



DECSAI

Departamento de Ciencias de la Computación e I.A.

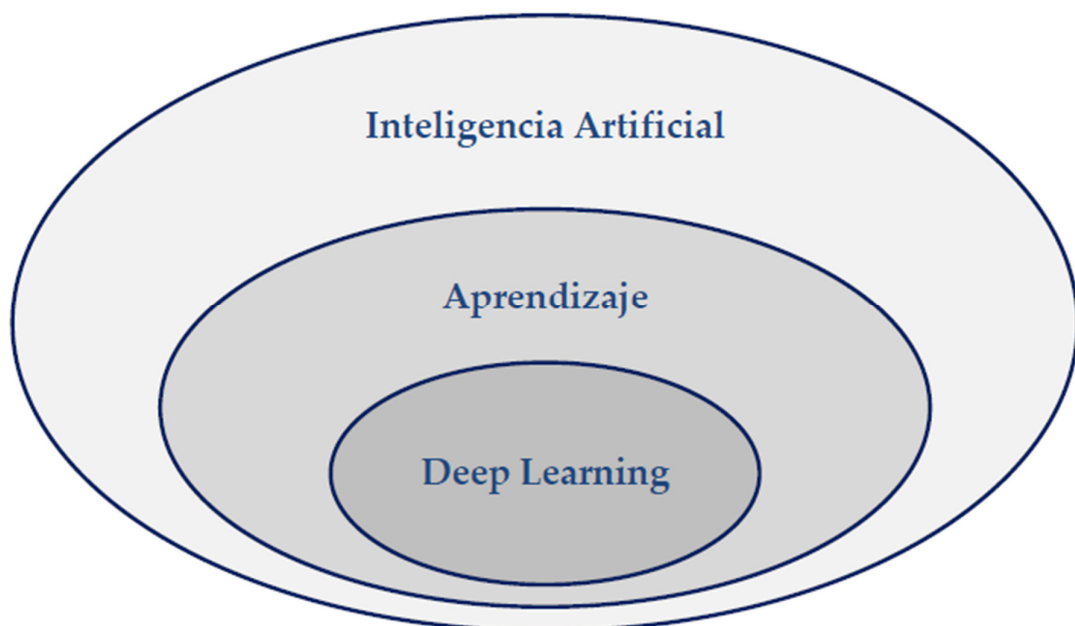
Universidad de Granada



Entrenamiento de redes neuronales

Fernando Berzal, berzal@acm.org

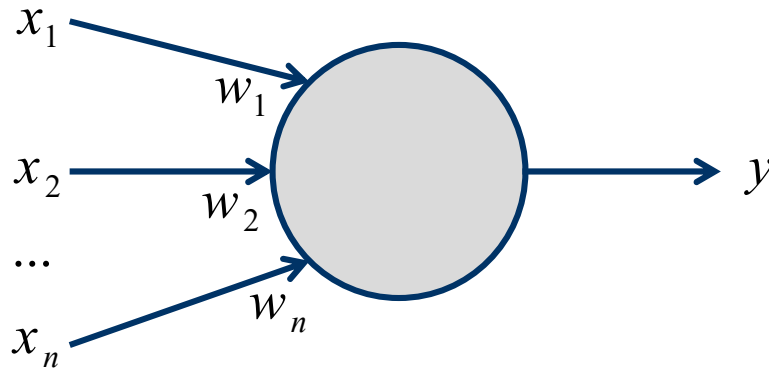
Deep Learning





Neuronas

El modelo computacional más simple de una neurona



$$y = \sum_i x_i w_i = x_1 w_1 + x_2 w_2 + \dots + x_n w_n$$

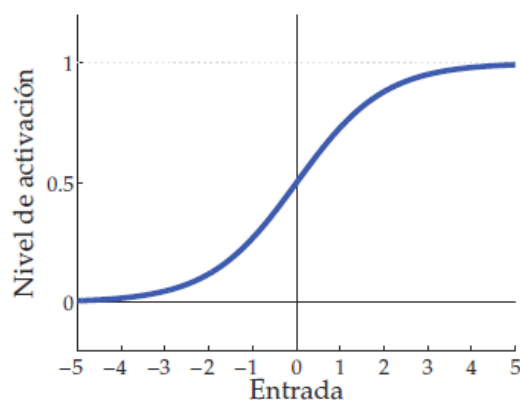


Modelo de neurona sigmoïdal

Función logística

$$z = \sum_i x_i w_i$$

$$y = \frac{1}{1 + e^{-z}}$$

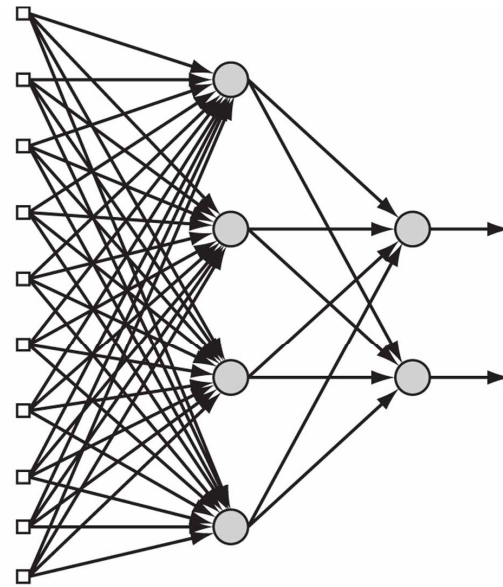




Redes neuronales multicapa [feed-forward neural networks]

Topología de red más habitual en la práctica.

Si hay más de una capa oculta, se denominan **"deep" neural networks**.



Input layer of source nodes

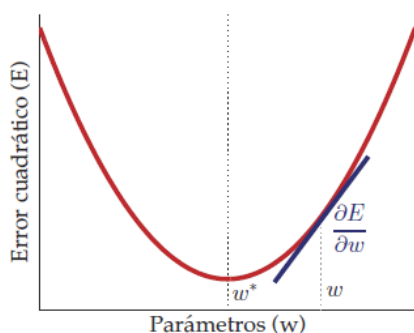
Layer of hidden neurons

Layer of output neurons

[Haykin: "Neural Networks and Learning Machines", 3rd edition]



Algoritmo de entrenamiento



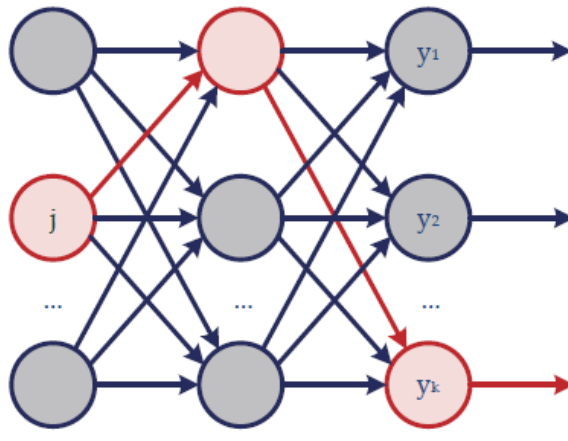
$$\Delta w_i = -\eta \frac{\partial E}{\partial w_i}$$



Backpropagation



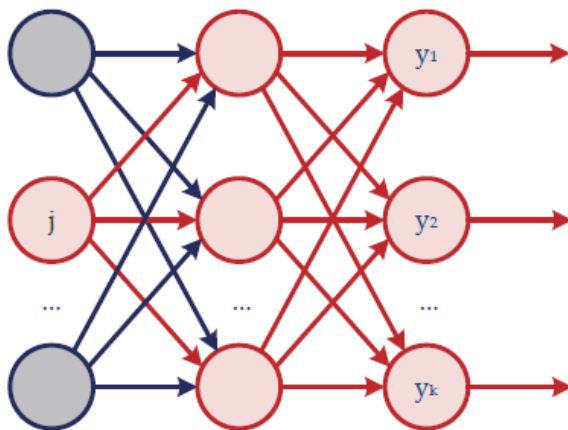
Las neuronas ocultas afectan a la salida de la red...



Backpropagation



... por múltiples caminos



Backpropagation



IDEA

No sabemos qué deben hacer las neuronas ocultas, pero podemos calcular cómo cambia el error cuando cambia su actividad.

- En vez de utilizar la salida deseada para entrenar las neuronas ocultas, usamos la derivada del error con respecto a sus actividades ($\delta E/\delta y$).
- La actividad de cada neurona oculta puede tener efectos en muchas neuronas de salida, por lo que debemos combinarlos.
- Una vez que tenemos las derivadas del error para todas las unidades ocultas, podemos calcular las derivadas del error para sus pesos de entrada.

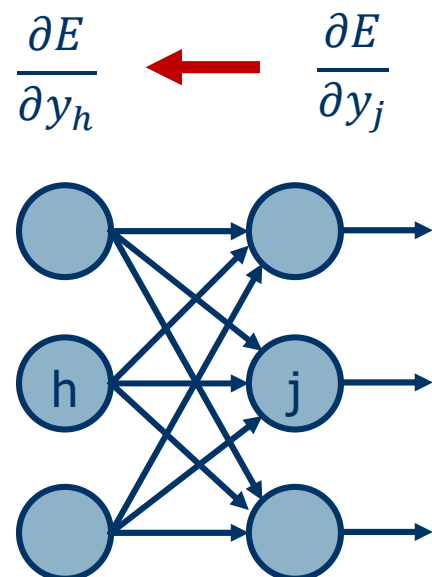


Backpropagation



Propagación de errores

- Convertimos la discrepancia entre la salida de la red y su salida deseada en una derivada del error ($\partial E/\partial y$).
- Calculamos las derivadas del error de cada capa oculta a partir de las derivadas del error de la capa superior.
- Usamos $\partial E/\partial y$ para obtener $\partial E/\partial w$.



