

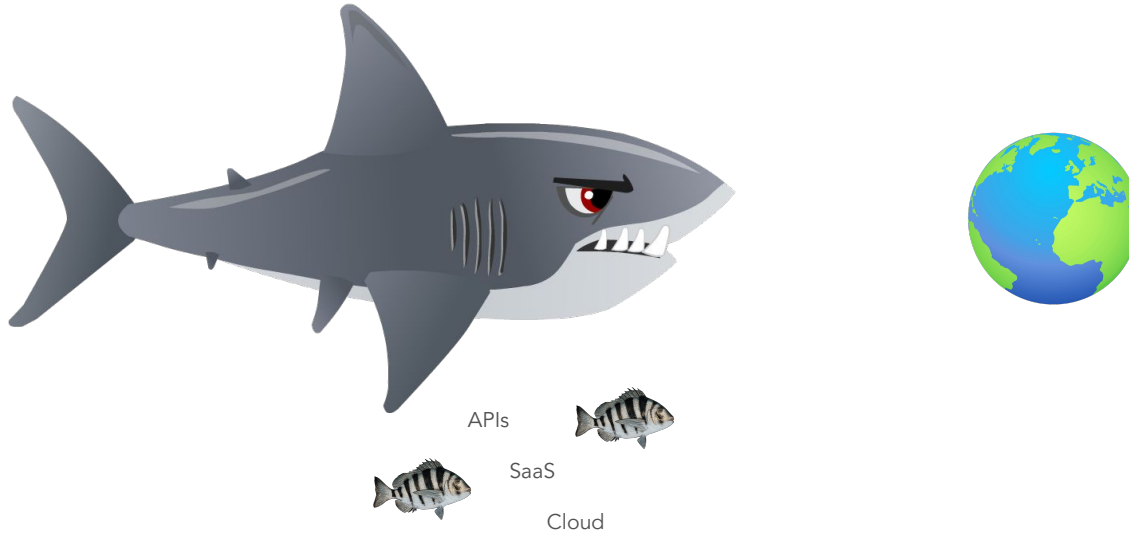


Generative AI at Hugging Face

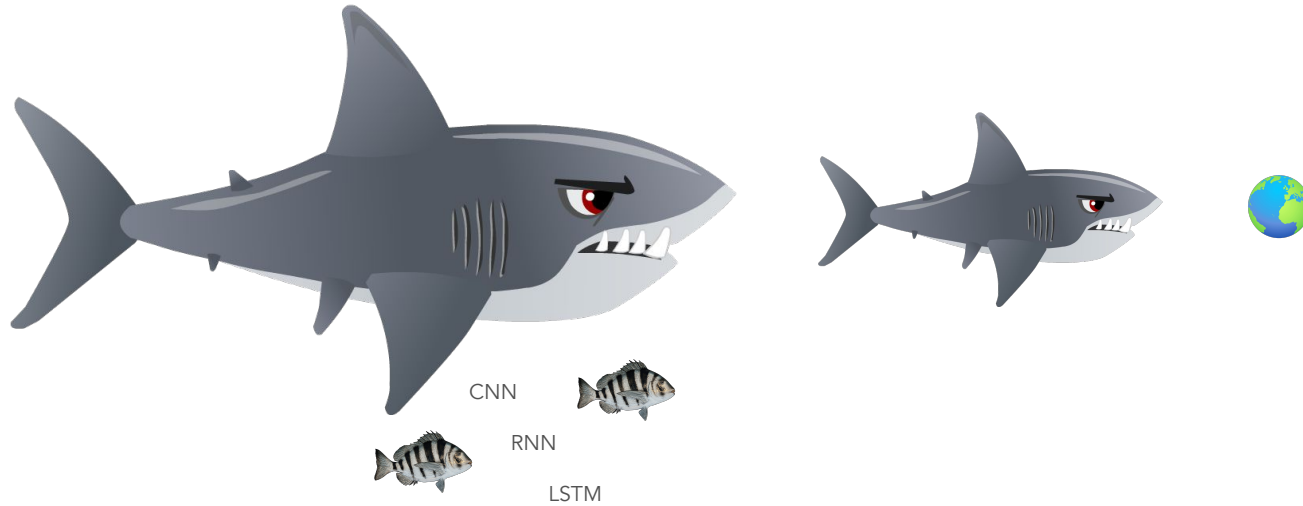
Thomas Wolf | Chief Science Officer @ Hugging Face

2011: Software is eating the world

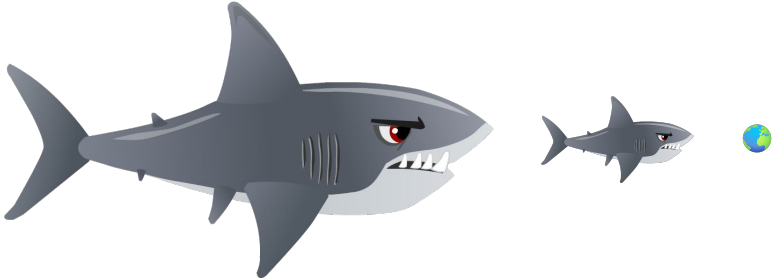
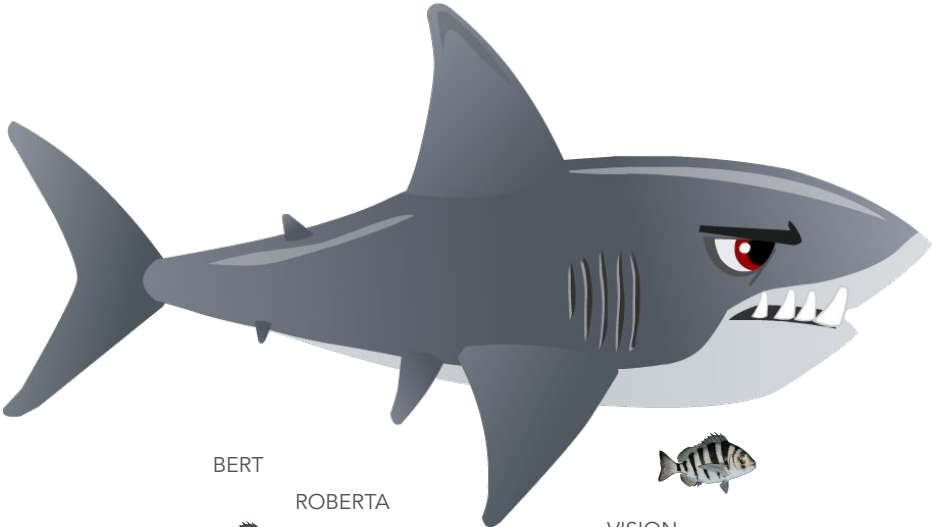
<https://a16z.com/2011/08/20/why-software-is-eating-the-world/>



