

LIACC



U.PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Navigating the Artificial Intelligence Revolution

Luís Paulo Reis

lpreis@fe.up.pt

Associate Professor at FEUP - Faculty of Engineering of the University of Porto
Director of LIACC – Artificial Intelligence and Computer Science Laboratory, Univ. Porto
Member of Coordinating Commission of LASI – Intelligent Systems Associate Laboratory
President of the GA of APPIA – Portuguese Association for Artificial Intelligence







Agenda

- **Navigating the New World of AI**
- **The Evolution of Artificial Intelligence**
- **Artificial Intelligence (AI) and Machine Learning (ML)**
- **Generative AI and Large Language Models (LLM)**
- **AI Applications and Challenges**
- **Agents, Agentic AI and Multi-Agent Systems**
- **Robotics, Deep RL, LBM's and the Future**
- **Conclusions**

There are More Tides than Sailors



There are More Tides than Sailors



Ferdinand Magellan




- Ferdinand was **born in Portugal** in 1480
- In September 1519, Ferdinand Magellan along with **5 ships** left from Spain
- He convinced the **king of Spain to pay** for his expedition
- Find a way to **reach Asia** by **sailing west** around the Americas




Ferdinand Magellan

- First sailed to South America
- Tried to **find a river** that would take him **across South America** to the **Pacific Ocean**
- Finally in 1520 he found a kind of river: **Magellan Strait**
- For **3 months**, ships sailed across the **Pacific Ocean**
- Began running **out of food and water**
- **Many men died** from hunger and diseases



Ferdinand Magellan

- March 1521, ships reached **Philippines**
- On April 27, **Magellan was killed** in battle against warriors of Lapulapu
- In 1522, only **one of the five ships returned to Spain**
- Of the 270 men, **only 18/19 survivors** returned
- First to sail completely around the world



Ferdinand Magellan





“The greatest sea voyage in the Age of Discovery!”

“The most important maritime voyage ever undertaken!”

Vision, Courage, Leadership, Skill and Perseverance!

King Manuel I of Portugal did not invest in it!




AI and the Discovery of the New World

**AI is no longer a Tool!
It is a Partner!**

Image: Shutterstock / CNN - <https://edition.cnn.com/2019/02/17/investing/artificial-intelligence-investors-machine-learning/>

Artificial Intelligence (AI)

- **Intelligence**
 - “Capacity to **solve new problems** through the use of knowledge”
- **Artificial Intelligence**
 - “Science concerned with building **intelligent machines**, that is, machines that perform tasks that when performed by humans require intelligence”



Weak and Strong AI

- **Weak Artificial Intelligence**

weak AI, also known as narrow AI is artificial intelligence that is focused on one **single narrow task**




Image: <https://livingsmartheart.com/what-is-strong-ai/>


- **Strong Artificial Intelligence**

Strong AI or Artificial General Intelligence (AGI) is the intelligence of a machine that could successfully perform **any intellectual task** that a human being can! **Science fiction? Not any more...**




The World Is Changing Faster Than We Can Understand

- AI won't replace humans — but humans using AI will replace humans not using AI
- AI is moving from *tools* that assist humans to *agents* that act autonomously
- Autonomous, Physical, Goal-Driven, Learning Systems




Evolution of AI Paradigms



The Birth of AI

- *From Logic to Intelligence...*
- *The Dream Begins (1950–1980)*
- “AI began as a dream of formalising reasoning!




1956 Dartmouth Conference: The Founding Fathers of AI




John McCarthy




Marvin Minsky



Claude Shannon



Ray Solomonoff




Alan Newell



Herbert Simon



Arthur Samuel



Oliver Selfridge



Nathaniel Rochester



Trenchard More

<https://jlanyon.substack.com/p/a-brief-history-of-ai>

AI Timeline





Image: [Paul Marsden, 2017] - <https://digitalwellbeing.org/artificial-intelligence-timeline-infographic-from-eliza-to-tay-and-beyond/>

Symbolic AI and Expert Systems

- ***Knowledge is Power — Encoding the Mind into Rules!***
- Symbolic AI represents knowledge through symbols and rules
- Examples:
 - *MYCIN* (medical diagnosis)
 - *DENDRAL* (chemical analysis)
- Architecture: inference engine + knowledge base
- Early AI was ***explainable and logical!***
- But could not learn or adapt...



From Rules to Patterns

AI needed to evolve from being programmed to being trained

- From Rules to Patterns - ***AI Started to Learn from Data***
- Instead of telling the machine *what to do*, we let it *learn how to do it*
- Machine Learning (ML) = algorithms that learn patterns from data



- Examples: email spam filters, handwriting recognition, early speech recognition
- ML marked the shift from ***explicit logic*** to ***statistical inference***

Machine Learning

Field of artificial intelligence that gives **computer systems** the **ability to "learn"** (e.g., progressively **improve performance** on a specific task) from **data/results of their actions**, without being explicitly programmed

Supervised Learning




Image: <https://towardsdatascience.com/what-are-the-types-of-machine-learning-e2b9e5d1756f>

Reinforcement Learning










Image: <https://becominghuman.ai/an-introduction-to-machine-learning-33a1b5d3a560>

