Clase 1

Introducción al aprendizaje profundo

Enzo Ferrante

eferrante@sinc.unl.edu.ar

@enzoferrante

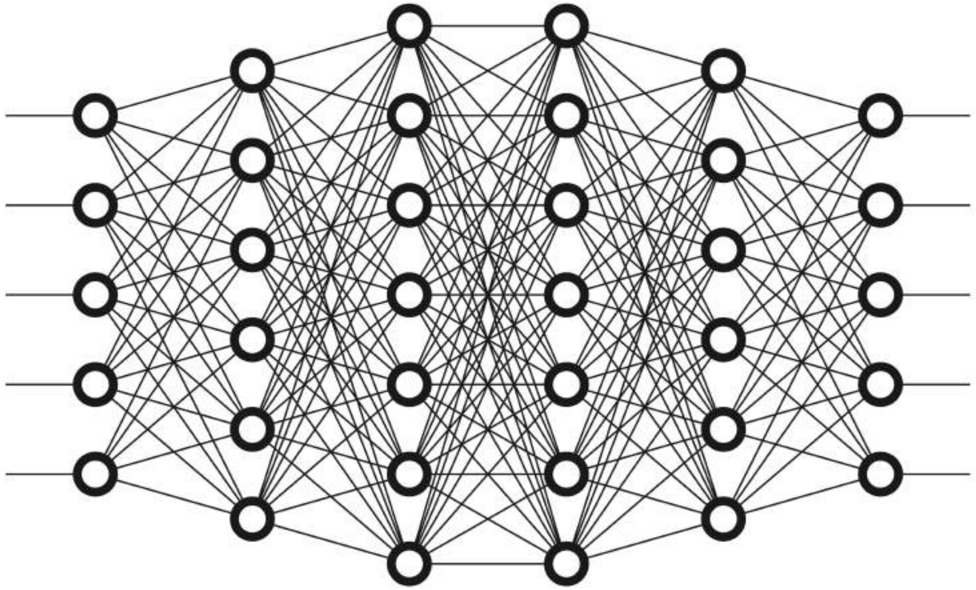

UNIVERSIDAD TORCUATO DI TELLA

Imperial College London

UNIVERSITY OF CAMBRIDGE

Stanford University

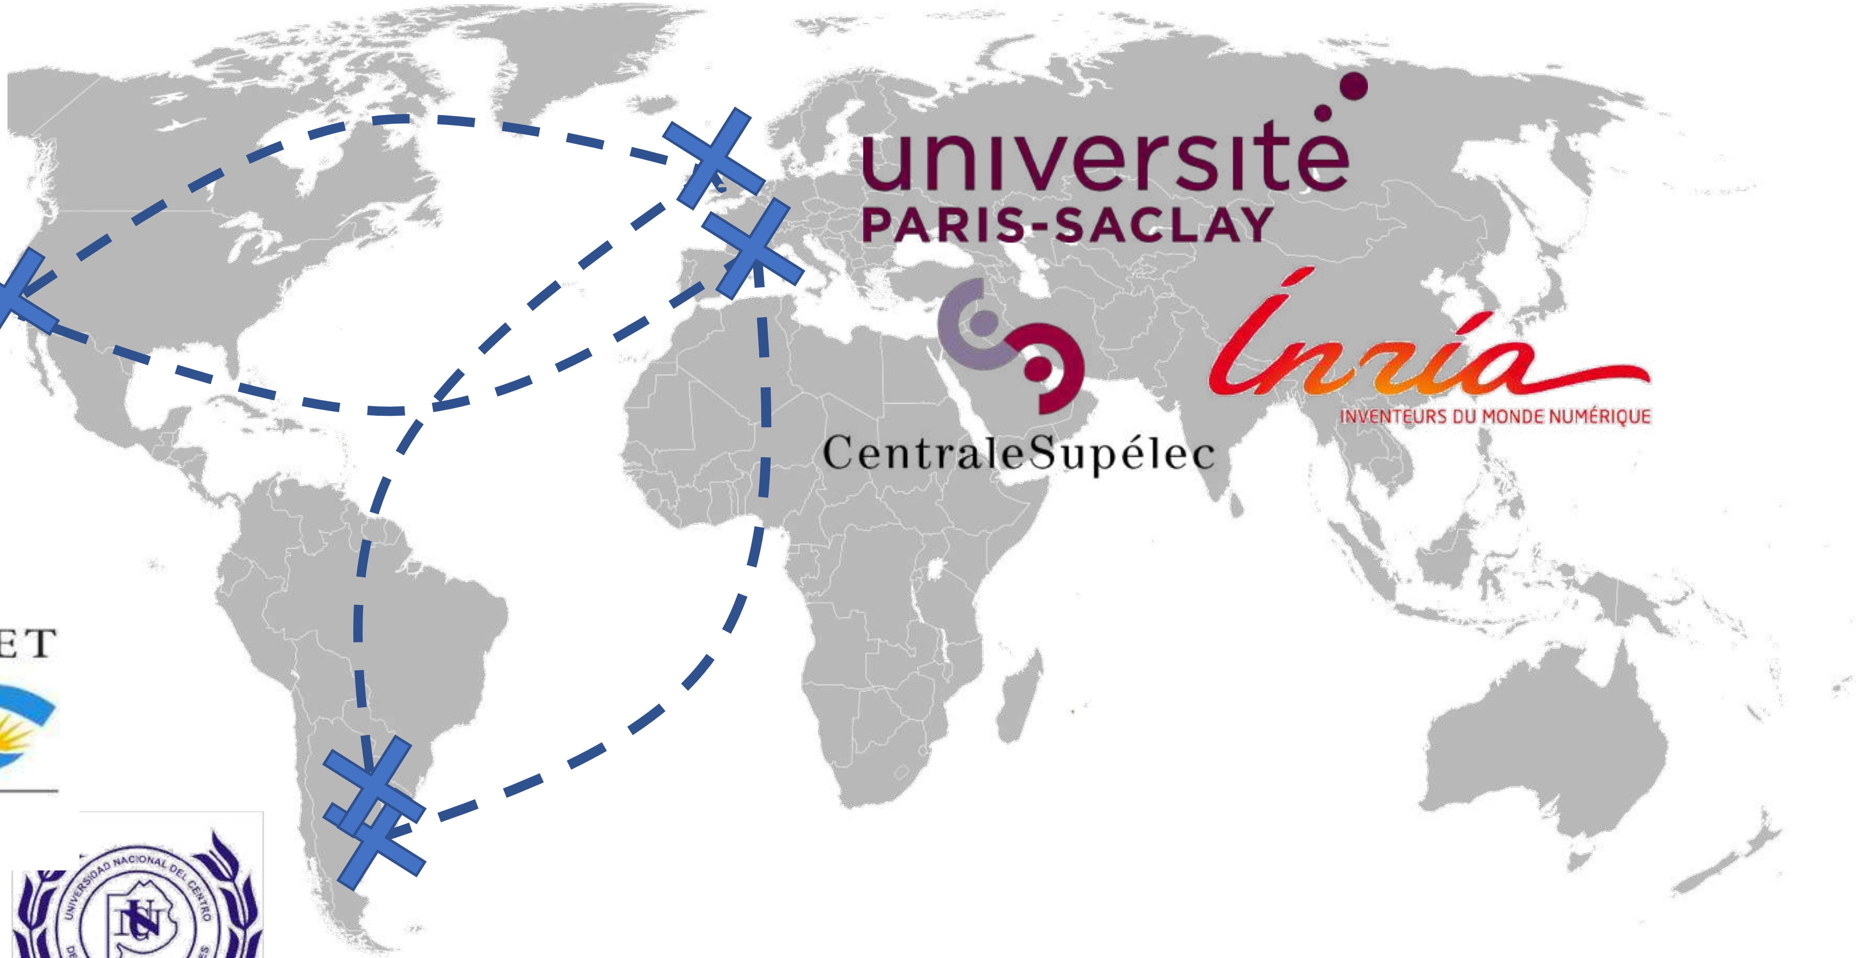

UNL

CONICET

sinc(i)

UNIVERSIDAD NACIONAL DEL CENTRO DE LA PROVINCIA DE BUENOS AIRES

sinc(i): Research institute for signals, systems and computational intelligence

UNL

CONICET

sinc(i)

Machine Learning for Biomedical Image Analysis

Bioing. Candelaria Mosquera

Ayudante en la materia

- **Bioingeniera**
Instituto Tecnológico de Buenos Aires (ITBA)

- **Estudiante del Doctorado en Ingeniería**
Universidad Tecnológica Nacional (UTN)

- **AI Engineer & Python developer**
Programa de Inteligencia Artificial en Salud del Hospital Italiano de Buenos Aires (PIASHIBA)

Video

https://www.youtube.com/watch?v=Dy0hJWltsyE&t=11s&ab_channel=NVIDIA

Aclarando conceptos:
IA, Aprendizaje Automático y Aprendizaje Profundo

Deep Learning

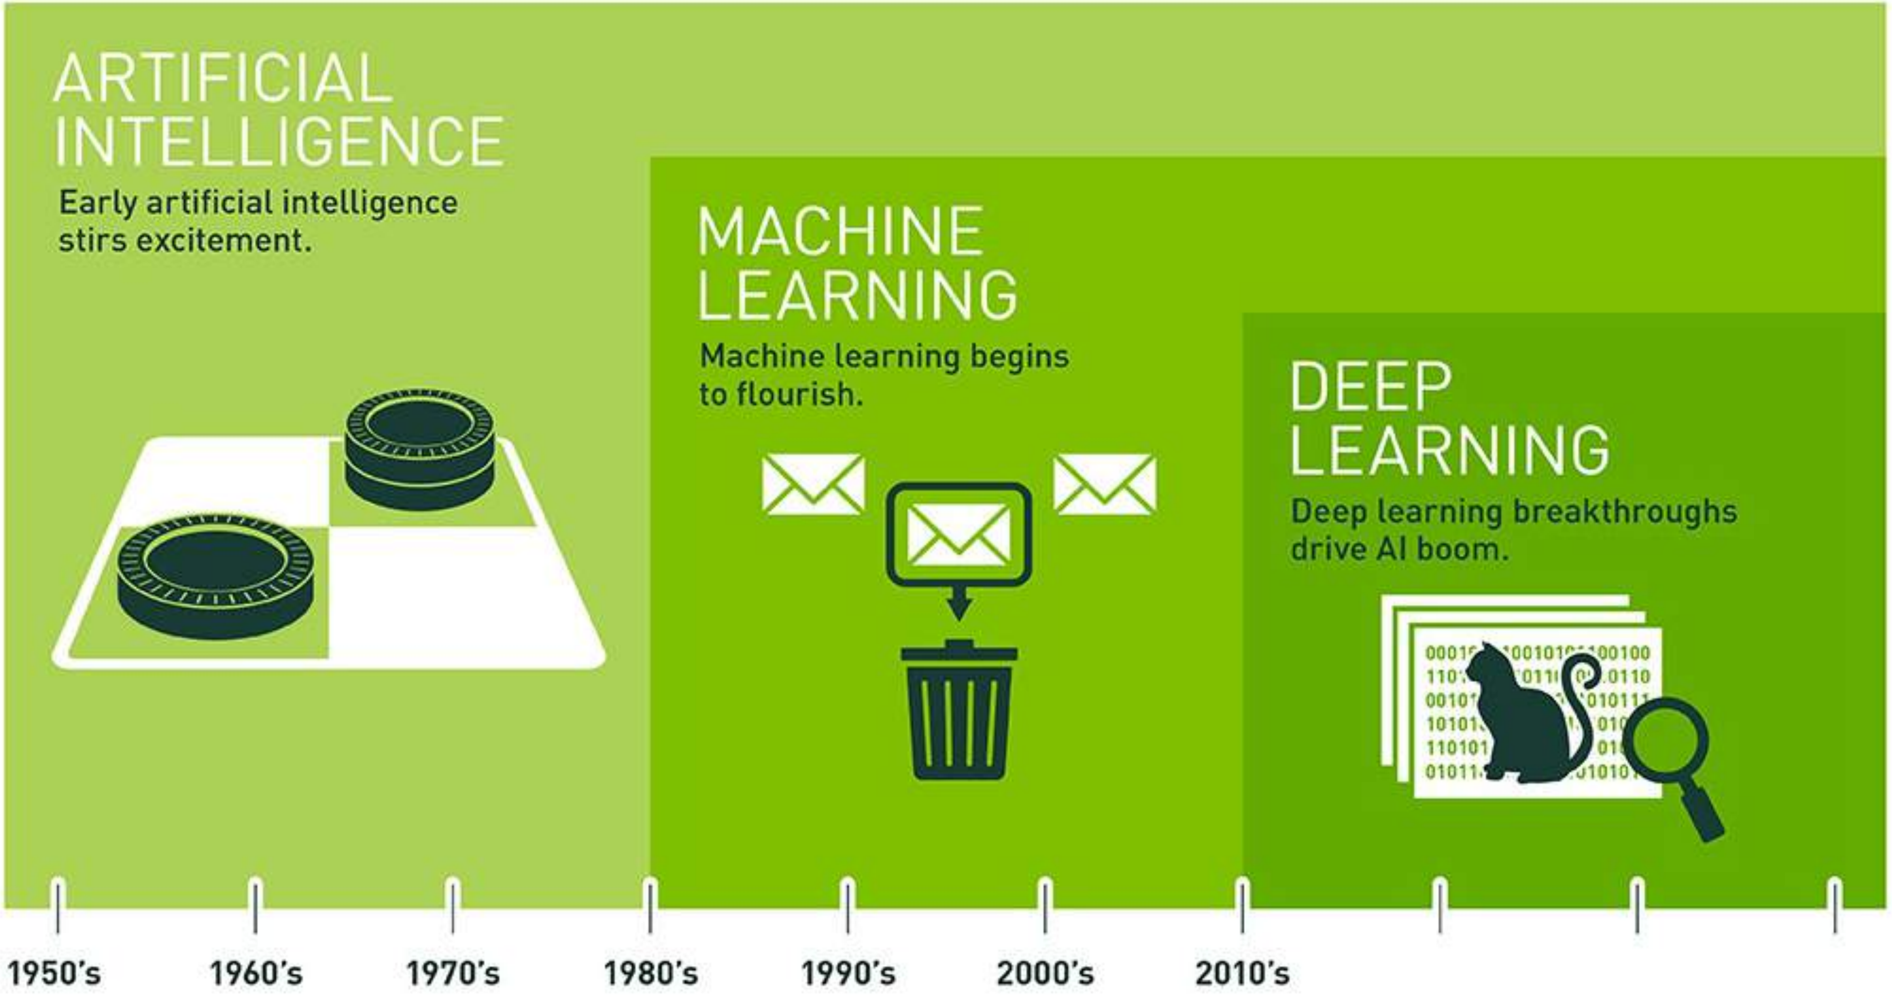

Imagen extraida de: https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

Inteligencia Artificial

Sus orígenes se consideran en la Conferencia de Inteligencia Artificial de Dartmouth en 1955

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories
August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate

Inteligencia Artificial

Comportamiento humano realizado por máquinas

inteligencia artificial.
1. f. Inform. Disciplina científica que se ocupa de crear programas informáticos que ejecutan operaciones comparables a las que realiza la mente humana, como el aprendizaje o el razonamiento lógico.

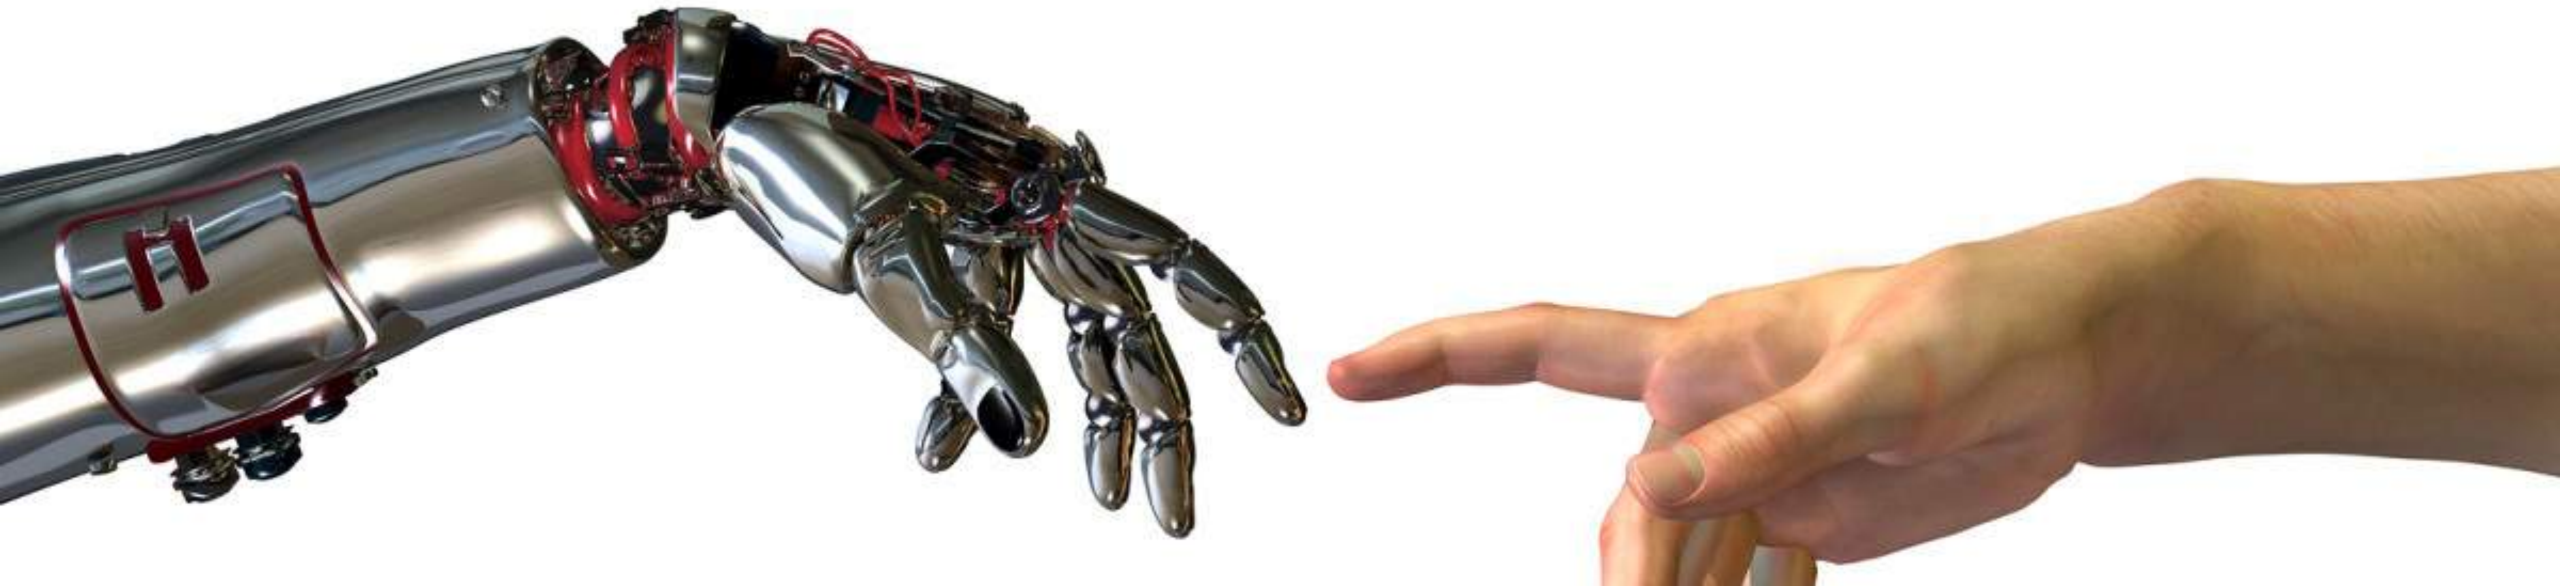

Aprendizaje Automático

Una estrategia para alcanzar la inteligencia artificial

Aprendizaje automático.

1. El estudio y construcción de algoritmos que pueden aprender de y hacer predicciones sobre datos.

Término acuñado en 1959 por Arthur Samuel (IBM)

Aprendizaje Automático

En la práctica, el aprendizaje maquinal puede resolver problemas de inteligencia artificial "estrechos" (Narrow AI).

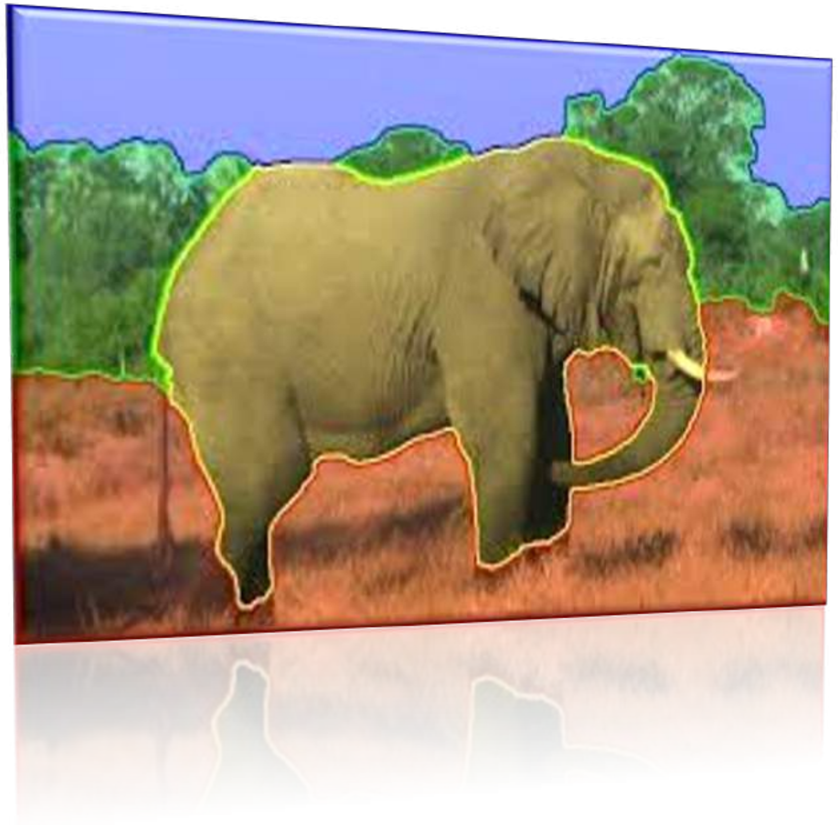

Segmentación de imágenes

Aprendizaje Automático

En la práctica, el aprendizaje automático puede resolver problemas de inteligencia artificial "estrechos" (Narrow AI).

Perro

Clasificación de imágenes

Aprendizaje Automático

En la práctica, el aprendizaje maquinal puede resolver problemas de inteligencia artificial "estrechos" (Narrow AI).

Chihuaua vs Muffin

Clasificación de imágenes

Aprendizaje Automático

En la práctica, el aprendizaje maquinal puede resolver problemas de inteligencia artificial "estrechos" (Narrow AI).

Reconocimiento de voz

Aprendizaje Automático

En la práctica, el aprendizaje maquinal puede resolver problemas de inteligencia artificial "estrechos" (Narrow AI).

- *direccion*: La dirección de la propiedad
- *ciudad*: La ciudad de la propiedad
- *provincia*: La provincia donde está localizada la propiedad
- *lat*: Latitud
- *lng*: Longitud
- *tipodepropiedad*: El tipo de propiedad (Casa, departamento, etc)
- *metrostotales*: Metros totales de la propiedad
- *metroscubiertos*: Metros cubiertos de la propiedad
- *antiguedad*: Antigüedad de la propiedad
- *habitaciones*: Cantidad de habitaciones
- *garages*: Cantidad de garages
- *banos*: Cantidad de baños
- *fecha*: Fecha de publicación
- *gimnasio*: Si el edificio o la propiedad tiene un gimnasio
- *usosmultiples*: Si el edificio o la propiedad tiene un SUM
- *piscina*: Si el edificio o la propiedad tiene un piscina
- *escuelascercanas*: Si la propiedad tiene escuelas cerca
- *centroscomercialescercanos*: Si la propiedad tiene centros comerciales cerca

→ Valor del inmueble

Aprendizaje Maquinal

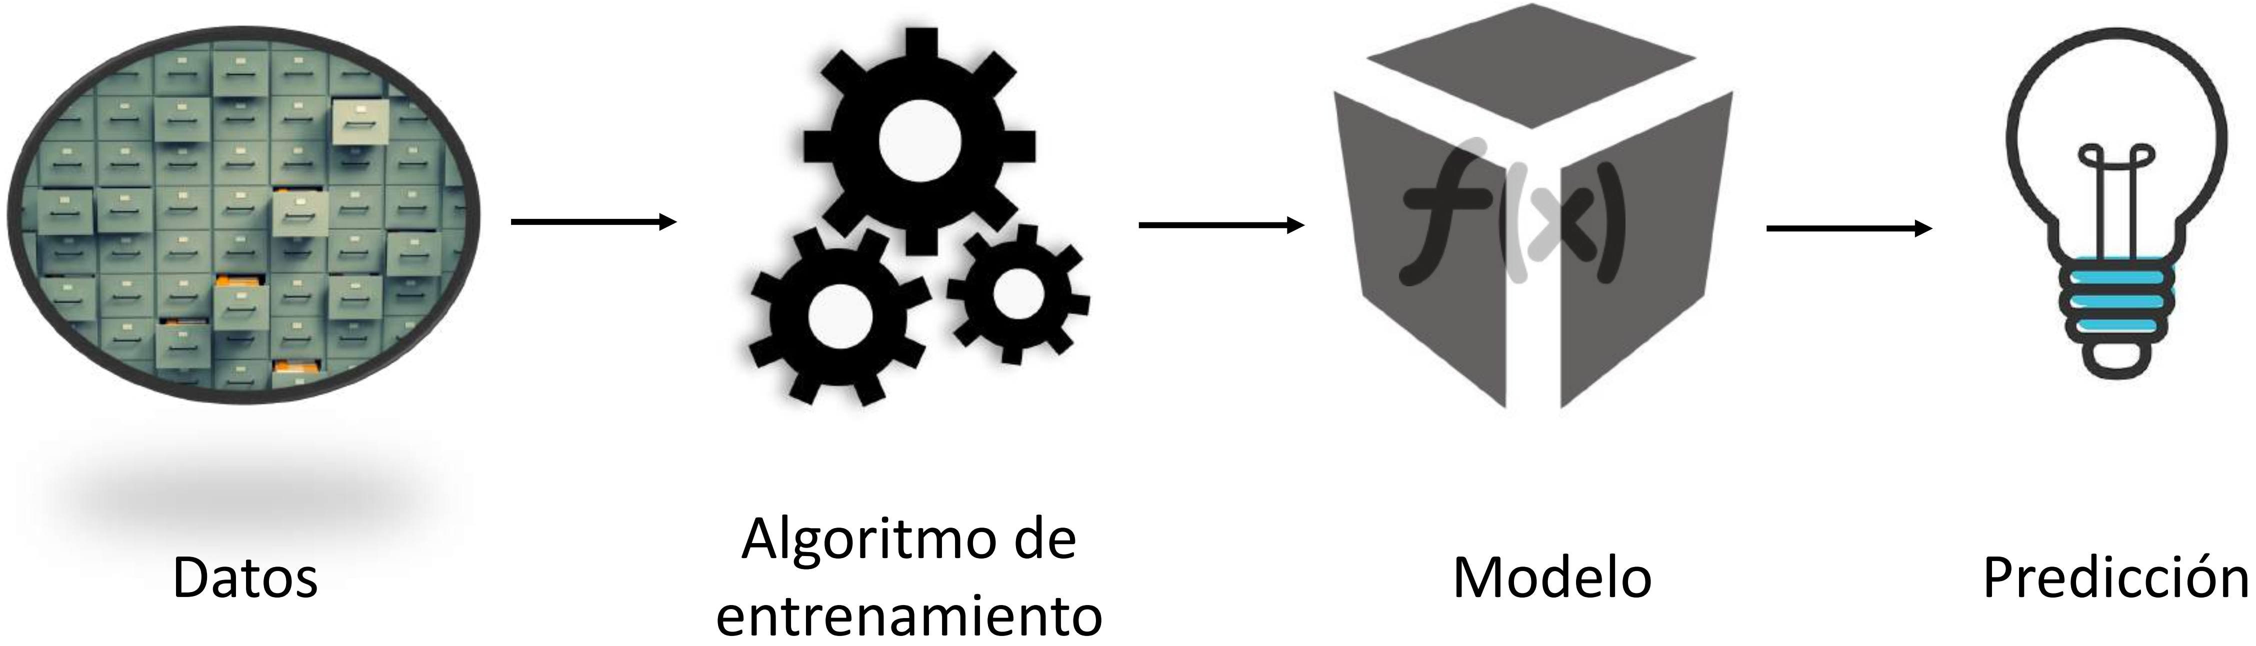

Aprendizaje supervisado

Datos de entrenamiento

Dato	Etiqueta
	Perro
	Perro
	Perro
	Perro

Dato	Etiqueta
	Gato
	Gato
	Gato
	Gato

→ Modelo entrenado

→ Modelo entrenado → Perro

Datos de test

Predicción

Aprendizaje no-supervisado

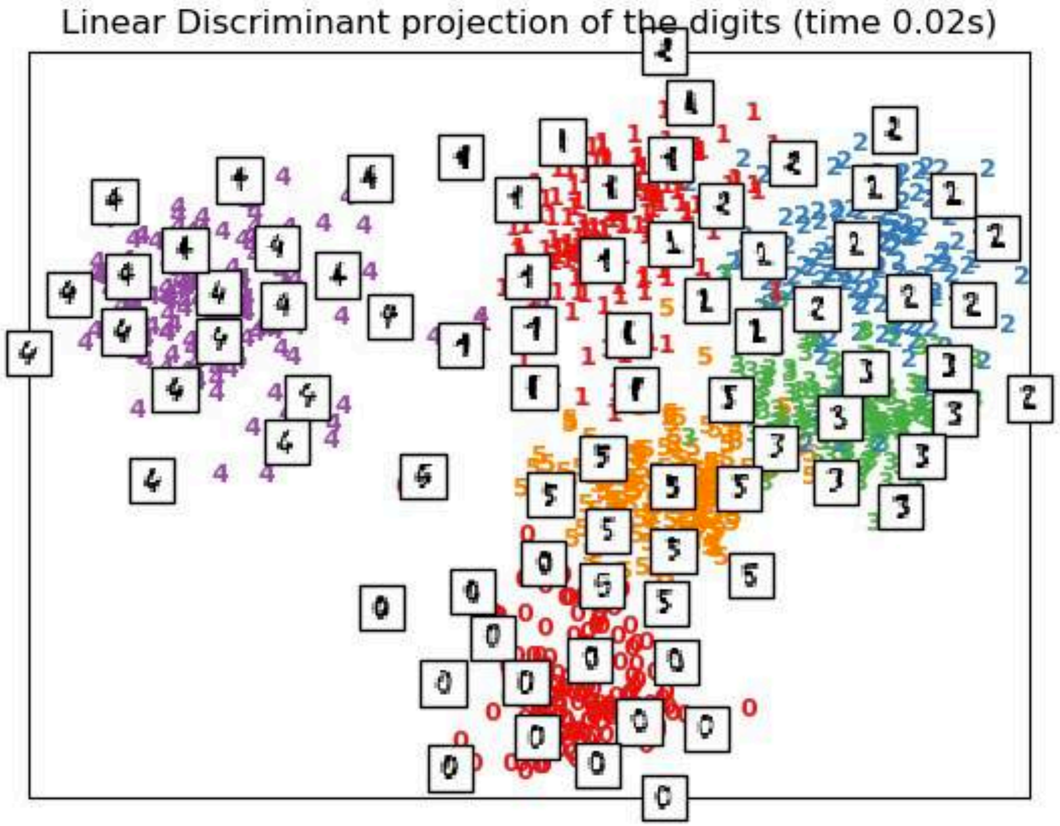

Fuente: http://scikit-learn.org/stable/modules/manifold.html

Aprendizaje semi-supervisado

Los alogrítmos de aprendizaje semi supervisado utilizan tanto datos etiquetados como no etiquetados durante el proceso de aprendizaje

