, efficient language models distilled from ChatGPT and trained on a large-scale dataset of 2.58M instructions. We explore different model architectures, sizes, and checkpoints, and

About

LaMini-LM: A Diverse Herd of Distilled Models from Large-Scale Instructions

Readme
526 stars
22 watching
24 forks
Report repository

Releases

No releases published

Packages

No packages published

Contributors 4

chiyuzhang94 chiyu
afaji Alham Fikri Aji
macabdul9 Abdul
minghao-wu Minghao Wu

Environments 1

LaMini-LM

A Diverse Herd of Distilled Models from Large-Scale (2.58M) instructions

Microsoft Designer

27th April 2023

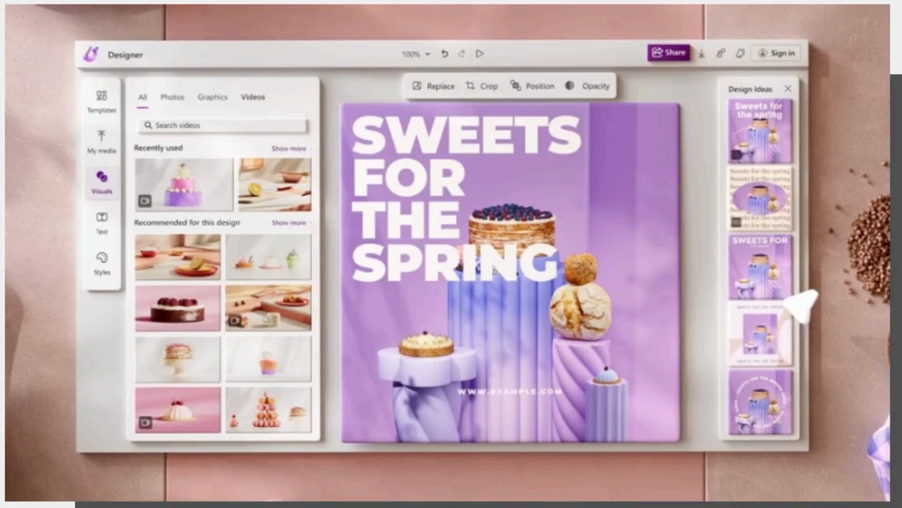

News Microsoft 365 Enterprise · 4 min read

Microsoft Designer expands preview with new AI design features

By Bryan Rognier, General Manager, Microsoft 365 Consumer

April 27, 2023

Edge

Personal and family

Small business

Creativity is more important to individuals than ever before. This reflects a trend that has added more than 165 million creators to the global creator economy in just the last three years.[1] As a result, people demand tools that help them to be both productive and creative. Microsoft 365 strives to empower individuals to achieve great things by constantly evolving our products to meet their changing needs. We continue to demonstrate this commitment with new tools that help unleash creativity and imagination by enabling any type of digital ideation and creation—no professional skills required. Today, we're excited to announce we're removing the waitlist and adding an expanded set of features to the **Microsoft Designer** preview. With new AI technology at the core, Microsoft Designer