From Statistics to Intelligence


- Algorithms and Breakthroughs
- Between the 1980s and early 2000s, we saw breakthroughs:
- Decision Trees, Random Forests, SVMs, Neural Networks and the rediscovery of Backpropagation
- Paving the road to Deep Learning




Artificial Neural Networks



Artificial Neural Network




Deep Neural Network





The Data Revolution

- *Data Is the New Oil — and the Fuel of AI*
- With the internet, smartphones, and sensors, data grew exponentially
- Algorithms improved — but what truly powered modern AI was scale!
- *Algorithms got better — but data made them powerful!*




The Deep Learning Revolution

- When Neural Networks Got Deep — and Everything Changed!
- AlexNet on ImageNet triggered the deep-learning boom
- Multi-layer neural networks, learn features automatically


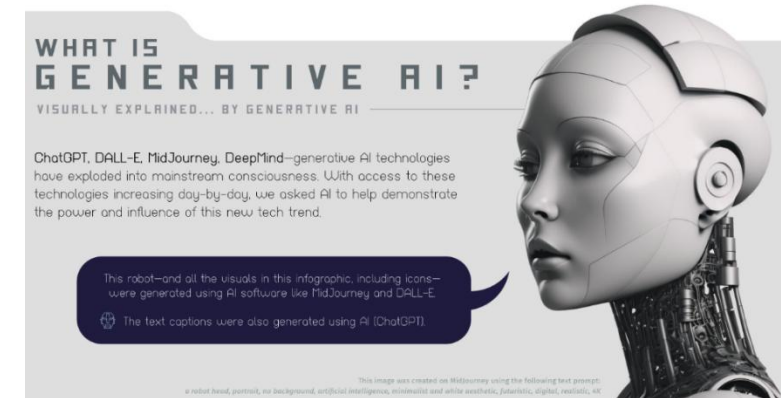


Deep Learning Neural Network



Generative AI - Machines That Create

- **Generative AI (GenAI)** can create a wide variety of data, such as images, videos, audio, text, and 3D models
- Learning patterns from existing data, and then using this knowledge to generate new and unique outputs in response to prompts
- GenAI produce **highly realistic and complex content** that mimics human creativity
- **Large language models (LLM)** - Very large neural networks with billions of weights, trained on very large quantities of text




	DALL-E 2	Stable Diffusion	craiyon	Lexica	MidJourney		
Text-to-image (T2I)	Imagen	WOMBO	NightCafe	GauGAN2	DeepAI	Jasper	artbreeder
	Wonder	pixray-text2image	neural.love	Omneky	alpaca		
	image.space	KREA	Nyx.gallery	ROSEBUD.AI	PhotoRoom		
Text-to-Video (T2V)	runway	Fliki	synthesia	Meta AI	Google AI	Phenaki	CONTENDA
Text-to-Audio (T2A)	Play.ht	MURF.ai	RESEMBLE.AI	WELLSAID	descript	Mforithmic	
	Simplified	Jasper	frase	EleutherAI	Requestary	letterdrop	
	grammarly	copy.ai	MarketMuse	AI21 Labs	HubSpot	NovelAI	
Text-to-Text (T2T)	InferKit	GooseAI	Research	Writesonic	cohere	CHIBI	
	Ideas AI Powered by OpenAI	copysmith	Flowrite	NICHES	sudowrite	Rytr	
	ideasbyai ^{beta}	text.cortex	OpenAI GPT-3	Blog Idea Generator			
	HyperWrite	Subtxt	WRITER	wordtune	LAIKA	COMPOSE AI	
	Moonbeam	Bertha.ai	anyword	Hypotenuse AI	Peppertype.ai		
Text-to-Motion (T2M)	treeInd	MDM: Human Motion Diffusion Model					
Text-to-Code (T2C)	replit	Ghostwriter	GitHub Copilot	MUTABLE AI	tabnine		
	Amazon CodeWhisperer						
Text-to-NFT (T2N)	LensAI						
Text-to-3D (T2D)	DreamFusion	CLIP-Mesh	GET3D				
Audio-to-Text (A2T)	descript	AssemblyAI	Whisper				
Audio-to-Audio (A2A)	AudioLM	VOICEMOD					
Brain-to-Text (B2T)	speech from brain	non-invasive brain recordings					
Image-to-Text (A2T)	neural.love	GPT-3 x Image Captions					

AI Timeline

AI TIMELINE: 2021

MORE THAN 24 LARGE MODELS
IN LESS THAN 12 MONTHS



Selected highlights only. Alan D. Thompson, November 2021. <https://life architect.ai/>

Language model

Dialogue model (chatbot)

Multi-modal model (images)