Aprendizaje por refuerzo

Aprendizaje a partir de la interacción con el ambiente

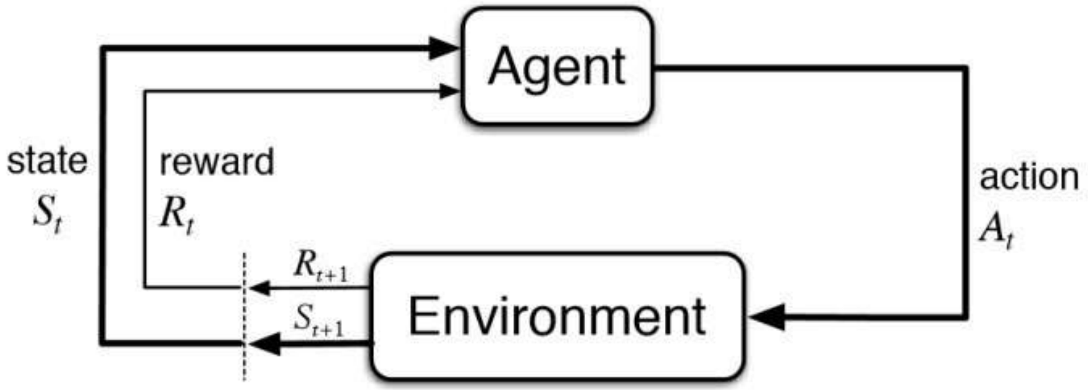

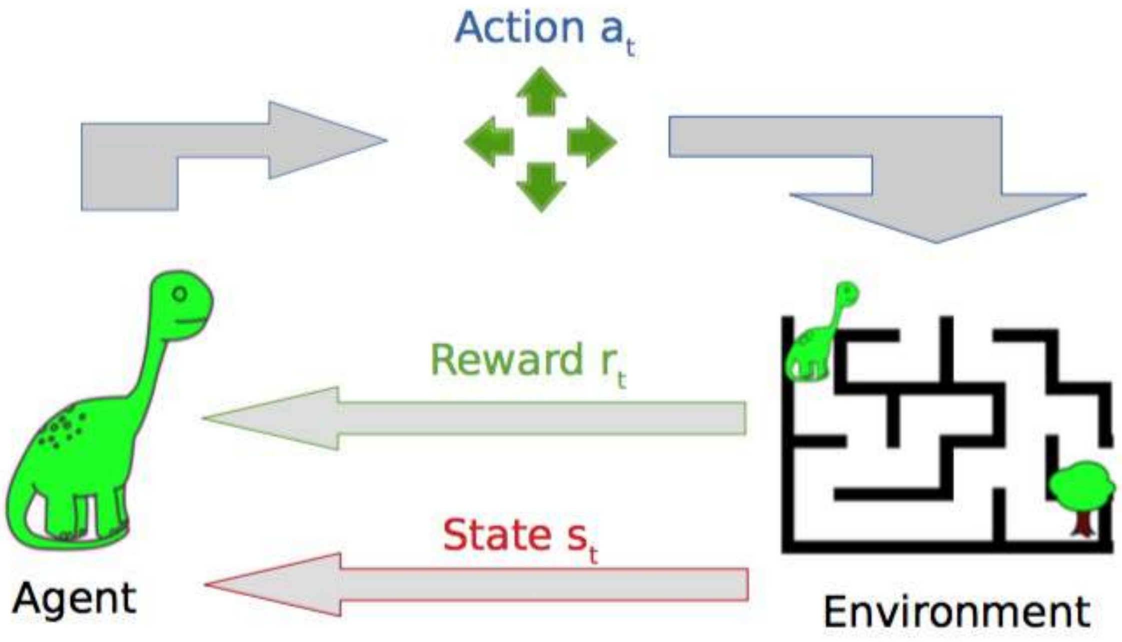

https://twimlai.com/whats-hot-in-ai-deep-reinforcement-learning-edition/

Aprendizaje automático

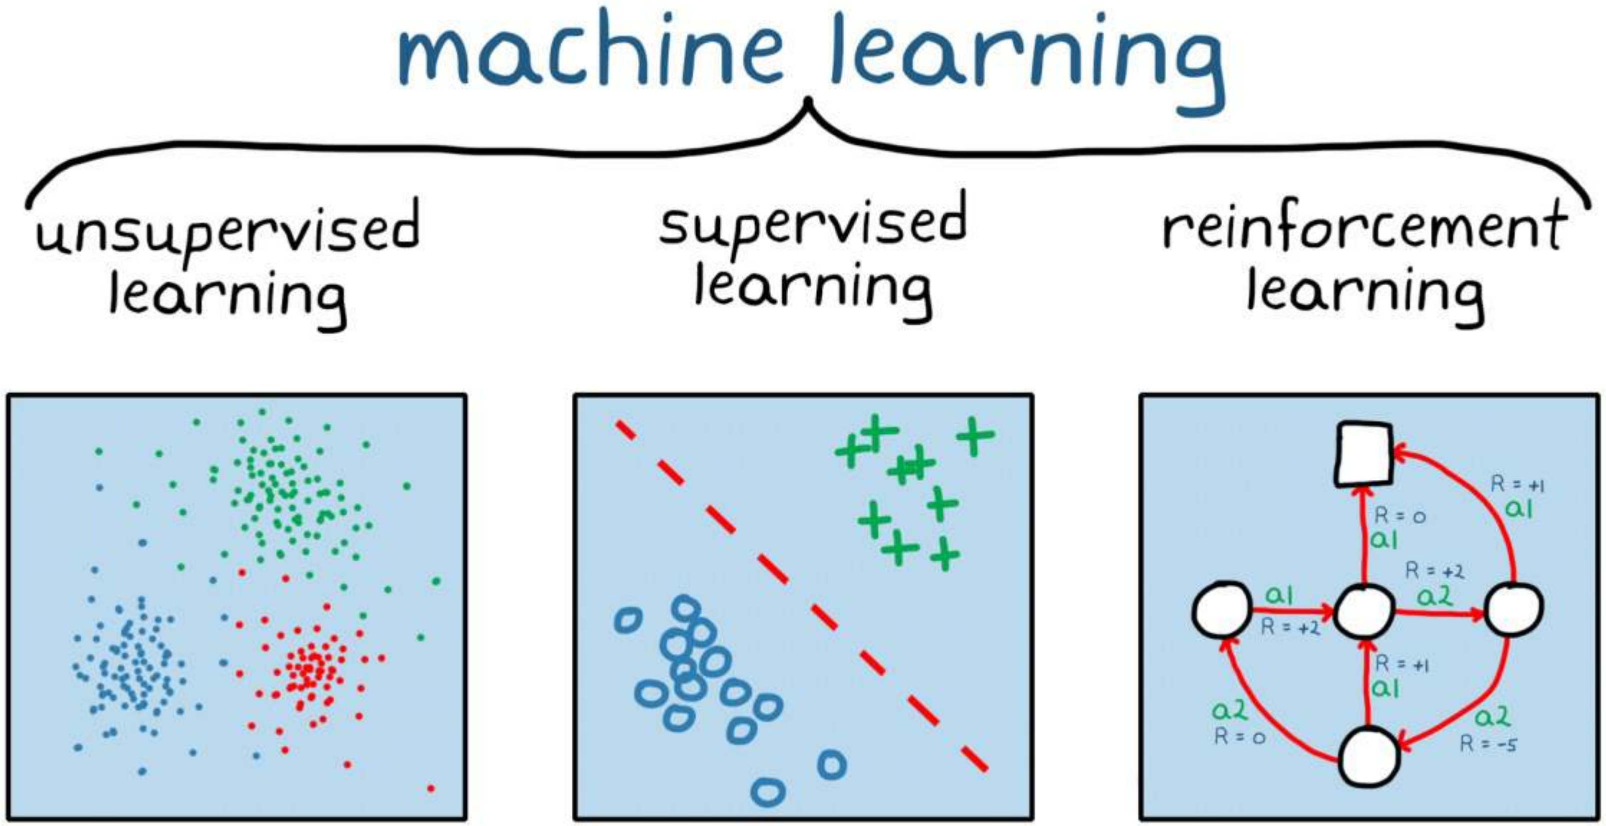

Fuente: https://es.mathworks.com/discovery/reinforcement-learning.html

Ejemplo: clasificación de imágenes

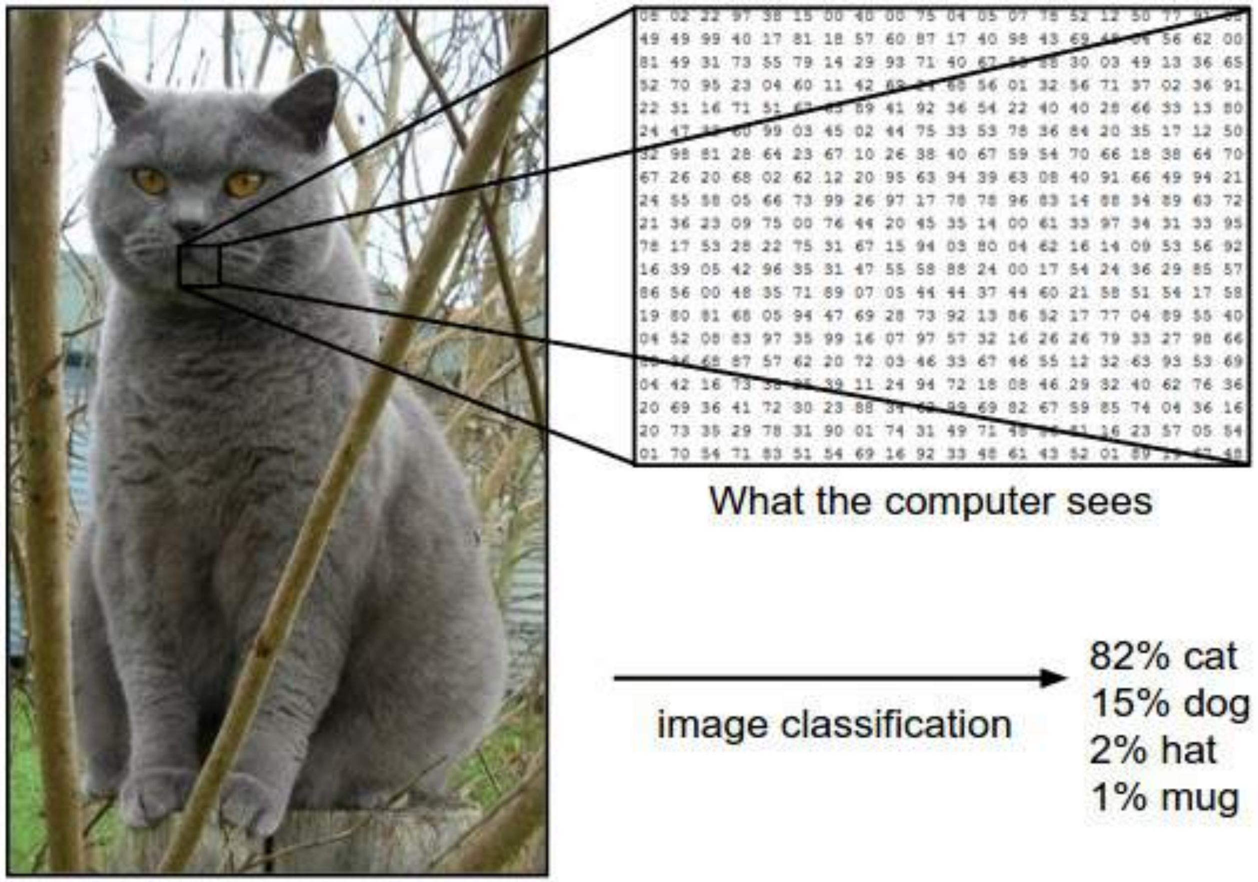

Fuente: http://cs231n.github.io/linear-classify/

Clasificación de imágenes

23	11	12	200	34	35	45	46	50	25
1	89	1	33	78	60	55	1	76	99
89	0	56	45	90	91	88	3	80	87
67	56	77	90	23	1	99	34	90	68
44	99	80	65	40	100	8	7	8	49
55	32	23	55	71	0	19	200	33	45
45	34	21	44	207	65	18	33	12	77
88	45	22	33	45	66	78	89	0	77
9	10	13	57	89	88	90	200	208	100
11	23	12	7	209	56	78	45	88	78

Tiene 4 patas

Tiene cola

Tiene dos orejas largas

Tiene ocico

Es bajito

Features

Perro salchicha!

Enfoque tradicional para la clasificación de imágenes
en aprendizaje maquinal supervisado

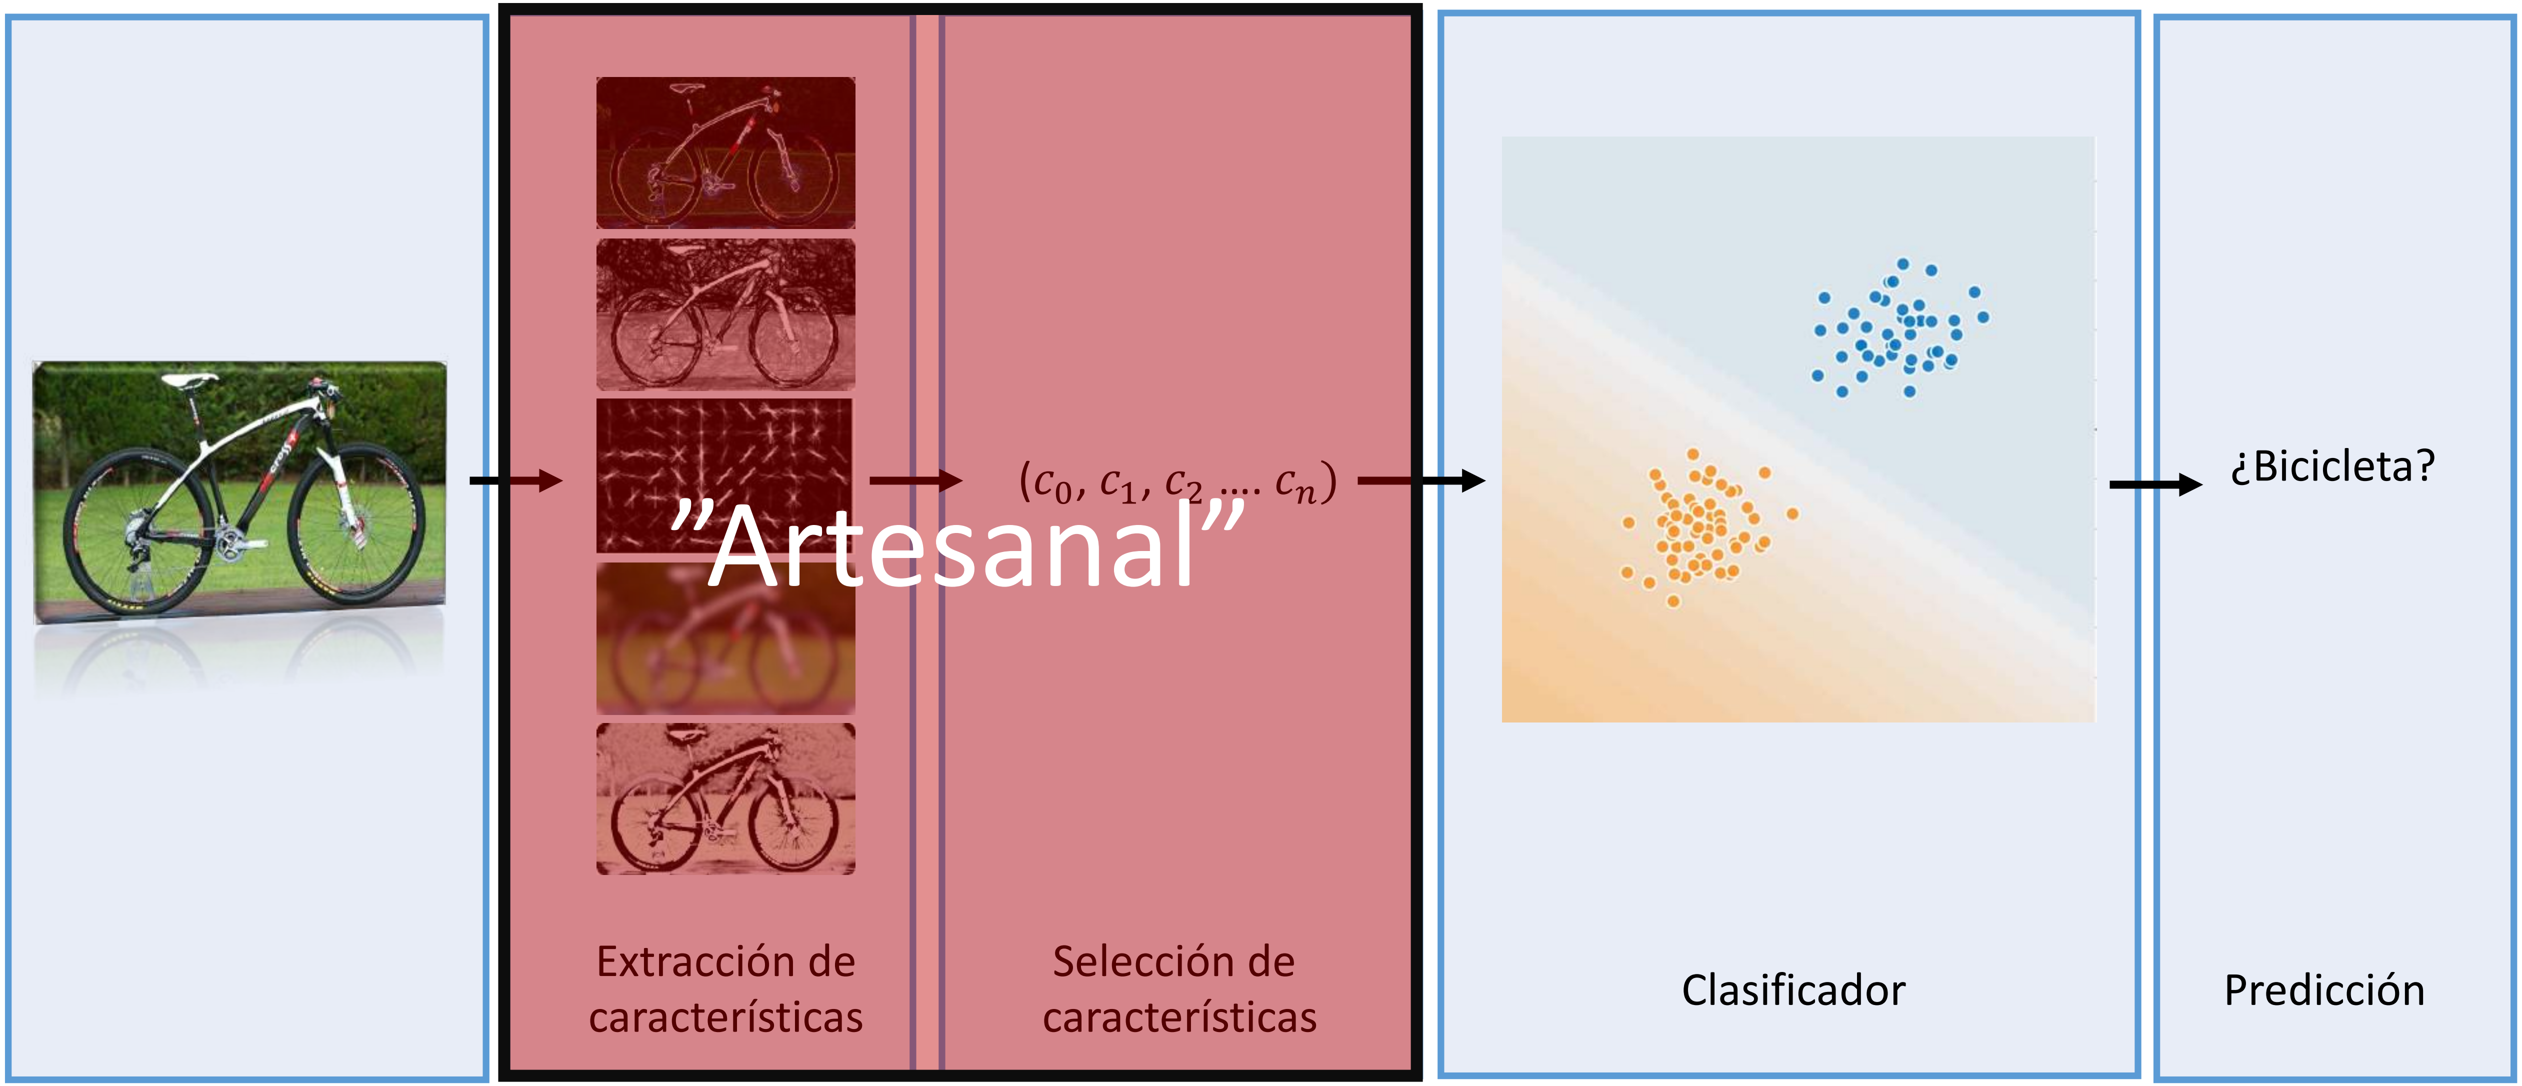

Espacio de características

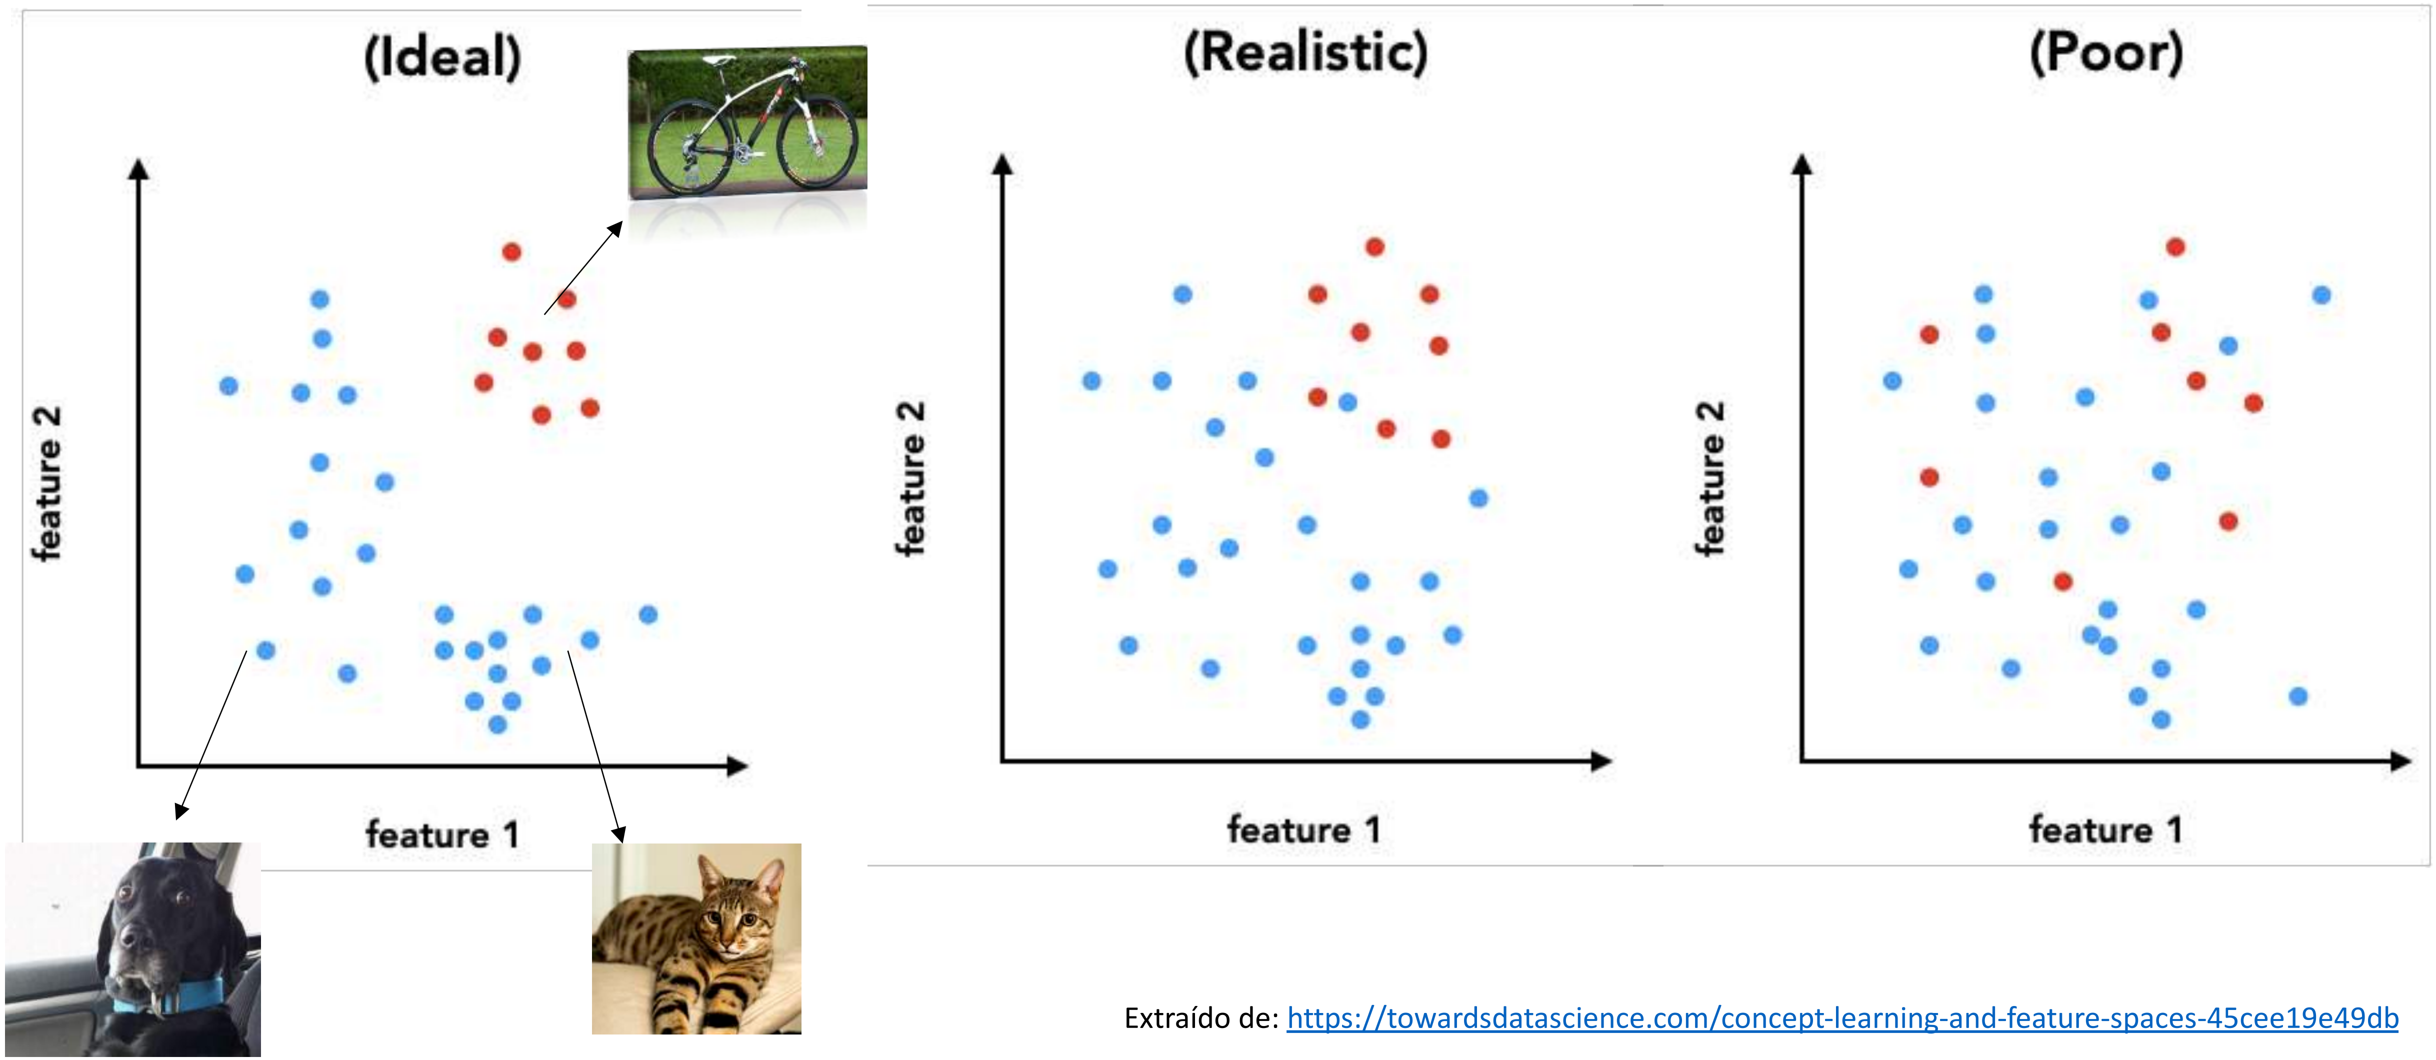

Extraído de: https://towardsdatascience.com/concept-learning-and-feature-spaces-45cee19e49db