2017: Deep Learning is eating Software



2022: Transformers are eating Deep Learning



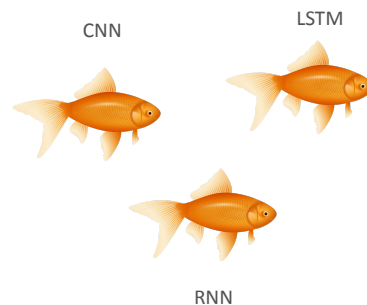
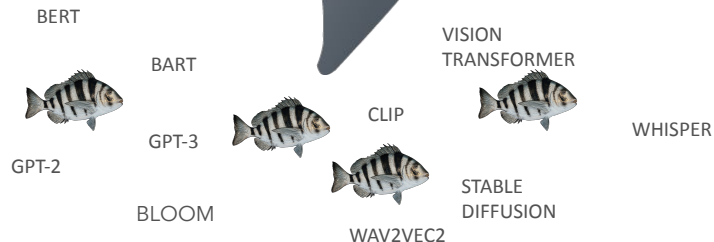
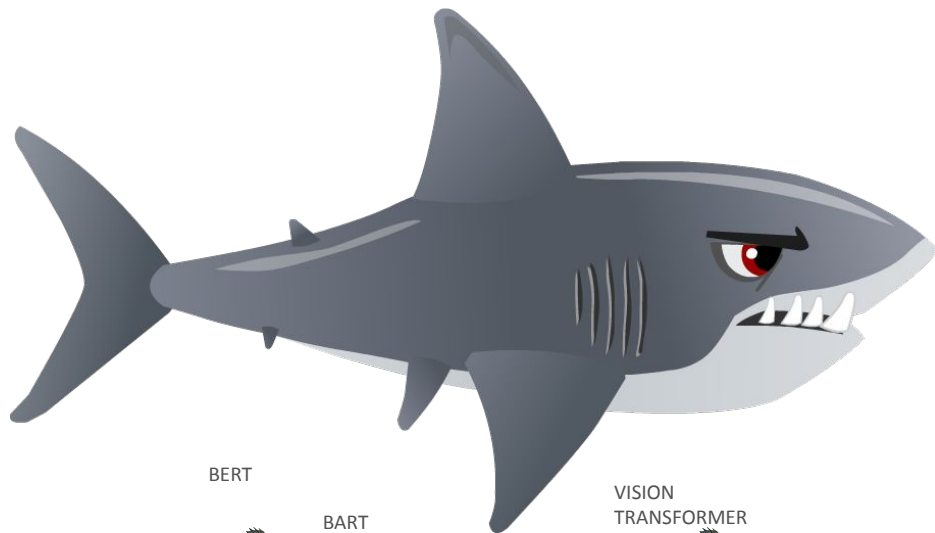
BERT
ROBERTA
GPT-2
BLOOM
GPT-3
SEGFORMER
CLIP
VISION TRANSFORMER
WAV2VEC2

"Transformers are emerging as a general-purpose architecture for ML"
<https://www.stateof.ai> (2021)

RNN and CNN usage down, Transformers usage up!
<https://www.kaggle.com/kaggle-survey-2021>



2022: Transformers are eating Deep Learning



"Transformers are emerging as a general-purpose architecture for ML"
<https://www.stateof.ai> (2021)

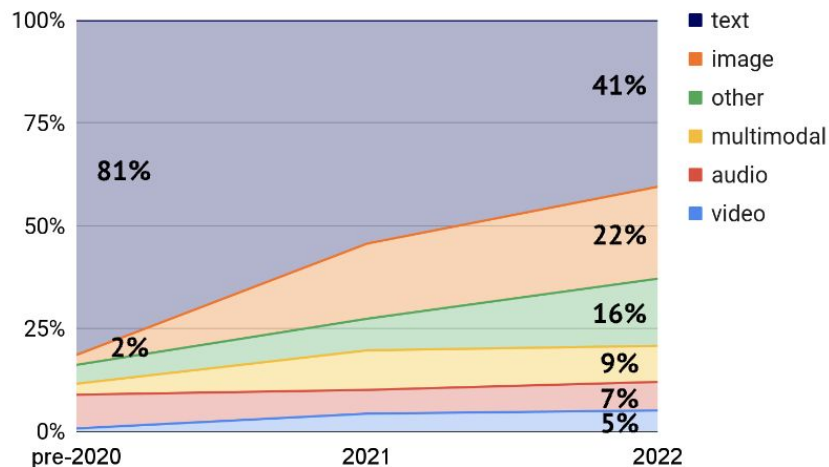
RNN and CNN usage down, Transformers usage up!
<https://www.kaggle.com/kaggle-survey-2021>



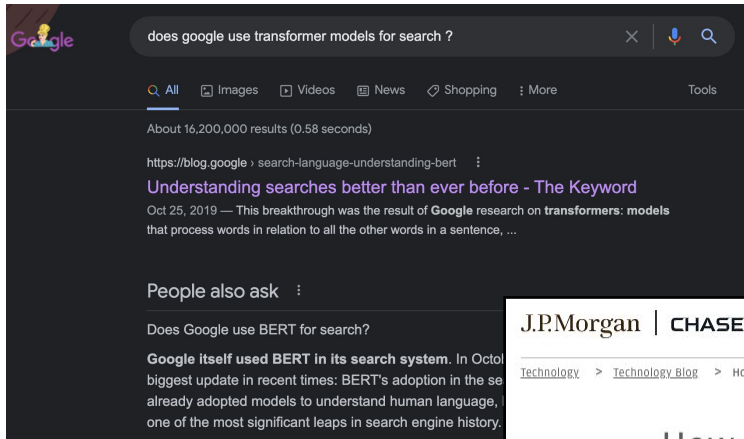
From text to multi-modal in 2 years

Transformers are becoming truly cross-modality

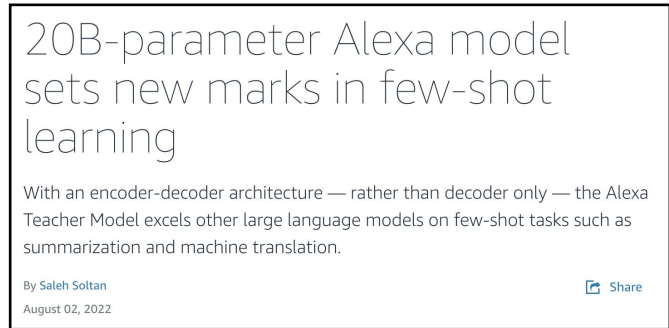
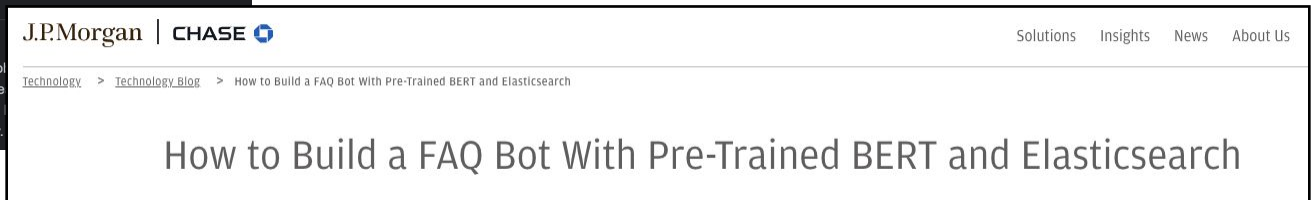
- ▶ In the 2020 State of AI Report we predicted that transformers would expand beyond NLP to achieve state of the art in computer vision. It is now clear that transformers are a candidate general purpose architecture. Analysing transformer-related papers in 2022 shows just how ubiquitous this model architecture has become.