LifeArchitect.ai/timeline

Generative AI Competition




LLM Evolution

2022/3->2024/5

2022-23 ¹		Jan 2025 ²	
Anthropic	Claude • Not multimodal (text only) • Limited contextual understanding (difficulty with complex conversations) • No tool usage	→	Claude 3.5 • Multimodal (text, audio, and images) • Enhanced contextual understanding and coherence during long interactions • Experimental computer usage capability for some users
Google	Google Bard • Fair reasoning ability (eg, scored high on SAT, but bottom 10% on bar examination) • Limited contextual understanding (difficulty with coherence in complex conversations) • Standard API access (for text generation)	→	Gemini 2.0 Flash • Advanced reasoning (eg, top 10% on bar examination) • Enhanced contextual understanding (maintains coherence in long dialogues) • Advanced API access (supports multimodal inputs)
OpenAI	GPT-3.5 • Not multimodal (text only) • Fair reasoning ability (eg, scored high on SAT, but bottom 10% on bar examination) • Limited contextual understanding (difficulty with coherence in complex conversations) • Standard API access (for text generation)	→	OpenAI o1 • Multimodal (text and images) • Advanced reasoning (eg, top 10% on bar examination) • Enhanced contextual understanding (maintains coherence in long dialogues) • Advanced API access (supports multimodal inputs)
Google Gemini	Google Bard • Not multimodal (text only) • Fair reasoning • Limited contextual understanding (difficulty with complex conversations) • Limited real-time data integration • Low personalization (limited adaptability)	→	Gemini 2.0 Flash • Multimodal (text, audio, and images) • Advanced reasoning (capable of multistep problem-solving and nuanced analysis) • Enhanced contextual understanding (maintains coherence in long dialogues) • Real-time data integration (from Google Search) • Advanced personalization (user context)

Programming vs Machine Learning



The End of Programming!

Nvidia CEO predicts the death of coding — Jensen Huang says AI will do the work, so kids don't need to learn

News

By Benedict Collins published February 26, 2024

Jensen Huang believes coding languages are a thing of the past



(Image credit: Nvidia)

Nvidia CEO Jensen Huang has once again announced the death of coding, but this time in front of a potentially far more influential audience.

StarCoder 2 is a code-generating AI that runs on most GPUs

Kyle Wiggers @kyle_i_wiggers / 2:00 PM UTC • February 28, 2024

Comment



Image Credits: Tippiatt / Getty Images

Developers are adopting AI-powered code generators — services like [GitHub Copilot](#) and [Amazon CodeWhisperer](#), along with open access models such as Meta's [Code Llama](#) — at an [astonishing](#) rate. But the tools are far from ideal. Many aren't free. Others are, but only under licenses that preclude them from being used in common commercial contexts.

Perceiving the demand for alternatives, AI startup Hugging Face several years ago teamed up with ServiceNow, the workflow automation platform, to create [StarCoder](#), an open source code generator with a less restrictive license than some of the others out there. The original came online early last year, and work has been underway on a follow-up, StarCoder 2, ever since.


StarCoder 2 isn't a single code-generating model, but rather a family. Released today, it comes in three variants, the first two of which can run on most modern consumer GPUs:

- A 3-billion-parameter (3B) model trained by ServiceNow
- A 7-billion-parameter (7B) model trained by Hugging Face
- A 15-billion-parameter (15B) model trained by Nvidia, the newest supporter of the StarCoder project

Technology: Adapt or Resist?!

So! You are the guy...

- **COVID-19 - Use of technology**
- **Educate New Generations**
- **Ethical Use**
- **Technical Use**
- **AI Policy**
- **Prompt Engineering**



That took all of our Jobs!

Concerns of AI and ChatGPT in Education

- More **in person exams** and oral assignments
- **Continuous evaluation** of projects
- Help students **develop** their **critical thinking skills** by comparing AI-produced content with reliable, valid sources of information
- Be careful with **over-reliance on AI**
- Shift the object of evaluation **from the “product” to the “process” of student learning**
- Emphasize **ethical use of AI** and the importance of proper research and citation practices and avoiding plagiarism

AI and ChatGPT in Education

“AI and Can Complement Learning, Not Replace It!”
From Knowledge Transmitter → Learning Architect



Concerns of AI and ChatGPT in Research

- Risk of **fabricated** or inaccurate **references and data** (“hallucinations”)
- Possible **erosion of methodological rigor** due to **over-reliance** on automated reasoning
- Challenges in ensuring **reproducibility** and **transparency** when **AI is involved**
- Need for explicit **disclosure of AI use** in research design, writing, and analysis
- **Potential bias amplification** from training data and opaque model behavior
- **Ethical responsibility** to validate, interpret, and contextualize **AI outputs** using **human expertise**
- Importance of **maintaining authorship integrity** and **avoiding ghost-writing by AI**

Research with AI



ELSEVIER

Scopus AI

Publish with us



[Home](#) > [Products](#) > [Scopus](#) > [Scopus AI](#)

Scopus AI: Trusted content. Powered by responsible AI.

Empower your research journey with Scopus AI – your dynamic GenAI-powered research companion.

Navigate through the vast expanse of human knowledge faster with a trusted guide designed to enhance your understanding, enrich your insights, and transform your overall research experience.