Espacio de características

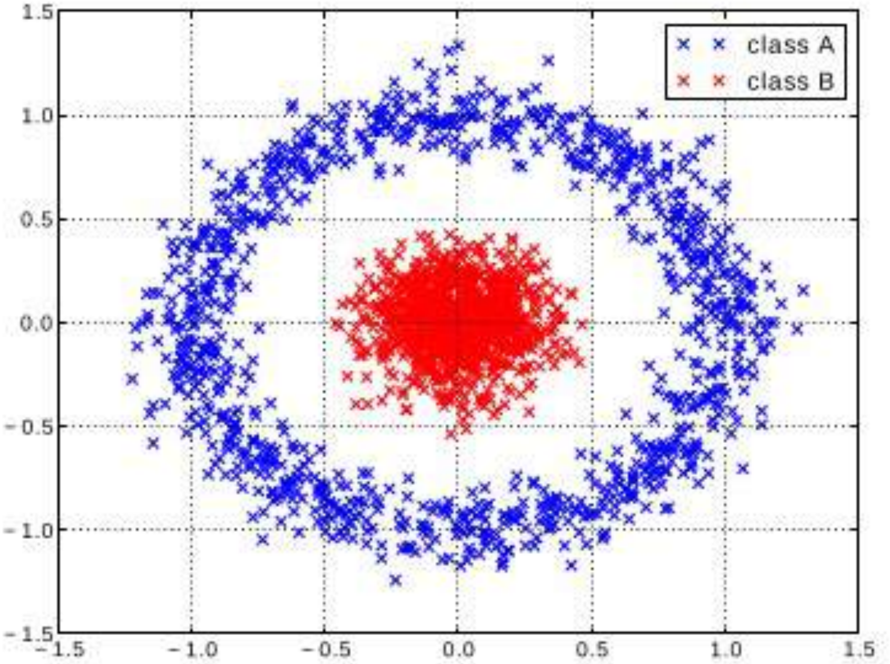

(a) A non linearly separable dataset.

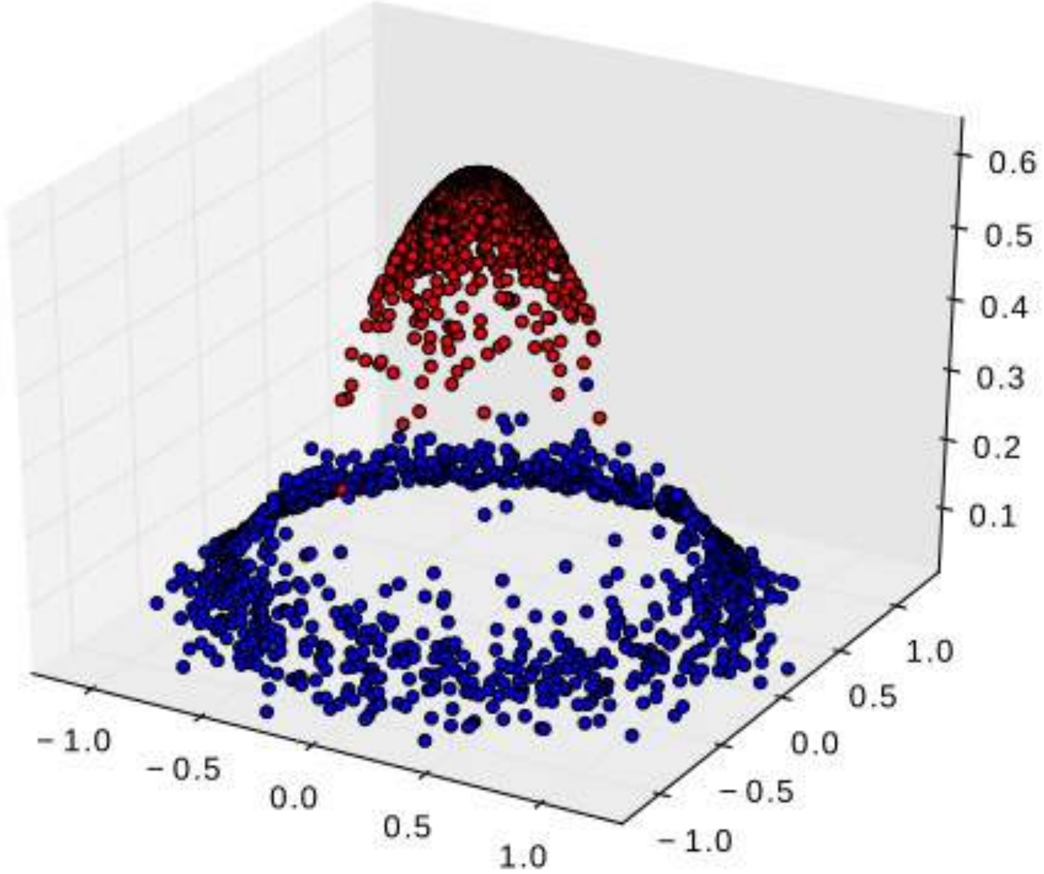

(b) Possible feature space representation.

Extraído de: http://ludovicarnold.altervista.org/teaching/optimization-machine-learning/changing-representations/

Deep Learning

Una técnica para implementar aprendizaje maquinal

Los modelos basados en deep learning son capaces de aprender **representaciones** de los datos de entrenamiento en **múltiples niveles de abstracción** (capas), componiendo módulos simples que sucesivamente transforman dichas representaciones en otras con mayor nivel de abstracción.

Deep Learning

Multiples niveles de abstracción

Único modelo entrenado ”end to end”

Bordes → Formas → Partes → Clasificador →

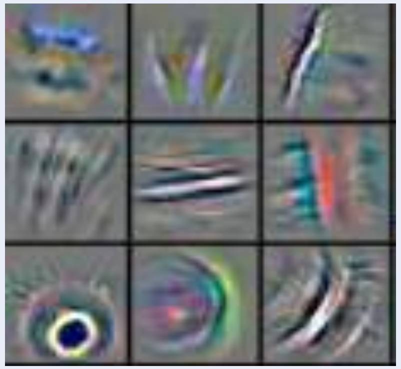

¿Bicicleta?

?

Redes Neuronales Artificiales

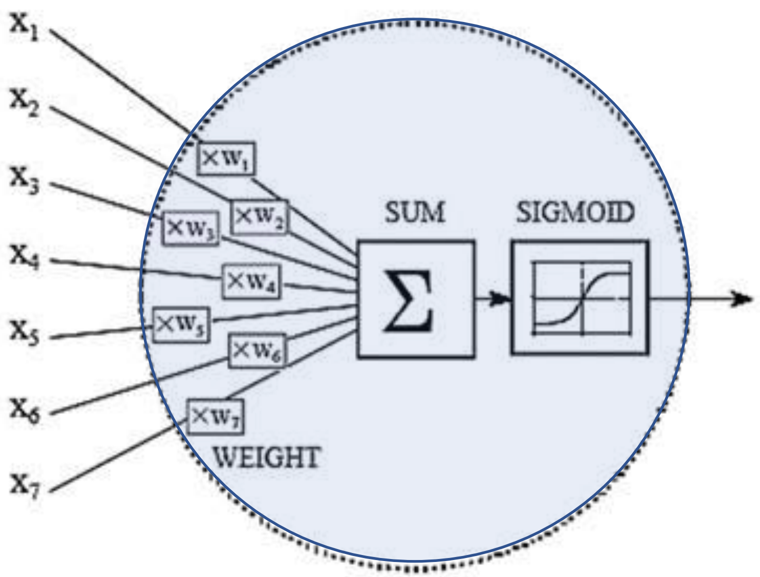

$$s\left(\sum x_i w_i\right)$$

Redes Neuronales Artificiales

Perceptrón simple

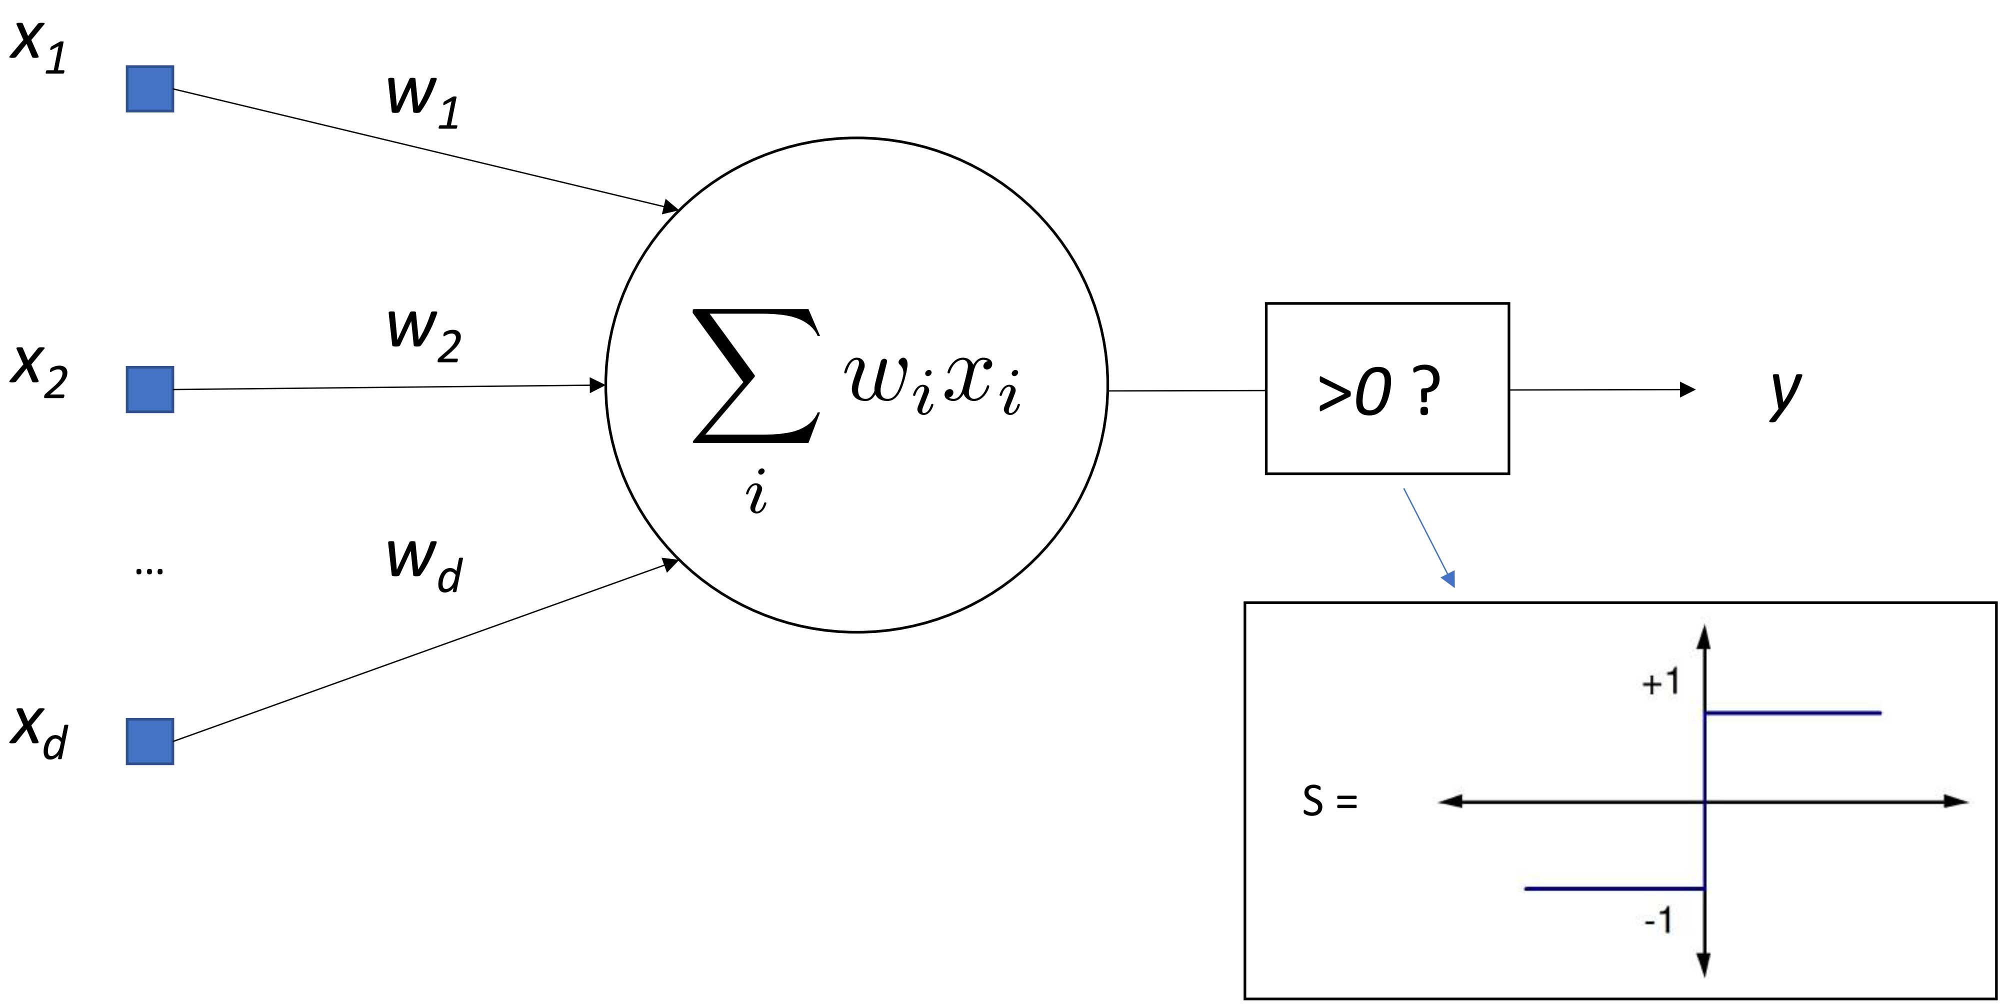

Redes Neuronales Artificiales

Perceptrón Multi Capa

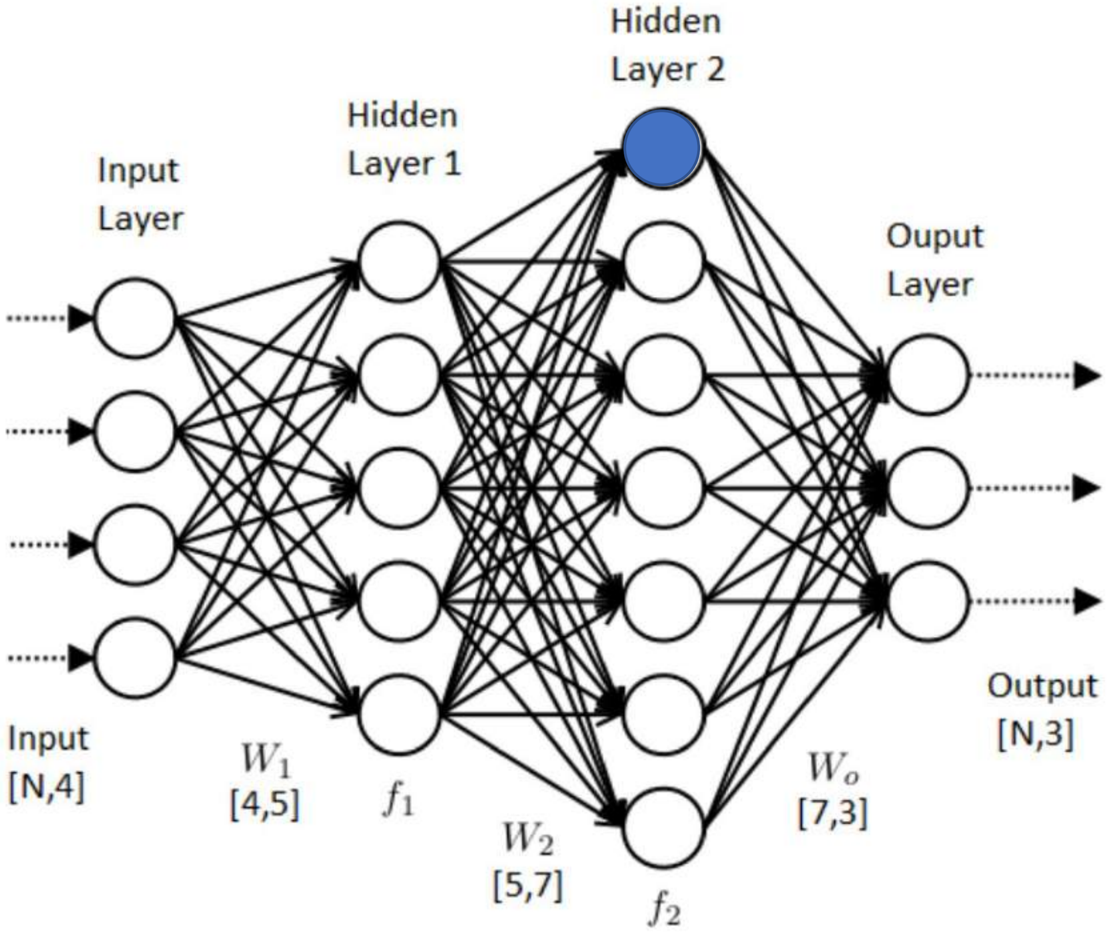

The New York Times

SUBSCRIBE NOW

Turing Award Won by 3 Pioneers in Artificial Intelligence

From left, Yann LeCun, Geoffrey Hinton and Yoshua Bengio. The researchers worked on key developments for neural networks, which are reshaping how computer systems are built.

From left, Facebook, via Associated Press; Aaron Vincent Elkaim for The New York Times; Chad Buchanan/Getty Images

By Cade Metz

March 27, 2019

Redes neuronales artificiales

Parámetros a aprender

Red neuronal artificial → $y = f(x, \boxed{w})$ → Ej: $f(x, w) = s(\sum x_i \boxed{w_i})$

Función de pérdida → $L(y, \bar{y})$ → Ej: $L(y, \bar{y}) = | y - \boxed{\bar{y}} |^2$

Etiqueta o Ground Truth

Cómo aprendemos *w*? → Gradiente descendiente

Gradiente descendente

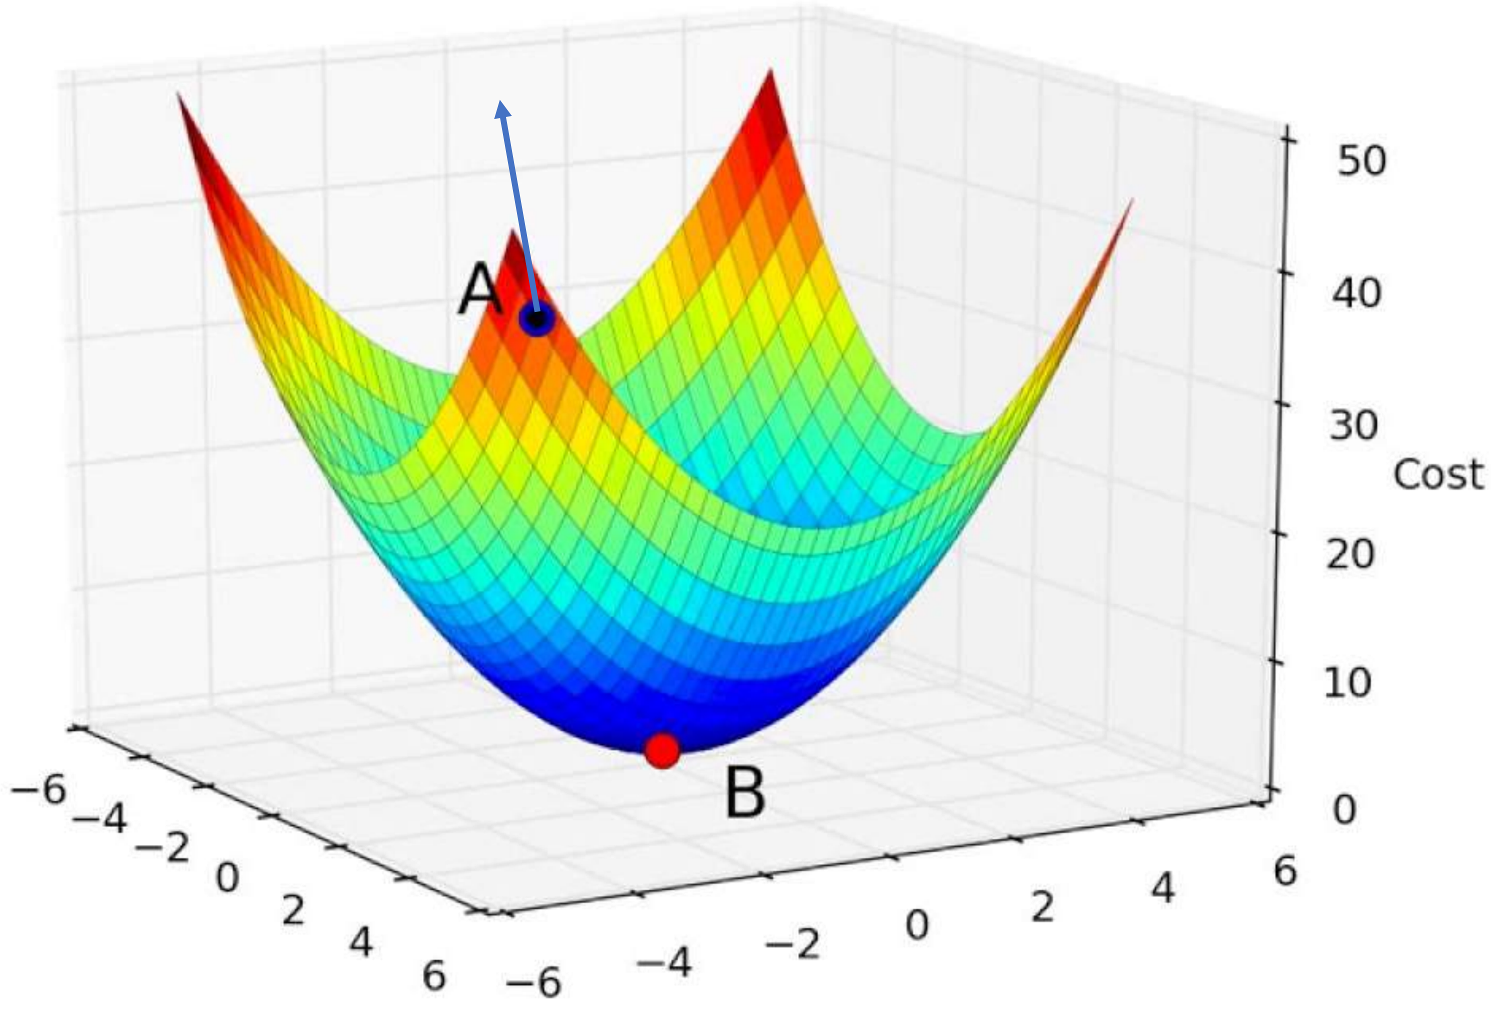

$$\nabla_{\mathbf{w}}\mathcal{L} = (\frac{\partial \mathcal{L}}{\partial w_1}, ..., \frac{\partial \mathcal{L}}{\partial w_D})$$

Gradiente descendente

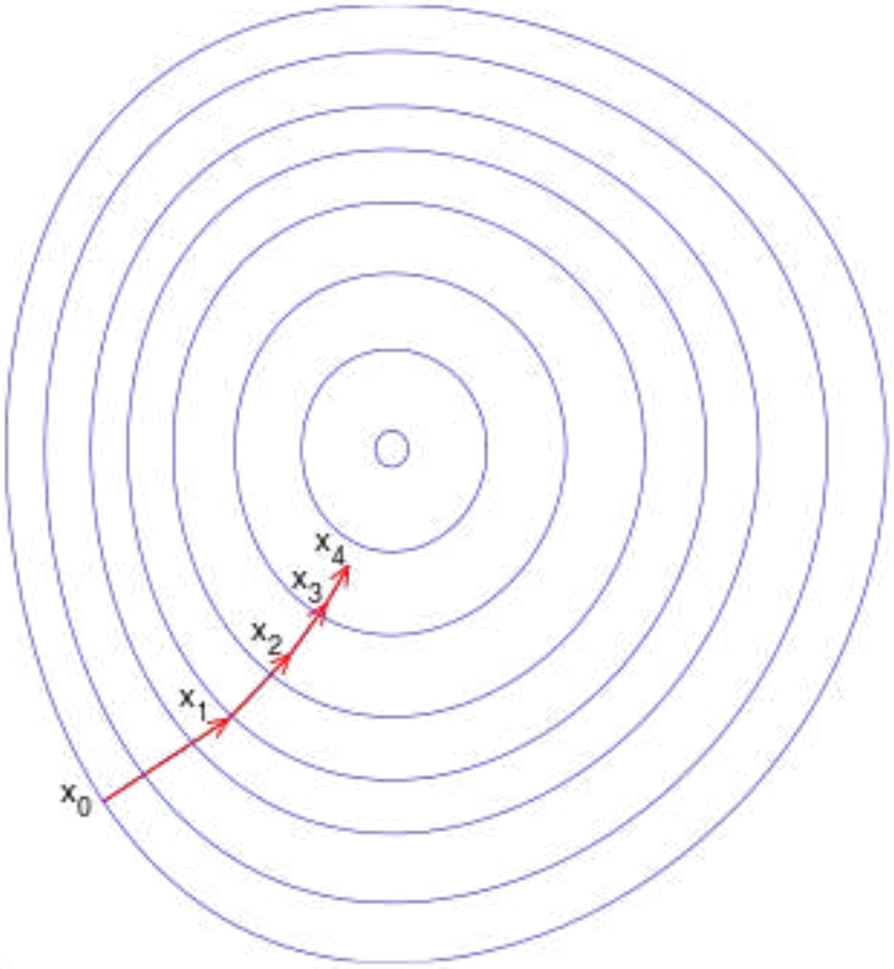

Cómo calcular el gradiente?

