RACCOONS PRESENTS:

Generative AI: the Potential and the Pitfalls

raccoons

Hi there!

Michiel Vandendriessche
Co-founder Raccoons
michiel@raccoons.be

Sam Hendrickx
Co-founder Raccoons
sam@raccoons.be

Hi there!

Matt Bytecaster
AI co-presenter Raccoons
hello@raccoons.be

Michiel Vandendriessche
Co-founder Raccoons
michiel@raccoons.be

Sam Hendrickx
Co-founder Raccoons
sam@raccoons.be

WHAT WE'LL COVER TODAY

What is Raccoons

An introduction to generative artificial intelligence

Use cases of generative AI

Generative AI for business

Limitations & considerations

What will the future bring for generative AI

NEXT UP:

What is Raccoons

We create agile teams tailored to your challenge

raccoons

What we do

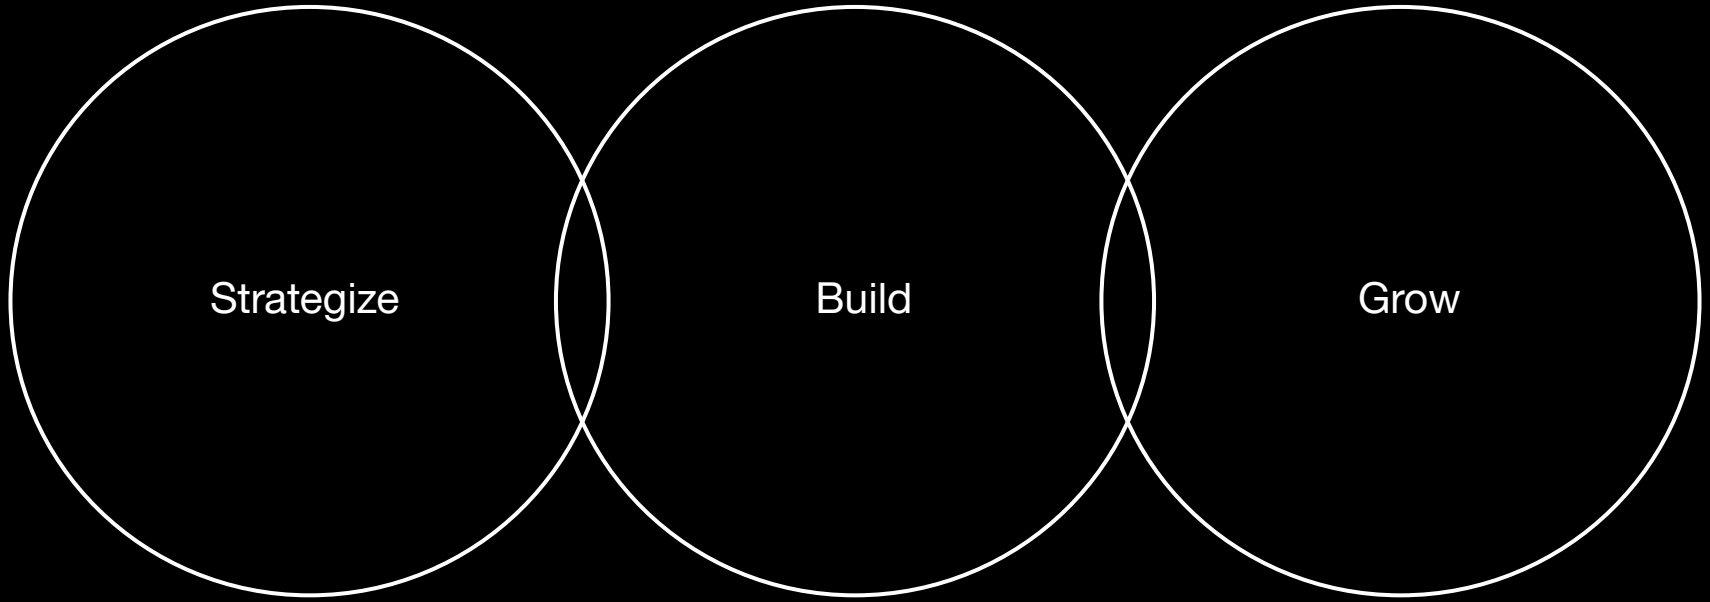

Trusted by industry leaders

NEXT UP:

An introduction to generative artificial intelligence

Artificial intelligence is the theory and development of computer systems able to perform tasks normally requiring human intelligence