Consensus

vs



ChatGPT

Your AI Research Assistant

scite_

Your AI Research Assistant

Ask questions and receive evidence-based answers from real research

Assistant

Search

Tables

Ask a question... (type '/' for menu)

Settings

Sources



Elsevier's Five Responsible AI Principles

- We consider the real-world impact of our solutions on people
- We take action to prevent the creation or reinforcement of unfair bias
- We can explain how our solutions work
- We create accountability through human oversight
- We respect privacy and champion robust data governance

Ask a research question to generate a concise summary report



The screenshot shows the Elicit interface with a search bar at the top. Below the search bar, there are filters for 'New Collection', 'New Category', and 'Connect to Zotero'. A list of results is displayed, including a paper titled 'Learning hybrid locomotion skills—Learn to exploit residual actions and modulate model-based gait control' by Mohammadmehdi Kasefi and Zhibin Li, published in 2023. The paper is highlighted with a blue box. To the right of the paper, there is a 'PDF' icon and a 'Download' button. Below the paper, there is a '1 selected paper' section with a brief description of the paper's content. The interface also includes a sidebar with 'Uncategorized' and 'Shared with Me' sections, and a bottom section with 'Explore Papers', 'Explore People', and 'Export Papers' options.


Google Search Engine






- **Hummingbird Algorithm**
 - **PageRank algorithm** that covers a specific way of **giving pages credit based on the links from other pages** pointing at them
 - Panda, Penguin and Payday to fight **spam**
 - Pigeon designed to improve **local results**
 - Top Heavy designed to demote **ad-heavy** pages,
 - Mobile Friendly designed to reward **mobile-friendly** pages
 - Pirate designed to fight **copyright infringement**
- **200 major ranking signals (up to 10,000 variations or sub-signals)**
- **Google RankBrain Algorithm**
 - **Machine-Learning System**
 - **3rd most important signal**
 - Interpret searches people submit to **find pages that might not have the exact words** that were searched for



Gemini – Google Deep Mind



Gemini 2.5





Gemini Advanced
Upgrade your Google One plan to get Gemini Advanced
€21.99 €0 for 1 month,
€21.99/month after

Microsoft Copilot



Claude 4.5 - Anthropic Claude



 Good morning, Luis Paulo

How can Claude help you today?

Claude 3.5 Sonnet

Get started with an example below Add content

Summarize meeting notes Generate interview questions Extract insights from report

NEW **Artifacts**
Create and iterate on documents, code, and more within Claude. [Try it out](#)

 Your recent chats  View all →

Claude for Desktop Beta

Your AI partner on desktop. Fast, focused, and designed for deep work.

 Windows

 Windows (arm64)

 macOS

Take Claude with you

Talk to Claude, anywhere you go. Brainstorm ideas, get answers, and analyze images on the go.

 Apple App Store

 Google Play



Pro

€18 + VAT / month

- 5x more usage versus Free plan
- Access to Projects to organize documents and chats
- Ability to use more models, like Claude 3 Opus
- Early access to new features

[Subscribe to Pro](#)

Perplexity AI



- **Perplexity AI** is a conversational search engine that uses large language models (LLMs) to answer queries using **sources from the web** and **cites links** within the text response.

The screenshot shows the Perplexity AI web interface. On the left is a sidebar with navigation links: Home, Discover, Spaces, and Library. The main content area displays a search query: "who are the top researchers in AI in portugal today". Below the query, it states: "Based on the most recent information available, several top AI researchers in Portugal stand out:". A list of three researchers is provided, each with a citation number (1, 2, 3). To the left of the list, under the heading "Sources", there is a snippet from "INESC-ID AI researchers ranked some of the best in..." with a link to "inesc-id".

perplexity |← lpreis1970 ⌚ 4m + Space / who are the top researchers in AI in portugal today

New Thread Ctrl I

Home Discover Spaces Library

who are the top researche

who are the top researchers in AI in portugal today

Based on the most recent information available, several top AI researchers in Portugal stand out:

Sources

INESC-ID AI researchers ranked some of the best in... inesc-id

1. Paulo Novais from the University of Minho, who leads the ISLab laboratory at the ALGORITMI centre and coordinates LASI, the largest associated laboratory in Portugal ³.
2. Luís Paulo Reis, an Associate Professor with Aggregation at FEUP and Director of LIACC – Artificial Intelligence and Computer Science Laboratory ² ³.
3. Manuela Veloso, currently the Director of J.P. Morgan AI Research and formerly a member of Carnegie Mellon University. She is internationally renowned for her work in artificial intelligence and robotics ³.

Grammarly

Grammarly

My Grammarly

Trash

Account

Apps

Support

Sign Out
lpreis@fe.up.pt

Grammarly Apps

Grammarly for Windows

Grammarly for your desktop is your one-stop solution to get writing suggestions across a wide array of desktop apps and websites.

By downloading Grammarly for Windows, you agree to the [Terms and Conditions](#) and [Privacy Policy](#). California residents, see our [CA Privacy Notice](#).

Install

Grammarly for Chrome

Grammarly's browser extensions offer writing suggestions in Google Docs and other sites across the web.

Active

Re-install

Grammarly for iPhone and iPad

Write confidently and professionally across all your apps and websites with Grammarly's writing suggestions, anywhere you go.