Backpropagation

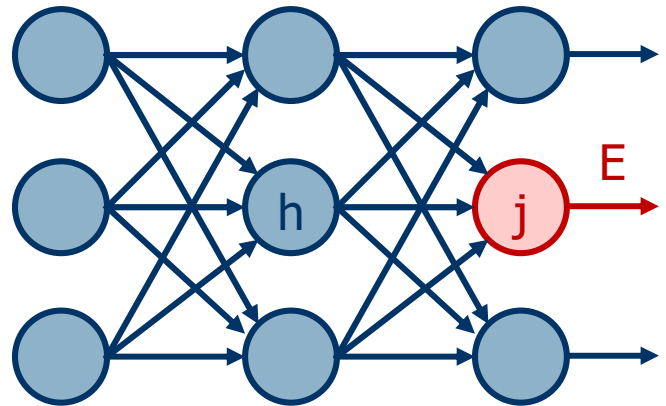


Propagación de errores $\partial E/\partial y$

Conocemos el error E en la capa de salida y cómo calcular su gradiente para las neuronas de salida j .

Podemos calcular cómo varía el error en función de la entrada neta de las neuronas (sigmoidales) de salida z_j :

$$\frac{\partial E}{\partial z_j} = \frac{\partial y_j}{\partial z_j} \frac{\partial E}{\partial y_j} = f'(z_j) \frac{\partial E}{\partial y_j} = y_j(1 - y_j) \frac{\partial E}{\partial y_j}$$



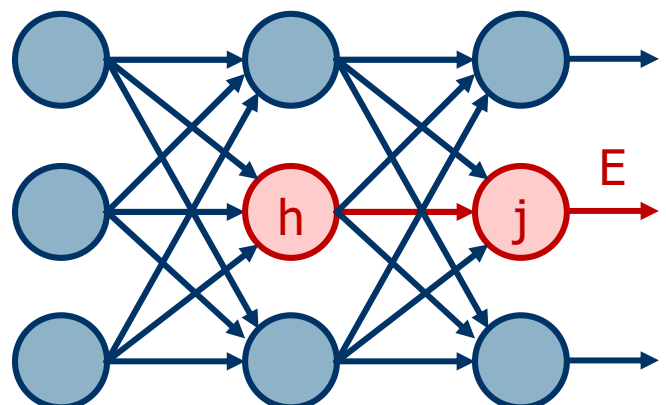
Backpropagation



Propagación de errores $\partial E/\partial y$

A continuación, podemos calcular cómo fluctúa el error en función de la salida y_h de la neurona oculta (h):

$$\frac{\partial E}{\partial y_h} = \frac{\partial z_j}{\partial y_h} \frac{\partial E}{\partial z_j} = w_{hj} \frac{\partial E}{\partial z_j}$$

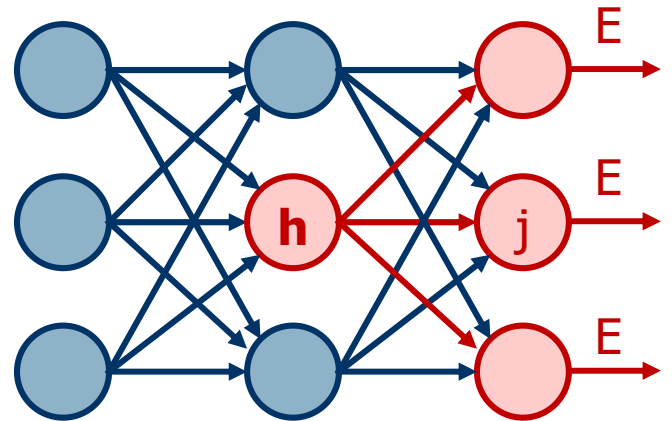


Backpropagation



Propagación de errores $\partial E/\partial y$

Pero tenemos que tener en cuenta que la neurona oculta influye en el error observado en todas las neuronas de la capa siguiente:



$$\frac{\partial E}{\partial y_h} = \sum_j \frac{\partial z_j}{\partial y_h} \frac{\partial E}{\partial z_j} = \sum_j w_{hj} \frac{\partial E}{\partial z_j}$$

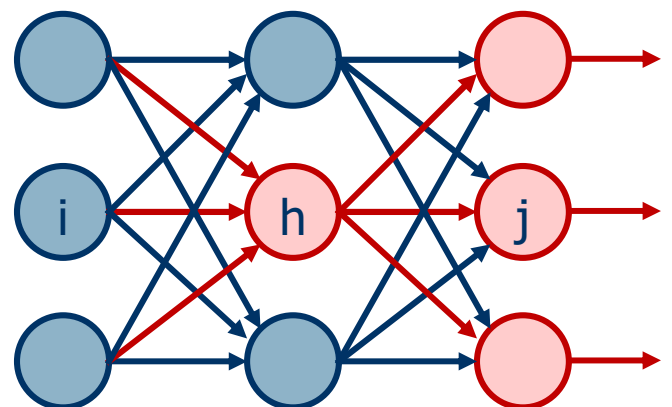


Backpropagation



Propagación de errores $\partial E/\partial y$

Por último, ya estamos en condiciones de calcular cómo fluctúa el error en función de los parámetros de la neurona oculta (h), los pesos w_{ij} :



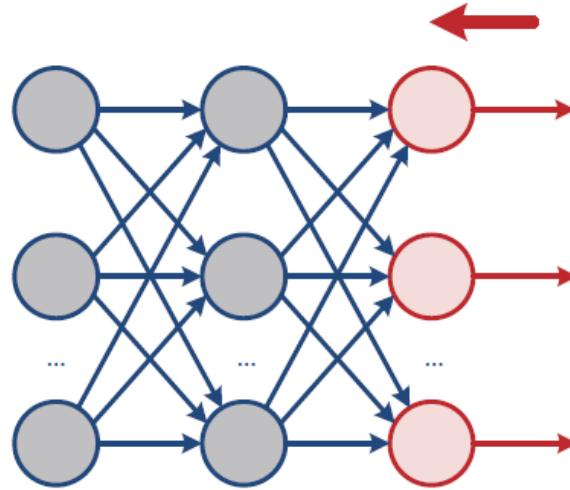
$$\frac{\partial E}{\partial w_{ih}} = \frac{\partial z_h}{\partial w_{ih}} \frac{\partial E}{\partial z_h} = x_{ih} \frac{\partial E}{\partial z_h} = y_i \frac{\partial E}{\partial z_h}$$



Backpropagation



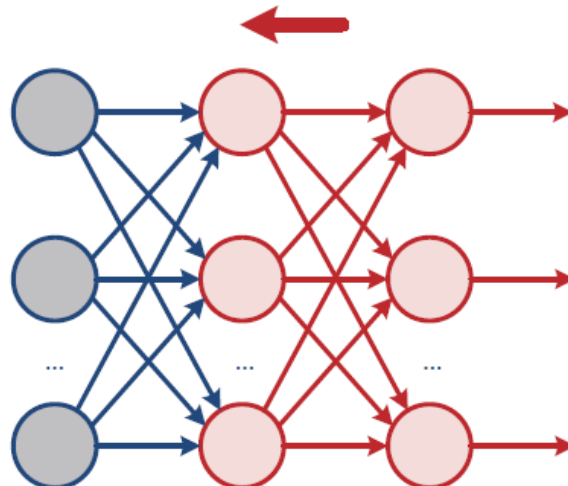
Propagación de errores $\delta E/\delta y$



Backpropagation



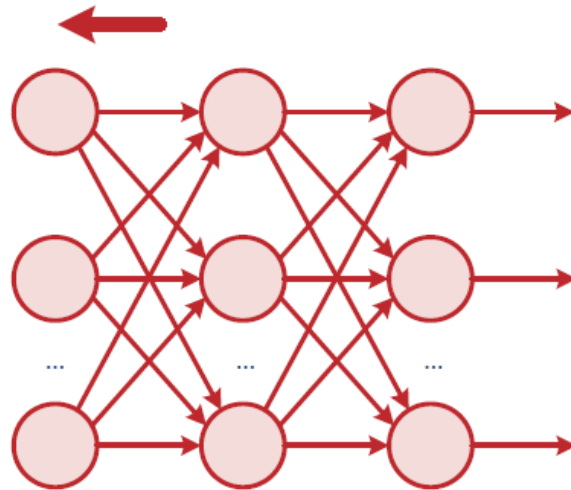
Propagación de errores $\delta E/\delta y$



Backpropagation



Propagación de errores $\delta E/\delta y$

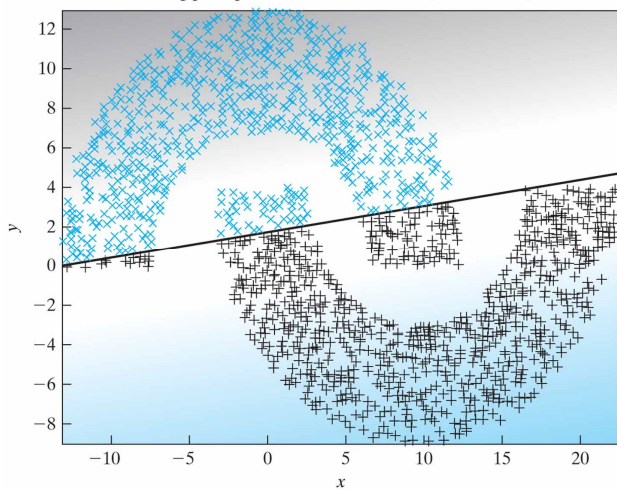


Backpropagation

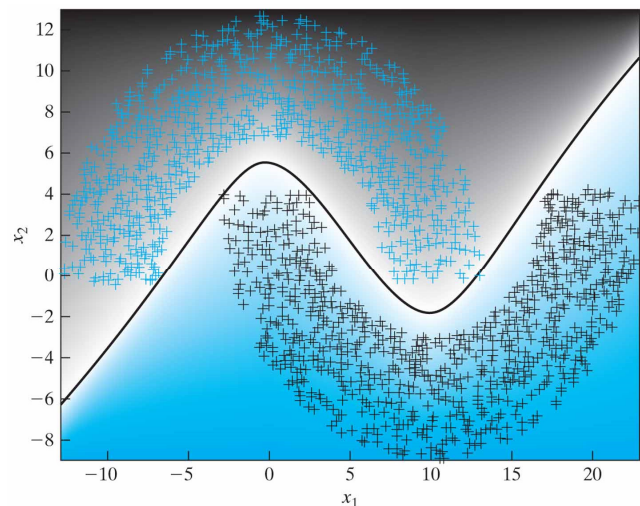


El resultado...