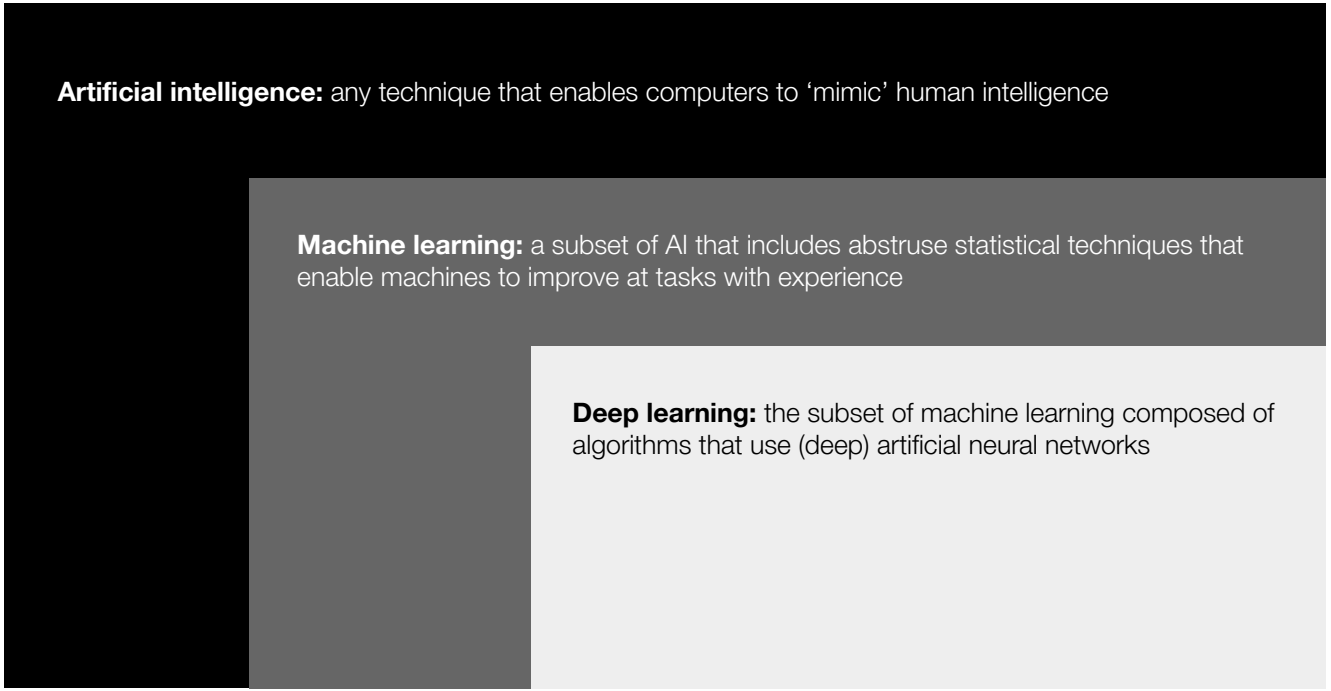
Artificial intelligence: any technique that enables computers to ‘mimic’ human intelligence
Machine learning: a subset of AI that includes abstruse statistical techniques that enable machines to improve at tasks with experience
Deep learning: the subset of machine learning composed of algorithms that use (deep) artificial neural networks