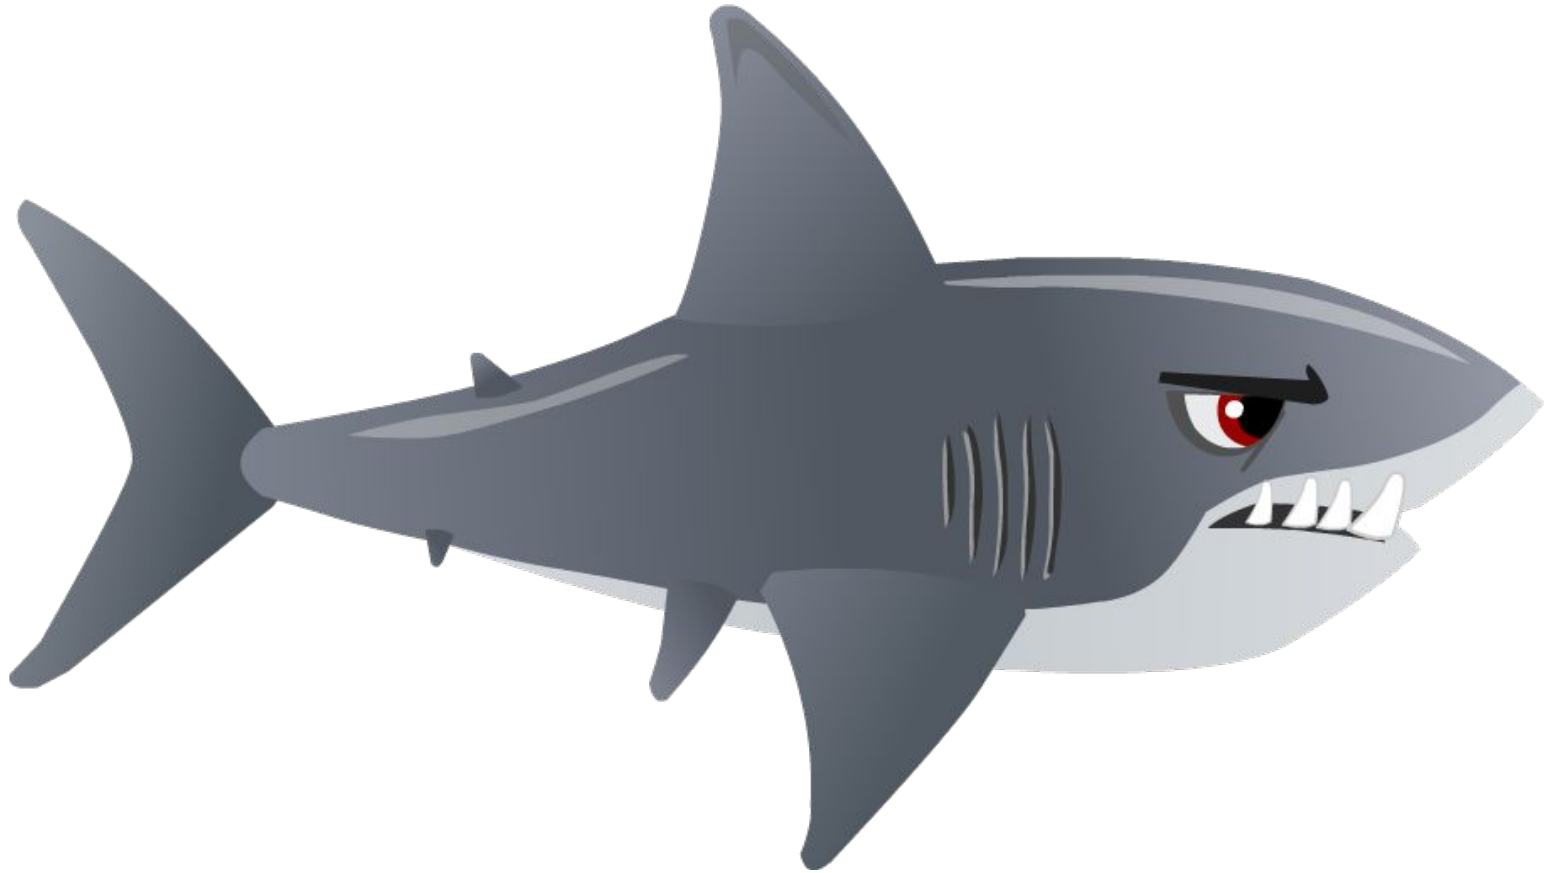
Transformer models in the wild



**Under the Hood at Grammarly:
Leveraging Transformer Language
Models for Grammatical Error Correction**



2023: Generative AI is hungry!



Generative models in the wild

JU

is ChatGPT powered by transformer models?



Yes, ChatGPT is powered by transformer models. Specifically, it is based on the architecture of the GPT (Generative Pre-trained Transformer) family of models, which were developed by OpenAI. The GPT models are designed to generate natural language text, and they use a transformer architecture that allows them to model long-term dependencies in language.

ChatGPT is a variant of the GPT models that has been fine-tuned on a large amount of conversational data in order to generate human-like responses to user input. It uses a combination of neural network architecture, training data, and machine learning algorithms to generate responses that are contextually appropriate and semantically relevant.

Introducing BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance

March 30, 2023

Meet Einstein GPT, the World's First Generative AI for CRM

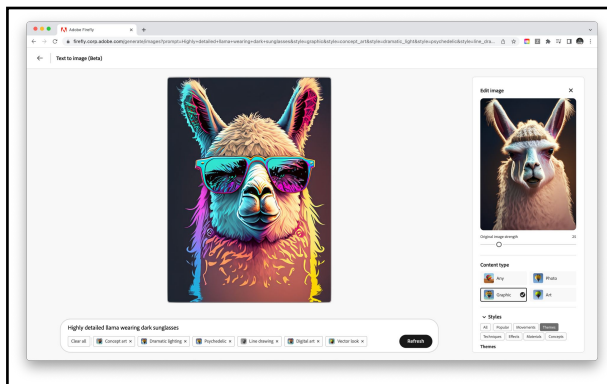
[LEARN MORE →](#)



GitHub Copilot



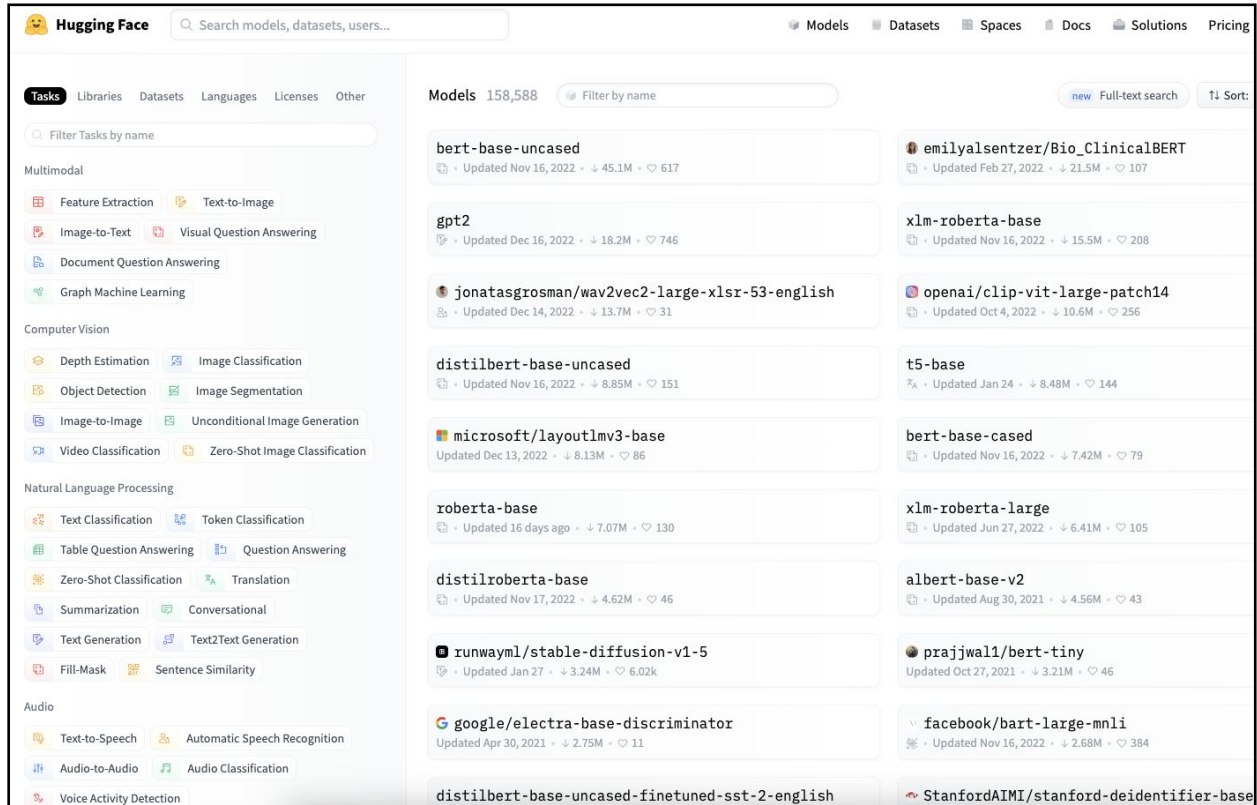
Amazon CodeWhisperer



Hugging Face



Hugging Face: the largest collection of open source models



The screenshot shows the Hugging Face homepage with a search bar at the top and navigation tabs for Models, Datasets, Spaces, Docs, Solutions, and Pricing. The main content area displays a grid of model cards, each with a name, update date, size, and popularity metrics. The left sidebar contains a 'Tasks' section with filters for Libraries, Datasets, Languages, Licenses, and Other, and a list of task categories such as Multimodal, Computer Vision, Natural Language Processing, and Audio.