- **1. Derivación analítica**: derivar a mano y escribir el código

- **2. Derivación numérica :** diferencias finitas

- **3. Derivación simbólica:** se realiza utilizando las reglas estudiadas en Análisis matemático pero automatizadas (ej: Maple, Mathematica)

Backpropagation

- **4. Derivación automática:**

  - Se definen las derivadas para las operaciones 'primitivas' (matemáticas y de control)

  - Se construye un grafo de operaciones y se deriva siguiendo la regla de la cadena.

Baydin et al. *"Automatic differentiation in machine learning: a survey"*. Journal of Machine Learning Research. 2017.

Frameworks que implementan diferenciación automática

TensorFlow

PYTORCH

Redes neuronales artificiales para el análisis de imágenes inspiradas en el sistema visual

Bases neuronales de la percepción visual

(Hubel & Wiesel, 1962 – Premio Nobel 1981)

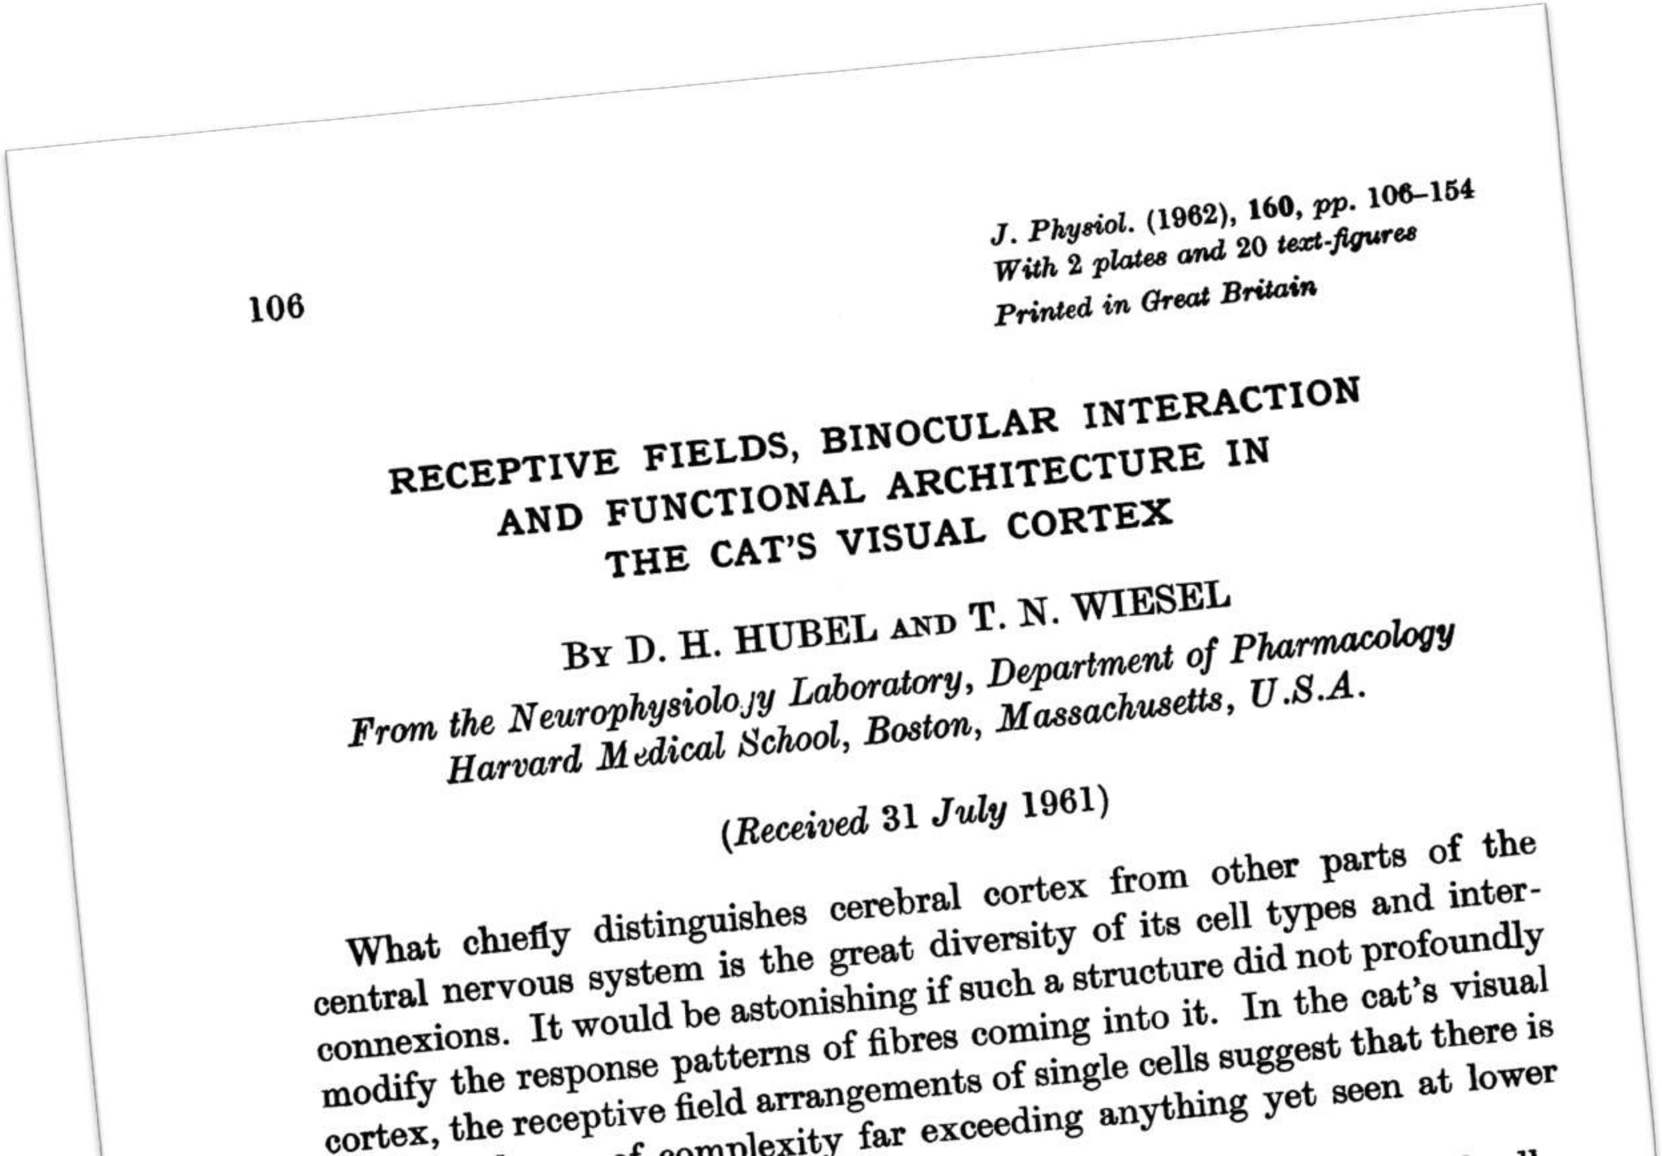

J. Physiol. (1962), **160**, *pp.* 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

By D. H. HUBEL and T. N. WIESEL

From the Neurophysiology Laboratory, Department of Pharmacology Harvard Medical School, Boston, Massachusetts, U.S.A.

(*Received* 31 *July* 1961)

What chiefly distinguishes cerebral cortex from other parts of the central nervous system is the great diversity of its cell types and interconnexions. It would be astonishing if such a structure did not profoundly modify the response patterns of fibres coming into it. In the cat's visual cortex, the receptive field arrangements of single cells suggest that there is ... of complexity far exceeding anything yet seen at lower

Redes neuronales inspiradas en el sistema visual

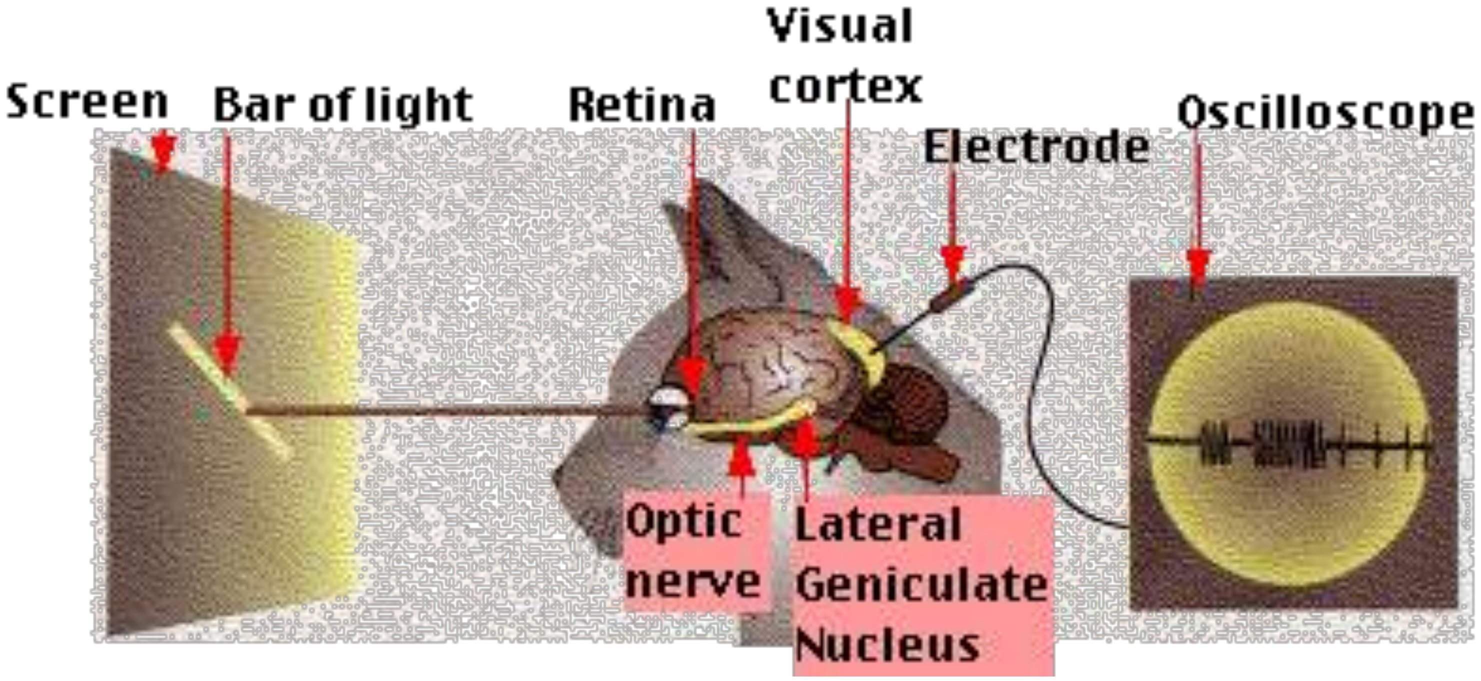

Redes neuronales inspiradas en el sistema visual

- Midieron la respuesta eléctrica en el cerebro de un gato estimulándolo con patrones simples en una pantalla.

- Encontraron que las **neuronas de la corteza visual** temprana se organizan de **forma jerárquica**, donde las **primeras** reaccionan a **patrones simples como líneas**, y las **posteriores** capas responden a **patrones más complejos** combinando las activaciones que reciben.

- En el modelo propuesto, las neuronas en las capas superiores tienen un mayor **campo receptivo** y son **menos sensibles a la posición desde la cual proviene dicho estímulo**.

Neocognitrón: red neuronal artificial inspirada en (Hubel & Wiesel, 1962)

(Fukushima 1980)

Biol. Cybernetics 36, 193–202 (1980)

Biological Cybernetics

Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position

Kunihiko Fukushima

NHK Broadcasting Science Research Laboratories, Kinuta, Setagaya, Tokyo, Japan

Abstract. A neural network model for a mechanism of visual pattern recognition is proposed in this paper. The network is self-organized by "learning without a teacher", and acquires an ability to recognize stimulus patterns based on the geometrical similarity (Gestalt) of their shapes without affected by their positions. This network is given a nickname "neocognitron". After completion of self-organization, the network has a structure similar to the hierarchy model of the visual nervous system proposed by Hubel and Wiesel. The network consists of an input layer (photoreceptor array) followed by a cascade connection of a number of modular structures, each of which is composed of two

reveal it only by conventional physiological experiments. So, we take a slightly different approach to this problem. If we could make a neural network model which has the same capability for pattern recognition as a human being, it would give us a powerful clue to the understanding of the neural mechanism in the brain. In this paper, we discuss how to synthesize a neural network model in order to endow it an ability of pattern recognition like a human being.

Several models were proposed with this intention (Rosenblatt, 1962; Kabrisky, 1966; Giebel, 1971; Fukushima, 1975). The response of most of these models, however, was severely affected by the shift in

Neocognitrón: red neuronal artificial inspirada en (Hubel & Wiesel, 1962)

(Fukushima 1980)

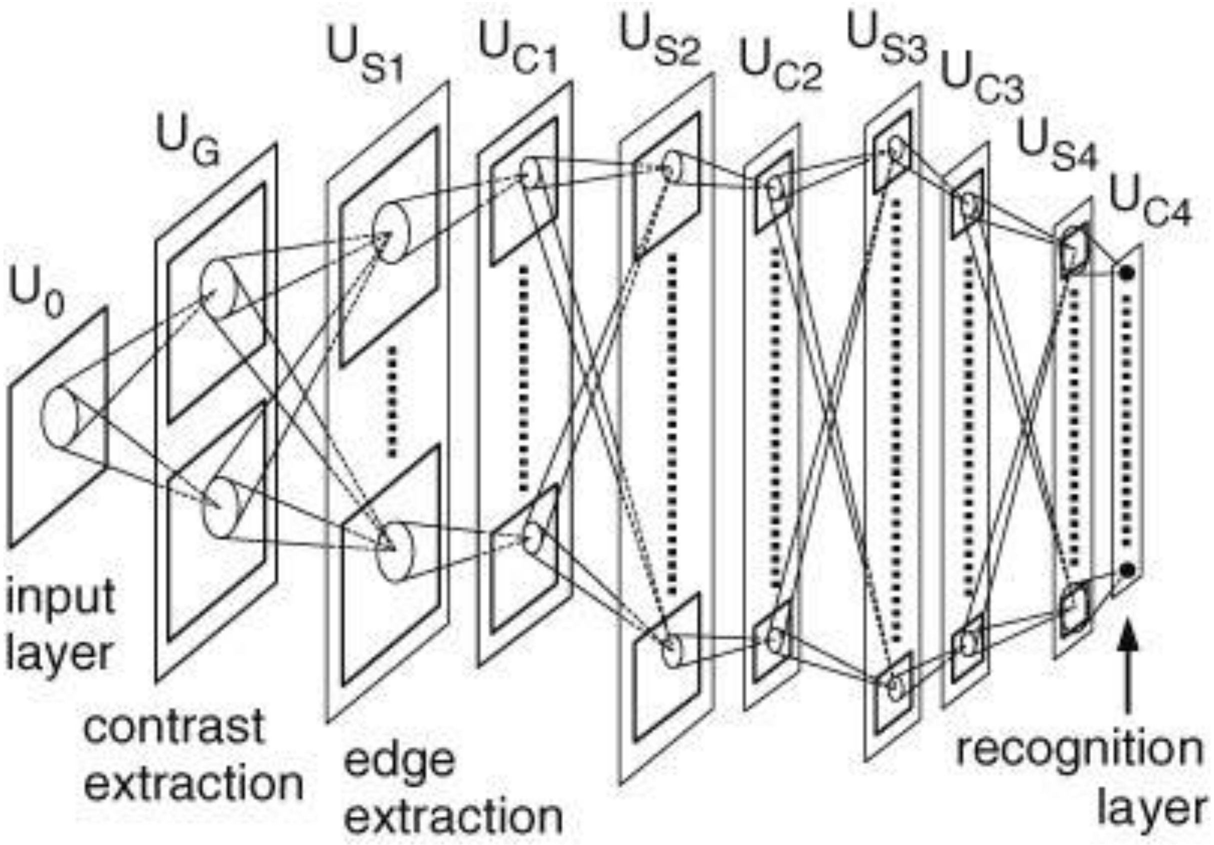

Neocognitrón: red neuronal artificial inspirada en (Hubel & Wiesel, 1962)

(Fukushima 1980)

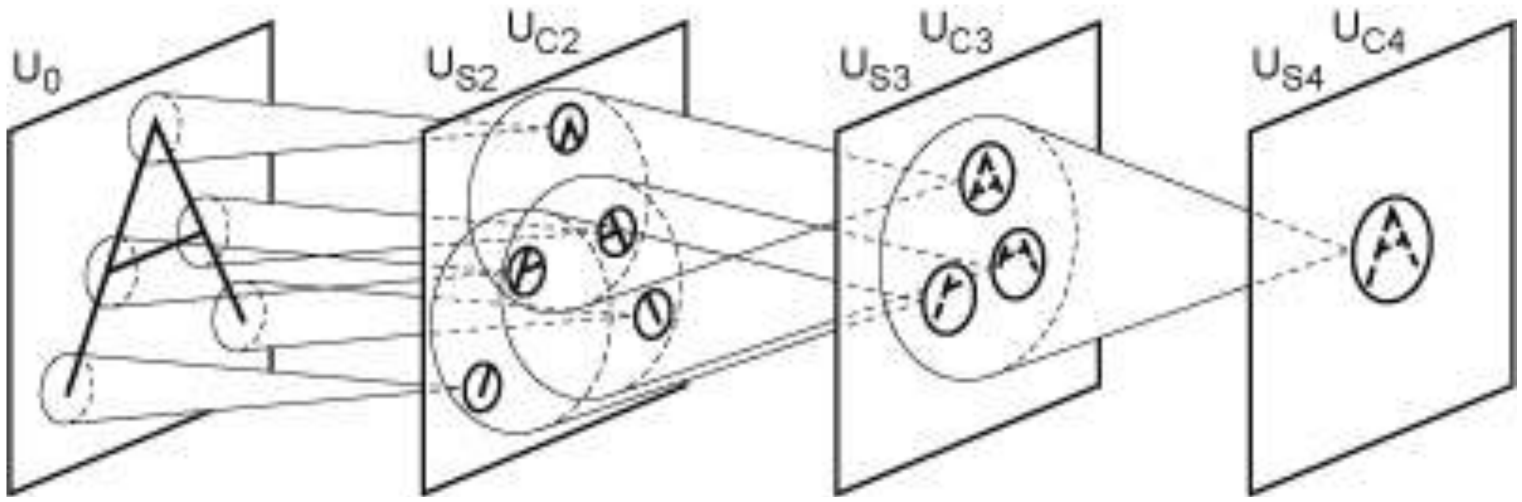

Redes neuronales convolucionales

(Lecun 1989)

Handwritten Digit Recognition with a Back-Propagation Network

Y. Le Cun, B. Boser, J. S. Denker, D. Henderson,
R. E. Howard, W. Hubbard, and L. D. Jackel
AT&T Bell Laboratories, Holmdel, N. J. 07733

ABSTRACT

We present an application of back-propagation networks to handwritten digit recognition. Minimal preprocessing of the data was required, but architecture of the network was highly constrained and specifically designed for the task. The input of the network consists of normalized images of isolated digits. The method has 1% error rate and about a 9% reject rate on zipcode digits provided