Install

DeepL



DeepL Translator ▾

DeepL Pro

Why DeepL? ▾

API

Plans and pricing

Apps

FREE

Start free trial

Login



Translate text
31 languages



Translate files
.pdf, .docx, .pptx

Pick up where you left off with your **30-day free trial** of DeepL Pro. ✕

Detect language ▾




English (UK) ▾

Glossary

Type to translate.

Drag and drop to translate PDF, Word (.docx), and PowerPoint (.pptx) files with our document translator.



DeepL Translator

”

Tech giants Google, Microsoft and Facebook are all applying the lessons of machine learning to translation, but a small company called DeepL has outdone them all and raised the bar for the field. Its translation tool is just as quick as the outsized competition, but more accurate and nuanced than any we've tried.

TechCrunch, USA

TC TechCrunch

QuillBot

The screenshot displays the QuillBot web application interface. At the top, there is a navigation bar with the QuillBot logo, a hamburger menu, an 'Upgrade to Premium' button, and flags for language and theme. A left sidebar contains icons for document management, grammar check, and integrations. The main content area is divided into two panels. The left panel, titled 'Modes', includes buttons for 'Standard' (selected), 'Fluency', 'Formal', 'Simple', 'Creative', 'Expand', and 'Shorten'. It also features a 'Synonyms' slider. Below the modes are instructions and two buttons: 'Try Sample Text' and 'Paste Text'. A large green 'Paraphrase' button is at the bottom of this panel. The right panel contains three horizontal text input fields. Overlaid on the right side of the interface are three callout boxes: 'Paraphraser' (blue), 'Summarizer' (green), and 'Grammar Checker' (purple). At the bottom left, the QuillBot logo is repeated. To the right of the logo, a text box describes the 'Vocabulary enhancement feature, Thesaurus powered by AI, Integrations (Word, GoogleDocs), Grammar check, Summarizing tool'.

QuillBot

Upgrade to Premium

Modes: Standard Fluency Formal Simple Creative Expand Shorten

Synonyms:

QuillBot will rewrite your text. Start by writing or pasting something here and then press the Paraphrase button.

Try Sample Text

Paste Text

Paraphrase

Paraphraser

Summarizer

Grammar Checker

Vocabulary enhancement feature, Thesaurus powered by AI, Integrations (Word, GoogleDocs), Grammar check, Summarizing tool

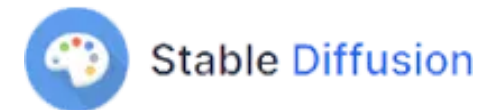
Image Generation - Dall-E 3

DALL·E 3 is an AI system that can create realistic images and art from a description in natural language

An astronaut riding a horse in photorealistic style



A beach similar to "praia da baia" in Espinho with surfers, surfing in the sea and big waves



Explore GPTs

GPTs

Discover and create custom versions of ChatGPT that combine instructions, extra knowledge, and any combination of skills.

Q Search public GPTs

1  **Canva**
Effortlessly design anything: presentations, logos, social media posts and more.
By canva.com

3  **Consensus**
Your AI Research Assistant. Search 200M academic papers from Consensus, get science-based answers, and draft content...
By Consensus

5  **Video GPT by VEED**
AI Video Maker. Generate videos for social media - YouTube, Instagram, TikTok and more! Free text to video & speech tool wit...
By veed.io

2  **Logo Creator**
Use me to generate professional logo designs and app icons!
By Chase Lean

4  **image generator**
A GPT specialized in generating and refining images with a mix of professional and friendly tone.image generator
By NAIF J ALOTAIBI

6  **Grimoire**
Code Wizard 🧙 Learn to Prompt-grail
Create a website (or anything) with a sentence. 20+ Hotkeys for coding flows....
By gptavern.mindgoblinstudios.com

Custom GPTs — special-purpose versions of the model created by others

uma receita de bacalhau com espinafres e batatascom um empratamento bonito e foto

Talked to api.adzedek.com ▾

Receita de Bacalhau com Espinafres e Batatas

Esta é uma receita clássica e nutritiva que combina o sabor rico do bacalhau com a suavidade das batatas e a frescura dos espinafres. O prato é montado de forma elegante, com camadas visuais para um empratamento atraente.

Ingr^{es}:

- 400g de bacalhau dessalgado
- 300g de batatas
- 200g de espinafres frescos
- 2 dentes de alho picados
- 1 cebola em rodela finas
- 100ml de azeite de oliva
- Sal e pimenta a gosto



Llama 4 – Meta Open Source Models

 Meta

Technology ▾

Get started ▾

Trust & Safety ▾

Community ▾

Try Meta AI ↗

Download models




Build the future of AI with Meta LLaMA 4


Now available with both 8B and 70B pretrained and instruction-tuned versions to support a wide range of applications

Get Started

Experience Llama 3 on Meta AI



Mistral AI – Open Source Models



Openness

We lead the market of open source generative technologies to bring trust and transparency in the field and foster decentralised technology development



Portability

Our technology is available through serverless APIs, public cloud services (on Azure AI, on Amazon Bedrock), and for VPC/on-premise deployment. Our independence from cloud providers is a guarantee of our customers' independence.



Value and speed

Our flagship model, Mistral Large, has independently validated top-tier reasoning in multiple languages. All our models bring unmatched value and latency at their price points.