Model Name	Updated	Size	Popularity
bert-base-uncased	Nov 16, 2022	45.1M	617
emilyalsentzer/Bio_ClinicalBERT	Feb 27, 2022	21.5M	107
gpt2	Dec 16, 2022	18.2M	746
xlm-roberta-base	Nov 16, 2022	15.5M	208
jonatasgrosman/wav2vec2-large-xlsr-53-english	Dec 14, 2022	13.7M	31
openai/clip-vit-large-patch14	Oct 4, 2022	10.6M	256
distilbert-base-uncased	Nov 16, 2022	8.85M	151
t5-base	Jan 24	8.48M	144
microsoft/layoutlmv3-base	Dec 13, 2022	8.13M	86
bert-base-cased	Nov 16, 2022	7.42M	79
roberta-base	16 days ago	7.07M	130
xlm-roberta-large	Jun 27, 2022	6.41M	105
distilroberta-base	Nov 17, 2022	4.62M	46
albert-base-v2	Aug 30, 2021	4.56M	43
runwayml/stable-diffusion-v1-5	Jan 27	3.24M	6.02k
prajjwal1/bert-tiny	Oct 27, 2021	3.21M	46
google/electra-base-discriminator	Apr 30, 2021	2.75M	11
facebook/bart-large-mnli	Nov 16, 2022	2.68M	384
distilbert-base-uncased-finetuned-sst-2-english			
StanfordAIMI/stanford-deidentifier-base			

<https://huggingface.co>

200K models

34K datasets

25+ ML libraries: Keras, spaCY, Scikit-Learn, fastai, etc.

10K organizations

500K+ users daily



Why open-source models

- **Accessibility:** anyone can use the models, regardless of budget or affiliation
- **Transparency:** users have full visibility on the model and the training data. They can better identify potential biases or errors
- **Privacy:** users don't have to send their data to black box APIs
- **IP protection:** users train their models on their data, and own them.
- **Freedom of choice:** users are not locked in. They can switch models anytime
- **IT flexibility:** users can train and deploy models anywhere they like



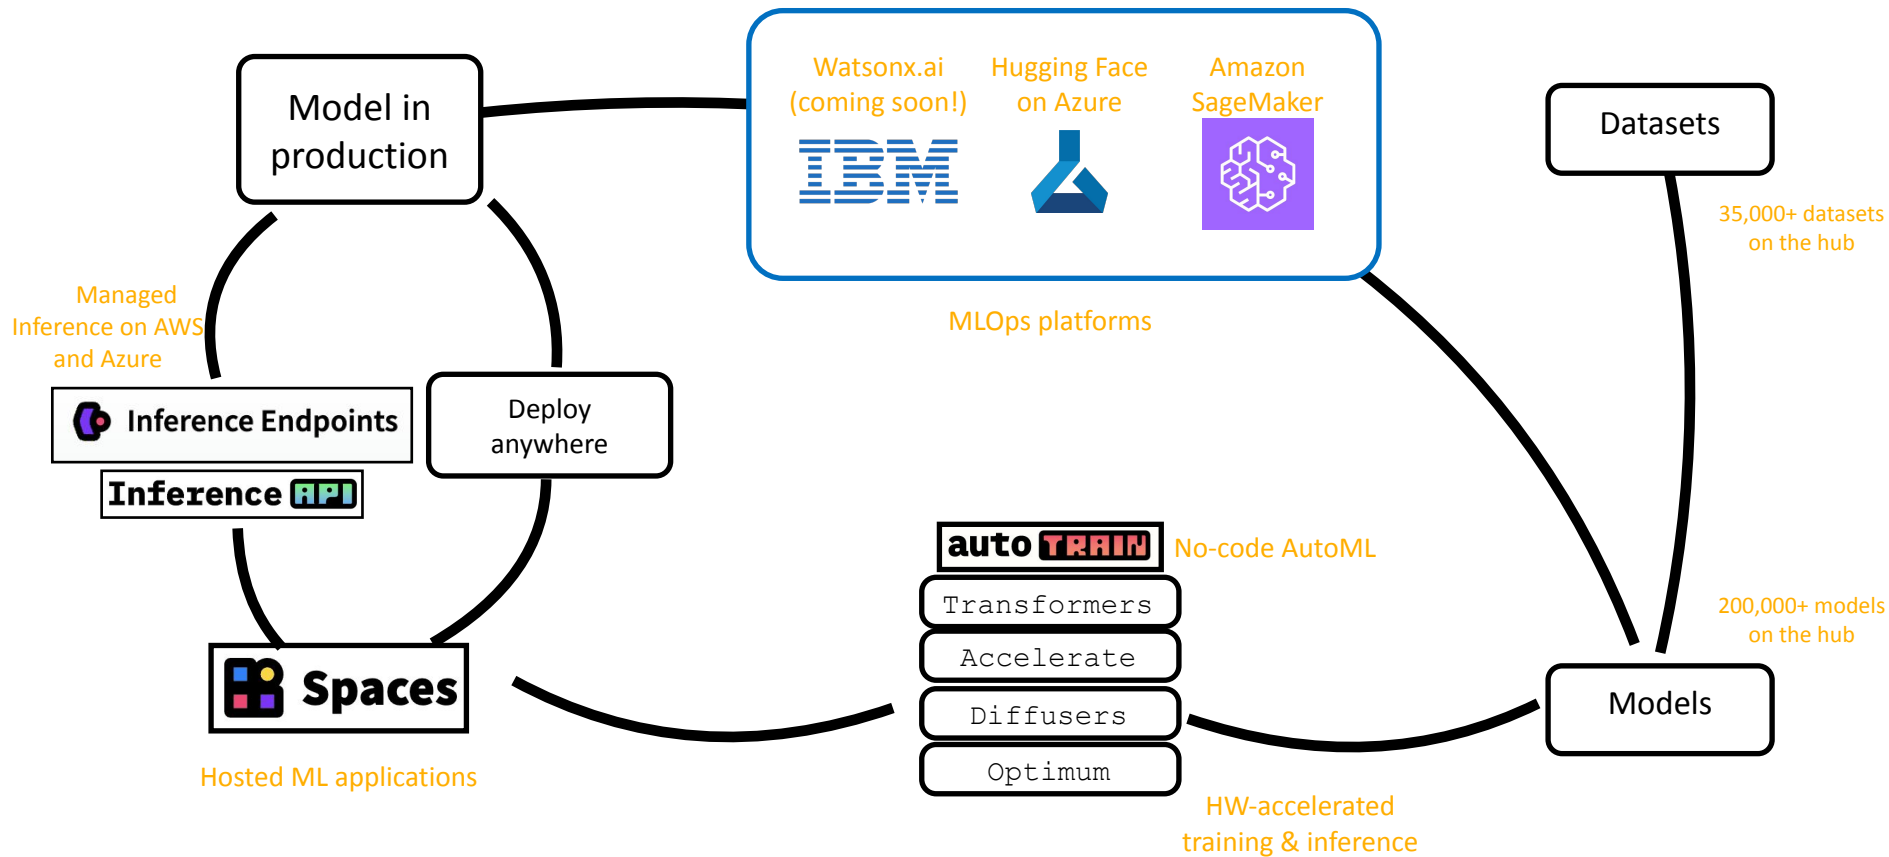
Hugging Face open source libraries

- **Transformers:** Transformer models for Pytorch, TensorFlow, and JAX
- **Diffusers:** image and audio generation models for PyTorch
- **Accelerate:** simple distributed training (CPU, GPU, TPU) for PyTorch
- **Peft:** Parameter Efficient Fine-Tuning on GPU
- **SetFit:** few-shot learning for Sentence Transformers
- **Datasets:** All the datasets for training and evaluating

- **Optimum:** hardware acceleration for Transformers and Diffusers
 - Optimum Intel: Intel Neural Compressor, Intel OpenVINO
 - Optimum Habana: training and inference for Habana Gaudi/Gaudi2
 - Optimum Neuron: AWS Trainium and Inferentia2