Redes neuronales convolucionales

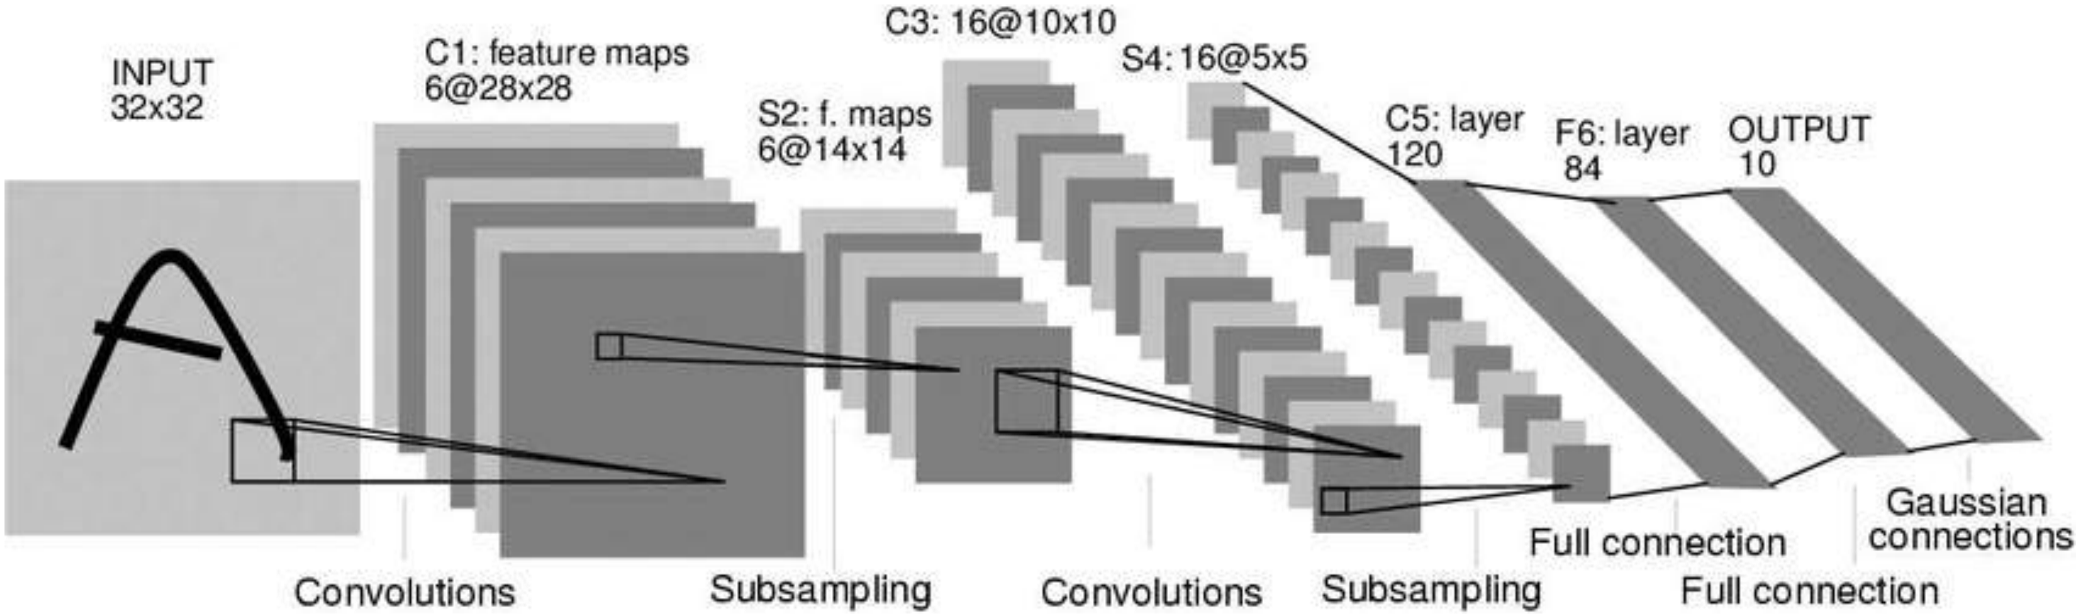

Convolución

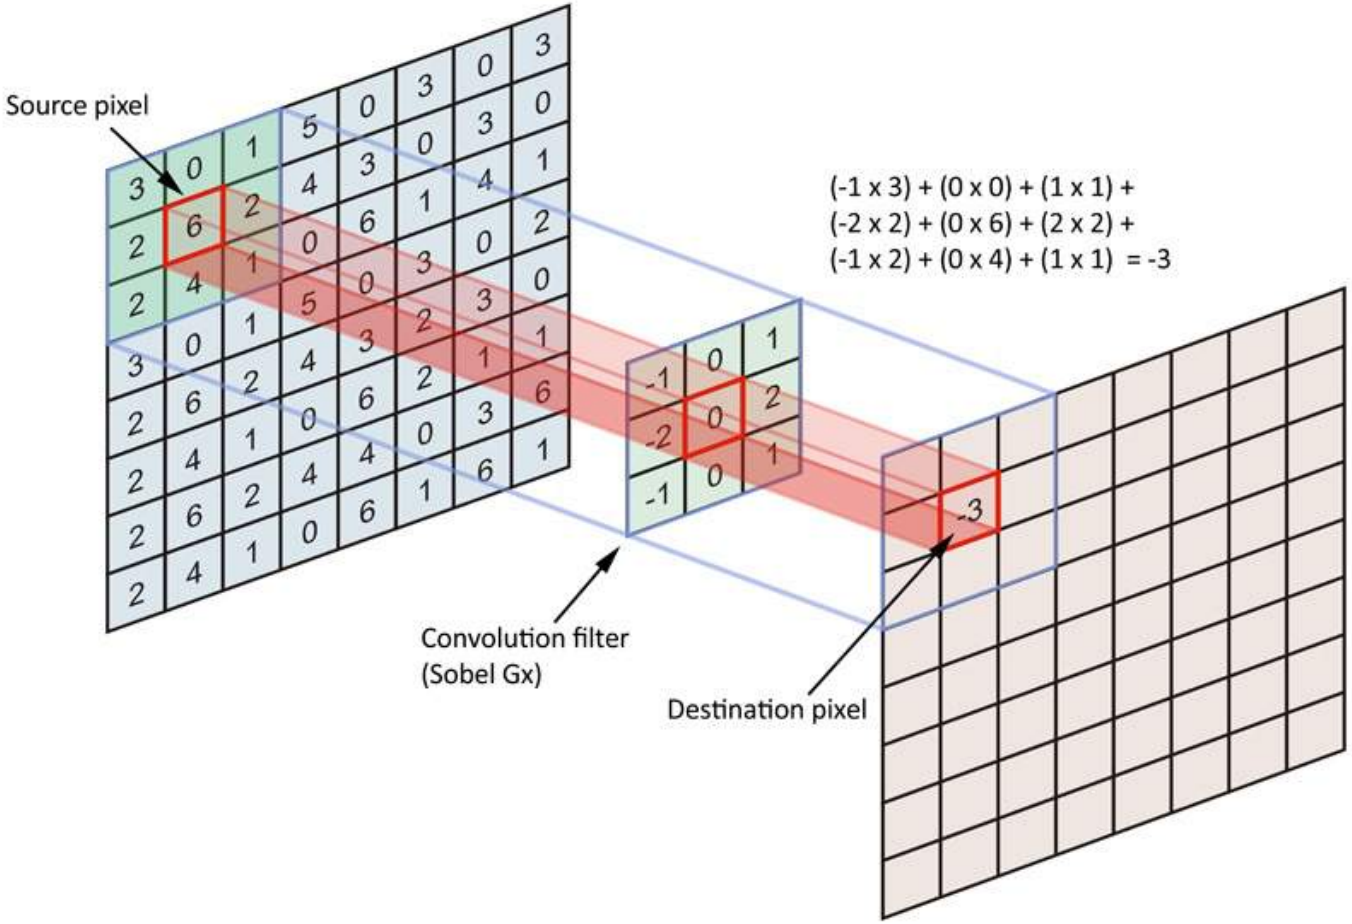

Redes neuronales convolucionales

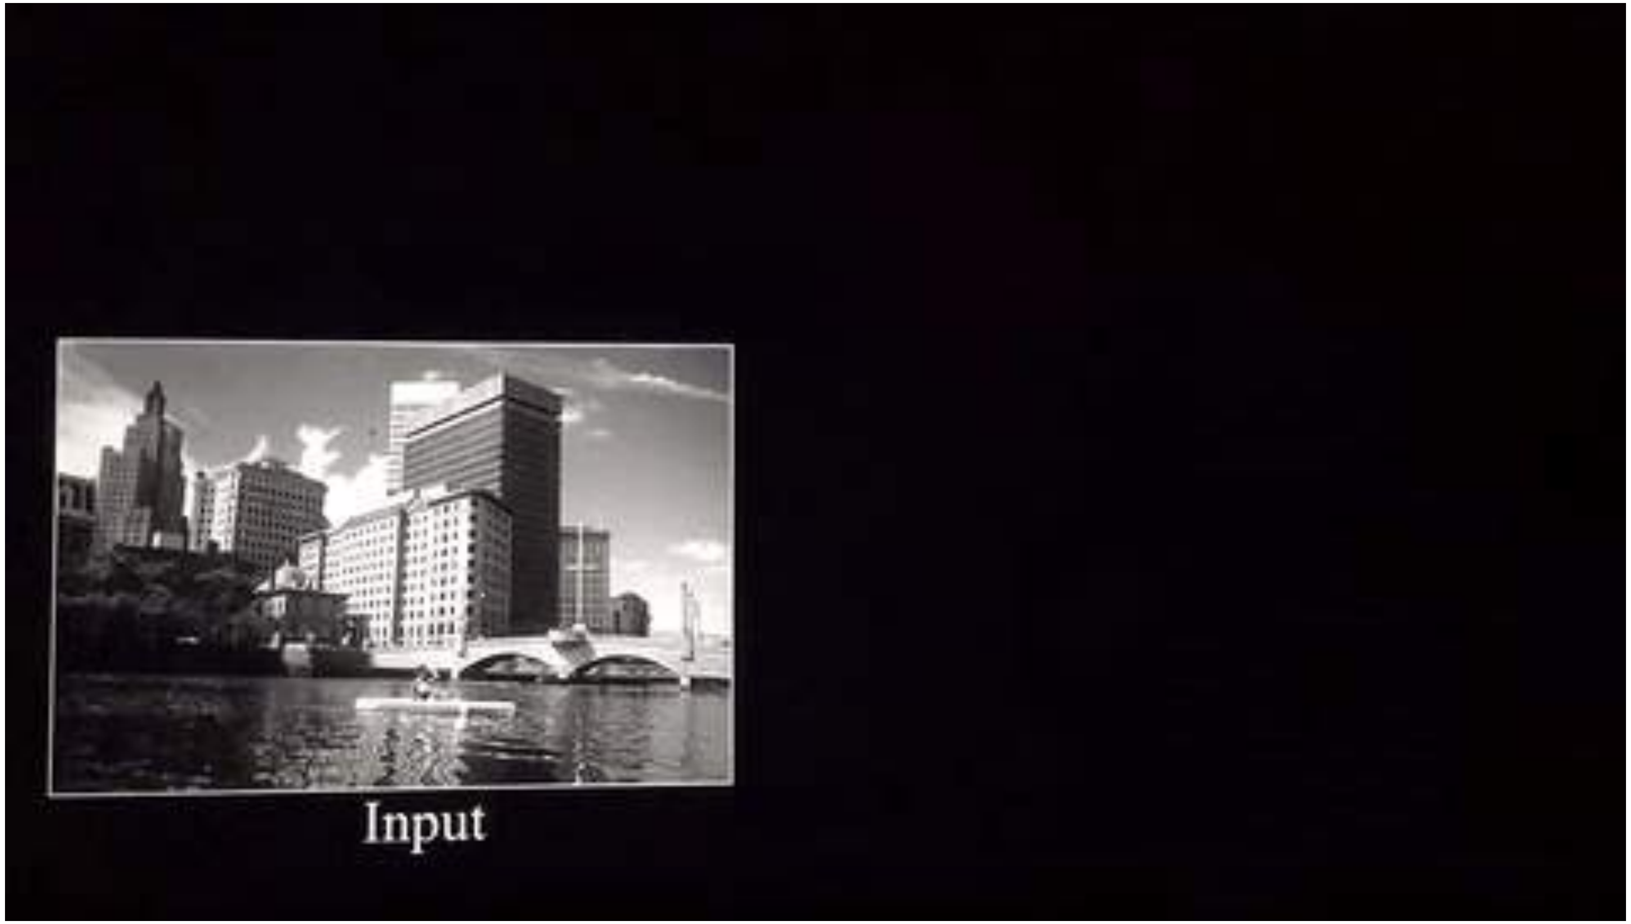

Redes Neuronales Convolucionales

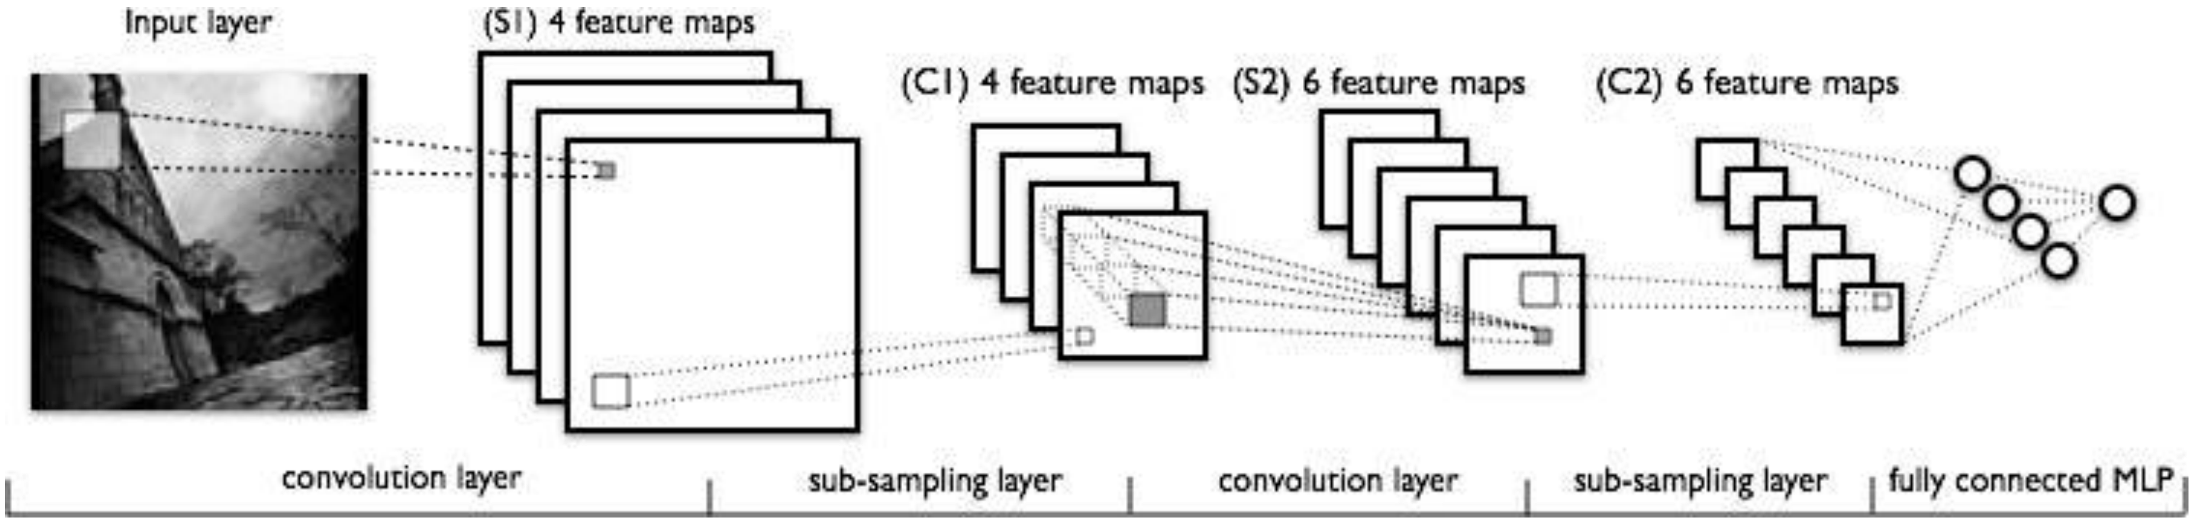

ConvNets para el análisis de imágenes: Ventajas

- Naturalmente adaptadas a la **estructura regular** de las imágenes (por medio de la operación de convolución)
- Invariantes respecto a traslaciones
- Aprendizaje **end-to-end**
- Bajos requerimientos de memoria: **weight sharing**
- **Eficientes** en test-time
- Buen grado de **generalización** si se entrenan con suficientes datos

ImageNet Challenge 2012
Clasificación de imágenes y detección de objetos


ImageNet
15.000.000 millones de imágenes, 1000 categorías

ImageNet Challenge

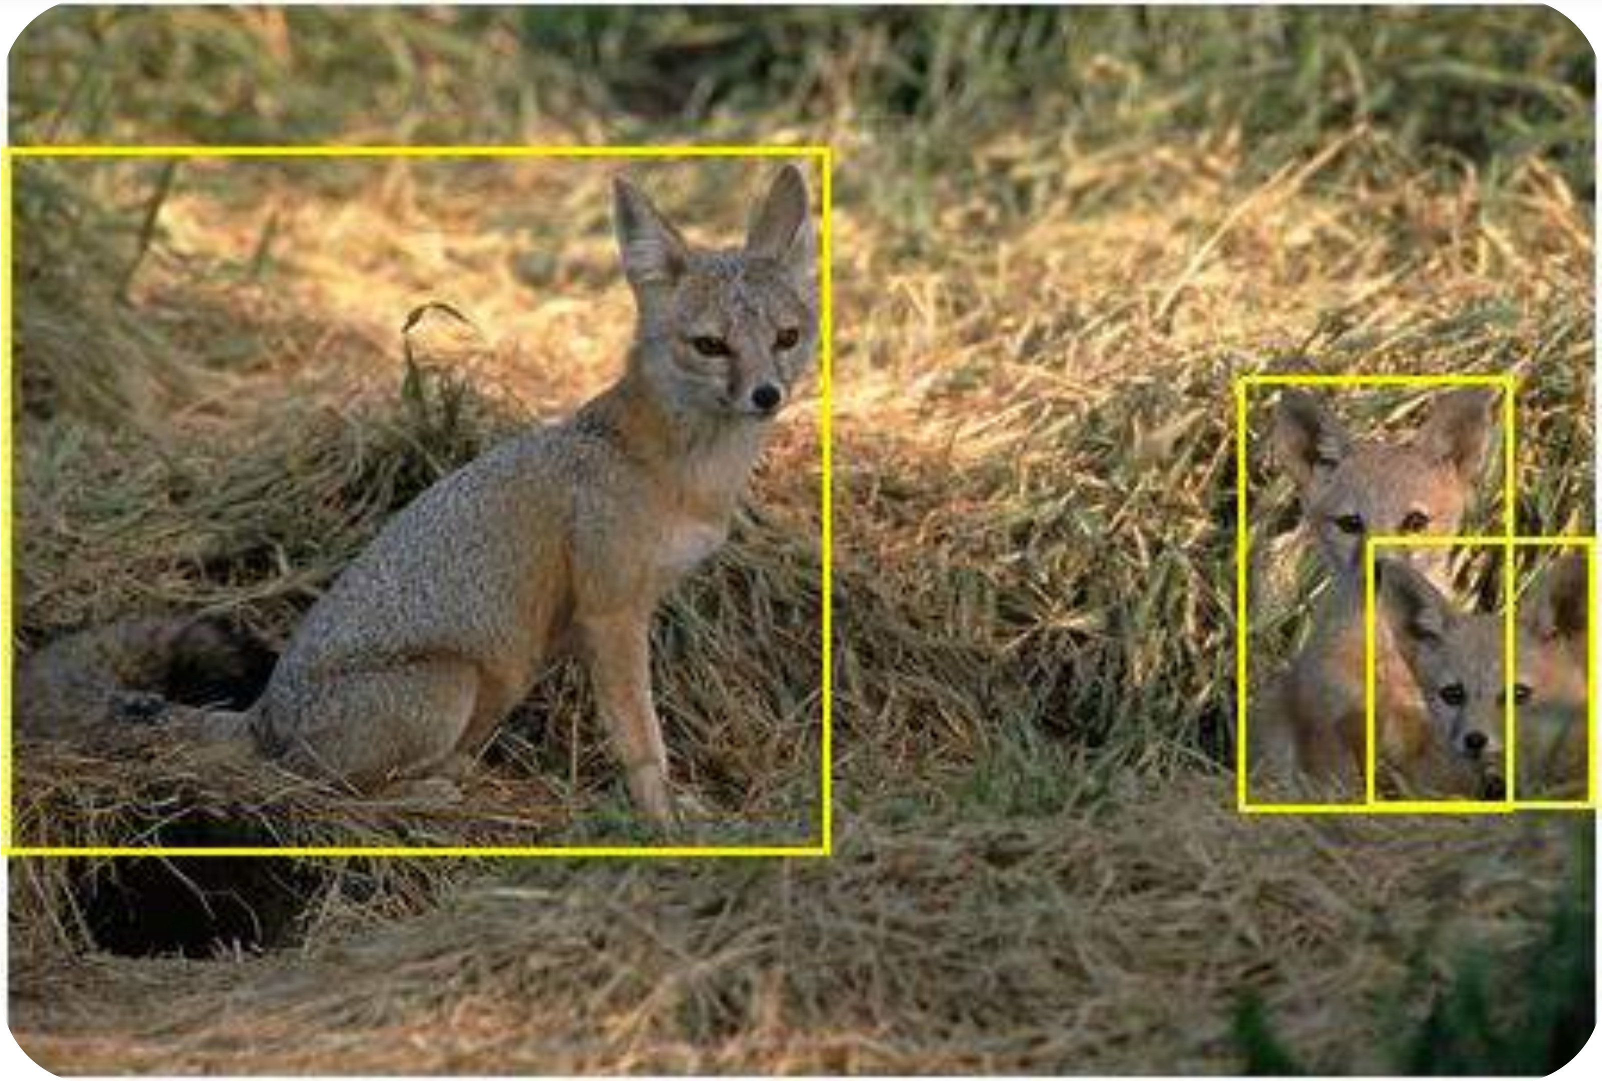

ImageNet Challenge

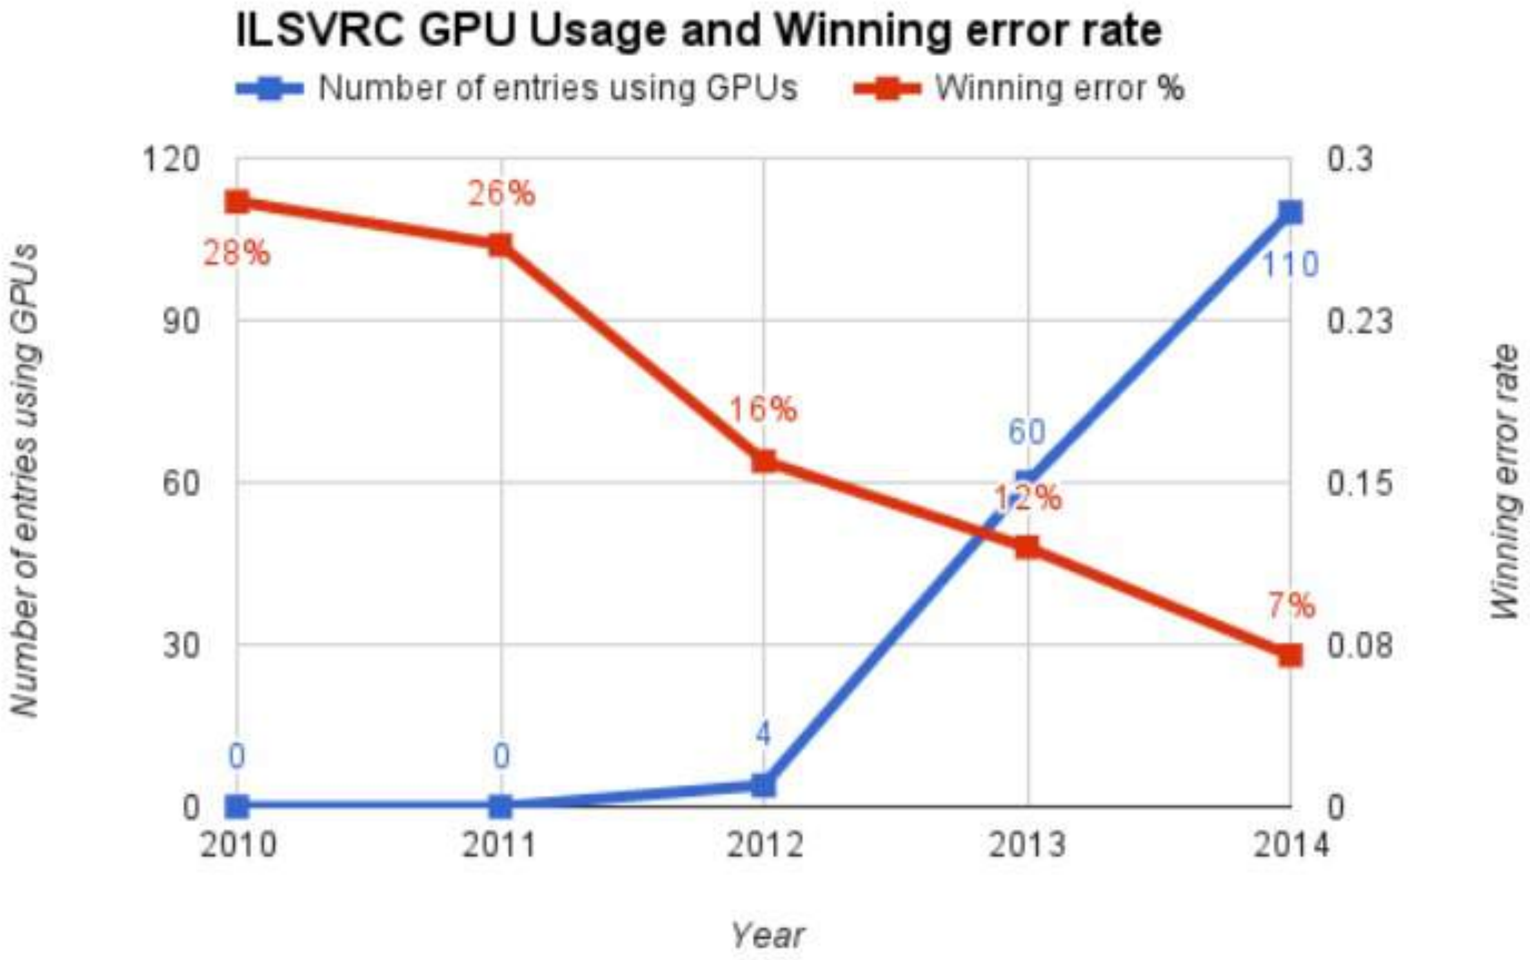

ImageNet

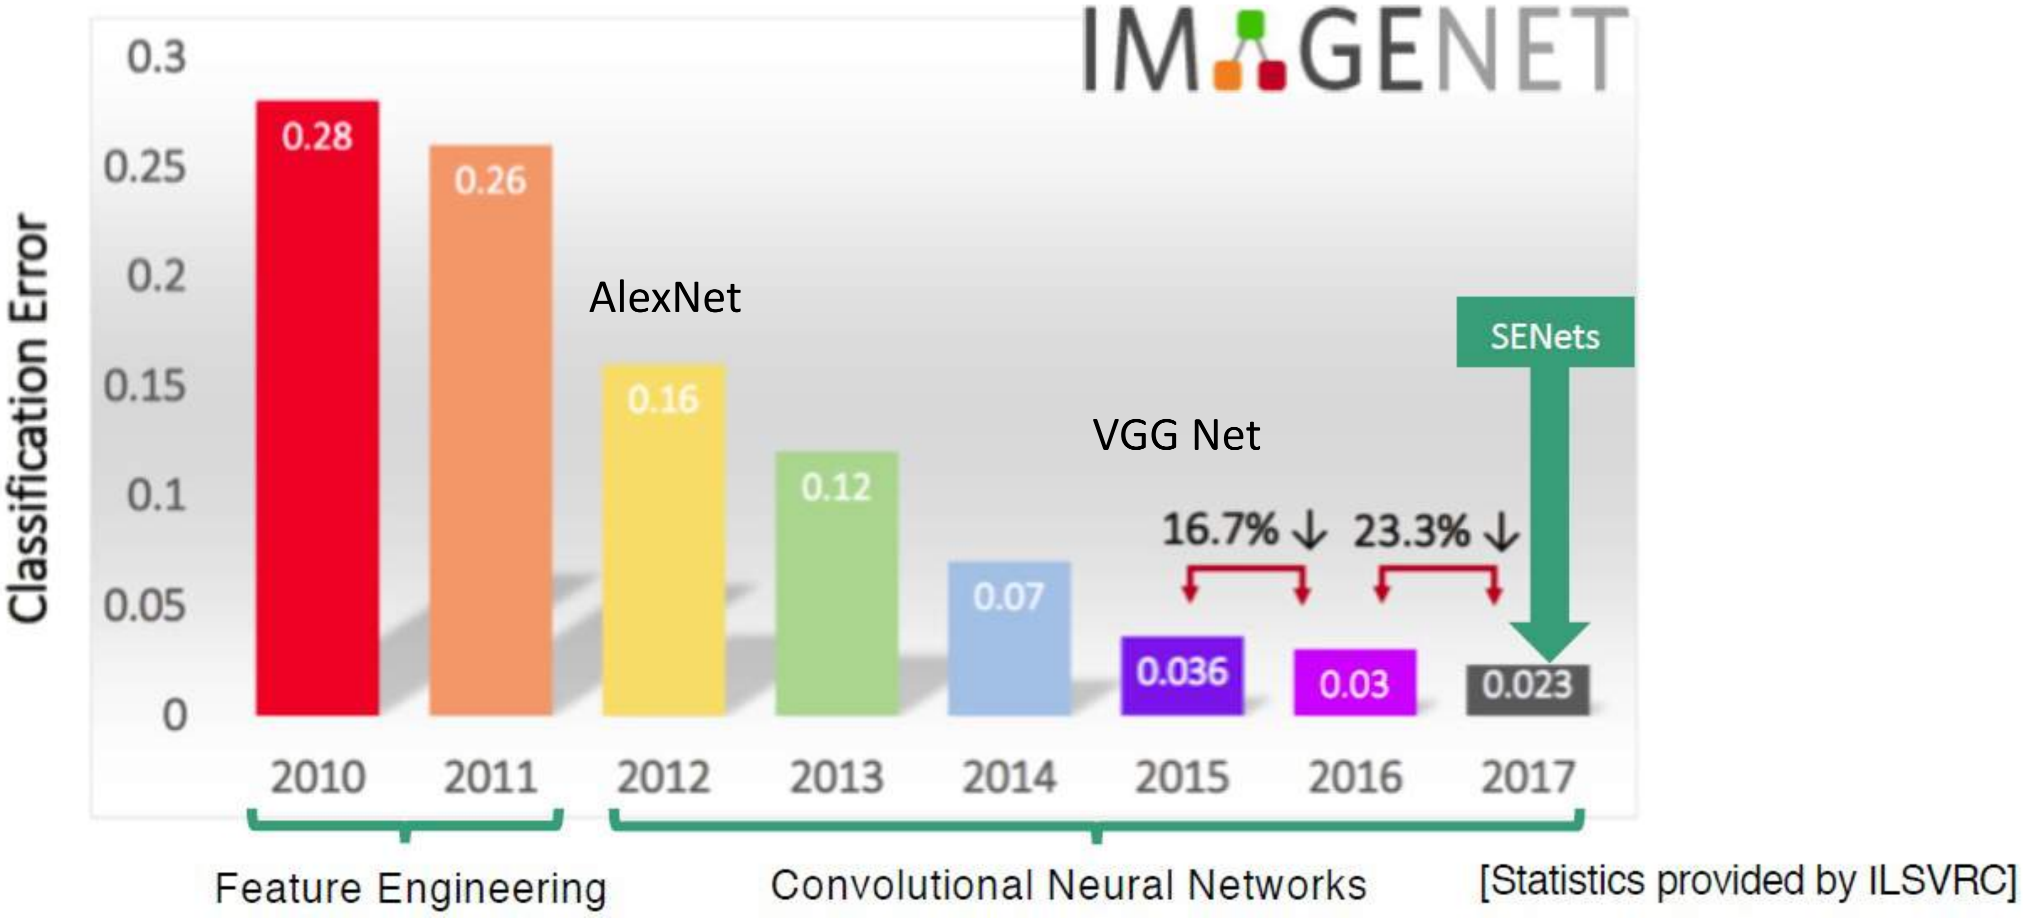

Image source: https://towardsdatascience.com/review-senet-squeeze-and-excitation-network-winner-of-ilsvrc-2017-image-classification-a887b98b2883