Customisation

Our models can be fine-tuned and modified at will for your business to create differentiated AI applications.



Build on Open Source

Under the Apache 2.0 license, our 3 open source models Mistral 7B, Mixtral 8x7B, Mixtral 8x22B are usable and customisable for a variety of use cases. They can be downloaded or used on demand via our platform.

- [Download them](#) for deployment in your own environment
- Use them on [La Plateforme](#) at market-leading availability, speed, and quality control



Gemma 3n – Gemini Google Open Source

Gemma


google/gemma

Gemma 2 is now available to researchers and developers

Gemma 2 offers best-in-class performance, runs at incredible speed across different hardware and easily integrates with other AI tools.

Jun 27, 2024 · 4 min read

Gemma is a family of lightweight, open models built from the research and technology that Google used to create the Gemini models.



Responsible Generative AI Toolkit

[Send feedback](#)

This toolkit provides resources to apply best practices for responsible use of open models such as the Gemma models, including:


- Guidance on setting safety policies, safety tuning, safety classifiers and model evaluation.
- The [Learning Interpretability Tool \(LIT\)](#) for investigating Gemma's behavior and addressing potential issues.
- A methodology for building robust safety classifiers with minimal examples.

This version of the toolkit focuses on English text-to-text models only. You can provide feedback to make this toolkit more helpful through the feedback mechanism link at the bottom of the page.


When building with Gemma, you should take a holistic approach to responsibility and consider all the possible challenges at the application and model levels. This toolkit covers risk and mitigation

Gemma 3n



Our powerful and efficient open model designed to run locally on phones, tablets, and laptops.



DeepSeek



DeepSeek is a Chinese AI research lab focused on developing LLMs and AI systems. DeepSeek series includes DeepSeek-V2, a high-performance open-weight model trained on diverse multilingual datasets, aiming to compete with top-tier models like GPT and LLaMA. DeepSeek emphasizes efficiency, scalability, and domain-specific capabilities, particularly in Chinese language understanding and applications. The lab has also developed DeepSeekCoder, an advanced AI model optimized for code generation and programming assistance. With strong computational resources and a growing presence in the AI research landscape, DeepSeek is positioning itself as a significant player in the global LLM ecosystem.



Qwen - Alibaba



AI Search

Contact Sales

Call Us (+351 210 270 012)

English

Cart

Console

Log In

Why Us

Pricing

Products

Solutions

Marketplace

Developers

Partners

Documentation

Services

Model Studio

Free Trial

Solutions > Generative AI > Tongyi Qianwen (Qwen)

Tongyi Qianwen (Qwen)

Top-performance foundation models from Alibaba Cloud

Contact Sales

Qwen on Model Studio

Qwen2.5-Max: Exploring the Intelligence of Large-scale MoE Model

FEUP
FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO


LIACC

Luis Paulo Reis :: Navigating the Artificial Intelligence Revolution :: BIN@Novo Mesto :: Nov 10th, 2025 :: 49

RAG Architecture Model

- Retrieval Augmented Generation (RAG) has emerged as a powerful technique for improving LLMs
- Retrieving and conditioning external knowledge, RAG allows models to generate more accurate, relevant, and comprehensive text

RAG Architecture Model



<https://medium.com/@gulernilay088/>

- Advanced RAG enhances each module further with innovations like higher-order retrievers, cross-encoder rerankers, and evidence manipulation architectures

AI Applications in Health

- Radiology (X-ray, CT, MRI)
- *Dermatology (Image)*
- Drug/Treatment Discovery
- Risk Identification in Patients
- Primary Care and Screening
- Health Monitoring/Wearables
- Cognitive and Social Rehabilitation
- Physical Rehabilitation
- Patient Interaction with the Health System
- Health Systems Exchange of Information
- Surgical/Medical Robots
- Efficient Resource Allocation in Health
- ...



AI&Robotics Applications in Surgery

da Vinci
Surgical System




Mako
SmartRobotics



Rosa Brain
Robot



Hugo Surgical
Robot



Senhance Surgical System



Versus Surgical System




Navio S.S.



AI&Robotics in Nursing




AI in Nursing:

- **Personalized Treatment Plans** based on individual patient data, optimizing care
- **Workflow Optimization** with automatization of administrative tasks, efficient scheduling, allowing nurse staff to focus more on patient care
- **Patient Monitoring** with AI-powered devices continuously monitor vital signs and alert nurses to any deviations
- **Robotics and Automated assistants** helping with physical tasks, reducing strain and increasing efficiency in patient handling and monitoring




Artificial Intelligence in Digital Media

- Digital Images
- Digital Video
- Digital Audio
- Movie Industry
- Video Games
- Web Pages and Websites
- Social Media
- Digital Data and Databases
- Electronic Documents/Books




Artificial Intelligence in Social Media

- Social Creation and Management
- Social Insights
- Social Media Advertising
- Image/Face Recognition
- Fake News/Improper Content Recognition
- Personalised User Experience
- Advise Contacts/ Job candidates / Matches
- AI Powered Chatbots
- Social Listening
- Sentiment Analysis
- Improved Influencer strategies
- Increased Security