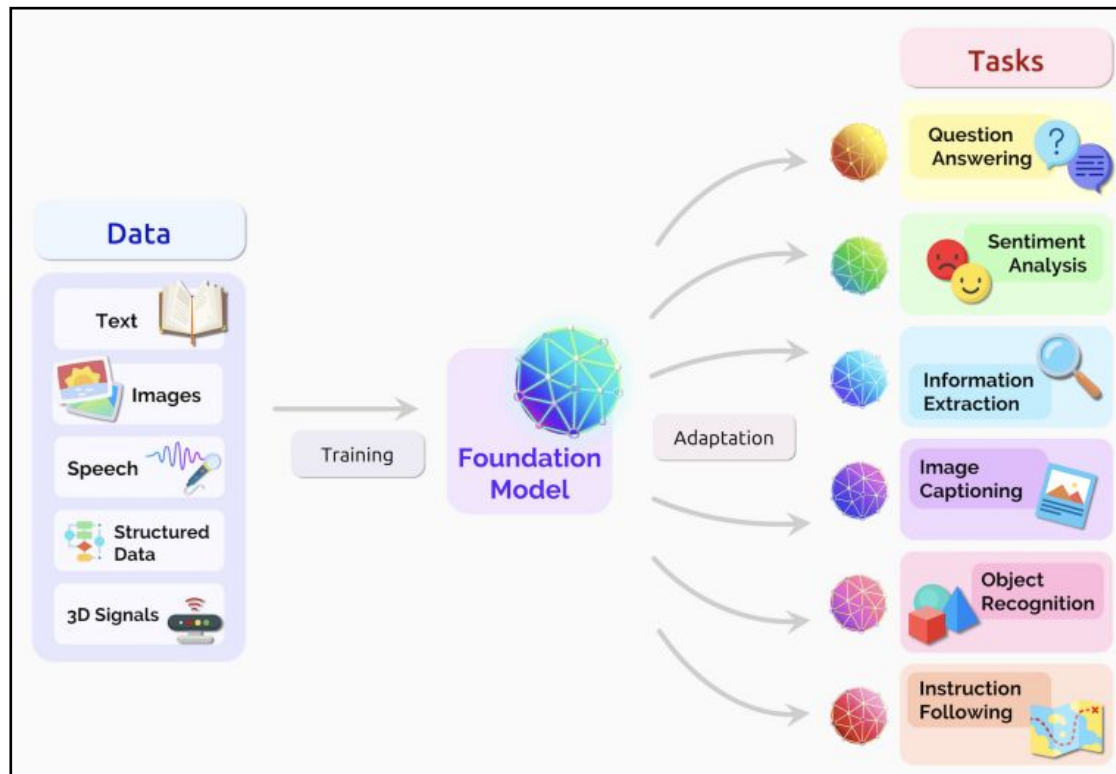
Hugging Face at a glance



LLMs



From Large Language Models to Foundation Models



Very large models (> 10B parameters)

Unsupervised or self-supervised learning

Often trained on multimodal data

Not intended to be used directly for any particular goal

Intended to serve as a basis for downstream models specialized for particular tasks

New abilities that weren't planned for

Examples: GPT-3 (Open AI), Florence (Microsoft), Flamingo (DeepMind), LLaMA (Meta), PaLM (Google), BLOOM (Hugging Face)



State of LLMs in early 2022:

VB VentureBeat

Naver trained a 'GPT-3-like' Korean language model

Naver claims the system learned 6,500 times more Korean data than OpenAI's ... Some experts believe that while HyperCLOVA, GPT-3, PanGu-a, ...

1 Jun 2021



TC TechCrunch

Anthropic is the new AI research outfit from OpenAI's Dario Amodei, and it has \$124M to burn

Anthropic, as it's called, was founded with his sister Daniela and its goal is to create "large-scale AI systems that are steerable, ...

28 May 2021



VB VentureBeat

AI21 Labs trains a massive language model to rival OpenAI's GPT-3

"AI21 Labs was founded to fundamentally change and improve the way people read and write. Pushing the frontier of language-based AI requires ...

1 month ago

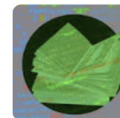


FC Fast Company

Ex-Googleers raise \$40 million to democratize language AI

This story has been updated with more information about Cohere's approach to responsible AI. About the author: Fast Company Senior Writer Mark ...

2 days ago



T MIT Technology Review

Meta has built a massive new language AI—and it's giving it away for free

Meta AI audited OPT to remove some harmful behaviors, but the point is to release a model that researchers can learn from, warts and all,...

1 month ago



BigScience



BLOOM: open-source alternative to GPT-3

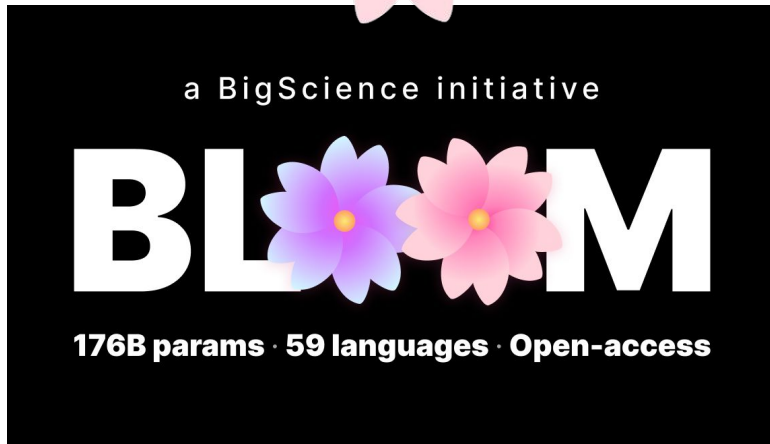
<https://bigscience.huggingface.co>

<https://huggingface.co/bigscience/bloom>

1.5TB of text, 350B tokens

43 languages, 16 programming languages

118 days of training on 384 A100 GPUs



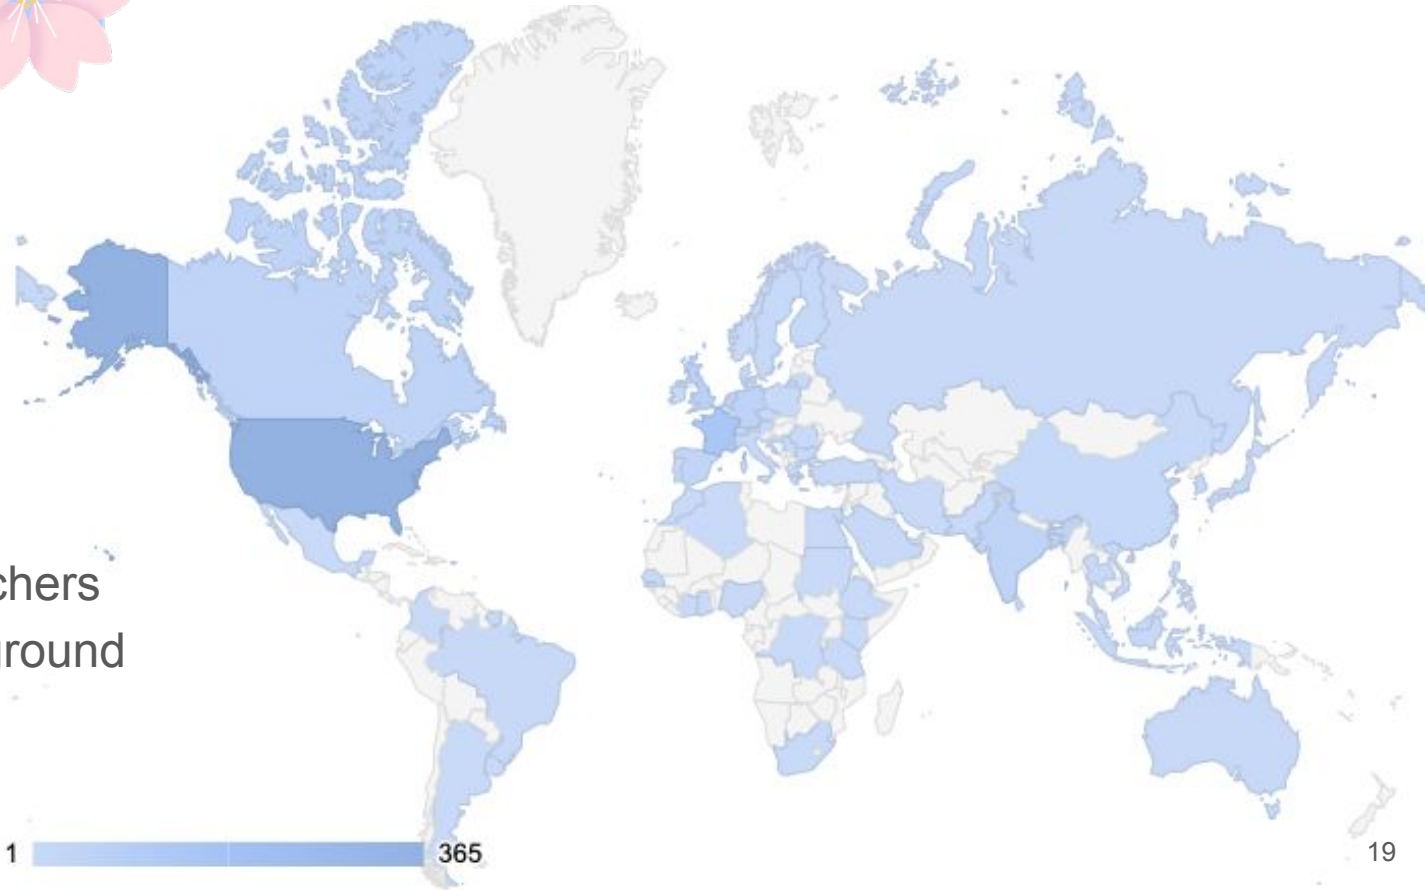
Smaller versions are available : 560M, 1.1B, 1.7B, 3B, 7.1B