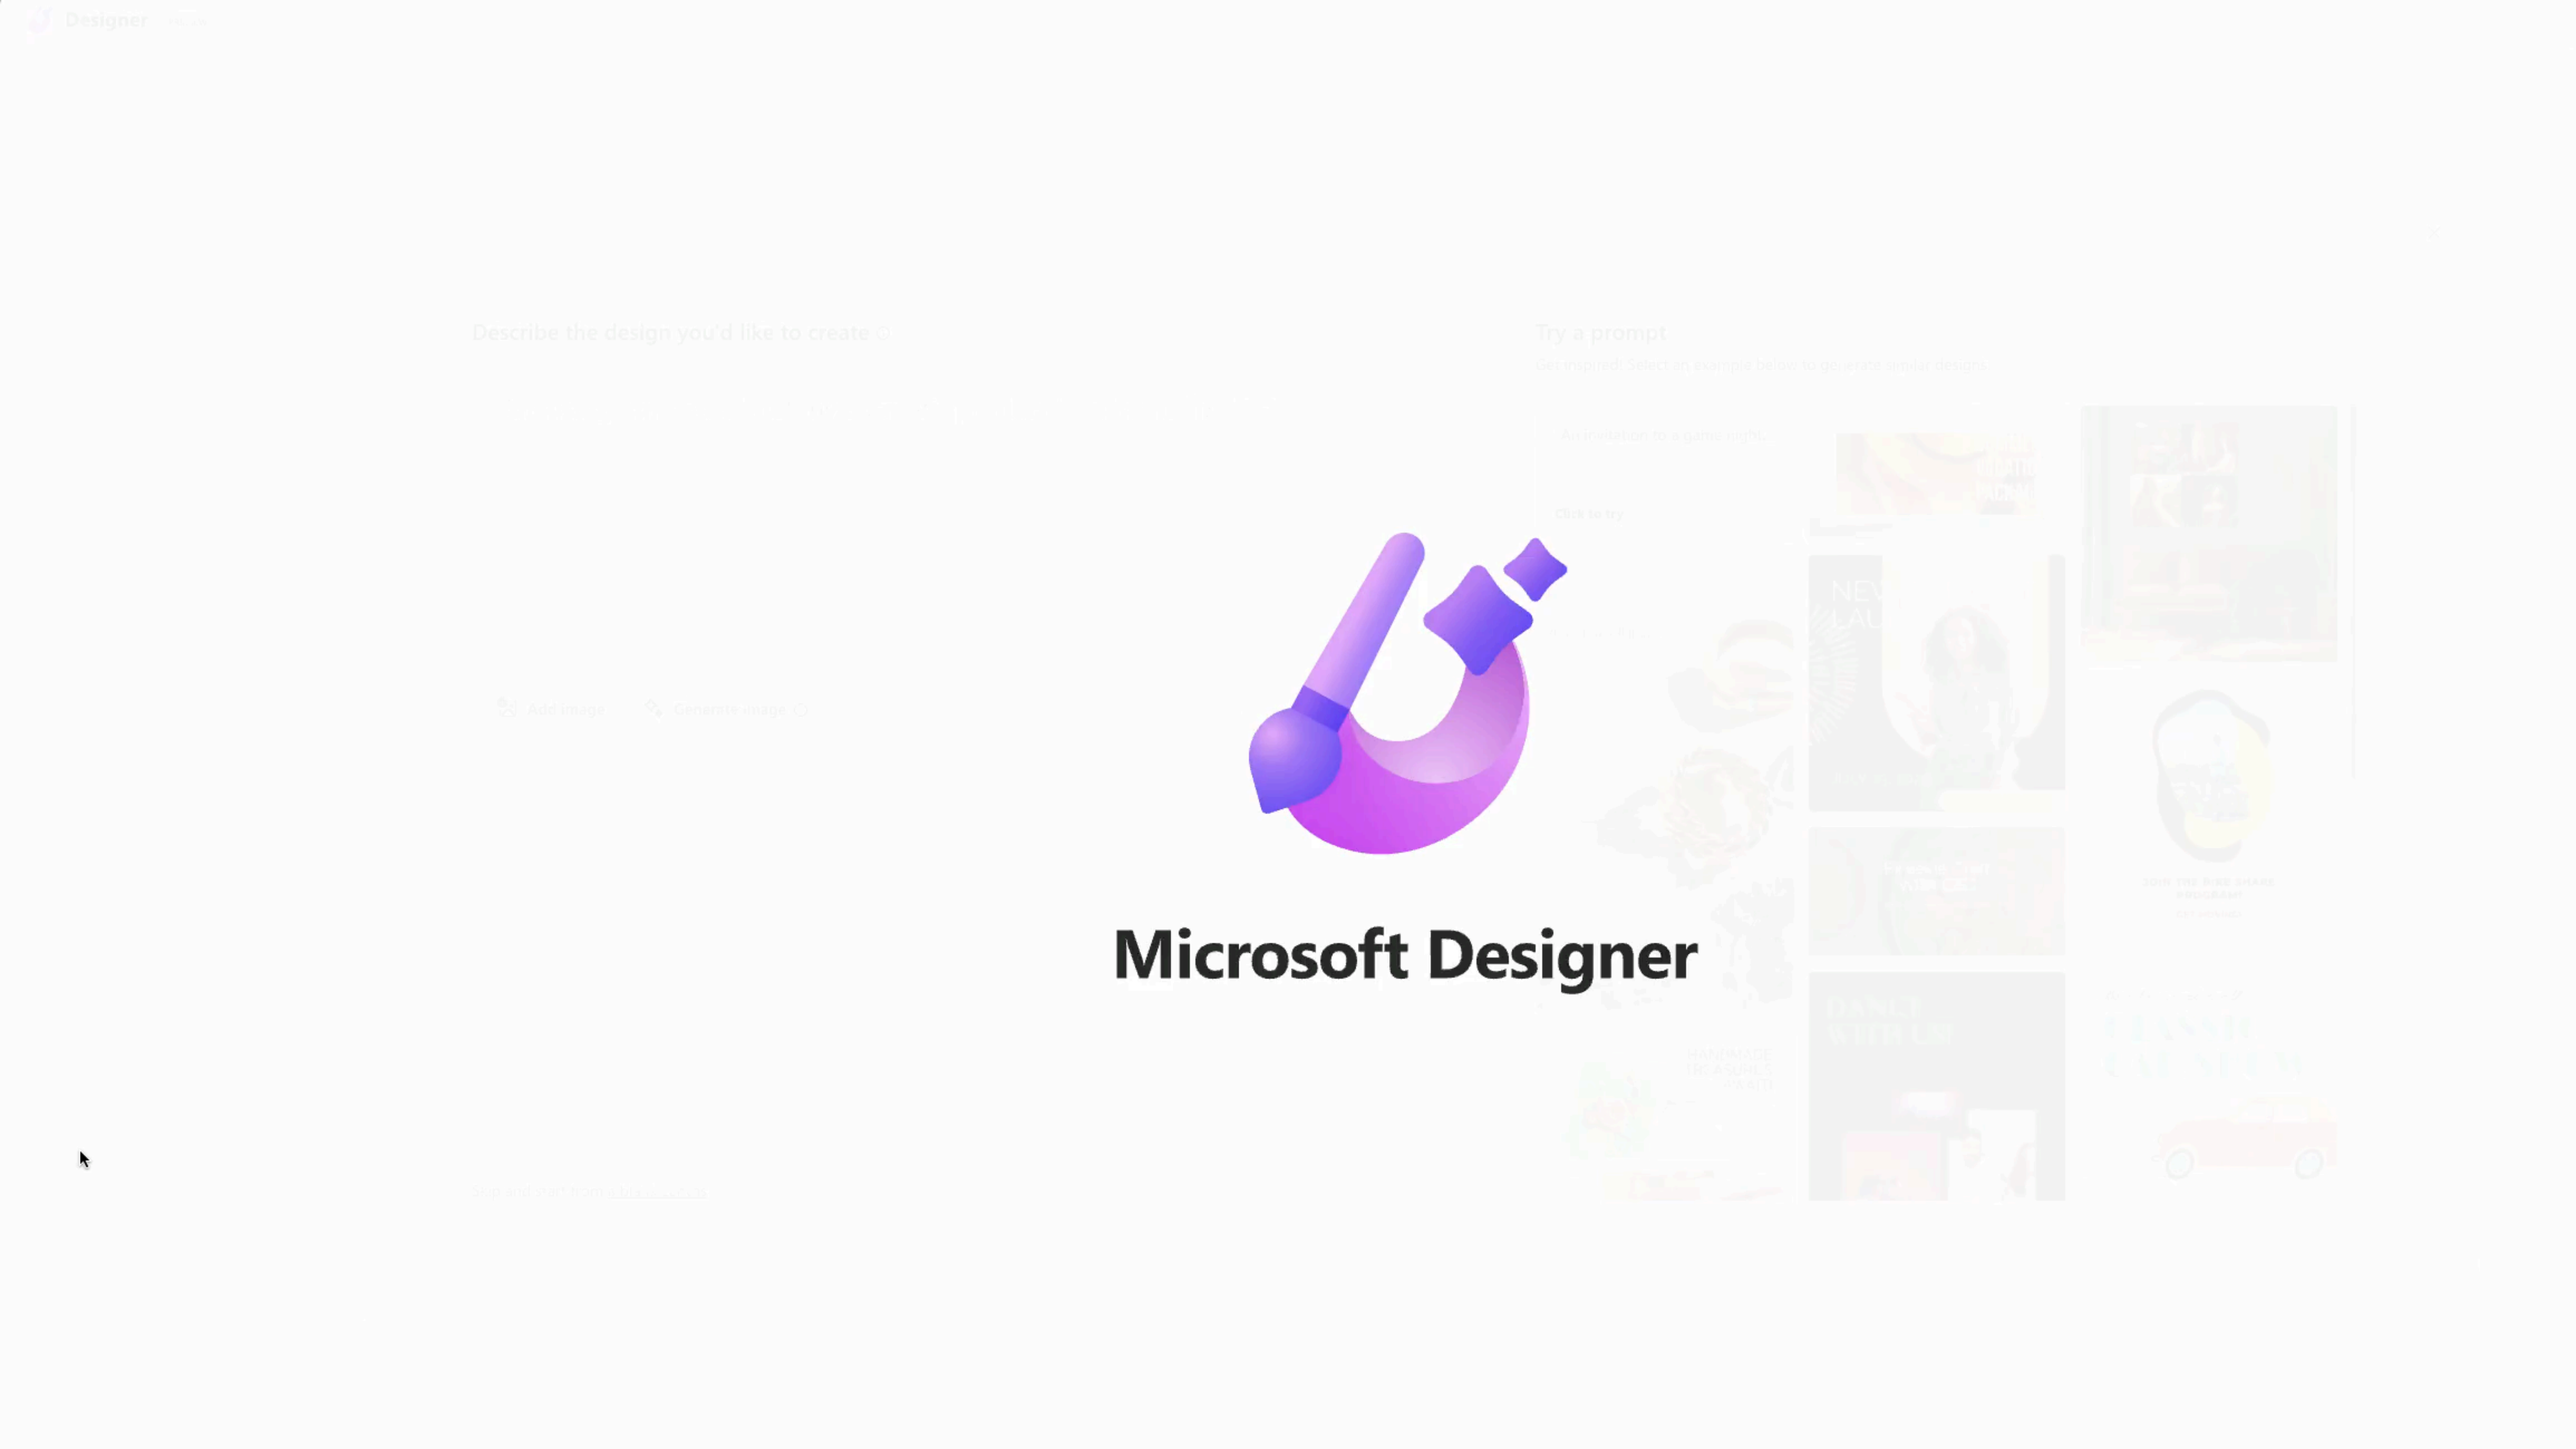
Microsoft Designer

Movie Trailer Generation