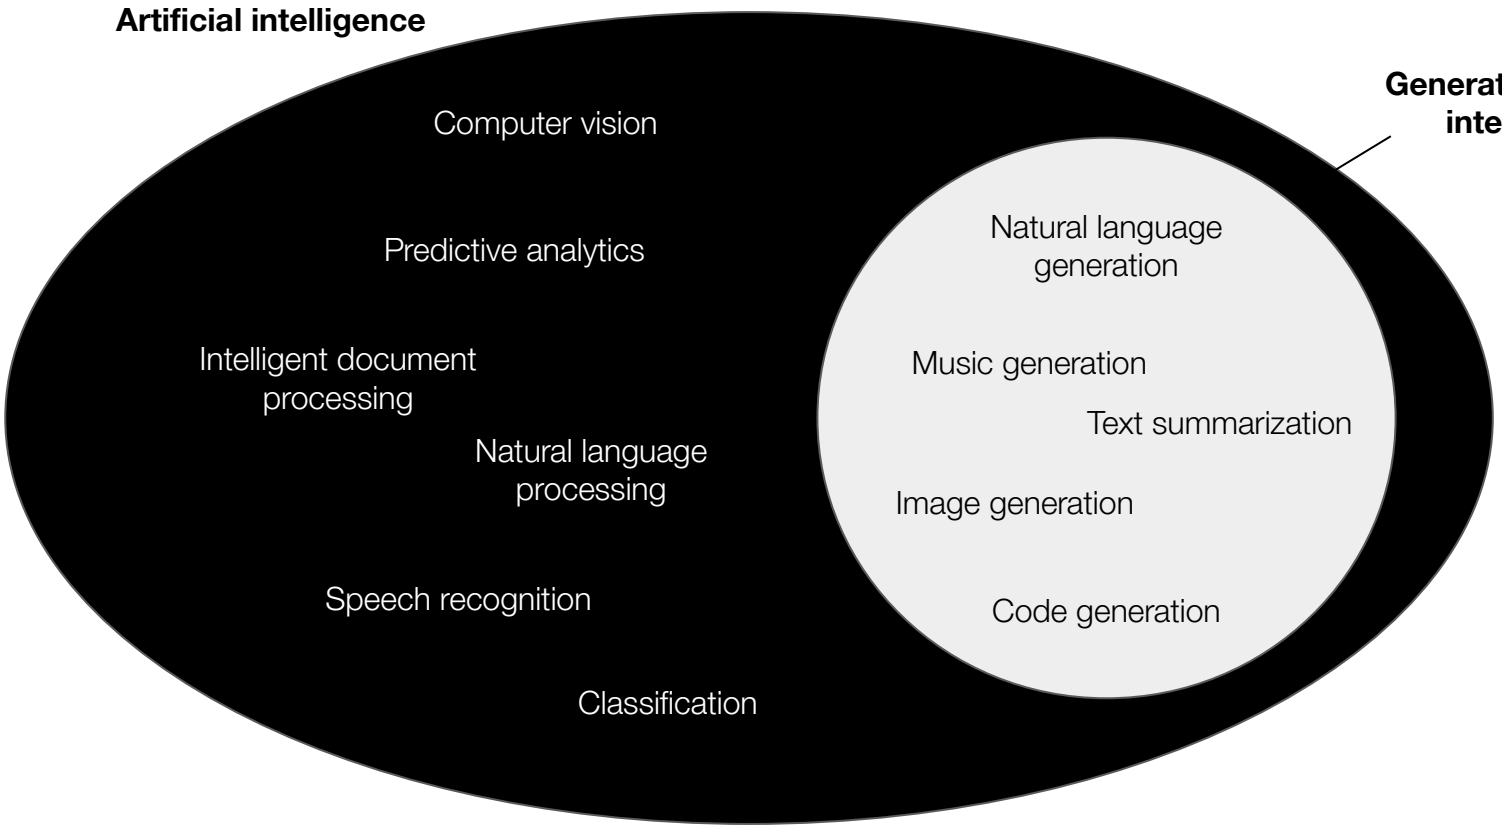
Artificial intelligence
Computer vision
Predictive analytics
Intelligent document processing
Natural language processing
Speech recognition
Classification
Generative artificial intelligence
Natural language generation
Music generation
Text summarization
Image generation
Code generation

Generative AI is not new, but in 2022 there has been a major shift towards releasing these systems to a broader audience.

Generative AI methods

The most prevalent techniques for generative AI are:

- Generative adversarial networks (GANs)
- Transformer networks
- Diffusion networks

NEXT UP:

Use cases of generative AI

Text

In the domain of Natural Language Processing (NLP), generative models are used to perform a wide range of tasks such as:

- **Text generation**: Generating new text that is similar to a given input text, such as writing a story or a news article.
- **Text summarization**: Generating a short summary of a given text, such as a news article or a book.
- **Dialogue systems**: Generating responses for chatbots or virtual assistants that can carry out a conversation with a human.

And many more...

ChatGPT

ChatGPT is a large language model developed by OpenAI that can generate human-like text and complete a wide range of language tasks with high accuracy.

- Generates **human-like text**
- Based on **GPT-4** LLM
- Can complete various **language tasks**
- **Fine-tunable** for domain knowledge
- **“Plus” version** available
- Rolling out **“Plugins”** marketplace

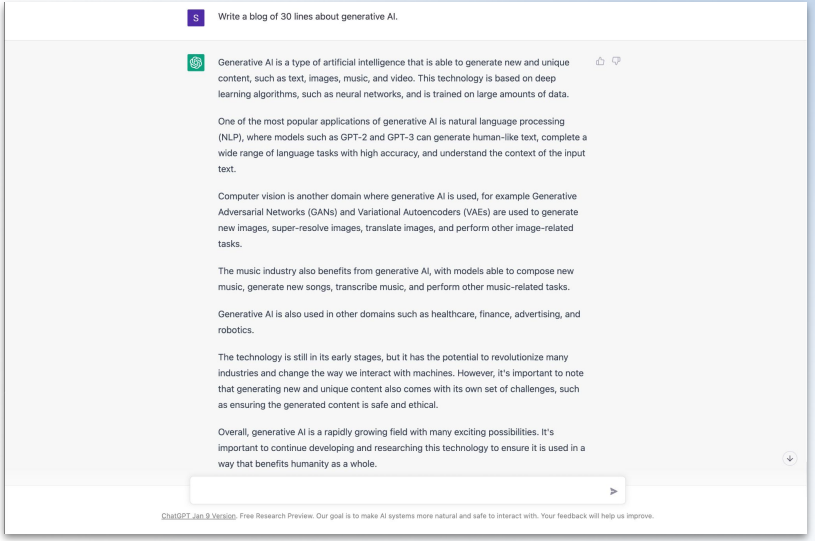

DEMO: chat.openai.com

Microsoft Bing

Microsoft released a new version of their search engine “Bing” that now uses a large language model to combine text generation with search.