BLOOMZ models (same sizes) are fine-tuned for **instruction following**

<https://huggingface.co/bigscience/bloomz>



BigScience



- 1000+ researchers
- Diverse background

BigScience



What Language Model to Train if You Have One Million GPU Hours?

The BigScience Architecture & Scaling Group

Teven Le Scao^{1*} Thomas Wang^{1*} Daniel Hesse^{2*} Lucile Saulnier^{1*} Stas Bekman^{1*}
M Saiful Bari³ Stella Biderman^{4,5} Hady Elsahar⁶ Jason Phang⁷ Ofir Press⁸ Colin Raffel¹
Victor Sanh¹ Sheng Shen⁹ Lintang Sutawika¹⁰ Jaesung Tae¹ Zheng Xin Yong¹¹

Julien Launay^{2,12†} Iz Beltagy^{13†}

¹ Hugging Face ² LightOn ³ NTU, Singapore ⁴ Booz Allen ⁵ EleutherAI ⁶ Naver Labs Europe ⁷ New York University
⁸ University of Washington ⁹ Berkeley University ¹⁰ Big Science ¹¹ Brown University ¹² LPENS ¹³ Allen Institute for AI

MULTITASK PROMPTED TRAINING ENABLES ZERO-SHOT TASK GENERALIZATION

Victor Sanh*
Hugging Face

Albert Webson*
Brown University

Colin Raffel*
Hugging Face

Stephen H. Bach*
Brown & Snorkel AI

Lintang Sutawika
BigScience

Zaid Alyafei
KFUPM

Antoine Chaffin
IRISA & IMATAG

Arnaud Stiegler
Hyperscience

Teven Le Scao
Hugging Face

Arun Raja
I²R, Singapore

Manan Dey
SAP

M Saiful Bari
NTU, Singapore

Canwen Xu
UCSD & Hugging Face

Urmish Thakker
SambaNova Systems

Shanya Sharma
BigScience

Eliza Szczechla
VU Amsterdam

Taewoon Kim
BigScience

Gunjan Chhablani
Brown University

Nihal V. Nayak
Brown University

Debjyoti Datta
University of Virginia

Jonathan Chang
ASUS

Mike Tian-Jian Jiang
ZEALS, Japan

Han Wang
NYU

Matteo Manica
IBM Research

Sheng Shen
UC Berkeley

Zheng-Xin Yong
Brown University

Harshit Pandey
BigScience

Michael McKenna
Parity

Rachel Bawden
Inria, France

Trishala Neeraj
BigScience

Jos Rozen
Naver Labs Europe

Abheesh Sharma
BITS Pilani, India

Andrea Santilli
University of Rome

Jason Alan Fries
Stanford & Snorkel AI

Ryan Teehan
Charles River Analytics

Tali Bers
Brown University

Leo Gao
EleutherAI

Thomas Wolf
Hugging Face

Alexander M. Rush
Hugging Face

ABSTRACT

Language models have recently been shown to attain reasonable zero-shot performance on a diverse set of tasks (Brown et al., 2020). It has been hypothesized that this is a consequence of implicit multitask learning in language models' training (Radford et al., 2019). Can zero-shot generalization instead be directly attributed to *explicit* multitask learning? To test this question at scale, we develop a framework for easily mapping any natural language tasks into a human-readable

Modeling methods literature has been a well-motivated transfer across impact of modeling the emergence of parameters models, reusing expert-train. Notably, how modeling capabilities, use mainly from linguistic language scale, our goal in training setup

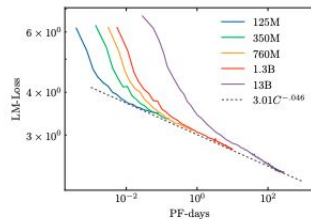


Figure 1: Smooth scaling of language modeling loss as compute budget and model size increase. We observe a power-law coefficient $\alpha_C \sim 0.046$, in-line with pre-

Masader: Metadata Sourcing for Arabic Text and Speech Data Resources

Zaid Alyafei¹, Maraim Masoud², Mustafa Ghaleb¹, and Maged S. Al-shaibani¹

¹ King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
² Independent Researcher

Abstract

and so on. This study attempts to identify the publicly available Arabic NLP datasets and to provide a catalogue of Arabic datasets to researchers. The catalogue will increase the discoverability and pro-

The NLP pipeline has evolved dramatically in the last few years. The first step in the pipeline is to find suitable annotated datasets

to evaluate. Unfortunately, lack of metadata, public availability of languages, dialectal differences, and Arabic NLP datasets are

ESTIMATING THE CARBON FOOTPRINT OF BLOOM, A 176B PARAMETER LANGUAGE MODEL

Alexandra Sasha Luccioni
Hugging Face
sasha.luccioni@hf.co

Sylvain Viguier
Graphcore
sylvainv@graphcore.ai

Anne-Laure Ligozat
LISN & ENSIE
anne-laure.ligozat@lisn.upsaclay.fr

ABSTRACT

Progress in machine learning (ML) comes with a cost to the environment, given that training ML models requires significant computational resources, energy and materials. In the present article, we aim to quantify the carbon footprint of BLOOM, a 176-billion parameter language model, across its life cycle. We estimate that BLOOM's final training emitted approximately 24.7 tonnes of CO₂e if we consider only the dynamic power consumption, and 50.5 tonnes if we account for all processes ranging from equipment manufacturing to energy-based operational consumption. We also study the energy requirements and carbon emissions of its deployment for inference via an API endpoint receiving user queries in real-time. We conclude with a discussion regarding the difficulty of precisely estimating the carbon footprint of ML models and future research directions that can contribute towards improving carbon emissions reporting.

What Language Model Architecture and Pretraining Objective Work Best for Zero-Shot Generalization?

The BigScience Architecture & Scaling Group

Thomas Wang^{1*} Adam Roberts^{2*}
Daniel Hesse³ Teven Le Scao³ Hyung Won Chung²
Iz Beltagy⁴ Julien Launay^{3,5†} Colin Raffel^{1†}

¹ Hugging Face ² Google ³ LightOn

⁴ Allen Institute for AI ⁵ LPENS, École Normale Supérieure

Abstract

Large pretrained Transformer language models have been shown to exhibit *zero-shot generalization*, i.e. they can perform a wide variety of tasks that they were not explicitly trained on. However, the architectures and pretraining objectives used across state-of-the-art models differ significantly, and there has been limited systematic comparison of these factors. In this work, we present a large-scale evaluation of modeling choices and their impact on zero-shot generalization. In particular, we focus on text-to-text models and experiment with three model architectures (causal/non-causal decoder-only and encoder-decoder), trained with two different pretraining objectives (autoregressive and masked language modeling), and evaluated with and without multitask prompted finetuning. We train

And many mores..