Classification using perceptron with distance = -4, radius = 10, and width = 6



(b) testing result



(b) Testing result

Perceptrón

Red multicapa



Aplicaciones



- Reconocimiento de voz
- Clasificación de imágenes
 - Imágenes médicas, p.ej. Radiografías
 - Señales de tráfico (vehículos autónomos)
 - Reconocimiento facial
 - ...
- Identificación y seguimiento de objetos, p.ej. Drones
- Traducción automática
- ...

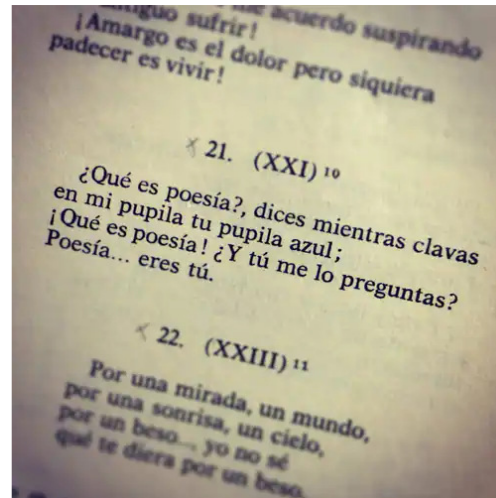


Pregunta



¿Puede un ordenador ser creativo?





IA Generativa - Textos

Fernando Berzal, berzal@acm.org

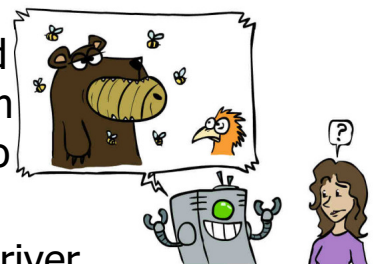
Capacidades de la I.A.

Historias "divertidas" (sin querer)

TALE-SPIN System, James Meehan, UC Irvine, 1976

http://en.wikipedia.org/wiki/Computational_creativity

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.
- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.
- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.



Berkeley CS188



Capacidades de la I.A.



¿Alguna mejora?

Agentes conversacionales



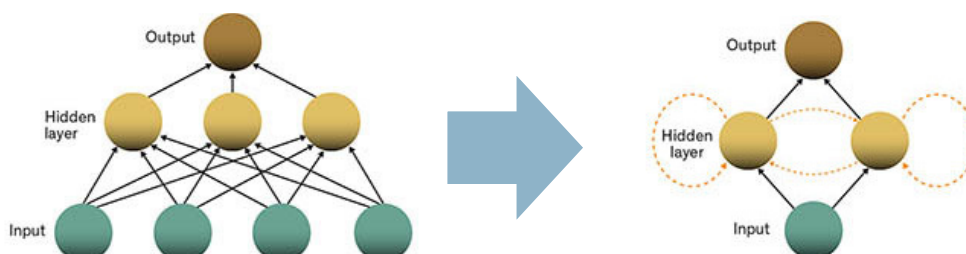
Capacidades de la I.A.



Redes neuronales recurrentes, 2016

<http://spectrum.ieee.org/computing/software/the-neural-network-that-remembers>

- "On tap at the brewpub. A nice dark red color with a nice head that left a lot of lace on the glass. Aroma is of raspberries and chocolate. Not much depth to speak of despite consisting of raspberries. The bourbon is pretty subtle as well. I really don't know that find a flavor this beer tastes like. I would prefer a little more carbonization to come through. It's pretty drinkable, but I wouldn't mind if this beer was available."
-- Review for a "fruit/vegetable beer".





2023

Model: GPT-4

BS Act as: Editor
Degree of revision: Substantial Revision
Type of edit: Enhance clarity and consistency
Change style to: Academic, PhD Work
Change tone to: Analytical
Change reader comprehension level to: advanced, assume extensive prior knowledge
Change length to: 1000 Words

My Text:

Brav.. hav u chaked the gpt-4 thing?? U can writ any branles drunken drivle in it and it changes it to any level even PhD hyper brain style.. trulli insane.. you, got to check it bro! Its a tru gift for all of us illiterate assholes around the the global flat plane! Now we can have universit jobs.. hhah crazy stuff...

Title: A Comprehensive Analysis of the GPT-4 Language Model and its Potential Impact on Scholarly Communication

Abstract: This article explores the functionality and capabilities of OpenAI's GPT-4, a sophisticated language model capable of generating coherent and semantically rich text. We examine the model's potential to transform written communication, particularly for those with limited literacy skills, by analyzing the model's capacity for tailoring text to various academic and linguistic requirements. Furthermore, we discuss the implications of the model's wide-ranging applications, including its potential to democratize access to higher education and professional opportunities.

Introduction

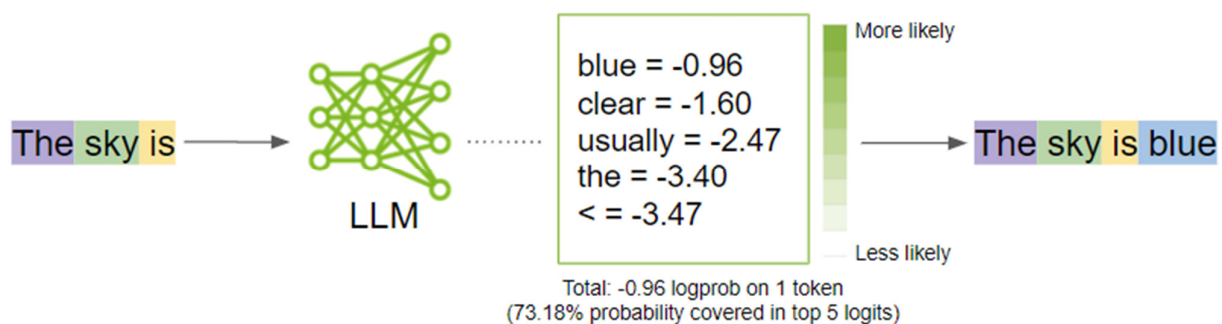


LLM [Large Language Models]



Red neuronal

Predicción de la siguiente palabra en una secuencia



IDEA

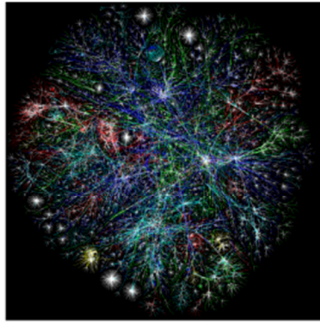
Aprender a predecir la siguiente palabra obliga a la red a aprender mucho acerca del mundo...



LLM [Large Language Models]

Entrenamiento del modelo base

El LLM "comprime" el contenido de Internet



Chunk of the internet,
~10TB of text



6,000 GPUs for 12 days, ~\$2M
~1e24 FLOPS



parameters.zip

~140GB file

p.ej. Llama 2 70B (Facebook), <https://ai.meta.com/llama/>

Andrej Karpathy: "Intro to LLMs", November 2023



LLM [Large Language Models]

Resultado

Una red que "sueña" documentos (alucina en el argot ;-)

```
package lecture06;

import java.util.*;

public class Field {

    public static Map<String, Card> cardIndex = new HashMap<String, Card>();

    public void addCard(Card card){
        this.cardIndex.put(card.getRegistrationNumber(),card);
    }

    public Card getCard(String registrationNumber) {
        return this.cardIndex.get(registrationNumber);
    }

    public int getCardCount() {
        return this.cardIndex.size();
    }
}
```

Java code dream

Hades Heroes

Author: Maureen Fergus

Product Details:

ISBN: 9780545946767

Format: Hardback

Pages: 176

Dims (mm): 140 x 223

Pub Date: 01-05-16

Pub Country: United States

Condition: NEW

Description:

Work hard. Love. Make your own luck. These are the keys to success in life at school. But when a student named Hercules wanders in and challenges him

Amazon product dream (?)

Blacknose dace

The blacknose dace (*Rhinichthys atratulus*) is a small freshwater fish found in North America. They are typically found in streams, rivers, and lakes, and are widely distributed throughout the eastern and central United States and Canada.

These fish have a slender body with a black stripe running horizontally along their sides, giving them their distinctive "blacknose" appearance. They have a flattened head with a pointed snout and small, round eyes. Their color can vary from dark brown to olive green on their back and sides, with a silvery-white underside.

Blacknose dace are omnivorous and feed on a variety of small aquatic insects, crustaceans, and plant matter. They are an important food source for larger predators such as bass,

Wikipedia article dream

Andrej Karpathy: "Intro to LLMs", November 2023



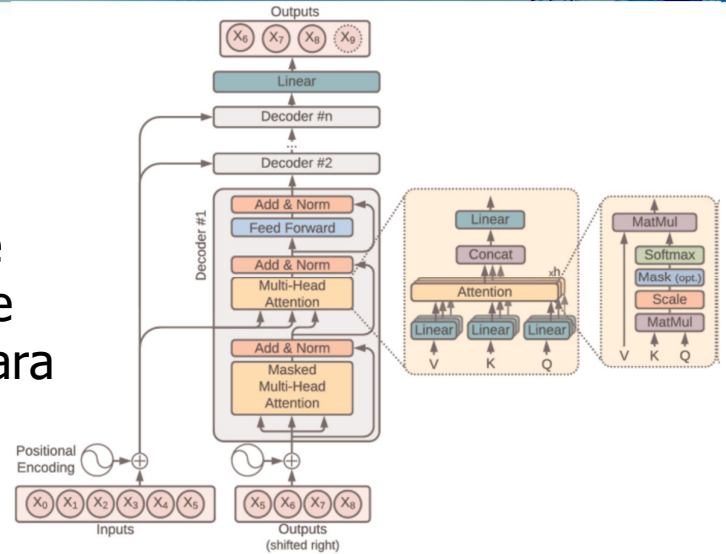
LLM [Large Language Models]

Resultado

¿Cómo funciona?