Computer Vision Revolution

Más datos

Mejores modelos

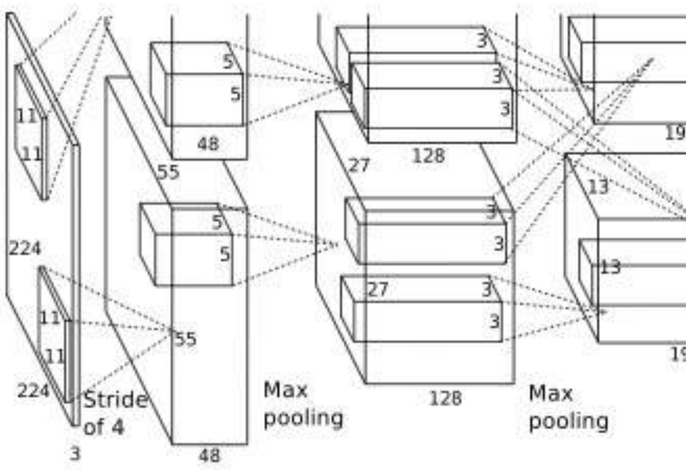

GPUs más poderosas

ImageNet Challenge

Algunas aplicaciones

2015: AlphaGO

AlphaGO

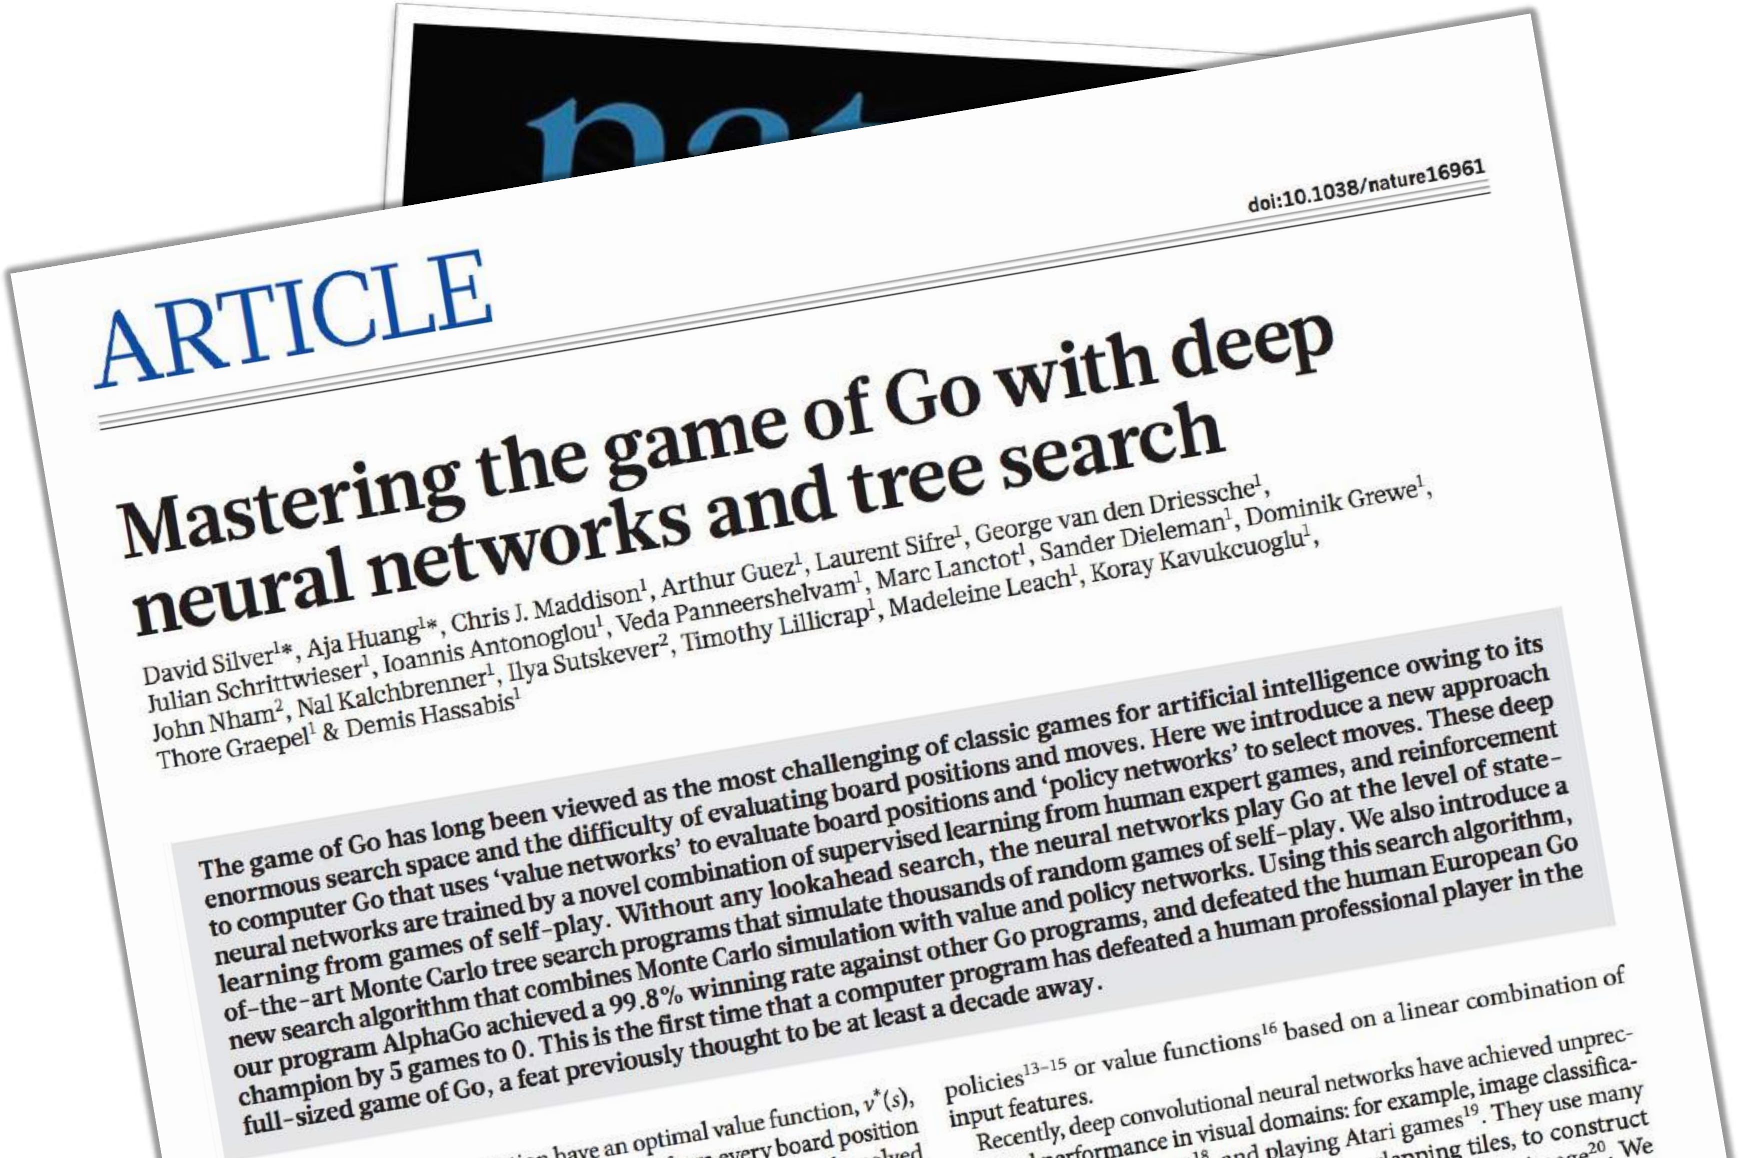

doi:10.1038/nature16961

ARTICLE

Mastering the game of Go with deep neural networks and tree search

David Silver[1]*, Aja Huang[1]*, Chris J. Maddison[1], Arthur Guez[1], Laurent Sifre[1], George van den Driessche[1], Julian Schrittwieser[1], Ioannis Antonoglou[1], Veda Panneershelvam[1], Marc Lanctot[1], Sander Dieleman[1], Dominik Grewe[1], John Nham[2], Nal Kalchbrenner[1], Ilya Sutskever[2], Timothy Lillicrap[1], Madeleine Leach[1], Koray Kavukcuoglu[1], Thore Graepel[1] & Demis Hassabis[1]

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

... have an optimal value function, $v^*(s)$, ...

policies[13–15] or value functions[16] based on a linear combination of input features.

Recently, deep convolutional neural networks have achieved unprecedented performance in visual domains: for example, image classification[17], ...[18], and playing Atari games[19]. They use many ... tiles, to construct ...[20]. We

AlphaGo

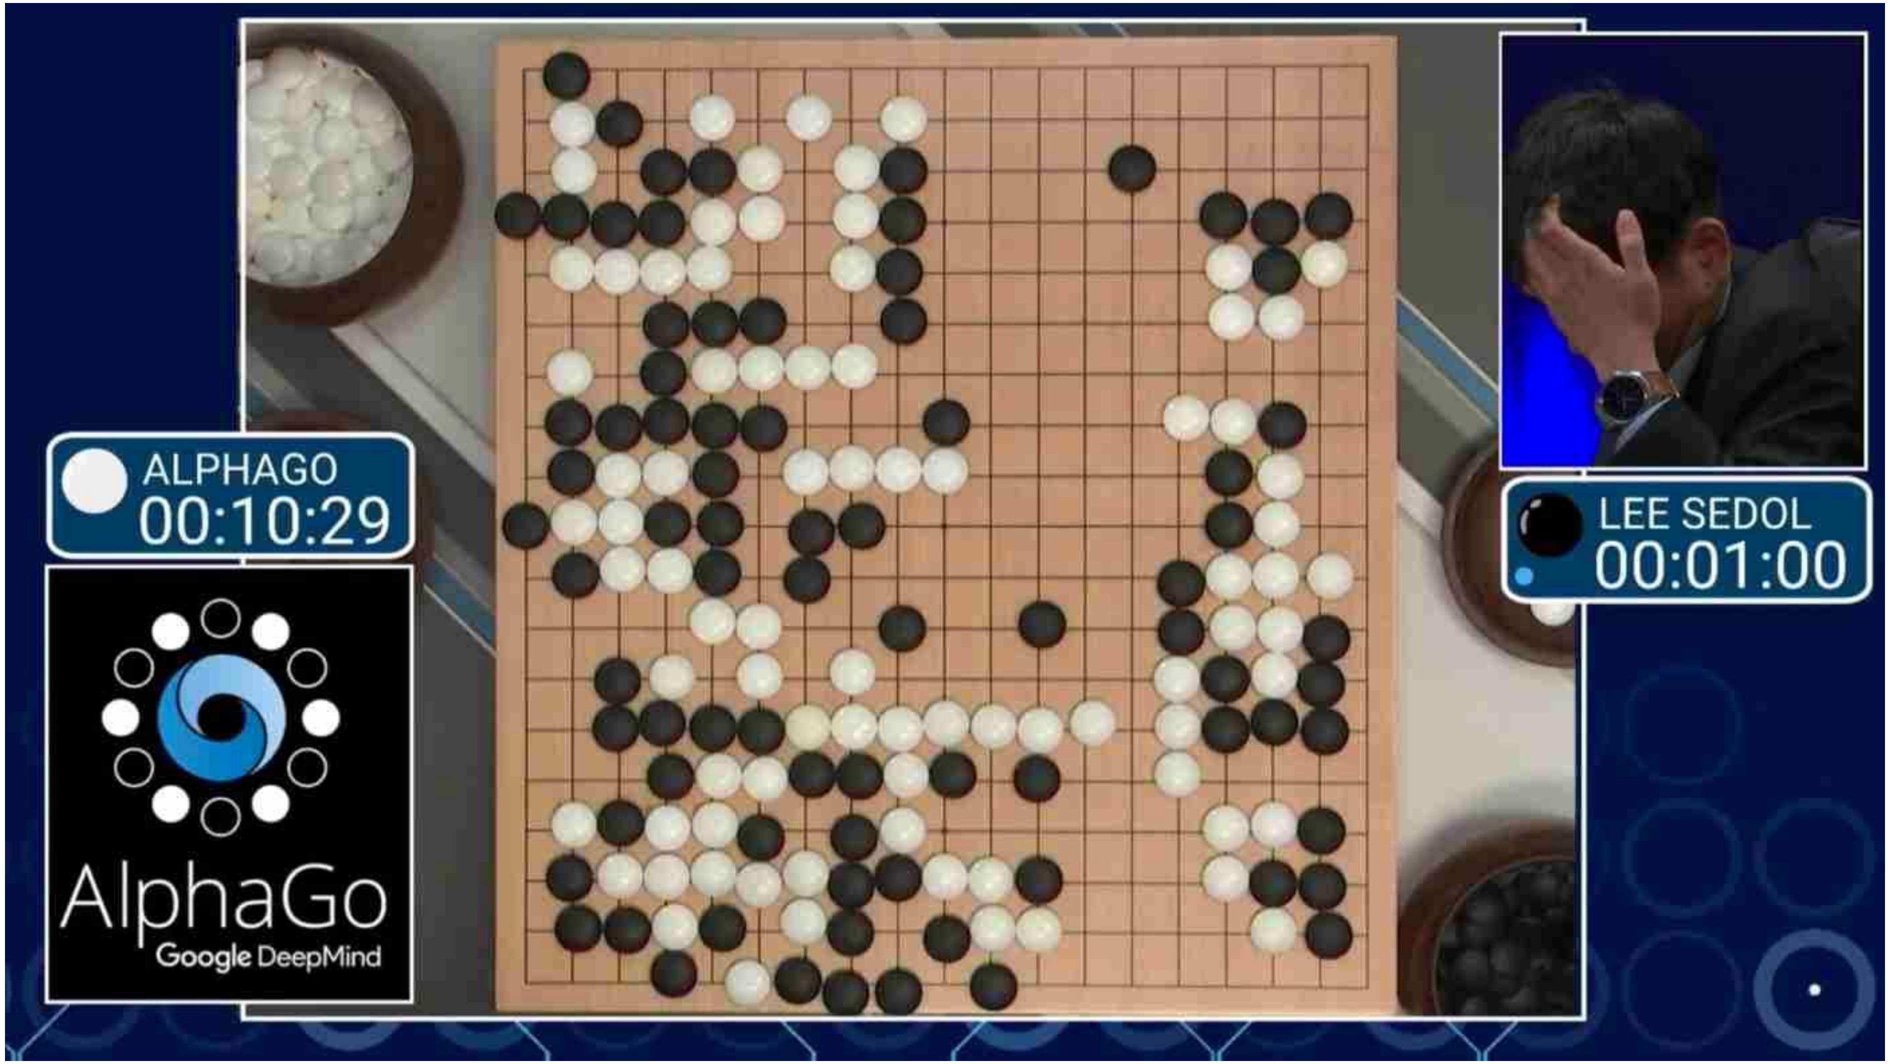

AlphaGo

- Entrenamiento inicial con 30 millones de tableros jugados por humanos expertos
- Luego refinado con entrenamiento basado en Aprendizaje por Refuerzo

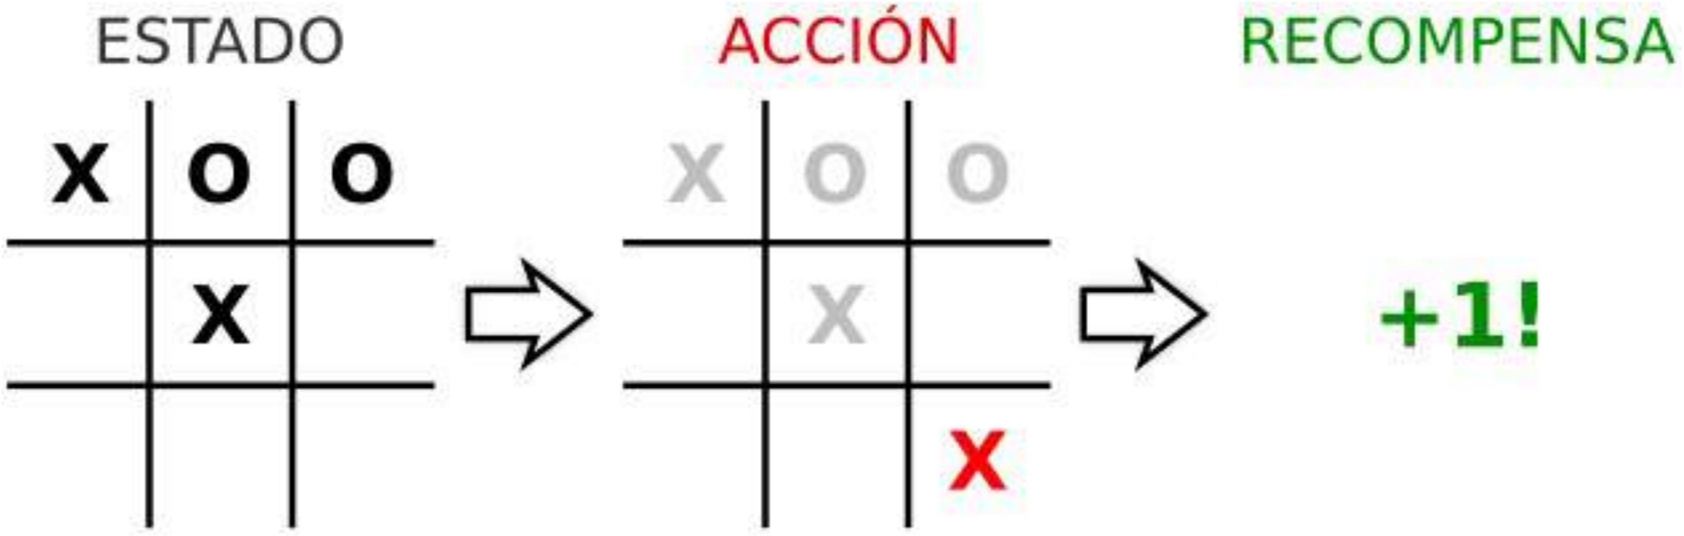

2017: AlphaGO Zero

AlphaGo Zero

ARTICLE

doi:10.1038/nature24270

Mastering the game of Go without human knowledge

David Silver[1]*, Julian Schrittwieser[1]*, Karen Simonyan[1]*, Ioannis Antonoglou[1], Aja Huang[1], Arthur Guez[1], Thomas Hubert[1], Lucas Baker[1], Matthew Lai[1], Adrian Bolton[1], Yutian Chen[1], Timothy Lillicrap[1], Fan Hui[1], Laurent Sifre[1], George van den Driessche[1], Thore Graepel[1] & Demis Hassabis[1]

A long-standing goal of artificial intelligence is an algorithm that learns, *tabula rasa*, superhuman proficiency in challenging domains. Recently, AlphaGo became the first program to defeat a world champion in the game of Go. The tree search in AlphaGo evaluated positions and selected moves using deep neural networks. These neural networks were trained by supervised learning from human expert moves, and by reinforcement learning from self-play. Here we introduce an algorithm based solely on reinforcement learning, without human data, guidance or domain knowledge beyond game rules. AlphaGo becomes its own teacher: a neural network is trained to predict AlphaGo's own move selections and also the winner of AlphaGo's games. This neural network improves the strength of the tree search, resulting in higher quality move selection and stronger self-play in the next iteration. Starting *tabula rasa*, our new program AlphaGo Zero achieved superhuman performance, winning 100–0 against the previously published, champion-defeating AlphaGo.

Much progress towards artificial intelligence has been made using supervised learning systems that are trained to replicate the decisions of human experts[1–4]. However, expert data sets are often expensive, unreliable or simply unavailable. Even when reliable data

trained solely by self-play reinforcement learning, starting from random play, without any supervision

AlphaGo Zero

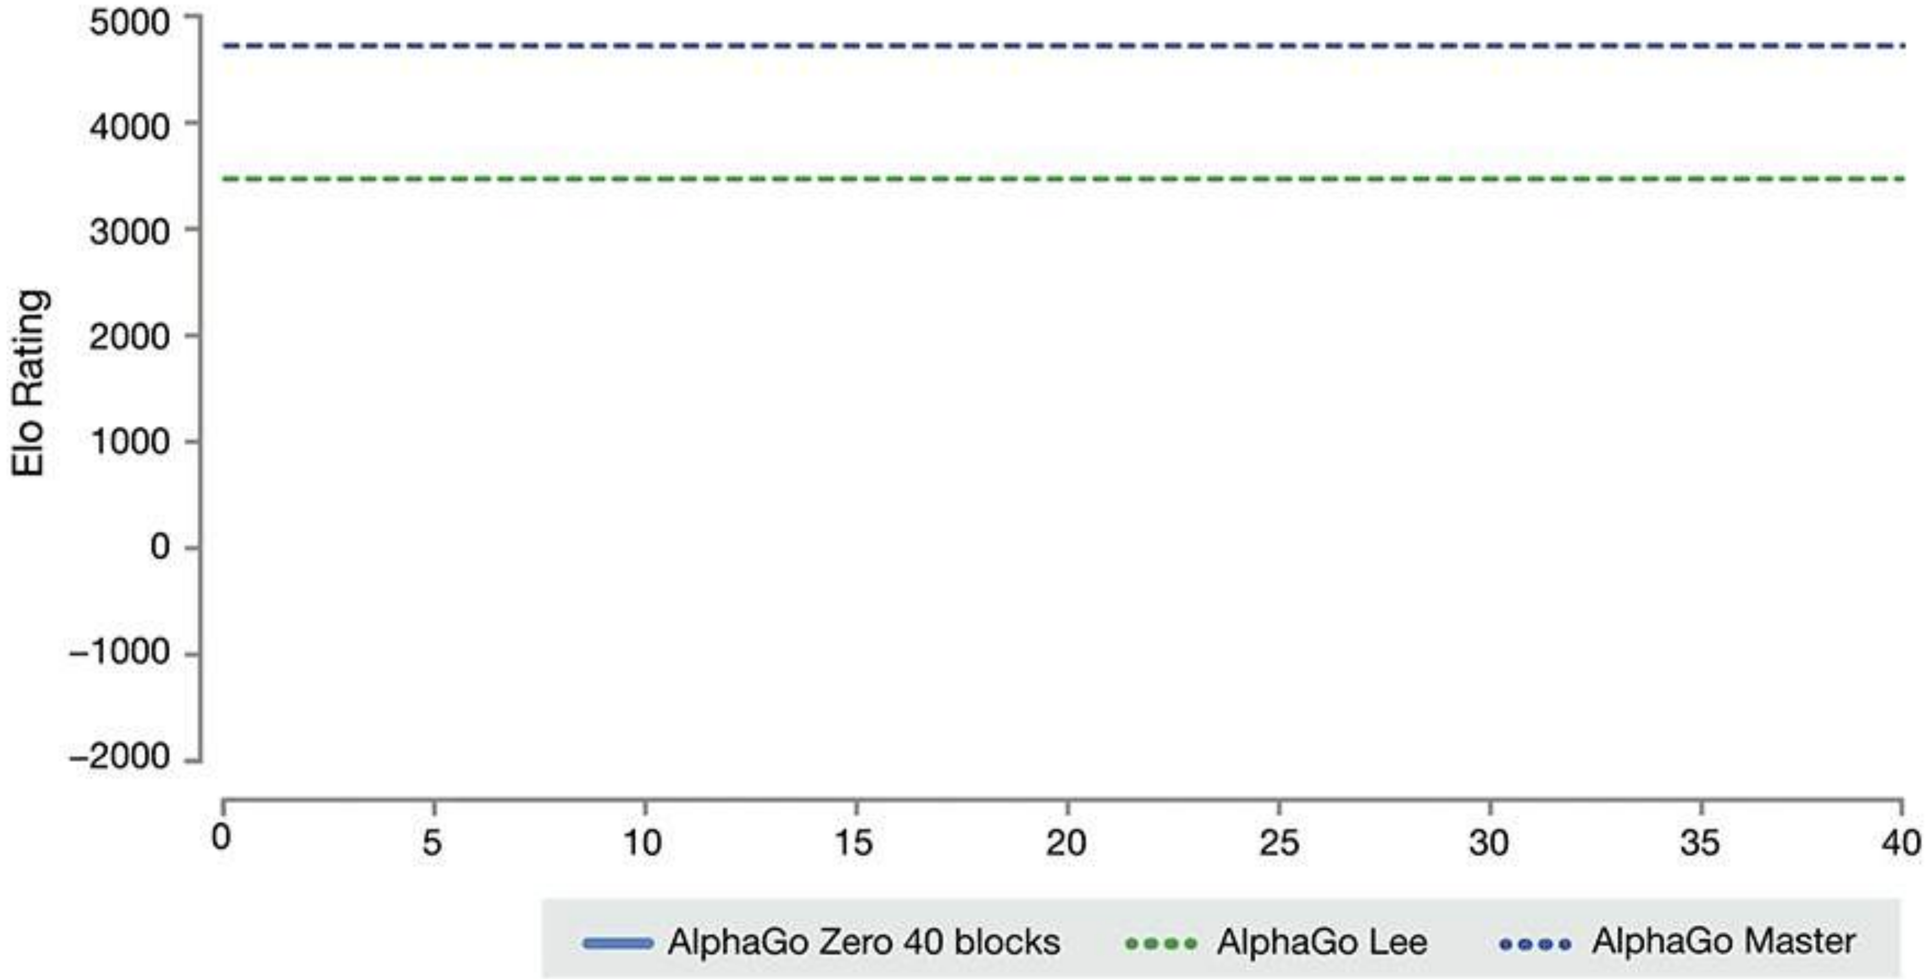

2015: Deep Reinforcement Learning en Atari

Deep Reinforcement Learning

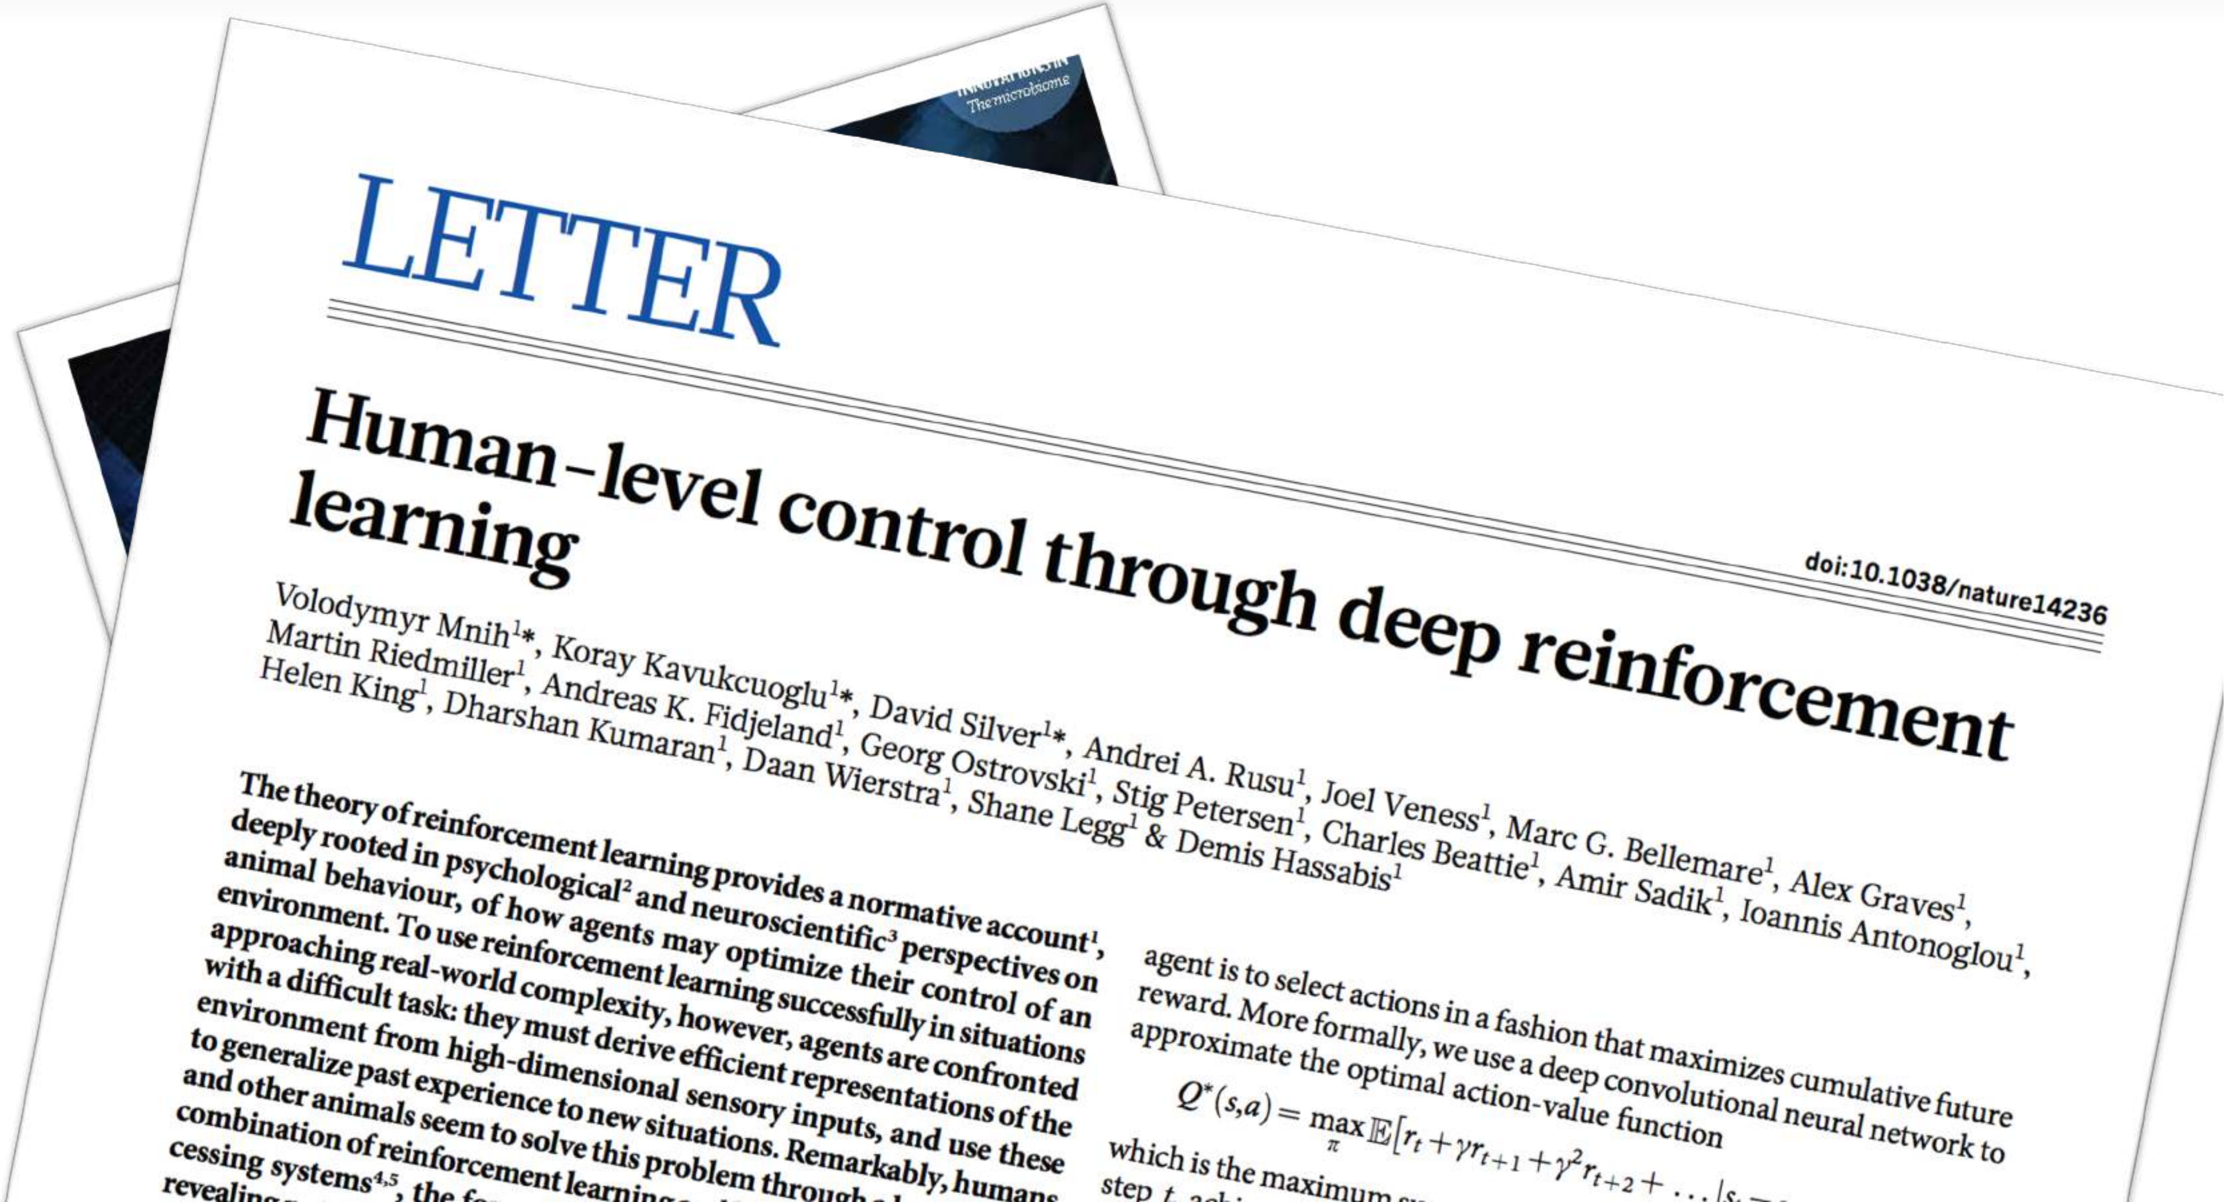

LETTER

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih[1]*, Koray Kavukcuoglu[1]*, David Silver[1]*, Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a ... combination of reinforcement learning ... cessing systems[4,5], the ... revealing ...

agent is to select actions in a fashion that maximizes cumulative future reward. More formally, we use a deep convolutional neural network to approximate the optimal action-value function

$$Q^*(s,a) = \max_{\pi} \mathbb{E}[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots | s_t = \ldots$$

which is the maximum ... step t, achi...