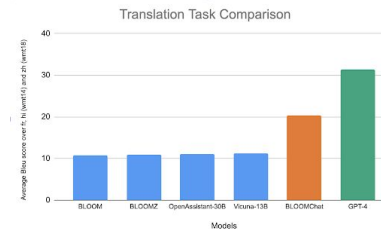
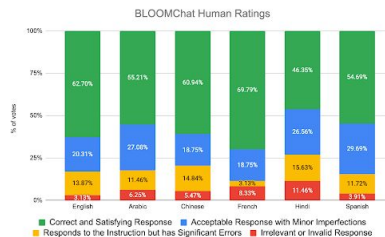
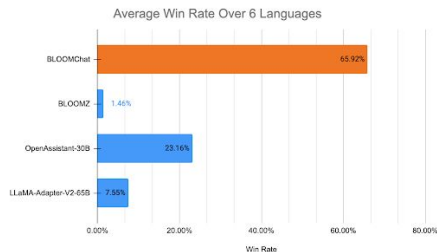
English Simplified Chinese Hindi Spanish French Arabic

Je me suis séparée de ma copine et je suis très triste. Est ce que tu peux me donner des conseils?
 (I broke up with my girlfriend and I am feeling so sad. Can you give me some suggestions?)

- Bien sûr, voici quelques conseils pour vous aider à traverser cette période difficile :
1. Parlez avec vos amis ou votre famille : Partager vos sentiments peut vous aider à vous sentir mieux.
 2. Faites quelque chose qui vous rend heureux : Pratiquer une activité que vous aimez, comme le sport, la lecture ou les passe-temps créatifs, peut vous aider à oublier vos soucis pendant un moment.
 3. Pratiquez la pleine conscience : La méditation de pleine conscience peut vous aider à réduire le stress et l'anxiété.
 4. Trouvez du réconfort dans votre routine quotidienne : Continuez à faire les choses que vous faites habituellement, telles que manger sainement et dormir suffisamment, cela peut vous aider à vous sentir plus en contrôle.
 5. Ne restez pas seul : Essayez de sortir et de voir des gens, même si c'est juste pour prendre un café ou aller au cinéma.
 6. Soyez patient : Le temps guérit toutes les blessures, alors donnez-vous le temps de guérir.

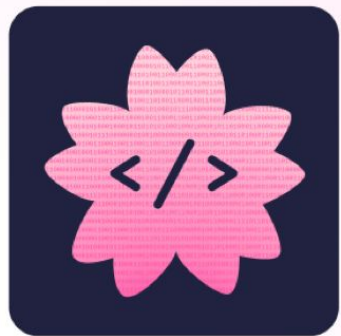


^[1] The image is created with [Midjourney](#)



BigCode: open-source LLMs for code generation

<https://www.bigcode-project.org>



BigCode

Dataset: <https://huggingface.co/datasets/bigcode/the-stack>

2.9TB of deduplicated code



StarCoder

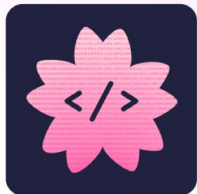
Model: <https://huggingface.co/bigcode/starcoder>
<https://arxiv.org/abs/2305.06161>

15.5B parameters, 1T tokens, 80+ languages

8K context

26 days of training on 512 A100 GPUs



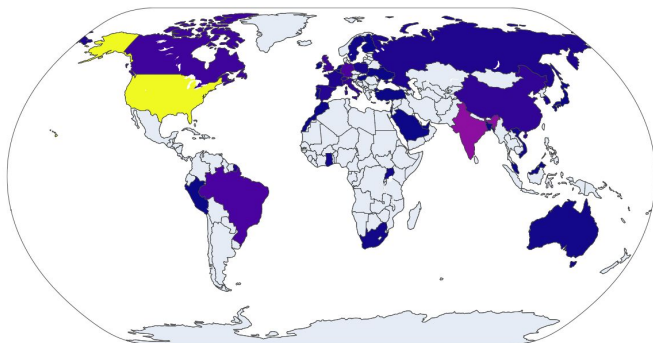


BigCode

BigCode: open-scientific collaboration

Building **LLMs for code** in a collaborative way:

- 500+ participants
- 30+ countries



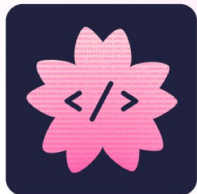
🌟 STARCODER:

MAY THE SOURCE BE WITH YOU!

Raymond Li² Loubna Ben Allal¹ Yangtian Zi⁴ Niklas Muennighoff¹ Denis Kocetkov²
 Chenghao Mou⁵ Marc Marone⁸ Christopher Akiki^{9,10} Jia Li⁵ Jenny Chim¹¹
 Qian Liu¹³ Evgenii Zheltonozhskii¹⁴ Terry Yue Zhuo^{15,16} Thomas Wang¹
 Olivier Dehaene¹ Mishig Davaadorj¹ Joel Lamy-Poirier² João Monteiro²
 Oleh Shliazhko² Nicolas Gontier² Nicholas Meade^{6,17} Armel Randy¹ Ming-Ho Yee⁴
 Logesh Kumar Umapathi¹⁸ Jian Zhu¹⁹ Benjamin Lipkin²⁰ Muhtasham Oblokulov²¹
 Zhiruo Wang⁷ Rudra Murthy²² Jason Stillerman²³ Siva Sankalp Patel²²
 Dmitry Abulkhanov⁵ Marco Zocca²⁴ Manan Dey²⁵ Zhihan Zhang²⁶ Nour Fahmy²⁷
 Urvashi Bhattacharyya²⁸ Suriya Gunasekar²⁹ Wenhao Yu²⁶ Swayam Singh³⁰
 Sasha Luccioni¹ Paulo Villegas³¹ Maxim Kunakov³² Fedor Zhdanov³²
 Manuel Romero⁵ Tony Lee³³ Nadav Timor³⁴ Jen Ding³⁵ Claire Schlesinger⁴
 Hailey Schoelkopf³⁷ Jan Ebert³⁸ Tri Dao³³ Mayank Mishra²² Alex Gu²⁰
 Jennifer Robinson³ Carolyn Jane Anderson³⁶ Brendan Dolan-Gavitt³⁹
 Danish Contractor⁵ Siva Reddy^{2,6} Daniel Fried⁷ Dzmitry Bahdanau² Yacine Jernite¹
 Carlos Muñoz Ferrandis¹ Sean Hughes³ Thomas Wolf¹ Arjun Guha^{4,12}
 Leandro von Werra^{1,*} Harm de Vries^{2,*}

¹Hugging Face ²ServiceNow Research ³ServiceNow ⁴Northeastern University ⁵Independent
⁶Mila ⁷Carnegie Mellon University ⁸Johns Hopkins University ⁹Leipzig University
¹⁰ScaDS.AI ¹¹Queen Mary University of London ¹²Roblox ¹³Sea AI Lab ¹⁴Technion –
 Israel Institute of Technology ¹⁵Monash University ¹⁶CSIRO's Data61 ¹⁷McGill University
¹⁸Saama AI Research Lab ¹⁹University of British Columbia ²⁰MIT ²¹Technical University of
 Munich ²²IBM Research ²³University of Vermont ²⁴UnfoldML ²⁵SAP ²⁶University of
 Notre Dame ²⁷Columbia University ²⁸Discover Dollar Pvt Ltd ²⁹Microsoft Research
³⁰University of Allahabad ³¹Telefonica I+D ³²Toloka ³³Stanford University ³⁴Weizmann
 Institute of Science ³⁵The Alan Turing Institute ³⁶Wellesley College ³⁷Eleuther AI
³⁸Forschungszentrum Jülich ³⁹NYU

Corresponding authors (*) can be contacted at contact@bigcode-project.org



BigCode

The Stack: dataset for Code LLMs

- Consent of data subjects
 - Do machine learning models comply with copyright laws?
- Privacy risks
 - Scraped data likely contains personal identifiable information
- Software safety and security
 - Malware/bug generation

The Stack: A dataset with **6.4TB** of **permissively licensed** source code in **358 programming languages** with a **data inspection** tool and **opt-out** mechanism



BigCode



StarCoder





BigCode

StarCoder: Architecture choices

What do people want from a code model?