Red con cientos de miles de millones de parámetros, que ajustamos iterativamente para mejorar las predicciones...

... no se sabe realmente cómo "colaboran" todos esos parámetros.



Andrej Karpathy: "Intro to LLMs", November 2023



28

LLM [Large Language Models]

Resultado

En cierto modo, almacena lo que ve, pero de manera imperfecta (compresión con pérdidas).

Recent viral example: "reversal curse"

Q: "Who is Tom Cruise's mother"?

A: Mary Lee Pfeiffer ✓

Q: "Who is Mary Lee Pfeiffer's son?"

A: I don't know ✗ 🤔

De ahí el "prompt engineering"...



Pedro Domingos ✓
@pmdomingos

Prompt engineering is like trying to shape someone's dreams by telling them things during the day.



Pedro Domingos ✓
@pmdomingos

With enough prompt engineering, you don't even need the LLM.



29

LLM [Large Language Models]

Refinamiento [fine-tuning]: Asistentes digitales

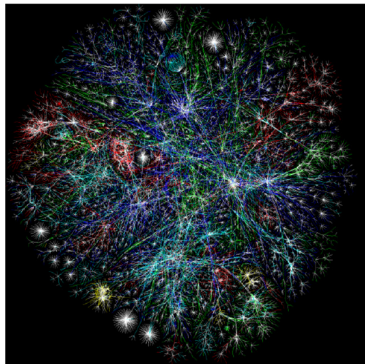
Se reemplaza el conjunto de entrenamiento (i.e. Internet) por uno preparado manualmente...

Training the Assistant

(Just swap the dataset, then continue training)



written by
people:



<USER>

Can you write a short introduction about the relevance of the term "monopsony" in economics? Please use examples related to potential monopsonies in the labour market and cite relevant research.

<ASSISTANT>

"Monopsony" refers to a market structure where there is only one buyer for a particular good or service. In economics, this term is particularly relevant in the labor market, where a monopsony employer has significant power over the wages and working conditions of their employees. The presence of a monopsony can result in lower wages and reduced employment opportunities for workers, as the employer has little incentive to increase wages or provide better working conditions...

~100K conversations

Andrej Karpathy: "Intro to LLMs", November 2023



LLM [Large Language Models]

Resultado: Asistentes digitales

Ya tenemos un chat bot...

<USER>

Can you help me with this code? It seems like there is a bug.
print("hello world")

Andrej Karpathy: "Intro to LLMs", November 2023



LLM [Large Language Models]

Resumen: Cómo crear tu propio ChatGPT



every
~year

Stage 1: Pretraining

1. Download ~10TB of text.
2. Get a cluster of ~6,000 GPUs.
3. Compress the text into a neural network, pay ~\$2M, wait ~12 days.
4. Obtain **base model**.

Knowledge



<USER>
Can you write a short introduction about the relevance of the term "monopsony" in economics? Please use examples related to potential monopsonies in the labour market and cite relevant research.
<ASSISTANT>
"Monopsony" refers to a market structure where there is only one buyer for a particular good or service. In economics, this term is particularly relevant in the labor market, where a monopsony employer has significant power over the wages and working conditions of their employees. The presence of a monopsony can result in lower wages and reduced employment opportunities for workers, as the employer has little incentive to increase wages or provide better working conditions...

every
~week

Stage 2: Finetuning

1. Write labeling instructions
2. Hire people (or use [scale.ai](#)!), collect 100K high quality ideal Q&A responses, and/or comparisons.
3. Finetune base model on this data, wait ~1 day.
4. Obtain **assistant model**.
5. Run a lot of evaluations.
6. Deploy.
7. Monitor, collect misbehaviors, go to step 1.

Alignment

Andrej Karpathy: "Intro to LLMs", November 2023



LLM [Large Language Models]

Refinamiento [fine-tuning]: Aprendizaje por refuerzo

Es mucho más sencillo comparar respuestas que crearlas...

You
Write a haiku about paperclips

ChatGPT
Steel curves intertwine,
Holding life's pages in line,
Quiet, small design.

You
Write a haiku about paperclips

ChatGPT
Silver arcs entwined,
Holding life's pages, aligned,
In steel we find, bind.

< 12 / 12 >

You
Write a haiku about paperclips

ChatGPT
Silver arcs entwined,
Humble joiner of pages,
Quiet strength, aligned.

< 2 / 2 >

e.g. OpenAI RLHF [Reinforcement Learning from Human Feedback]

Andrej Karpathy: "Intro to LLMs", November 2023



LLM [Large Language Models]

Ranking ELO

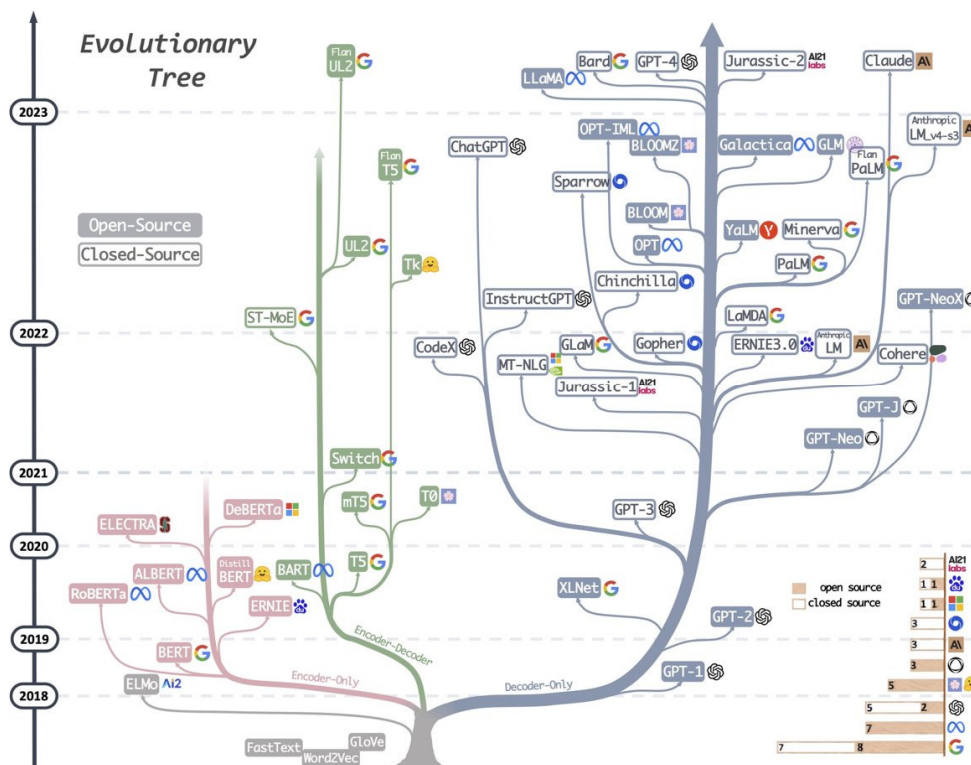
<https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

Model	Arena Elo	95% CI	Votes	Organization	License
GPT-4o-2024-05-13	1289	+6/-5	12808	OpenAI	Proprietary
GPT-4-1106-preview	1250	+3/-3	81015	OpenAI	Proprietary
GPT-4-Turbo-2024-04-09	1252	+4/-2	59114	OpenAI	Proprietary
Gemini 1.5 Pro API-0409-Preview	1248	+3/-3	62929	Google	Proprietary
Claude 3 Opus	1246	+2/-2	116342	Anthropic	Proprietary
GPT-4-0125-preview	1244	+3/-2	75206	OpenAI	Proprietary
Bard (Gemini Pro)	1208	+5/-5	12387	Google	Proprietary
Llama 3 70B Instruct	1203	+3/-2	118358	Meta	Llama 3 Community
Claude 3 Sonnet	1199	+3/-3	94050	Anthropic	Proprietary
Reka-Coxe-20240501	1195	+4/-3	33503	Reka AI	Proprietary
Qwen-Max-0428	1187	+5/-4	20794	Alibaba	Proprietary
GPT-4-0314	1189	+3/-3	54137	OpenAI	Proprietary
Command R+	1189	+3/-3	60230	Cohere	CC-BY-NC-4.0
Claude 3 Haiku	1181	+3/-2	84423	Anthropic	Proprietary



LLM [Large Language Models]

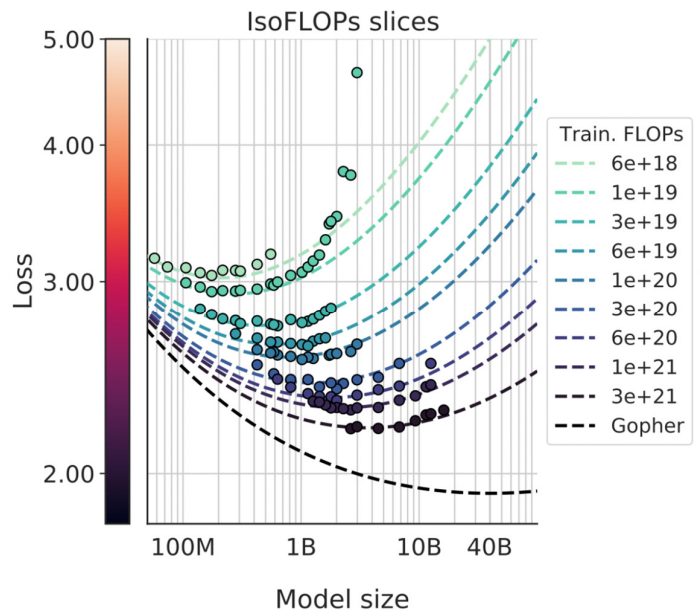
Evolución



LLM [Large Language Models]

Evolución

Rendimiento observado en función del número de parámetros (N) y de la cantidad de texto (D) usada en su entrenamiento.