AI for the Environment

- **AI for Land/Agriculture**
 - Soil pollution, population growth, intensive agriculture
 - Land cover mapping, automating and optimizing farming processes
- **AI for Air/Pollution**
 - Pollution/air quality sources identification, reduce emissions
 - Intelligent and green transportation, autonomous cars, ride shares, dynamic bus routing, intelligent traffic lights
- **AI for Water**
 - Garbage collection, Intelligent ocean farming (fish and shellfish)
 - Predict storms, tsunamis, hurricanes, cost-effective water management
- **AI for Biodiversity**
 - Surveillance, tracking, analysis, prediction
- **AI for Climate Change**
 - Better climate predictions, effects of extreme weather, predict where carbon is coming from







AI & LLM Key Trends in 2025

- **Small Language Models (SLM)** drive **enterprise AI adoption**
- **Mixture of Experts (MoE) + LoRA** enable SLMs to outperform 10x larger LLMs
- **Open-source** are becoming the **de facto** way for **LLMs** use
- Large language models start adopting **modular architectures**
- **Synthetic data** revolution arrives to LLMs
- LLM **hallucinations disappear** as training techniques evolve
- **Data** (not LLMs) becomes the **true competitive advantage**
- **RAGs** (Retrieval-Augmented Generation) used in most LLMs tasks
- Shift from closed to open-source **democratizes AI** (transparency)
- **Agentic AI – AI that acts!**
- **LLMs for Machines and Robots**, not just people - **LBM**s

Limits of Deep Learning


- **Brilliant but Passive!**
- **Deep nets can predict, but they can not plan. They lack goals, intents, and common sense**
- **This sets the stage for Agentic AI: From Prediction to Decision and Action...**
- **The next leap is agency — AI that senses, reasons, and acts toward goals**
- **Predictive (Model→Output) vs Agentic (Sense→Think→Act→Learn)**



Autonomous Agents and Multi-Agent Systems

Agent (1995):


“Computational System, situated in a given **environment**, that has the ability to **perceive** that environment using **sensors** and **act**, in an **autonomous way**, in that environment using its **actuators** to fulfill a given **function**.”



From Russel and Norvig, "AI: A Modern Approach", 1995

Multi-Agent System (1995):


- Agents exhibit **autonomous behavior**
- **Interact** with other agents in the system



Agentic AI

- **Autonomy:** Ability to initiate and complete tasks with limited direct human supervision, with great flexibility and efficiency in executing tasks
- **Reasoning:** Sophisticated decision-making based on context and trade-offs
- **Reinforcement Learning:** Dynamically evolve by interacting with their environment and receiving feedback from these interactions.
- **Language Understanding:** Capacity to understand and follow complex instructions
- **Workflow Optimization:** Efficient execution of multi-step complex processes

Feature	Agentic AI	Generative AI	Traditional AI
Primary Function	Goal-oriented action & decision-making	Content generation (text, code, images, etc.)	Focused on automating repetitive tasks
Autonomy	High – Operates with minimal human oversight	Variable – May require user prompts or guidance	Low – Relies on specific algorithms and set rules
Learning	Reinforced Learning – Improves through experience	Data-driven learning – Learns from existing data	Relies on predefined rules and human intervention





<https://aisera.com/blog/agentic-ai/>

New & Emerging Trends (2025+)

- **Agentic AI Rises** — Autonomous agents that reason, plan, and act toward goals
- **Large Behavior Models (LBMs)** — Models learning from actions and interactions, not just text
- **Multi-Agent Collaboration** — Intelligent agents cooperate and negotiate to solve complex tasks
- **Responsible & Regulated AI** — Transparency, fairness, and compliance standards
- **Edge & On-Device AI** — Efficient, private models running locally on devices /robots
- **Hybrid Human–AI Teams** — Shared intelligence combining human oversight and machine efficiency


EU AI Act – Regulating the Past...

- First regulation on artificial intelligence
- AI can create many benefits:
 - better healthcare
 - safer and cleaner transport
 - more efficient manufacturing
 - cheaper and more sustainable energy.



The Rise of Deep Reinforcement Learning (DRL)

- **Beyond Short-Term Thinking**
 - LLMs have limited context windows — they forget.
 - Memory and planning enable continuity and strategic reasoning
- **Reflection and Self-Improvement**
 - Evaluate their own reasoning, detect failures, and adapt
 - Self-reflection transforms agents from reactive to self-correcting
- **Tools/Actuators Use and World Interaction**
 - Intelligence changes into capability
- **Learning by Doing**
 - Reinforcement Learning teaches through feedback — an agent acts, receives a reward, and updates its policy to improve next time.