- Bing runs on a **next-generation OpenAI LLM** customized for search
- Microsoft claims that it is **more capable** and accurate than ChatGPT or GPT-3.5
- Combines text generation with **search**
- Also incorporates **paid ads** in generated content

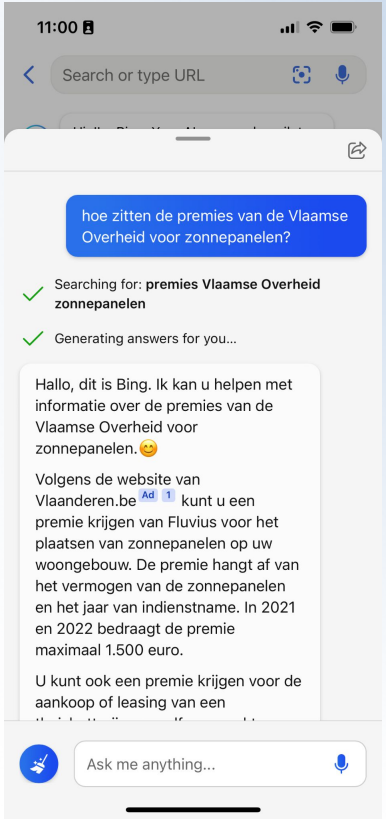

DEMO: bing.com
(limited availability today)

Google Bard

Google released their own version of ChatGPT/Bing called **“BARD”**. It seemed that Google rushed to the game a little too soon since the demo wasn’t received that well.

- Also **combines search with text generation**
- Google developed their **own AI models (LaMDA)** for Bard
- More to come in the next upcoming months

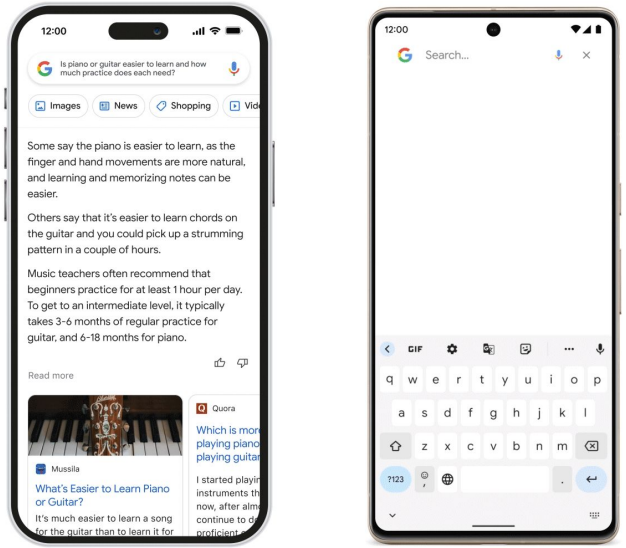

DEMO: Not available at the moment

OpenAI API

OpenAI launched a suite of API's that allow developers to integrate a wide spectrum of different language generation models into their own applications.

- **Text completion**: based on the GPT-3 model, with fine-tuning possibilities
- **Text editing**: based on GPT-3
- **Text insertion**: based on GPT-3
- **Chatting**: based on the newer GPT-3.5 turbo and GPT-4 language models, that power ChatGPT

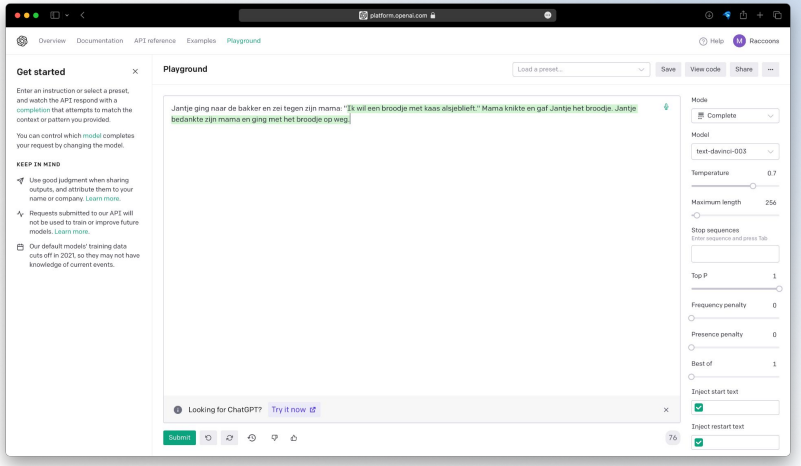

DEMO: platform.openai.com

Images & video

In the domain of Computer Vision, generative models are used to perform a wide range of tasks such as:

- **Image synthesis**: Generating new images that are similar to a given input image, such as creating a new picture of a person's face or a landscape.
- **Super-resolution**: Generating high-resolution images from low-resolution images, such as increasing the resolution of an image taken with a low-quality camera.
- **Image-to-image translation**: Translating an image from one domain to another, such as converting a picture of a day scene to a night scene or a sketch to a photograph.