LLM Scaling

"Training Compute-Optimal Large Language Models"
arXiv, 2022, <https://arxiv.org/abs/2203.15556>



LLM [Large Language Models]

Regulación...

Grading Foundation Model Providers' Compliance with the Draft EU AI Act

Source: Stanford Research on Foundation Models (CRFM), Institute for Human-Centered Artificial Intelligence (HAI)

	OpenAI	cohere	stability.ai	ANTHROPIC	Google	BigScience	Meta	AI21 labs	ALPHA ALPHA	EleutherAI	Totals
Draft AI Act Requirements	GPT-4	Cohere Command	Stable Diffusion v2	Claude	PaLM 2	BLOOM	LLaMA	Jurassic-2	Luminous	GPT-NeoX	
Data sources	●○○○	●●●○	●●●●	○○○○	●●○○	●●●●	●●●●	○○○○	○○○○	●●●●	22
Data governance	●●○○	●●○○	●●○○	○○○○	●●○○	●●●●	●●○○	○○○○	○○○○	●●●●	19
Copyrighted data	○○○○	○○○○	○○○○	○○○○	○○○○	●●●●	○○○○	○○○○	○○○○	●●●●	7
Compute	○○○○	○○○○	●●●●	○○○○	○○○○	●●●●	●●●●	○○○○	●○○○	●●●●	17
Energy	○○○○	●○○○	●●●●	○○○○	○○○○	●●●●	●●●●	○○○○	○○○○	●●●●	16
Capabilities & limitations	●●●●	●●●○	●●●●	○○○○	●●●●	●●●○	●●○○	●○○○	●○○○	●●●○	27
Risks & mitigations	●●●●	●●○○	○○○○	●○○○	●●●●	●●○○	●○○○	●○○○	○○○○	●○○○	16
Evaluations	●●●●	●●○○	○○○○	○○○○	●●○○	●●○○	●○○○	○○○○	●○○○	●○○○	15
Testing	●●●○	●●○○	○○○○	○○○○	●●○○	●●○○	○○○○	●○○○	○○○○	○○○○	10
Machine-generated content	●●●●	●●●○	○○○○	●●●●	●●○○	●●●●	○○○○	●●●●	●○○○	●●○○	21
Member states	●●●●	○○○○	○○○○	○○○○	●●●●	○○○○	○○○○	○○○○	●○○○	○○○○	9
Downstream documentation	●●●●	●●●●	●●●●	○○○○	●●○○	●●○○	○○○○	○○○○	○○○○	●●●●	24
Totals	25 / 48	23 / 48	22 / 48	7 / 48	27 / 48	36 / 48	21 / 48	8 / 48	5 / 48	29 / 48	



ChatGPT



ChatGPT

GPT-3.5

175B parámetros

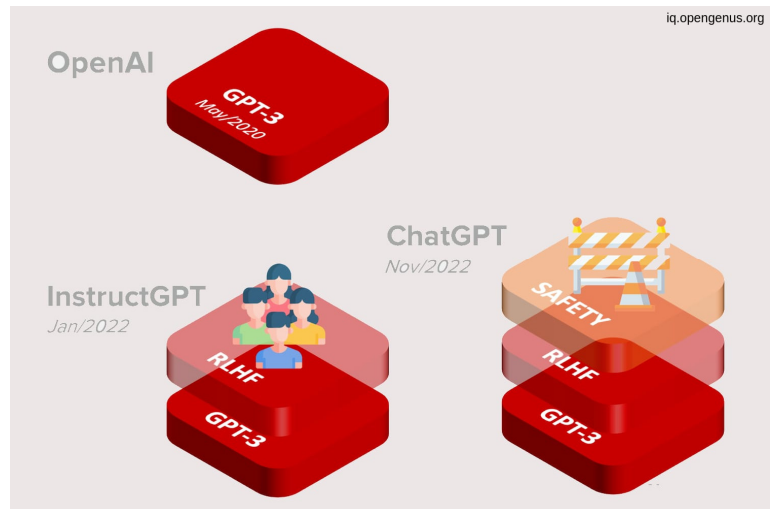
+

RLHF

Reinforcement Learning
from Human Feedback

+

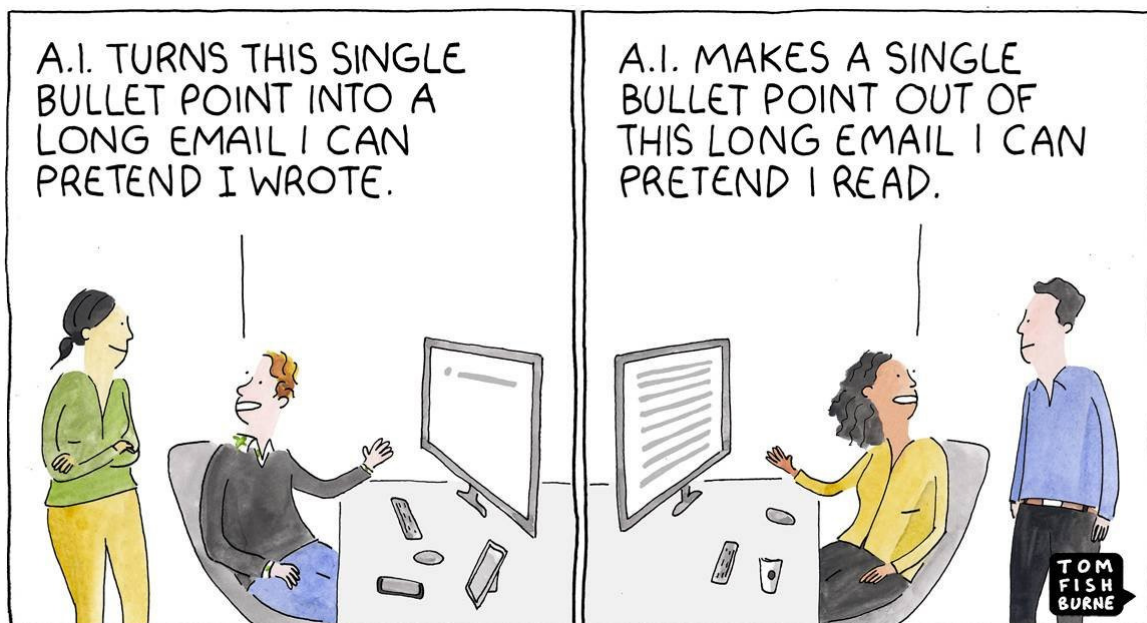
Safety features



2022



IA Generativa



Análisis de sentimientos

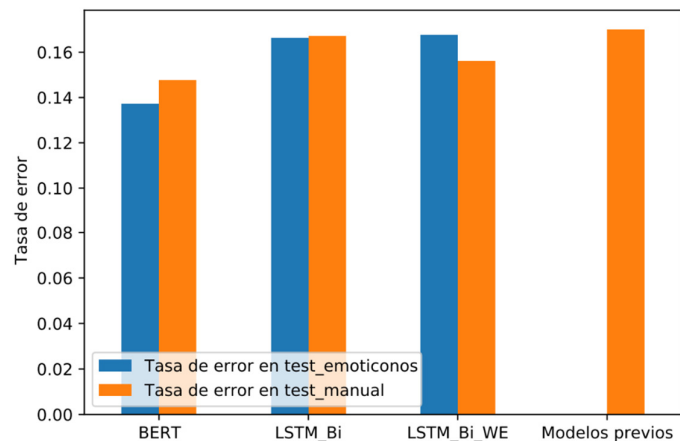
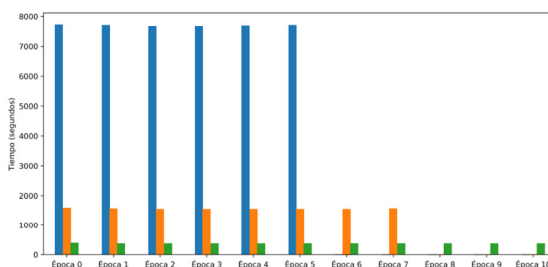


Análisis de sentimientos



Análisis de sentimientos con transformers / LLMs

Modelo	Nº de parámetros total	Nº de parámetros ajustables	Nº de parámetros "congelados"
Modelo 1: LSTM_Bi	6,242,945	6,242,945	0
Modelo 2: LSTM_Bi_WE	4,884,609	84,609	4,800,000
Modelo 3: BERT	108,311,810	108,311,810	0





DECSAI

Departamento de Ciencias de la Computación e I.A.

Universidad de Granada



IA Generativa - Imágenes

Fernando Berzal, berzal@acm.org

Síntesis de imágenes

Transferencia de estilos



Leon A. Gatys, Alexander S. Ecker & Matthias Bethge:

A Neural Algorithm of Artistic Style

arXiv, 2015. <http://arxiv.org/abs/1508.06576>



Síntesis de imágenes



<https://thispersondoesnotexist.com/>



StyleGAN <https://arxiv.org/abs/1812.04948> CVPR'2019



Síntesis de imágenes



"You sketch, the AI paints"



GauGAN, NVIDIA, CVPR'2019



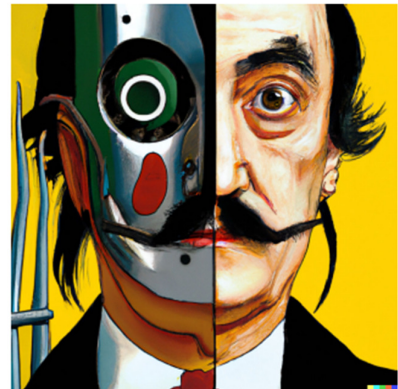
Síntesis de imágenes



DALL-E



Versión modificada de GPT-3 para generar imágenes de 256x256 a partir de descripciones textuales [12B parámetros]



vibrant portrait painting of Salvador Dalí with a robotic half face

DALL-E 2 (2022) 1024x1024

@ Bing Creator

<https://www.bing.com/create>



a dolohin in an astronaut suit on saturn. artstation

2021



Síntesis de imágenes



Midjourney

<https://www.midjourney.com/>



2022



Síntesis de imágenes



Modelos multimodales...