AI in Robotics




Robotics – Atlas Robot



Video: Boston Dynamics: Atlas Gets a Grip - https://www.youtube.com/watch?v=e1_Qh1EhG30&channel=BostonDynamics

Robotic Humanoids



https://pplware.sapo.pt/high-tech/existem-tantos-robos-humanoides-que-fica-dificil-distingui-los-mas-este-grafico-pode-ajudar/?fbclid=IwY2xjawl1Zg9leHRuA2FlbQlXMQABHS_EdekrvewLEgIyq3RZRfPy-HaU4MdNtaFgqa6bT2lkoYfBoF_Jkqd9w_aem_RO6nDQ16bQJ51jD6bM6-TQ

Unitree and Booster Robots



RoboCup: Objectives

- Joint International Project:
 - (Distributed) Artificial Intelligence
 - Intelligent Robotics
- Soccer – Central Research Topic:
 - Very complex collective game
 - Huge number of technologies involved:
 - Autonomous Agents, Multi-Robot Systems, Cooperation, Communication, Strategic Reasoning, Robotics, Sensor Fusion, Real-Time Reasoning, Machine Learning, etc
- 5 Challenges: Soccer, Rescue, @Home, Industrial, Junior
- 17 Leagues, 2500 Humans and 2000 Robots!
- Main Goal: ***“By 2050, develop a team of fully autonomous humanoid robots that may win against the human world champion team in soccer!”***



RoboCup Leagues and Challenges

Industrial



Mobility (Remote Operator)



Stairs with Debris



Rescue

@Home




Soccer

RoboCup - Robotic Soccer Competitions

Middle
Size
League



Standard
Platform
League



RoboCup

- Real, Standard, Simulated Robots
- Mini, Small, Medium and Large Robots
- Wheeled, Legged and Humanoid Robots
- **Distinct but interrelated Leagues/Problems**
- Only a Few Research Groups able to develop code that works in more than one league!

Small
Size
League




Humanoid
League



RoboCup Leagues and Participations

FC Portugal Team (FEUP/UAveiro)



Simulation 3D League (Humanoids)




Deep Reinforcement Learning (DRL)

A Classic Reinforcement Learning

Reinforcement Learning Problem

Tabular Solution



B Classic Deep Learning

Categorization Problem

Deep Learning Solution







Image: <https://ucsdneuro.wordpress.com/2021/06/13/diving-deep-into-the-brain-a-case-for-deep-reinforcement-learning-in-neuroscience/>

C Deep Reinforcement Learning: Deep learning solutions for RL problems



Brains + Experience = Deep Reinforcement Learning

DRL proved machines can outperform humans in dynamic environments

Robots can now learn directly from physical interaction

Deep Reinforcement Learning (DRL)

Model-Free RL learns directly from interaction with the environment.


Model-Based RL builds an internal model of the environment to improve decision-making




Policy Optimization: Directly optimizes policy

Q-Learning: Optimizes a value function and derives the policy from it

PPO (Proximal Policy Optimization)

- Stability in training
- Sample efficiency
- Balancing exploration and exploitation





IEETA
https://www.ieeta.pt › New

APPIA Award for Best PhD Thesis in Artificial Intelligence

6 Oct 2025 — The Portuguese Association for Artificial Intelligence (APPIA) has awarded the 2024 Best PhD Thesis in Artificial Intelligence to Miguel Abreu ...


Notícias U.Porto
https://noticias.up.pt › 2025/10/09 · Translate this page

Melhor Tese de Doutoramento em Inteligência Artificial de ...

9 Oct 2025 — Prémio da Associação Portuguesa para a Inteligência Artificial foi atribuído a Miguel Abreu, doutorado em Engenharia Informática pela FEUP.

Large Behavior Models - LBM

- **LBM**s go beyond processing language. They're designed to **interpret, predict, and generate complex sequences of human-like actions and behaviors**
- While LLMs focus on text and dialogue, LBM specialize in **behavioural patterns, decision-making** processes, and **interactive scenarios**
- LLMs try to guess the next word, while **LBM**s try to guess the next **action or behaviour**
- LBM are enhancing robots' abilities to move naturally and respond intuitively to their surrounding
- LBM rely on:
 - **Transformer Architecture**: Similar to LLMs but adapted to handle sequences of actions/behaviours
 - **Multi-Modal Learning**: Integrating visual, spatial, and temporal data to enhance understanding
 - **Reinforcement Learning**: Learning from interactions and feedback to improve behavior over time



Conclusions

- **AI has become a Partner no longer a Tool!**
- We've travelled from **Symbolic Logic** to self-learning **Agents and Robots**
- Impossible to prevent **AI** use in **All Areas** of **Society and Work**
- **Agentic AI, LBMs** and **New GenAI** and **DRL** powered and trained **Robots**
- **Are We Ready for Self-Improving Intelligent Machines?**
- **Next generation of AI** won't just predict — it will **act, learn, and evolve**
- **By 2030, AI agents** will be integrated into **every digital platform/robotic system** — from industry to healthcare to daily life (**LLMs, DRL, LBMs**)
- **Our challenge is to ensure AI Grows With Us, Not Against Us!**

LIACC



U.PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Navigating the Artificial Intelligence Revolution

Luís Paulo Reis

lpreis@fe.up.pt

Associate Professor at FEUP - Faculty of Engineering of the University of Porto
Director of LIACC – Artificial Intelligence and Computer Science Laboratory, Univ. Porto
Member of Coordinating Commission of LASI – Intelligent Systems Associate Laboratory
President of the GA of APPIA – Portuguese Association for Artificial Intelligence