And many more…

DALL·E 2

DALL·E 2 is an AI model created by OpenAI that generates images from natural language text, and can also complete text based on a given prompt.

- Fine-tuned version of DALL·E, which was released in December 2020.
- Able to **generate a wide range of images from text**, from simple drawings to more complex and realistic images.
- Also able to complete text prompts in a creative and interesting way, making it a versatile tool for creative tasks.

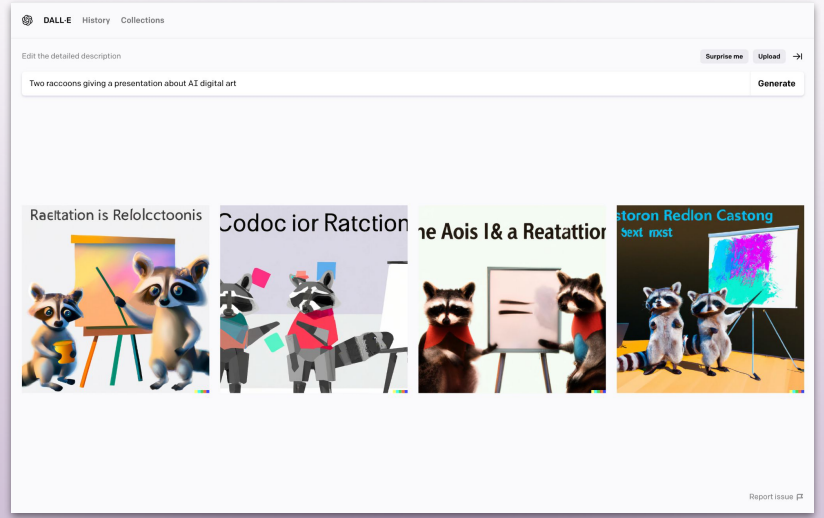

DEMO: openai.com/dall-e-2/

Stable Diffusion

Stable Diffusion is a latent text-to-image diffusion model capable of generating photo-realistic images given any text input, cultivates autonomous freedom to produce incredible imagery.

- Latent text-to-image diffusion model that generates photo-realistic images given text input
- Cultivates autonomous freedom to produce high-quality imagery
- Empowers users to create stunning art within seconds.

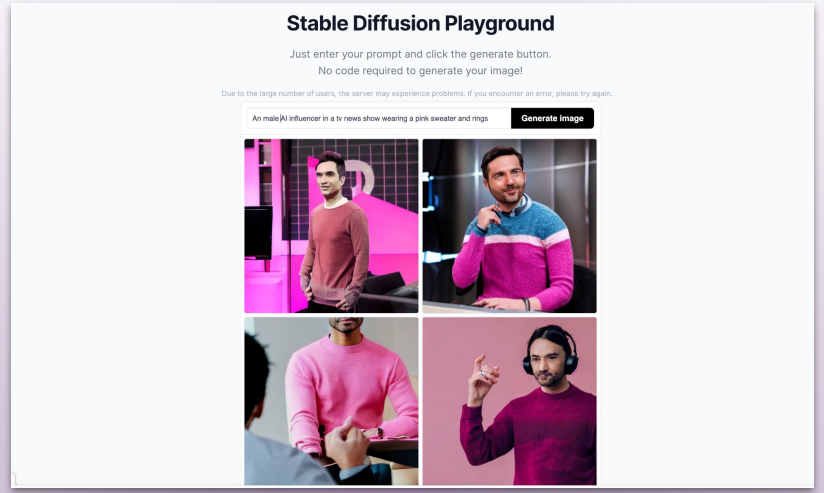

DEMO: stablediffusionweb.com/#demo

Midjourney

Midjourney is an independent research lab exploring new mediums of thought and expanding the imaginative powers of the human species.

- Use natural language to generate images
- If often used for character design, web design…
- Accessible via Discord

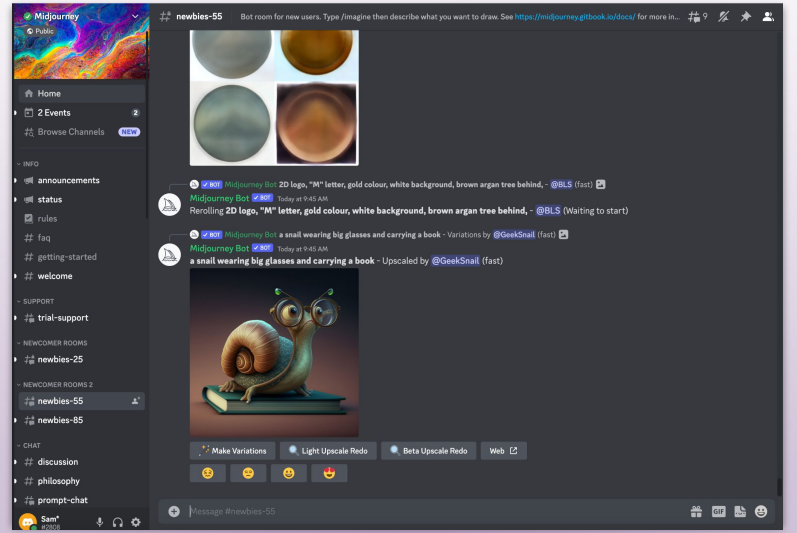

DEMO: midjourney.com

Bing Image Creator

Bing Image Creator is the image generator from Microsoft based on Dalle·E-2. The results are equally impressive as Midjourney.