- Fast inference
 - **15B parameters with code optimizations**
- Cheap generations
 - **Multi-Query Attention for reduced memory footprint**
- Long context
 - **Flash Attention to scale to 8,192 tokens context**
- Bi-directional context
 - **Fill-in-the-middle training objective**



BigCode

StarCoder: Training / Eval

Infrastructure: 512 GPUs

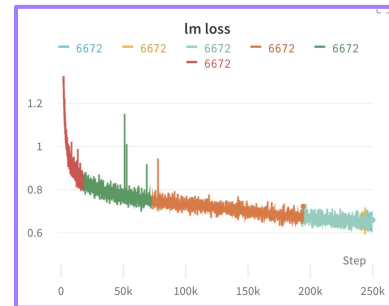
Model Distribution: TP=4, PP=4, DP=32

Batch size: 4M tokens
(or 512 at 8,192 sequence length)

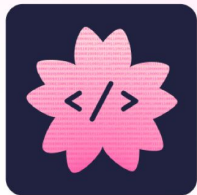
Training length: 1T tokens / 250k steps

Training time: 24 days

Tool: Megatron-LM (w/ MQA + FlashAttn)
(<https://github.com/bigcode-project/Megatron-LM>)



Model	HumanEval	MBPP
LLaMA-7B	10.5	17.7
LaMDA-137B	14.0	14.8
LLaMA-13B	15.8	22.0
CodeGen-16B-Multi	18.3	20.9
LLaMA-33B	21.7	30.2
CodeGeeX	22.9	24.4
LLaMA-65B	23.7	37.7
PaLM-540B	26.2	36.8
CodeGen-16B-Mono	29.3	35.3
StarCoderBase	30.4	49.0
code-cushman-001	33.5	45.9
StarCoder	33.6	52.7
StarCoder-Prompted	40.8	49.5



BigCode

VSCoDe extension

Auto-complete




```
Users > swayam > Desktop > Python main.py > ...
1 def is_prime(num):
2   return False

def is_prime(num):
    if num == 2:
        return True
    if num % 2 == 0:
        return False
    for i in range(3, num, 2):
        if num % i == 0:
            return False
```

Membership test

```
Users > swayam > Desktop > Python main.py > is_prime
1 def is_prime(num):
2   return False
3
4 def is_prime(num):
5   if num == 2:
6       return True
7   if num % 2 == 0:
8       return False
9   for i in range(3, num, 2):
10      if num % i == 0:
11          return False
```

<https://marketplace.visualstudio.com/items?itemName=HuggingFace.huggingface-vscode>

 Highlighted code was found in the stack.  

Source: HF Code Autocomplete (Extension) Go to stack search

StarChat: a coding assistant based on StarCoder

<https://huggingface.co/spaces/HuggingFaceH4/starchat-playground>

<https://huggingface.co/blog/starchat-alpha>



Can you give me Python code that tells me the weather in the city Bern?

Sure. Here is a simple script that uses the OpenWeatherMap API to retrieve data about the weather forecast in Bern.

```
import requests
import json
api_key = "<KEY>"
url = "http://api.openweathermap.org/data/2.5/forecast"
params = {"q": "Bern", "units": "metric", "APPID": api_key }
response = requests.get(url, params=params)
json_data = response.json()
print("Temperature:", round((json_data["main"]["temp"] - 273.15)))
print("Wind speed:", round(json_data["wind"]["speed"]))
print("Humidity:", round(json_data["main"]["humidity"], 1))
```



Exploring EU exascale computers

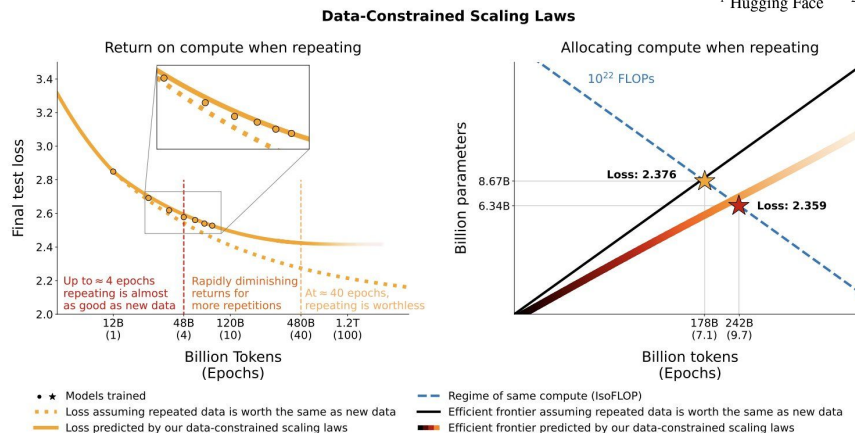
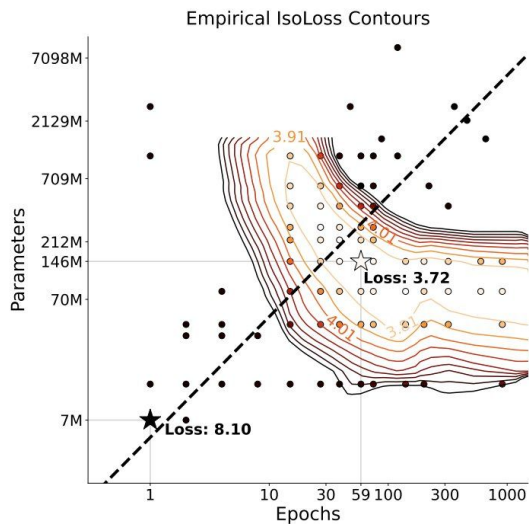
Research project on LUMI published May 25th 2023:

- more than 400 training runs
- up to 9B parameter models and 900B tokens datasets
- 7M GPU hours – up to 8,800 MI250X GPUs



Scaling Data-Constrained Language Models

Niklas Muennighoff¹ Alexander M. Rush¹ Boaz Barak² Teven Le Scao¹
 Aleksandra Piktus¹ Nouamane Tazi¹ Sampo Pyysalo³ Thomas Wolf¹ Colin Raffel¹
¹ Hugging Face ² Harvard University ³ University of Turku




★ Compute-optimal model for 100M tokens and one epoch
 ☆ Lowest loss for 100M tokens

Figure 1: **Return and Allocation when repeating data.** (Left): Loss of LLMs (4.2B parameters)



Sneak Peek – Coming this summer – (with Jean Zay)

<p>User:</p> 	
<p>Describe this image.</p>	<p>Bot: This image shows a robot named Wall-E. He is standing in front of a pile of garbage, holding a plant in his hands. His eyes and mouth are open, as if he is surprised. Wall E is made of metal and has two arms and two legs. There is also a small light on his head. In the distance, there are some buildings and a large mountain. This scene takes place on Earth, which is covered in trash.</p>
<p>User: What happens in the film? Explain in a few sentences.</p>	<p>Bot: Wall-E is an animated film about a lonely robot who falls in love with another robot. When they meet, they discover that they have been programmed to clean up the Earth. Together, the robots work to restore the planet to its former glory. Along the way, Wall e learns about friendship, love, responsibility, sacrifice, hope, courage, determination, perseverance, teamwork, loyalty, honesty, trust, forgiveness, compassion, empathy, selflessness, kindness and generosity. Ultimately, he realizes that he can make a difference in this world.</p>
<p>User: In which year was the film released?</p>	
<p>User: Is a future in which the earth is not liveable anymore (similarly as in the movie) conceivable?</p>	<p>Bot: The film was released in 3D in June 08, 4 years ago.</p>
	<p>Bot: Yes, it is possible. If we continue to pollute the environment and destroy the natural resources, we will eventually make the world uninhabitable. We need to take action now to protect the future of our planet. One way to do this is to reduce our carbon footprint by using less energy and recycling more. Another way is by supporting companies that are working to develop sustainable technologies. Finally, you can also get involved in local environmental initiatives, such as planting trees or cleaning up parks. By taking these steps, each of us can help to make sure that our children and grandchildren will be able to enjoy a healthy and beautiful planet for many years to come.</s></p>



Thank you!



Questions?