DALL·E

<https://openai.com/dall-e-3>

Midjourney

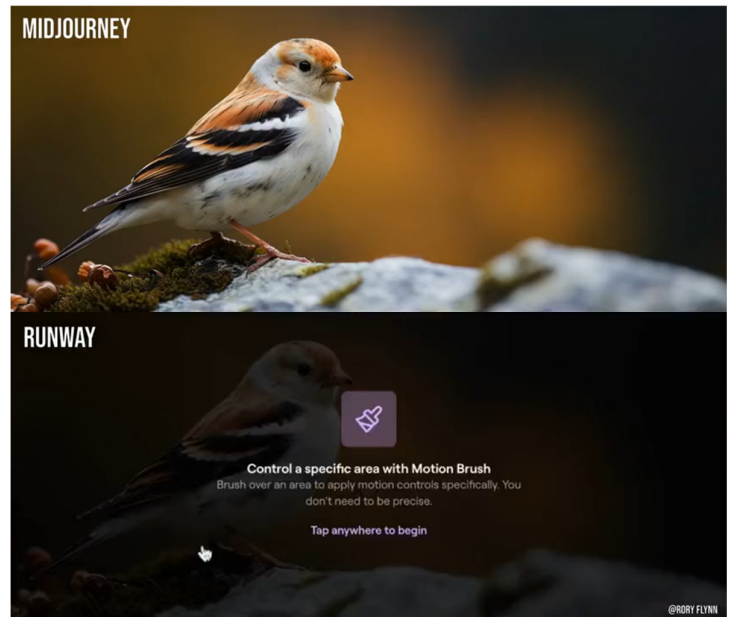
<https://www.midjourney.com/>

Runway

<https://runwayml.com/>

Pika

<https://pika.art/>

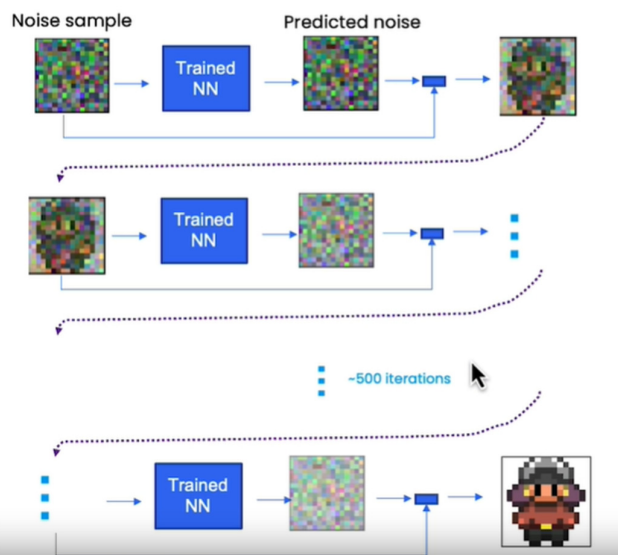


Modelos de difusión

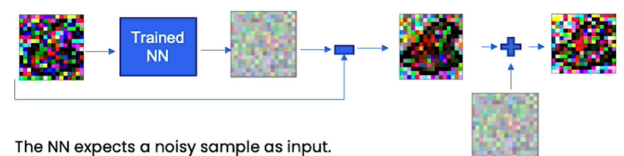


Sampling

NN tries to fully predict the noise at each step. Realistically, it's just a prediction. You need multiple steps to get high quality sprites.



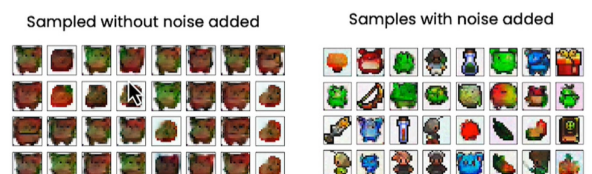
Sampling Iteration Details



The NN expects a noisy sample as input.

You can add in additional noise before it gets passed to the next step.

Empirically, this stabilizes the NN so it doesn't collapse to something closer to the average of the dataset.

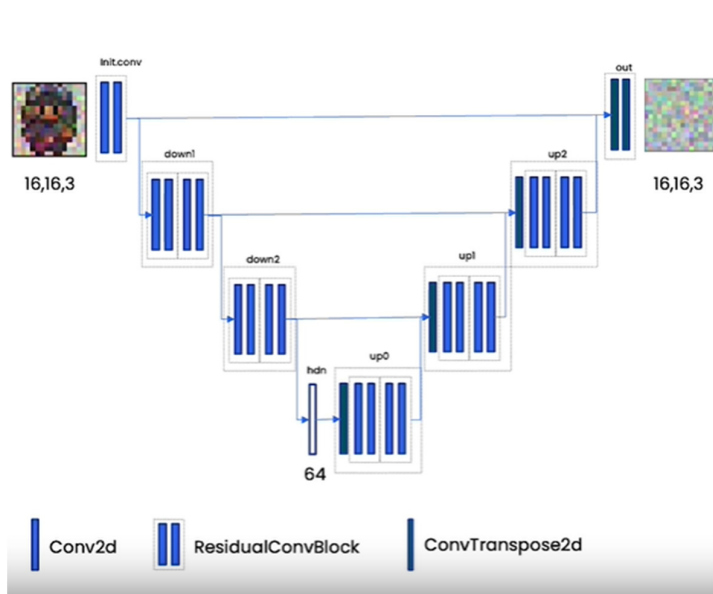


DeepLearning.AI: "How Diffusion Models Work," 2023

<https://www.deeplearning.ai/short-courses/how-diffusion-models-work/>



Modelos de difusión



DeepLearning.AI: "How Diffusion Models Work," 2023

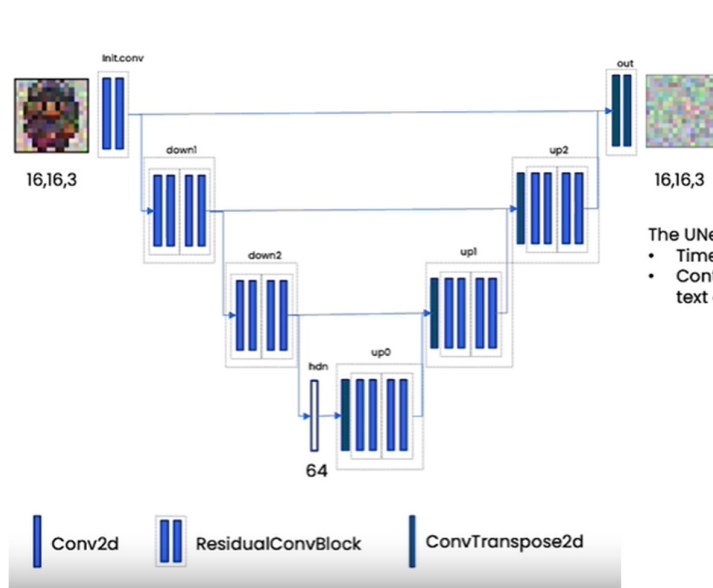
<https://www.deeplearning.ai/short-courses/how-diffusion-models-work/>



Modelos de difusión

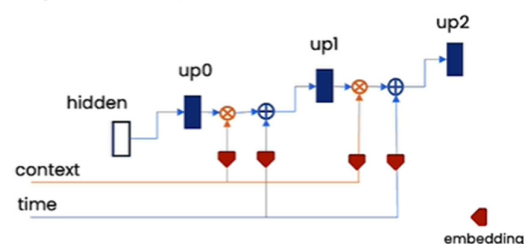


Generación condicionada de imágenes... p.ej. LLM/estilo



The UNet can take in more information in the form of embeddings

- Time embedding: related to the timestep and noise level.
- Context embedding: related to controlling the generation, e.g. text description or factor (more later).



DeepLearning.AI: "How Diffusion Models Work," 2023

<https://www.deeplearning.ai/short-courses/how-diffusion-models-work/>



Modelos de difusión



Estilo: Definición de categorías...

Context

- Context is a vector for controlling generation.
- Context can be text embeddings, e.g. > 1000 in length.
- Context can also be categories, e.g. 5 in length.



DeepLearning.AI: "How Diffusion Models Work," 2023

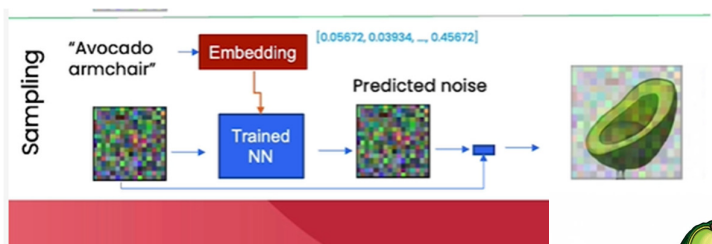
<https://www.deeplearning.ai/short-courses/how-diffusion-models-work/>



Modelos de difusión



LLM: Descripción textual de una escena...