- Use natural language to generate images
- Based on Dalle·E-2 model by OpenAI
- Integrated into Bing chat experience and Microsoft Edge browser

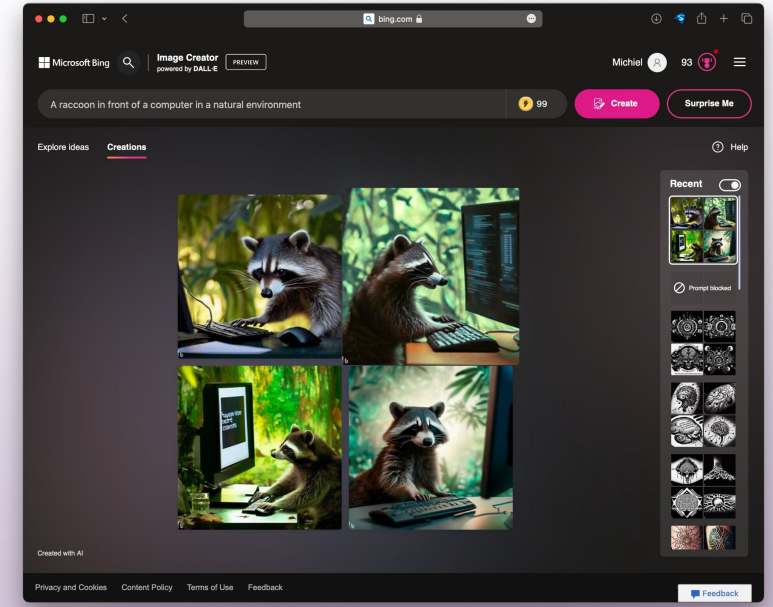

DEMO: www.bing.com/images/create/

Make-A-Video

Make-A-Video is a generative AI model that allows users to create videos automatically from text, images, or other multimedia inputs.

- Automates the video creation process
- Can use text, images, or other multimedia inputs
- Outputs a video as the final product

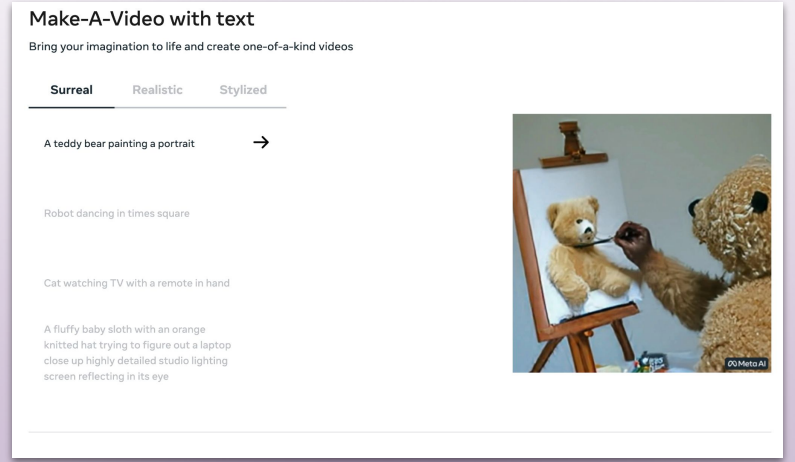

DEMO: makeavideo.studio/

D-ID

D-ID is a video generation tool that uses images from people to create a video of that person speaking a provided audio file of text message.

- Input an image and an audio file or text message
- Generates a video from the static image
- API for developers available

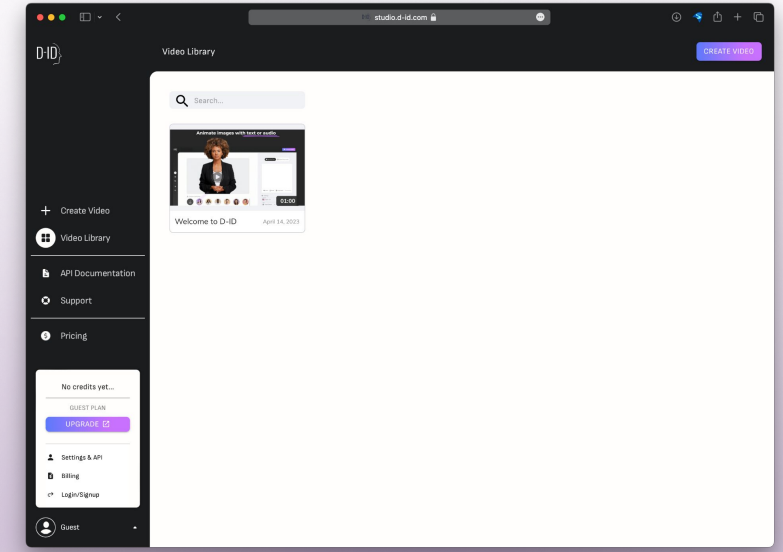

DEMO: d-id.com

Audio & music

In the domain of audio, generative models are used to perform a wide range of tasks such as:

- **Voice cloning:** generate a text-to-speech engine based on someone's voice
- **Music generation**: Creating a new piece of music from scratch, such as a new classical piece or a new pop song
- **Music style transfer**: Applying the style of one piece of music to another, such as converting a rock song to a classical piece

And many more…

ElevenLabs

ElevenLabs launched their VoiceLab which is their one-stop-shop for voice cloning and creating new voices. Samples uploaded to VoiceLab are instantly available for use in Speech Synthesis.

- **Voice cloning** based on very short example
- Instantly available afterwards as a **text-to-speech** engine
- **New voice generation** also possible *(without voice cloning)*
- Used to be a **very hard problem** to solve *(large dataset required)*

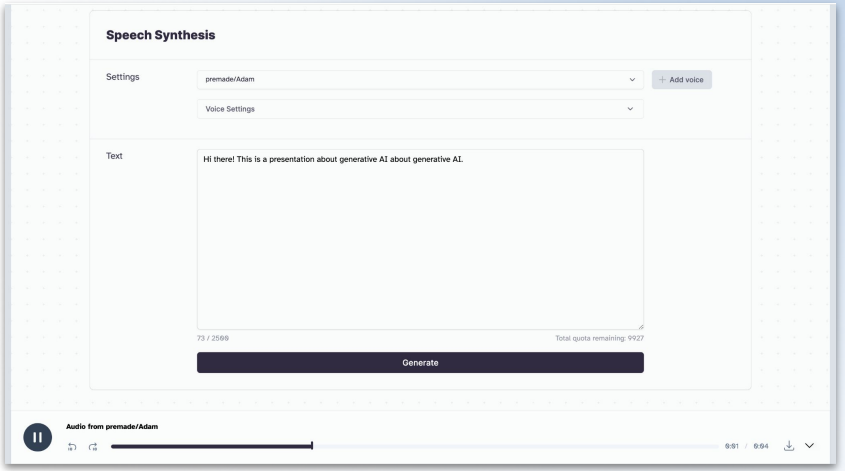

DEMO: elevenlabs.io

Riffusion

Riffusion is a text-to-music model. It creates an audio track from a description of music.

- Diffusion model based on text-to-image models.
- Able to **generate a wide range of audio samples from text**.

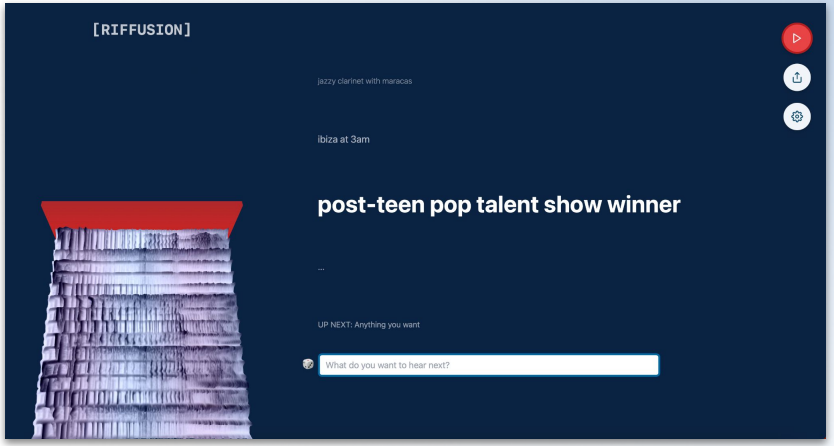

DEMO: riffusion.com

Coding

In the domain of code generation, generative models are used to perform a wide range of tasks such as:

- **Code generation**: Generating code that performs a specific task or solves a specific problem, such as sorting an array or finding the shortest path in a graph.
- **Code completion**: Generating code snippets that complete a given code, such as completing a missing function or a missing loop.
- **Code summarization**: Generating a summary of a given code, such as identifying the main methods or classes of a program.
- **Code explanation**: Generating an explanation of how a given piece of code works.

And many more…

GitHub Copilot

GitHub Copilot is an AI-powered autocomplete tool for developers, which suggests code snippets and helps complete tasks within GitHub.

- GitHub Copilot uses machine learning to understand the context of a developer's code and **suggests** relevant code snippets.
- It can help developers save time and reduce errors by **automating** repetitive tasks.
- Integrated with GitHub, so developers can access it directly within the platform.