Arcades

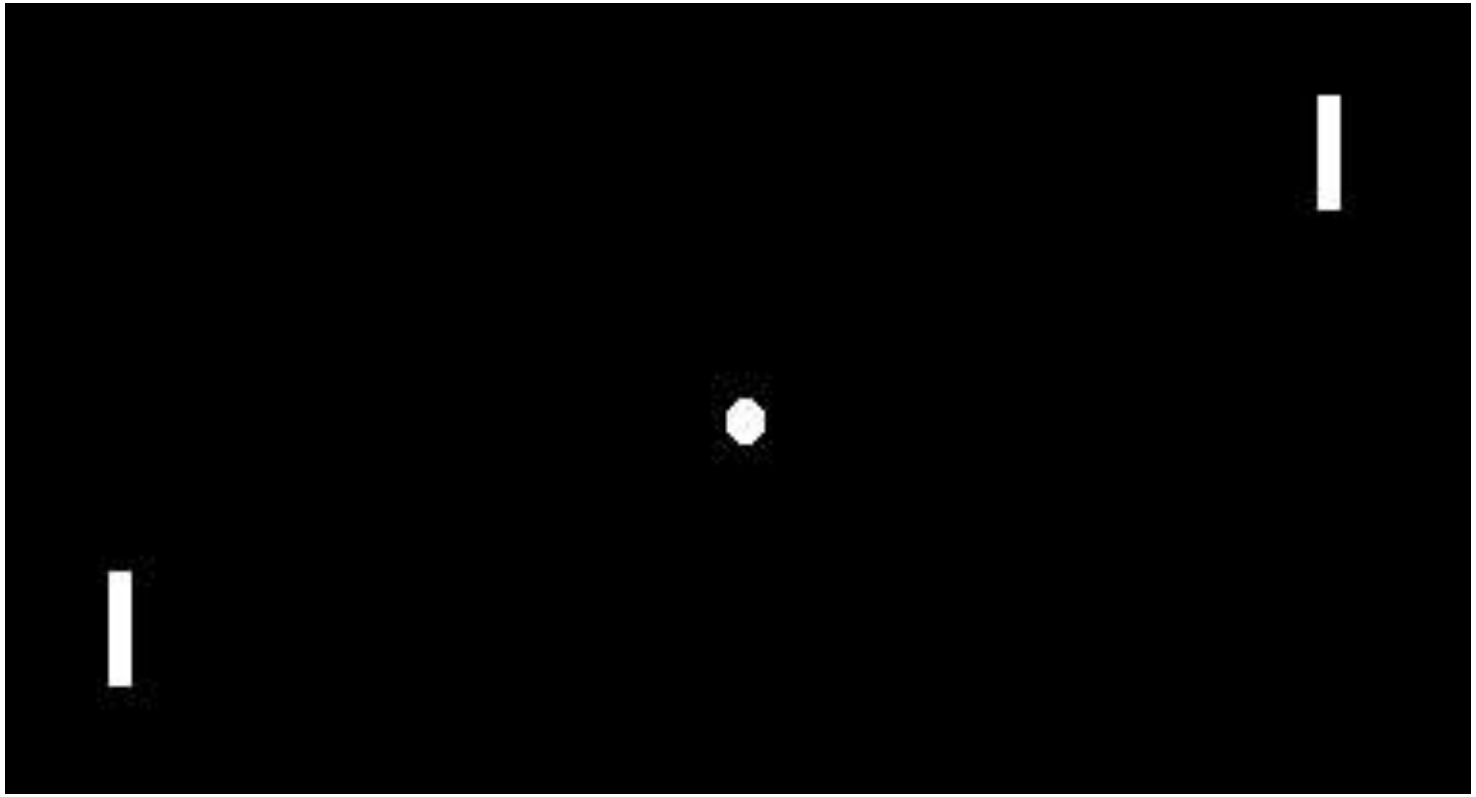

Arcades

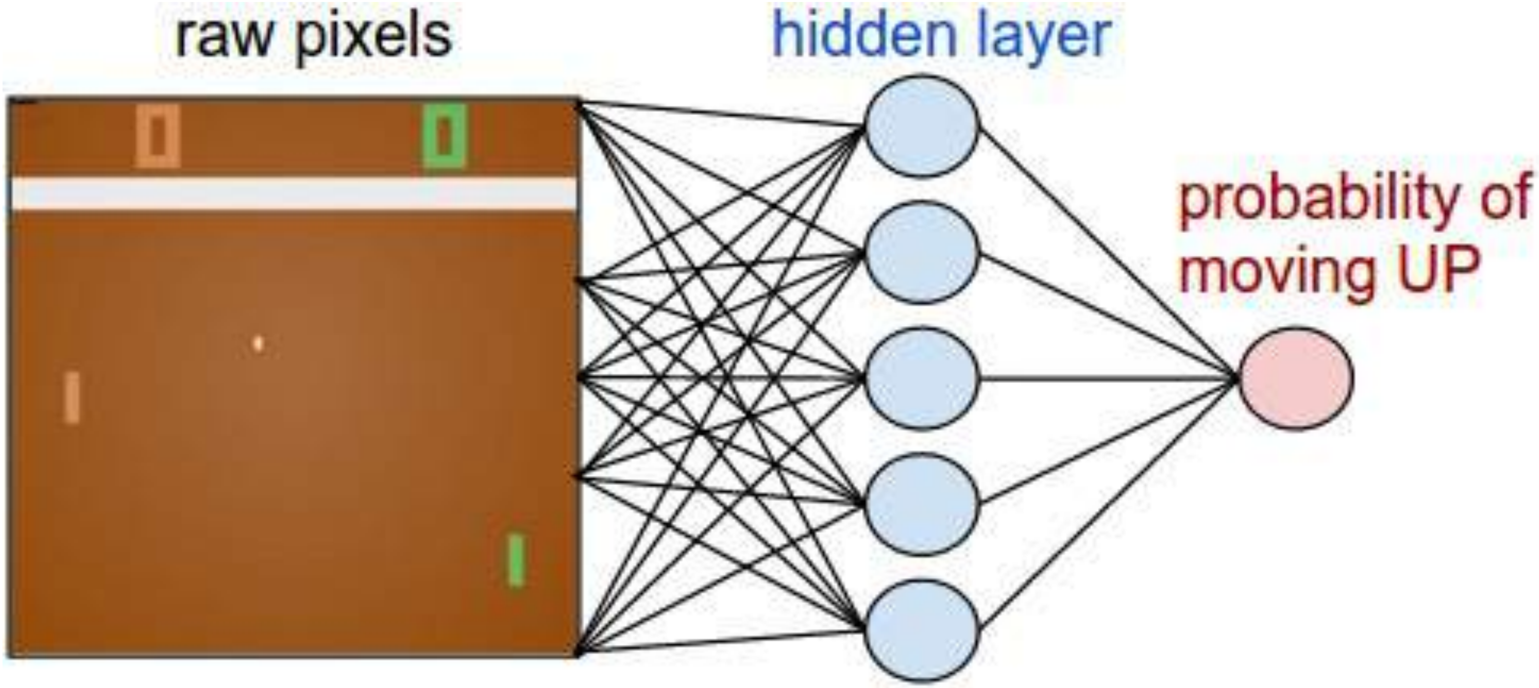

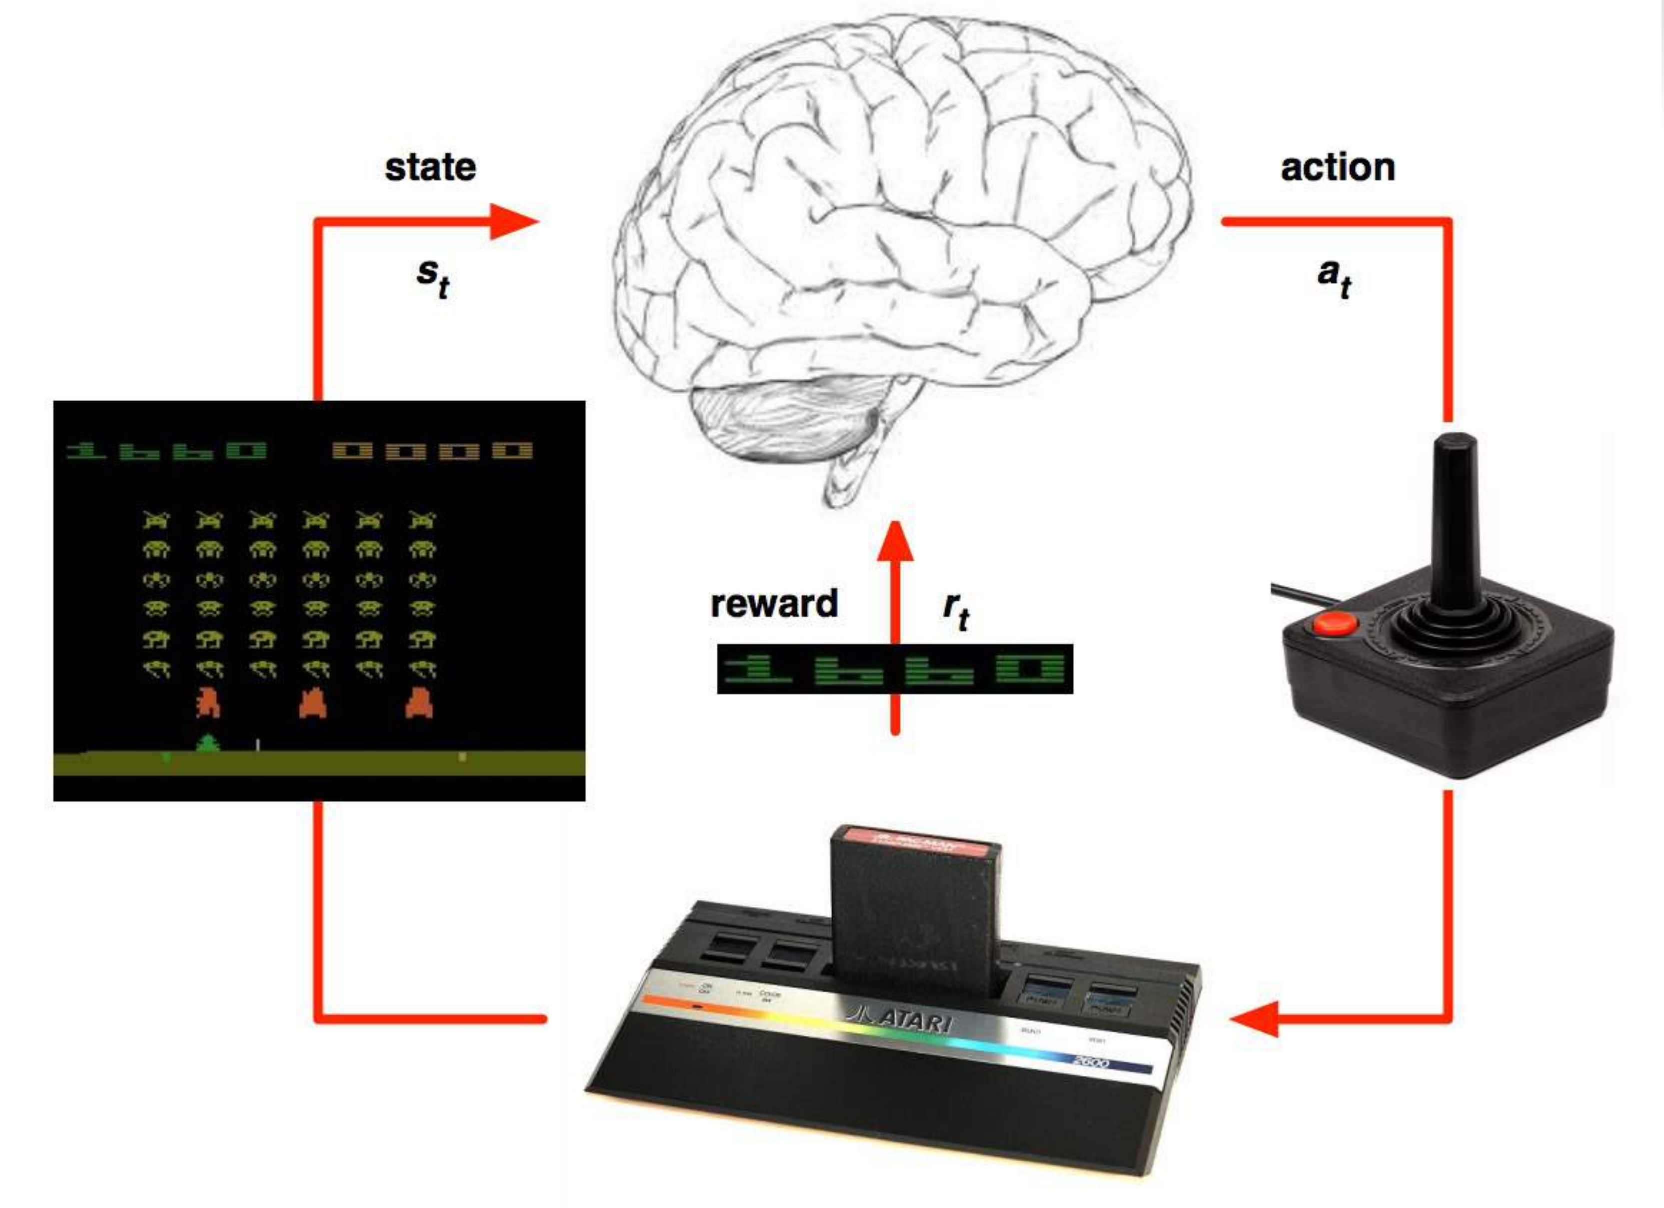
state
s_t
action
a_t
reward
r_t

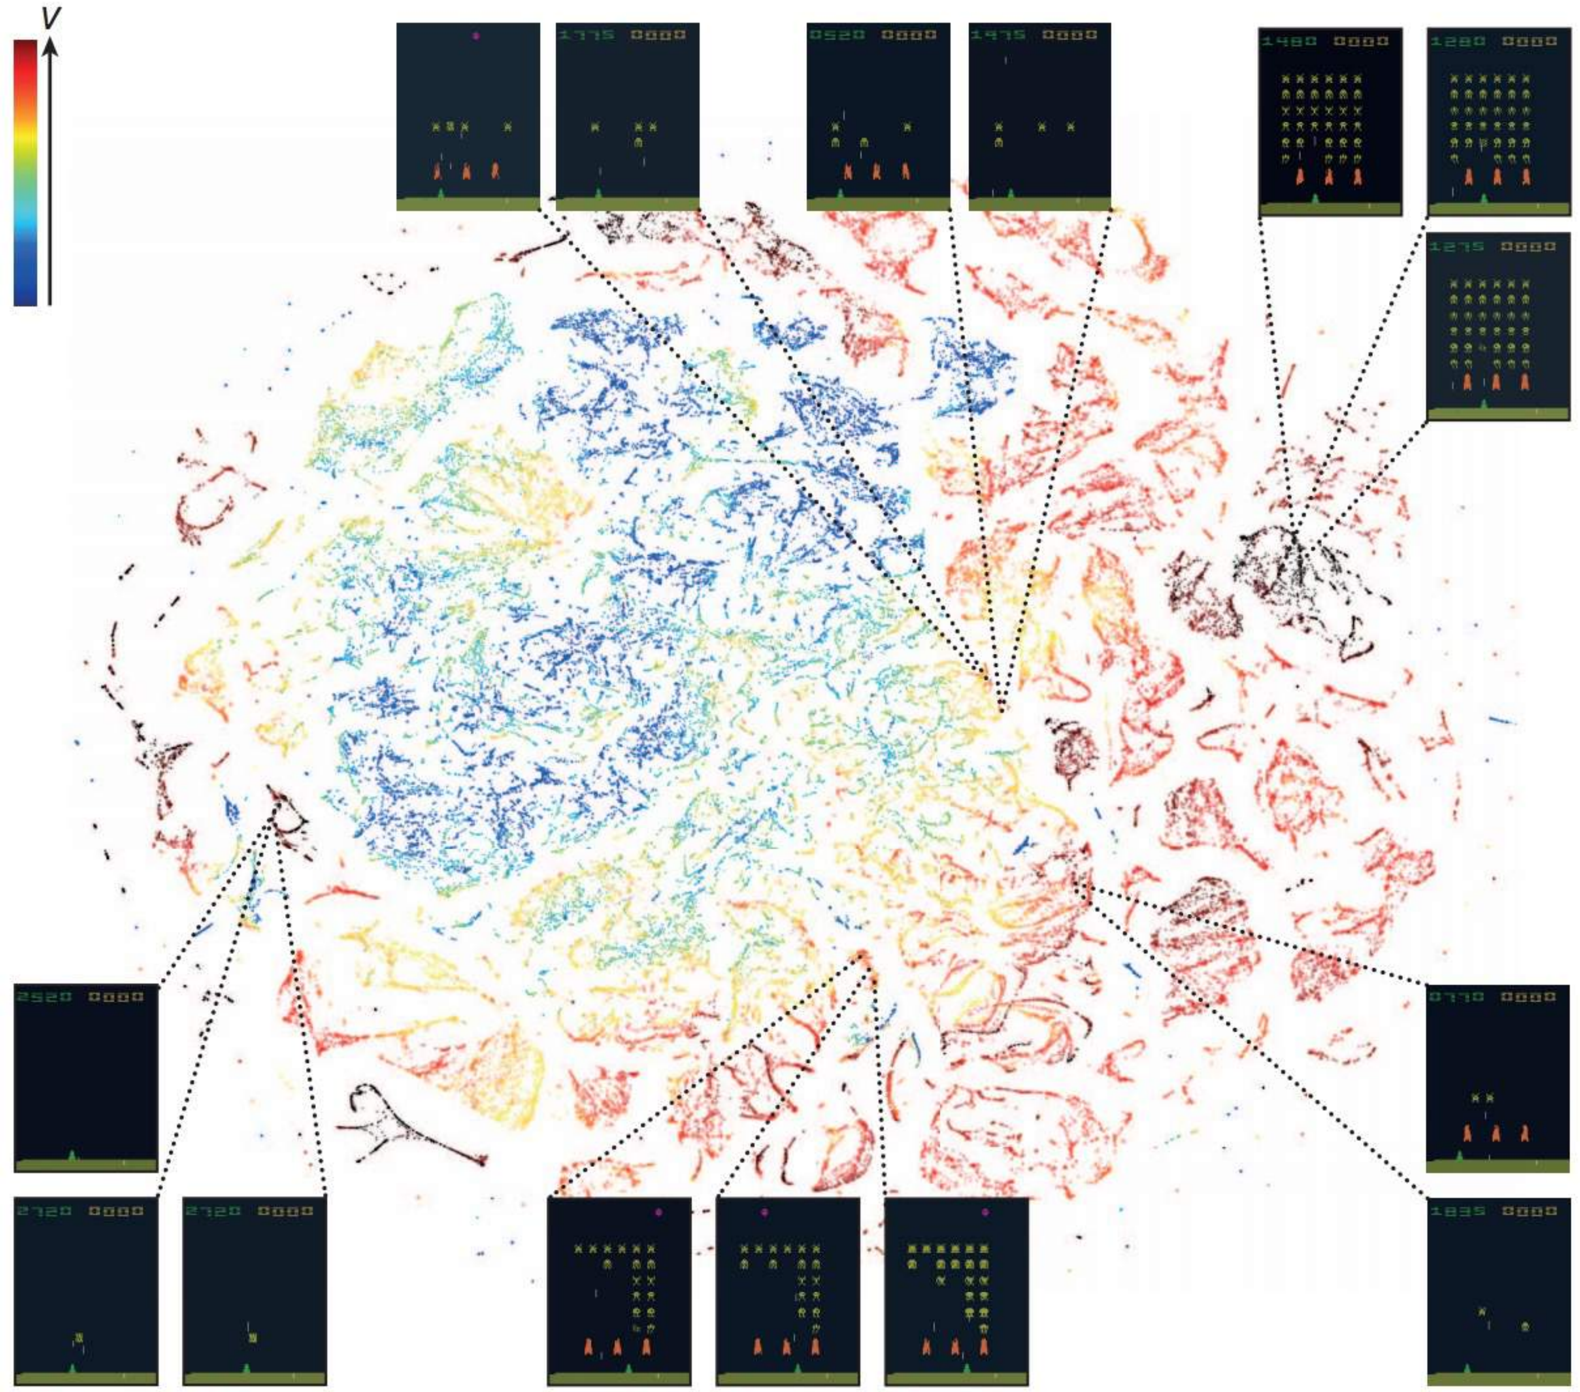
V

Arcades

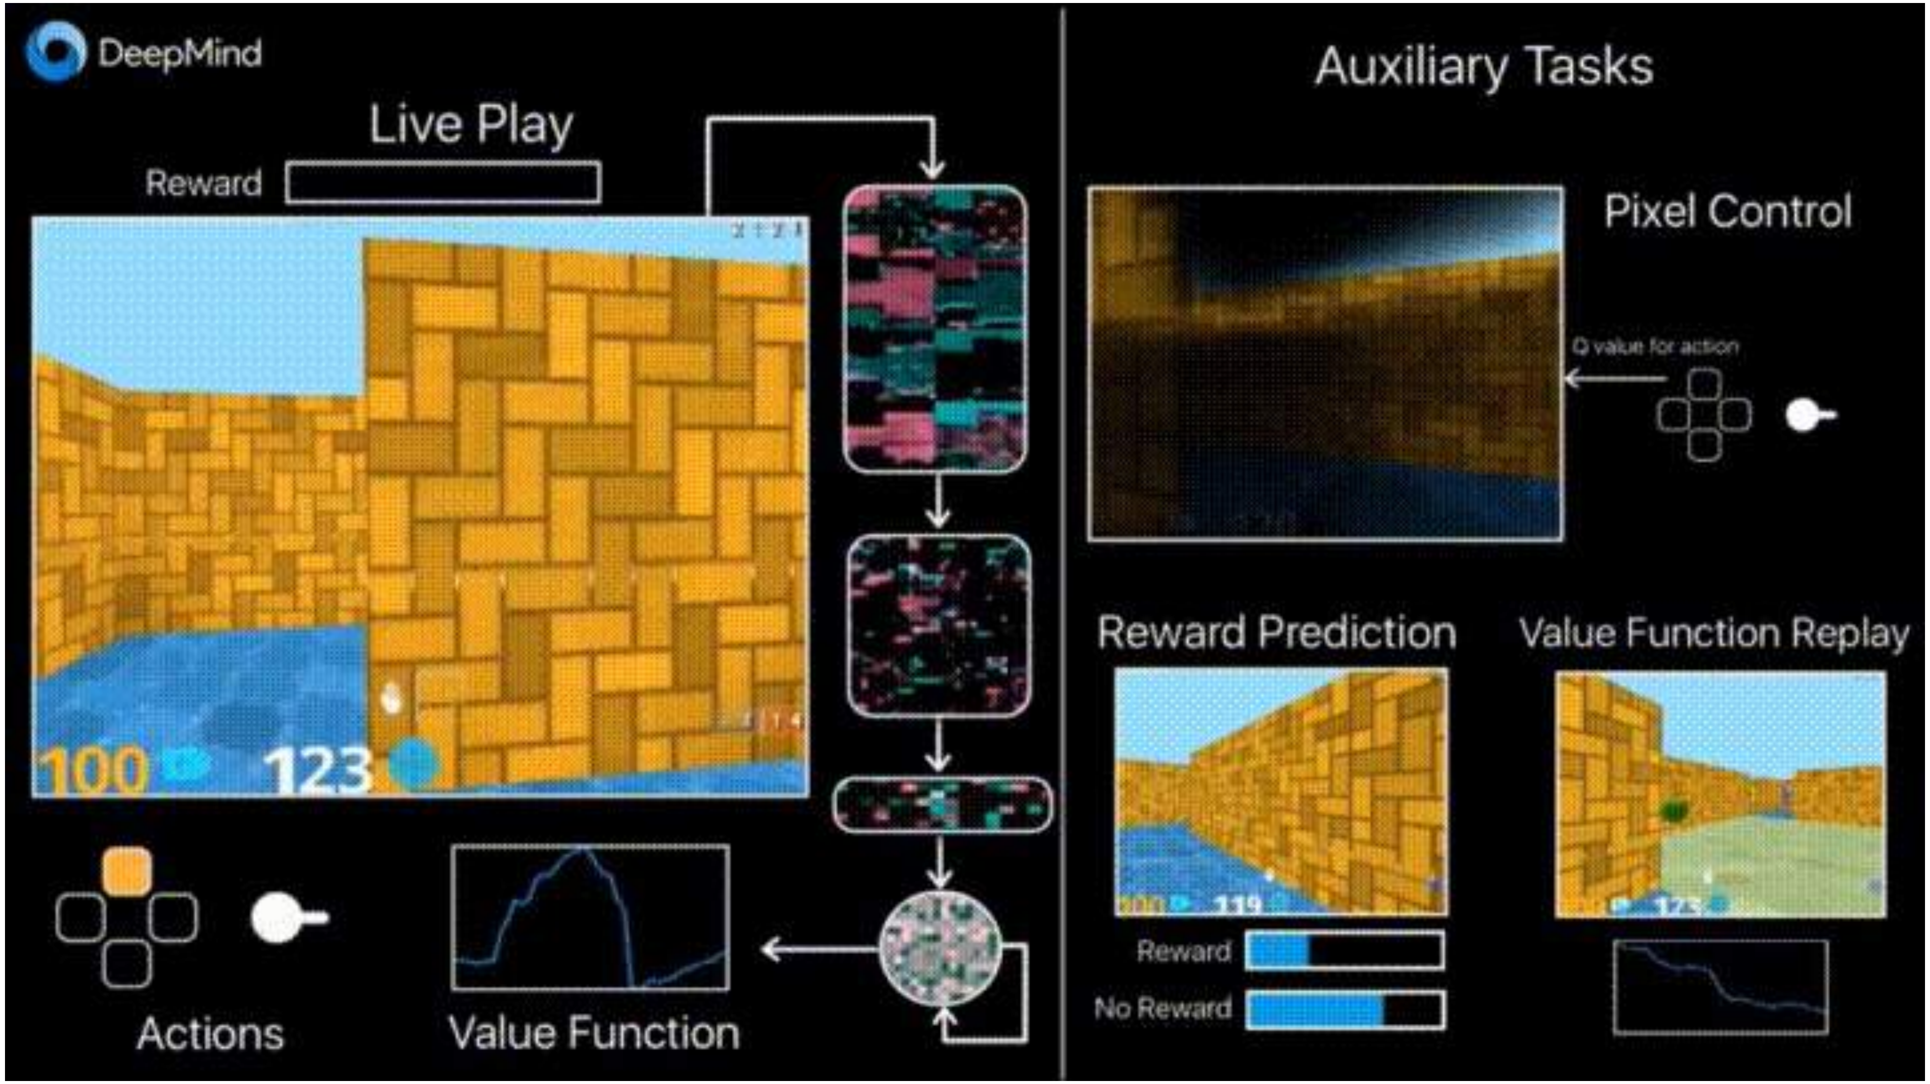

DeepMind
Live Play
Reward
100
123
Actions
Value Function
Auxiliary Tasks
Pixel Control
Q value for action
Reward Prediction
Value Function Replay
Reward
No Reward