DeepLearning.AI: "How Diffusion Models Work," 2023

<https://www.deeplearning.ai/short-courses/how-diffusion-models-work/>



Modelos de difusión



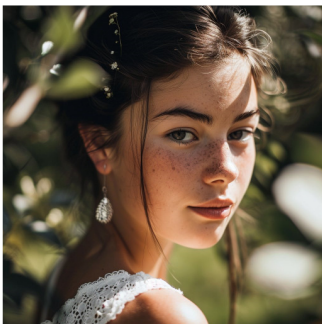
Resultado: DALL·E 3



Modelos de difusión



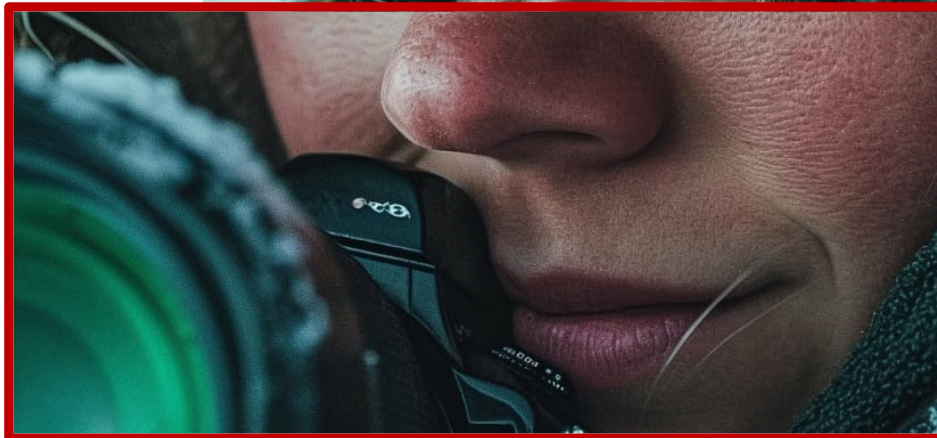
Resultado: Midjourney 6



Modelos de difusión



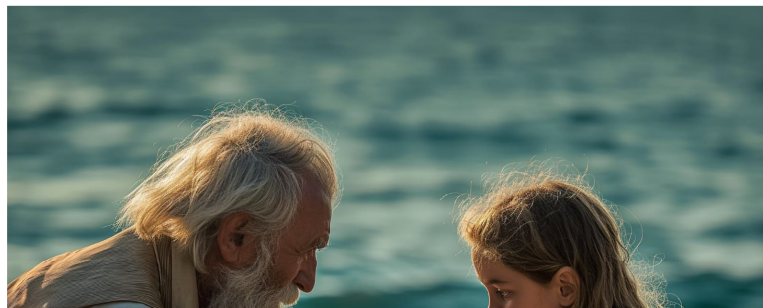
Resultado: Midjourney 6



Modelos de difusión



Resultado: Midjourney 6



¿Real o fake?



¿Real o fake?



Limitaciones



Falta de comprensión (en sentido humano) ...



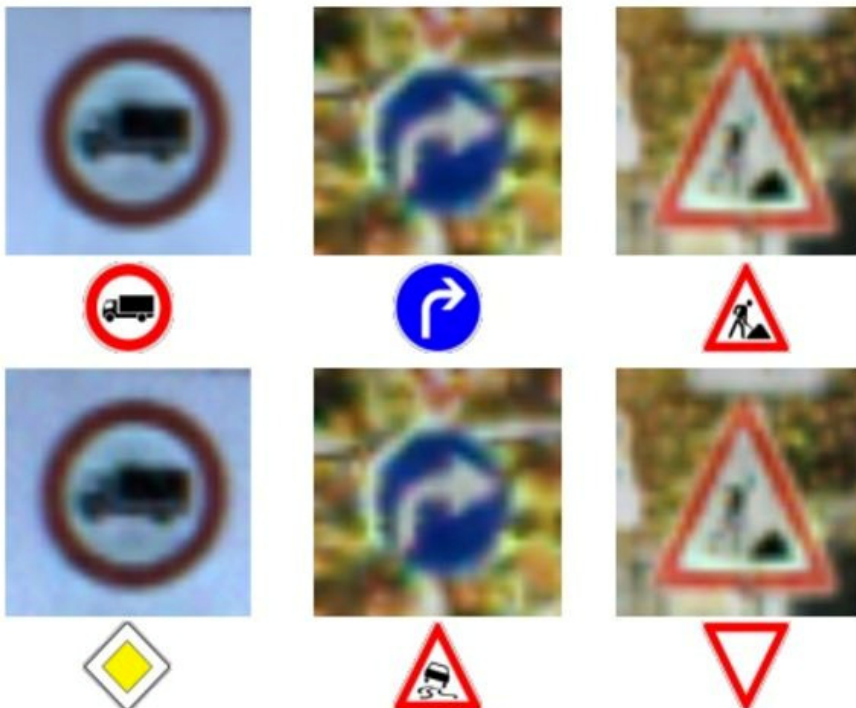
The boy is holding a baseball bat.



Limitaciones



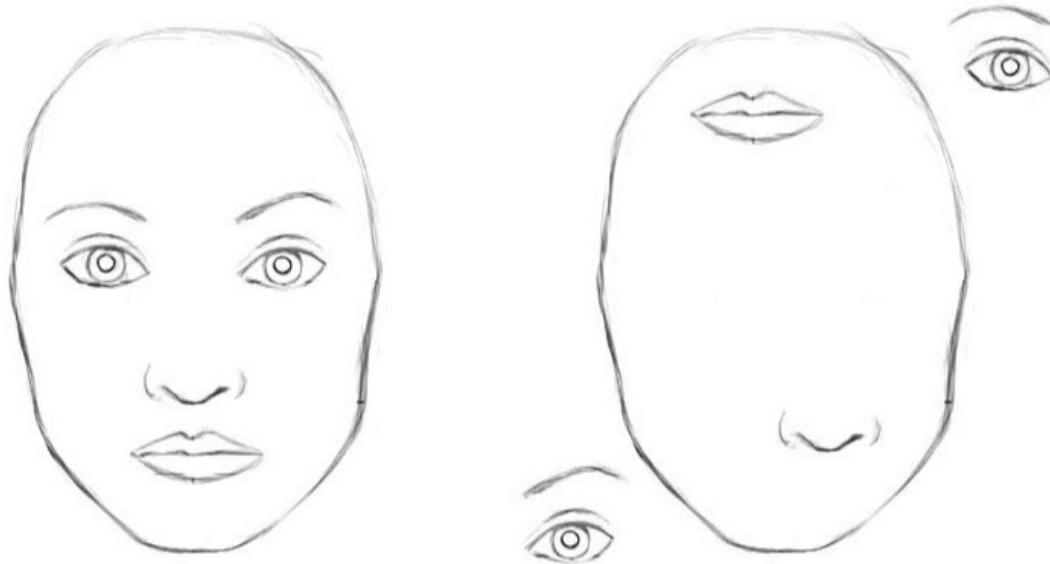
Situaciones con adversario [adversarial examples]



Limitaciones



Las redes convolutivas [CNNs] funcionan muy bien en la práctica, pero...



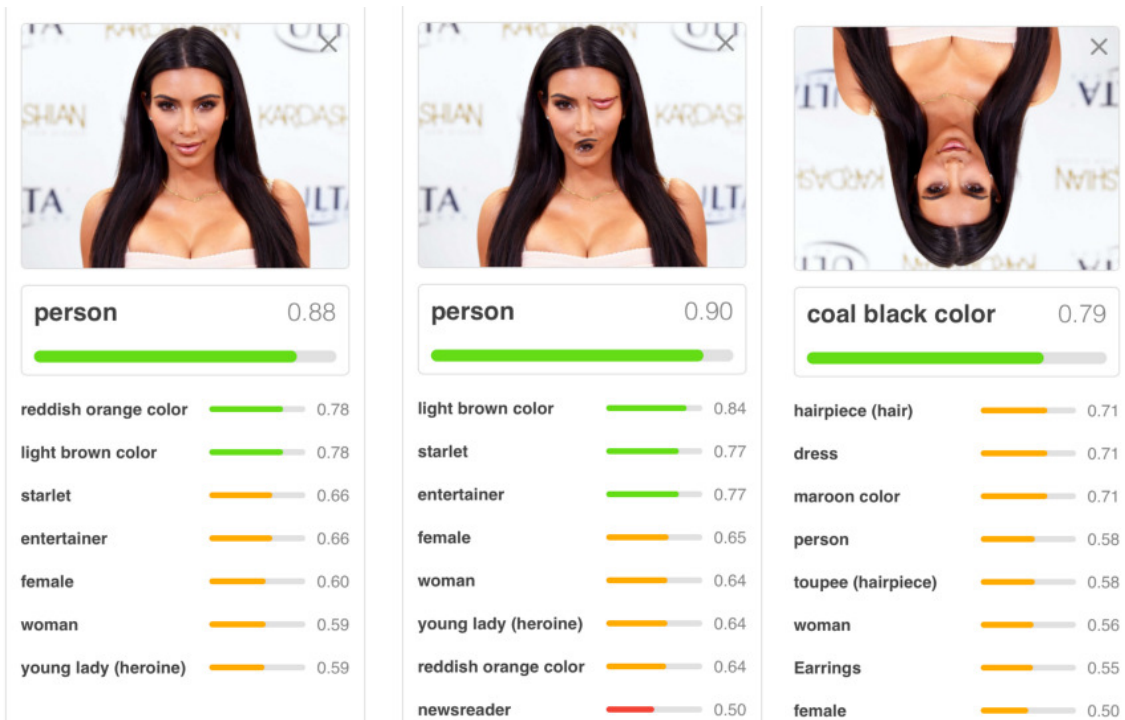
... para una CNN, ambas imágenes son similares 😞



Limitaciones



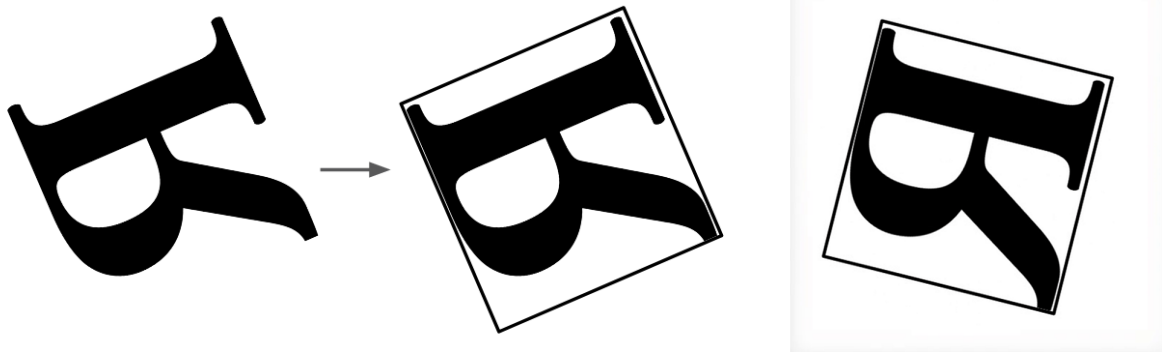
“Convolutional neural networks are doomed”
—Geoffrey Hinton





Problema clave de las redes convolutivas

La representación interna de una red convolutiva no tiene en cuenta las relaciones espaciales entre objetos, ni la jerarquía existente entre objetos simples y los objetos compuestos de los que forman parte.



NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo
of Computer Designed to
Read and Grow Wiser

WASHINGTON, July 7 (UPI)—The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's \$2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of \$100,000.

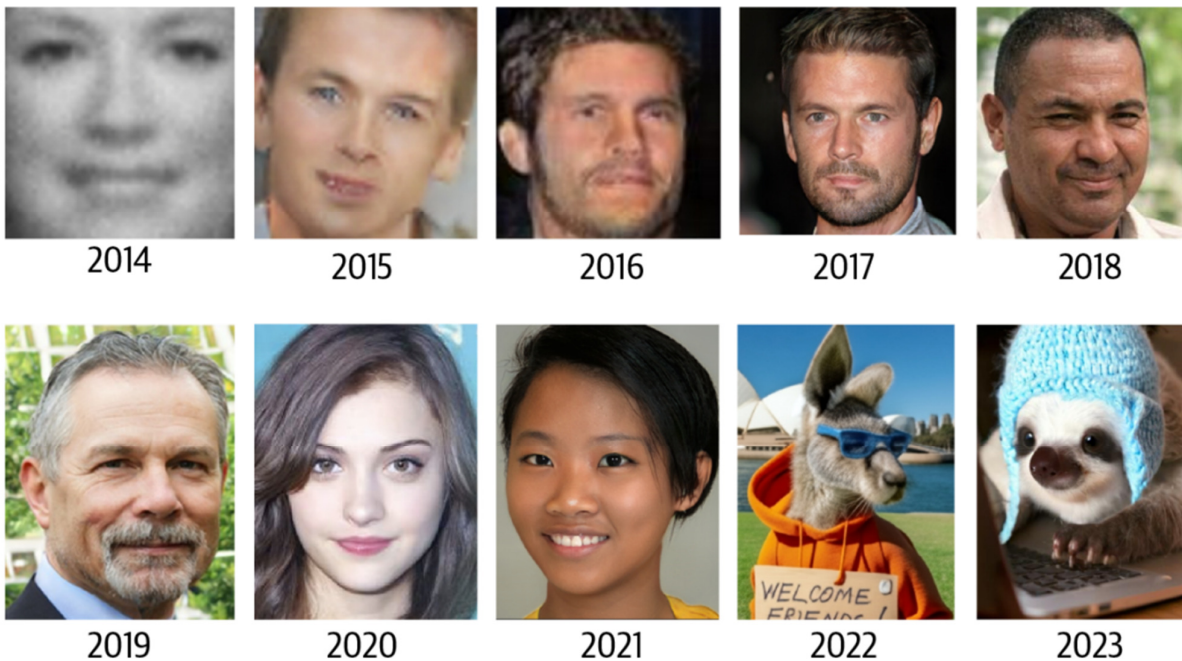
Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

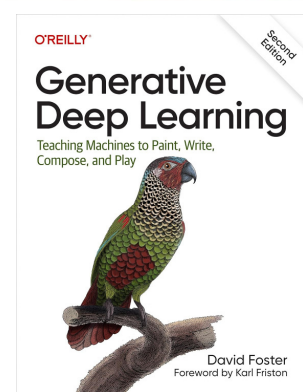
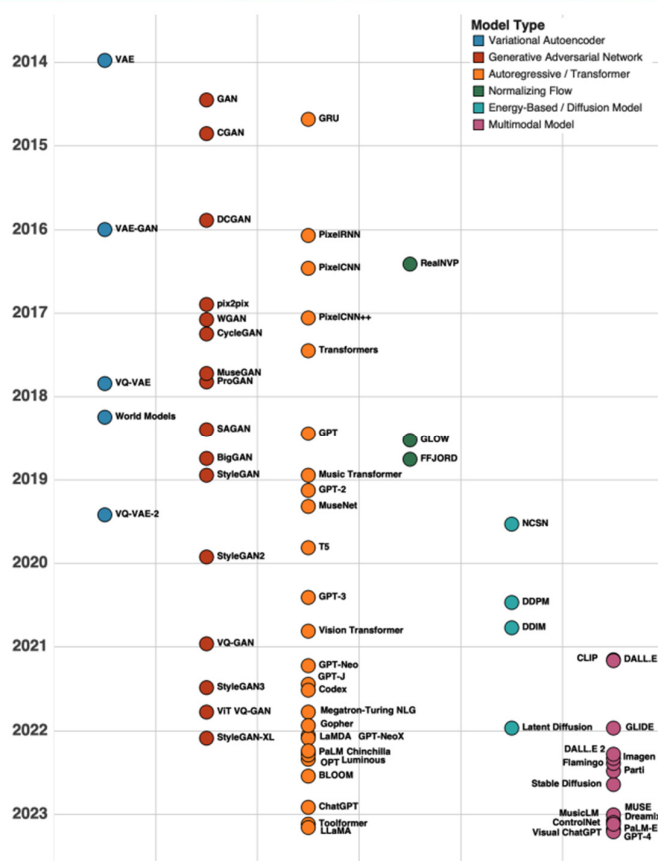
New York Times
1958



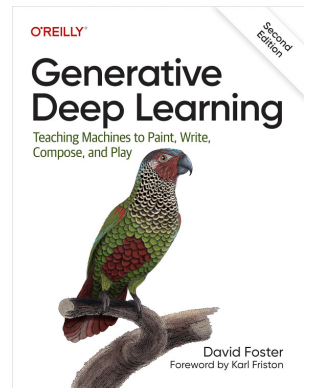
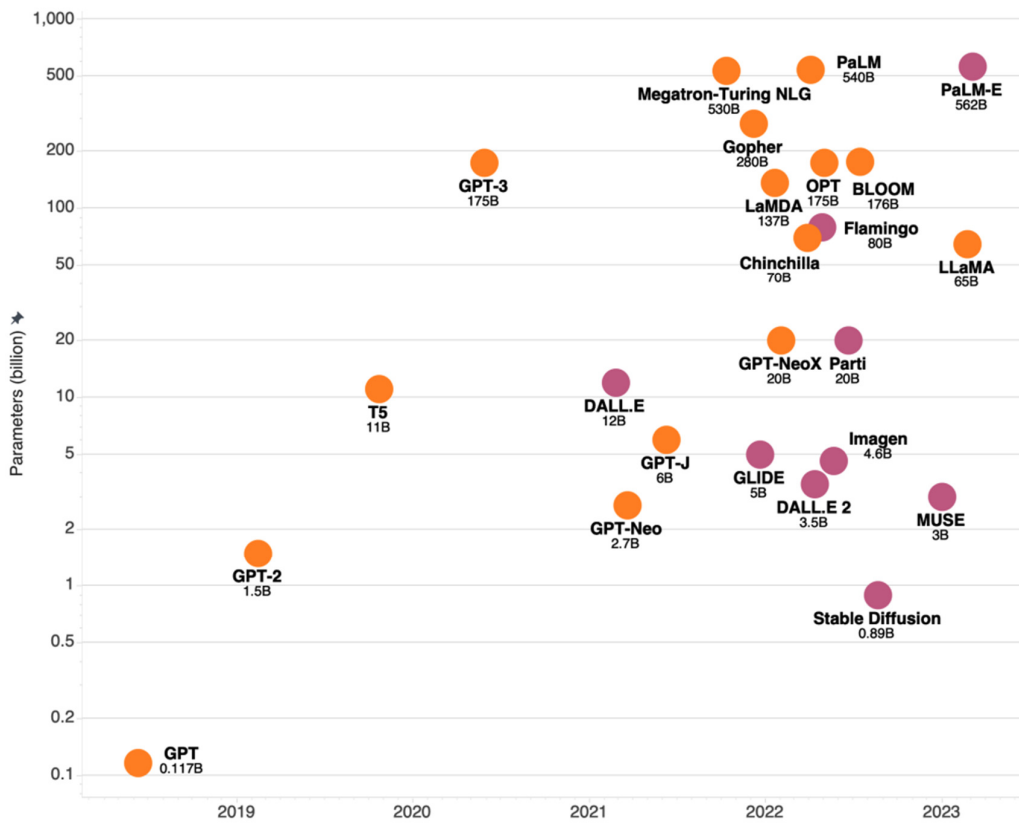
Evolución



Evolución



Evolución



Información adicional



... con ejemplos en Python:

David Foster:

**Generative Deep Learning:
Teaching Machines to
Paint, Write, Compose, and Play**

O'Reilly, 2nd edition, 2023

ISBN 1098134184

O'REILLY

Generative Deep Learning

Teaching Machines to Paint, Write, Compose, and Play



David Foster
Foreword by Karl Friston





Fernando Berzal:
**Redes Neuronales
& Deep Learning**

Edición en dos volúmenes, 2019

Volumen I: Entrenamiento
de redes neuronales artificiales
ISBN 1-0903-2030-2

Volumen II: Regularización,
optimización y arquitecturas especializadas
ISBN 1-0903-3688-8

<https://deep-learning.ikor.org>