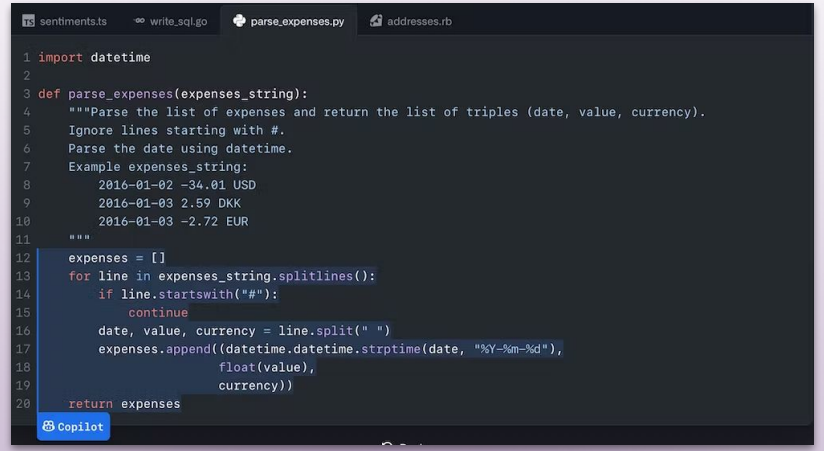

DEMO: github.com/features/copilot

Let’s bring it all together

Demo created by Sam Hendrickx, Michiel Vandendriessche and ChatGPT

NEXT UP:

Generative AI for business

Applications for businesses

Generative AI has the potential to revolutionize the way businesses operate by automating repetitive tasks, creating new products and designs, personalizing content and experiences and making data-driven decisions.

- **Content creation**: Generative AI can be used to generate new content, such as images, text, or audio.
  Example: a media company could use it to generate new headlines or article summaries.

- **Insight gathering**: Gather insights from large (un)structured datasets
  Example: get information about your best selling product based on transaction history in Excel

- **Product design and prototyping**: Generative AI can be used to generate new product designs and help businesses quickly prototype and test new ideas.
  Example: For example, a furniture company could use it to generate new furniture designs.

GuacMobile

Mash up the road in style

Sustainable fuel source: The GuacMobile runs on avocado oil, a renewable and sustainable fuel source that reduces dependence on fossil fuels and helps to reduce the car's environmental impact.
Low emissions: Avocado oil is a clean-burning fuel that generates low emissions, making the GuacMobile an environmentally friendly mode of transportation.
Cost-effective: Avocado oil is a relatively inexpensive fuel source, which can help to reduce the overall cost of owning and operating the GuacMobile.
Unique design: The GuacMobile's avocado-inspired design sets it apart from other cars on the market, making it a stylish and eye-catching choice for consumers.
Innovative technology: The GuacMobile is equipped with innovative technology, such as a custom-designed engine that can efficiently run on avocado oil, to enhance the driving experience and ensure a safe journey.

Custom domain knowledge in generative AI

The number one question we get is “can we incorporate our specific domain knowledge into existing general generative AI models?”. There are multiple ways to address this problem and it depends on the specific use case.

- **Prompt engineering**: By incorporating custom domain knowledge in prompts given to generative models, you can provide information that the model will take into account in the answers.

  Example: Our cheapest product costs 5 euros and has free shipping. Generate a response to this customer’s question…

- **Retraining existing models**: This depends on the fact if the model is open or closed sourced (however some closed models also allow fine-tuning).

A new era for AI and Google Workspace

NEXT UP:

Limitations & considerations

As with every innovation, with excitement, so comes risks. There are still ethical and practical considerations, maybe even risks, that you must consider when calling upon generative AI.

- **Generative AI can be wrong**
- **Systemic bias**
- **Using the right tool for the right reason**
- **Intellectual property rights**
- **Sustainability**
- **Safety & security**
- **Always put a human in the loop**
- **Models are often trained on historic data *(e.g. ChatGPT up until 2021)***

What about job displacement?

EXPERT VS TOOL?

EXPERT + TOOL > EXPERT

NEXT UP:

What will the future bring for generative AI

We believe there is a bright future ahead of us for generative AI. Maybe we will think back about the present time and see that this is definitely a new “iPhone moment”? What do you think the future will look like? This is our guess:

- **Newer and better models (GPT-4, GPT-5, Google PaLM, Google LaMDA)**
- **Almost all big-tech companies are stepping their game regarding generative AI (massive resources allocated)**
- **More legislation**
- **More and competitive generative AI tools for consumers: Google/Bing/…**
- **Integration into cloud platforms for developers: Microsoft Azure, Google Cloud…**
- **…**

The hype around generative AI does not mean we’re closer to artificial general intelligence.

FREE WHITEPAPER

How generative AI will disrupt Belgian businesses

raccoons


WHITEPAPER
How generative AI will disrupt Belgian businesses

raccoons