Photo Inpainting

Photo Inpainting

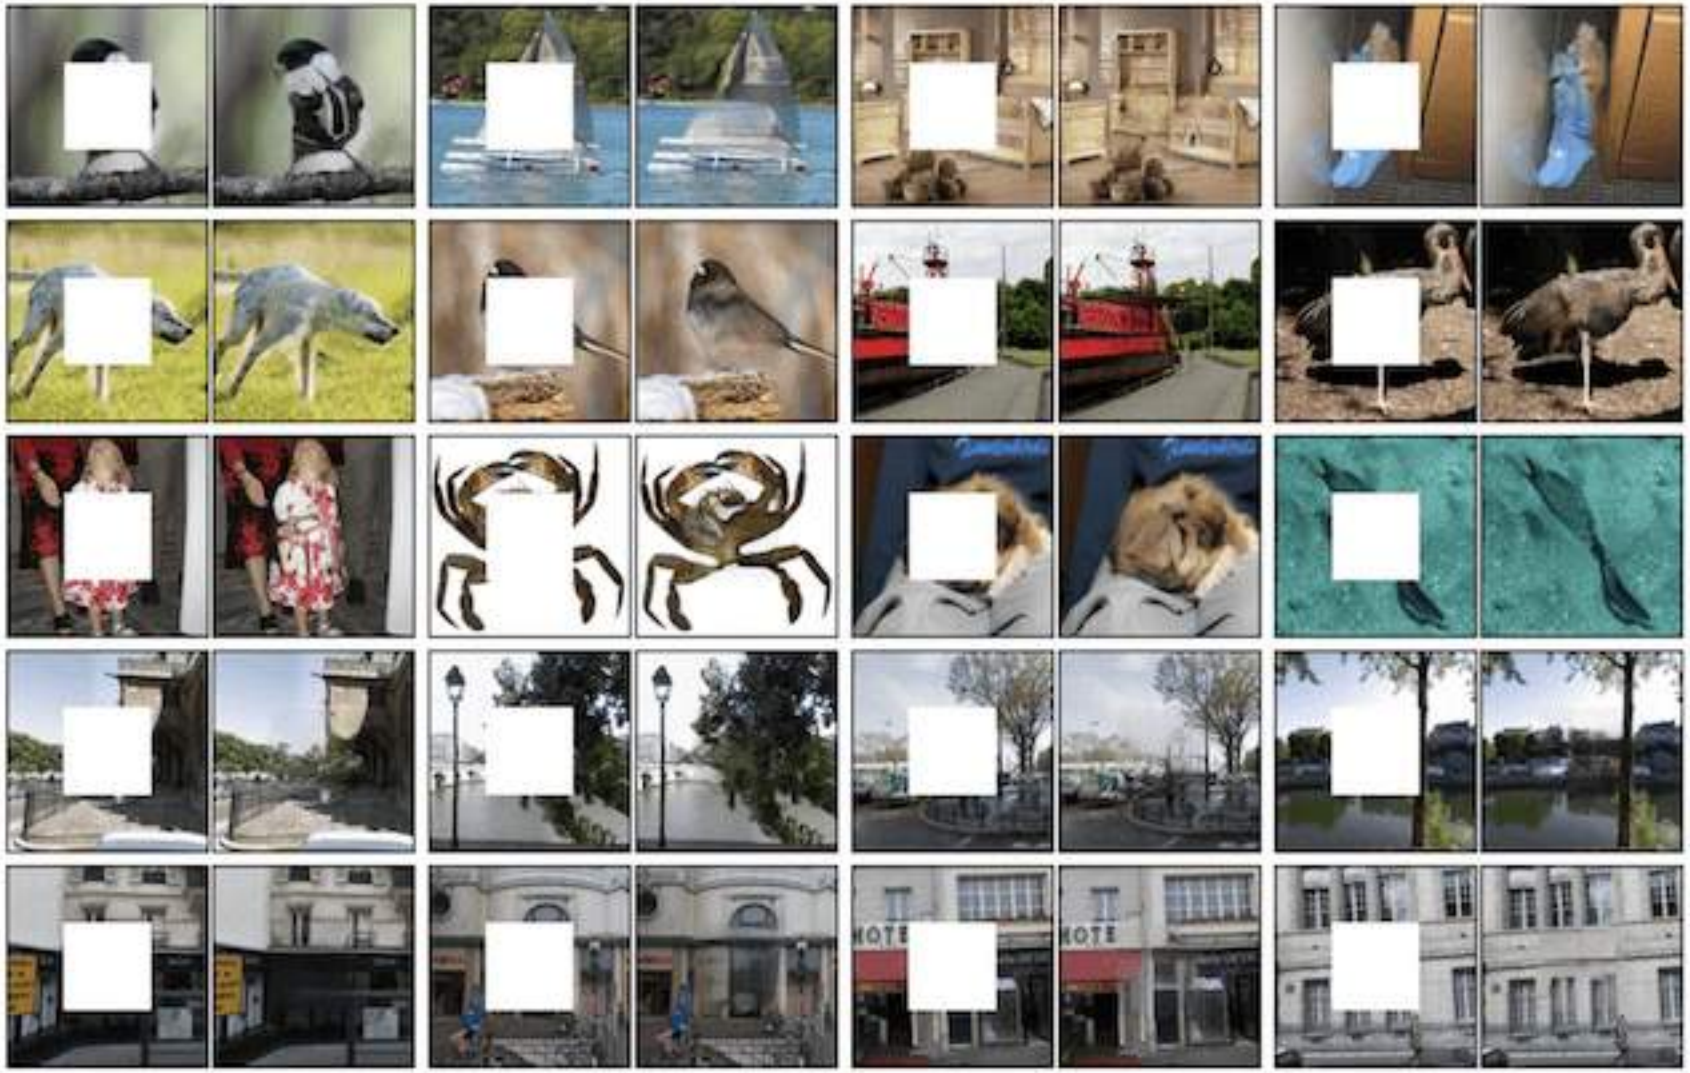

Modelos de lenguaje

Generación de código asistida

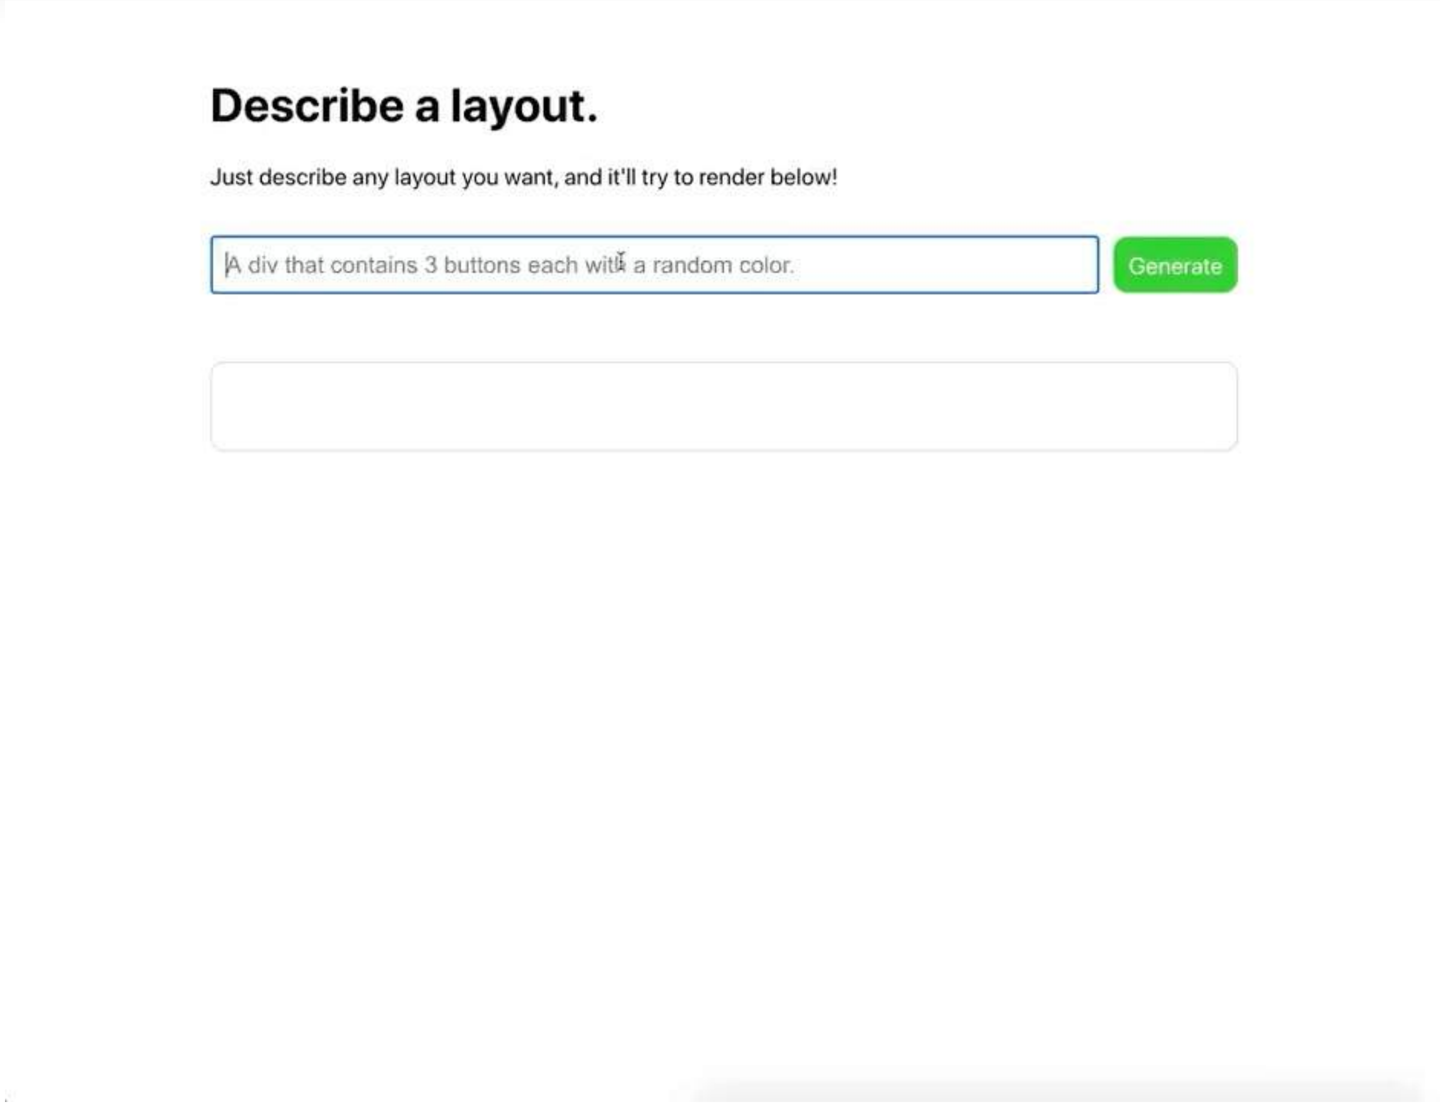

Fuente: https://twitter.com/sharifshameem/status/1282676454690451457

Generación de código asistida

```
// Here are the 2 description:code pairs used to give GPT-3
some context for how to provide a response

// sample 1
description: a red button that says stop
code: <button style={{color: 'white', backgroundColor:
'red'}}>Stop</button>

//sample 2
description: a blue box that contains 3 yellow circles with
red borders
code: <div style={{backgroundColor: 'blue', padding: 20}}><div
style={{backgroundColor: 'yellow', border: '5px solid red',
borderRadius: '50%', padding: 20, width: 100, height: 100}}>
</div><div style={{backgroundColor: 'yellow', borderWidth: 1,
border: '5px solid red', borderRadius: '50%', padding: 20,
width: 100, height: 100}}></div><div style={{backgroundColor:
'yellow', border: '5px solid red', borderRadius: '50%',
padding: 20, width: 100, height: 100}}></div></div>
```

Fuente: https://twitter.com/sharifshameem/status/1282676454690451457

Predicción de apariencia futura

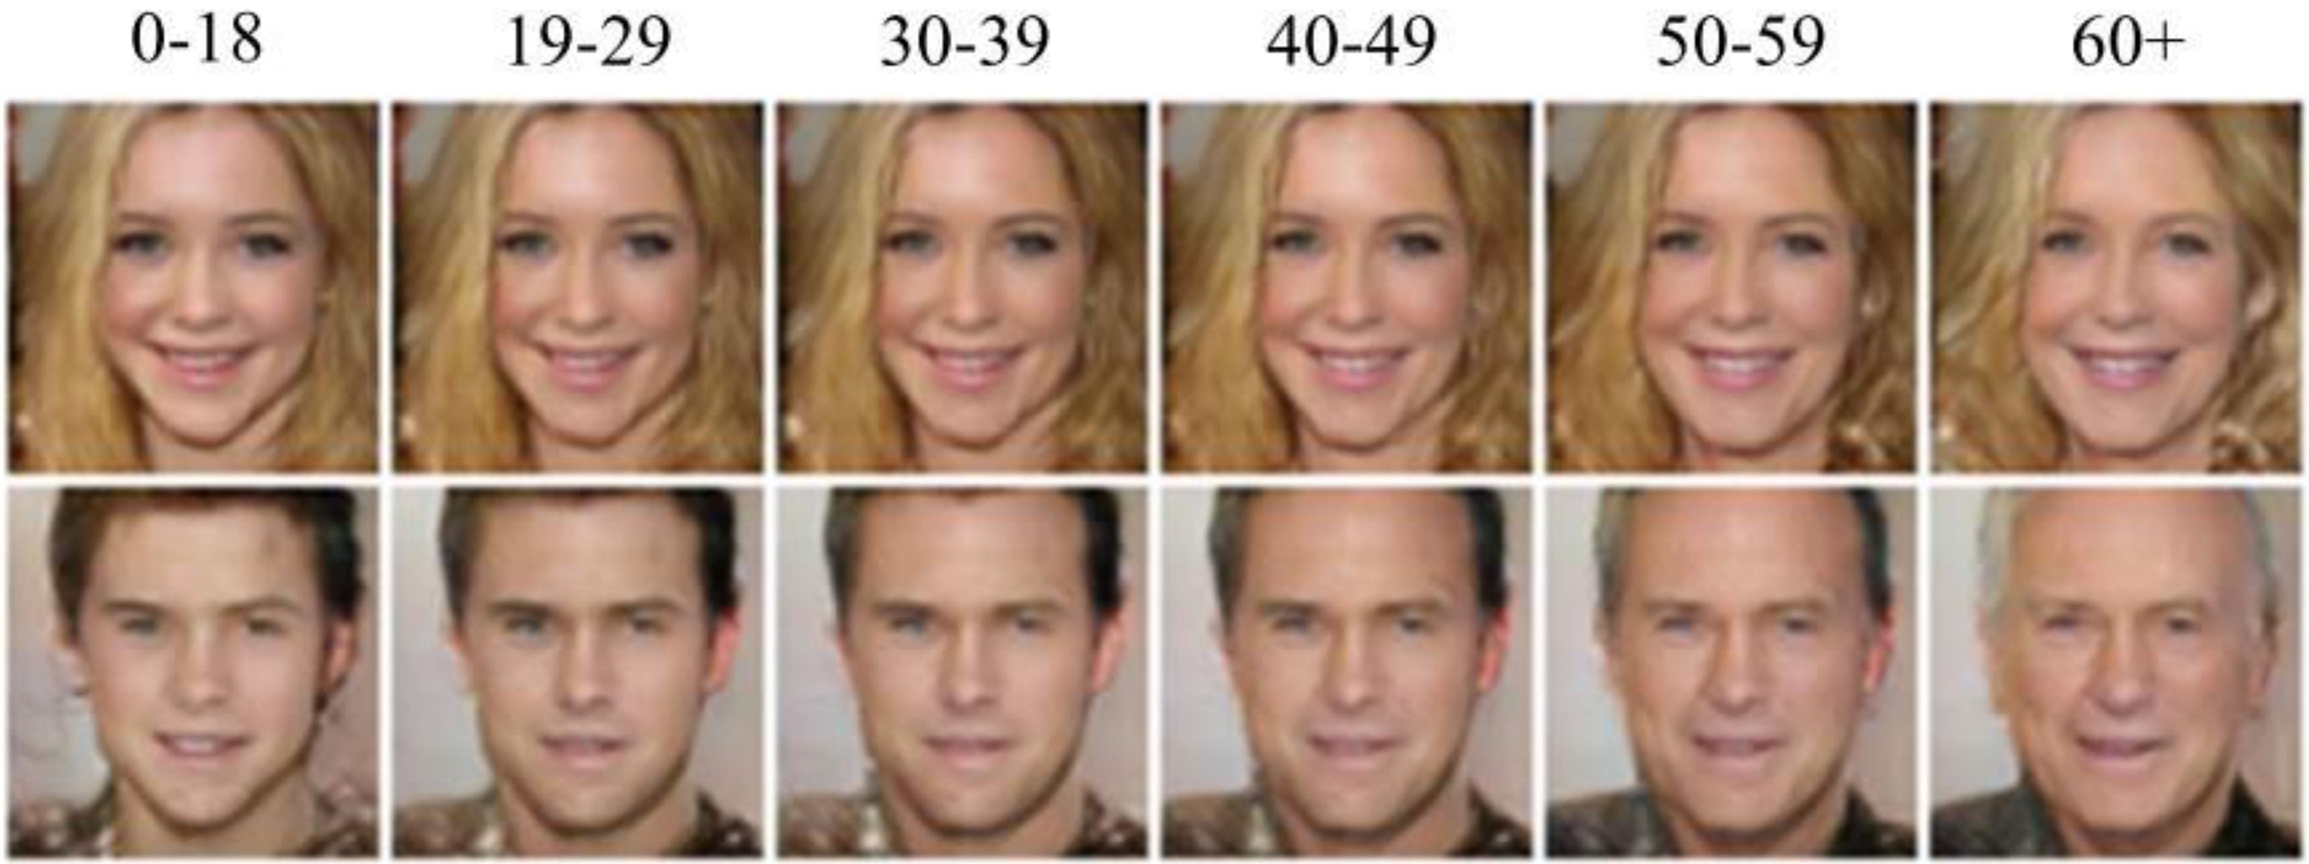
0-18
19-29
30-39
40-49
50-59
60+

Traducción de imágenes

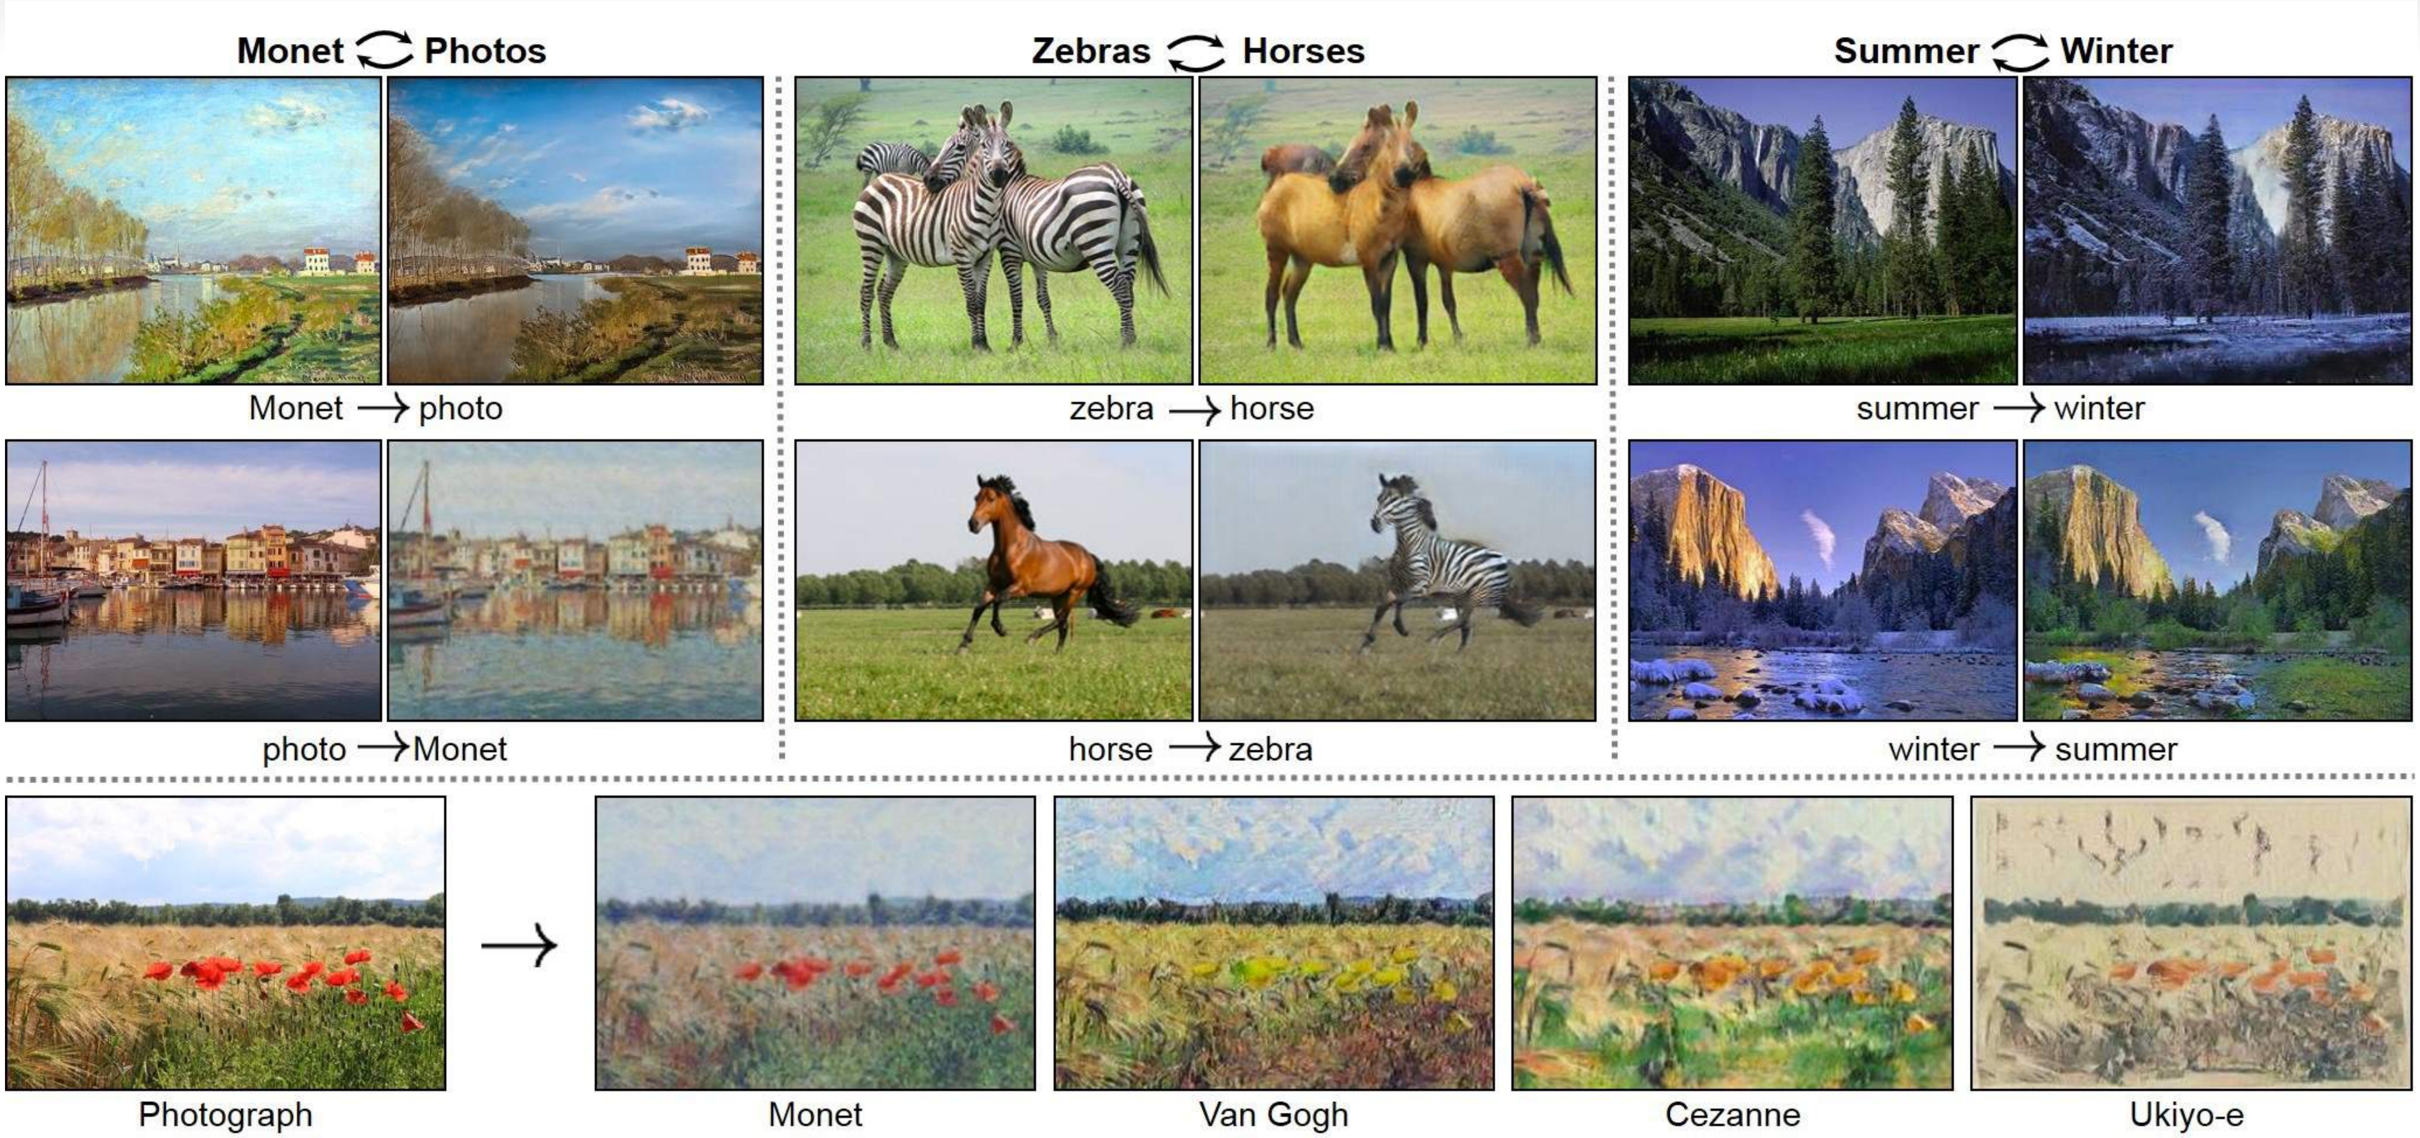
Monet ⇄ Photos
Zebras ⇄ Horses
Summer ⇄ Winter
Monet → photo
zebra → horse
summer → winter
photo → Monet
horse → zebra
winter → summer
Photograph
Monet
Van Gogh
Cezanne
Ukiyo-e


Input
Output
Input
Output

Aerial photo to map　　Map to aerial photo

input　output　input　output

Traducción de imágenes con redes adversarias

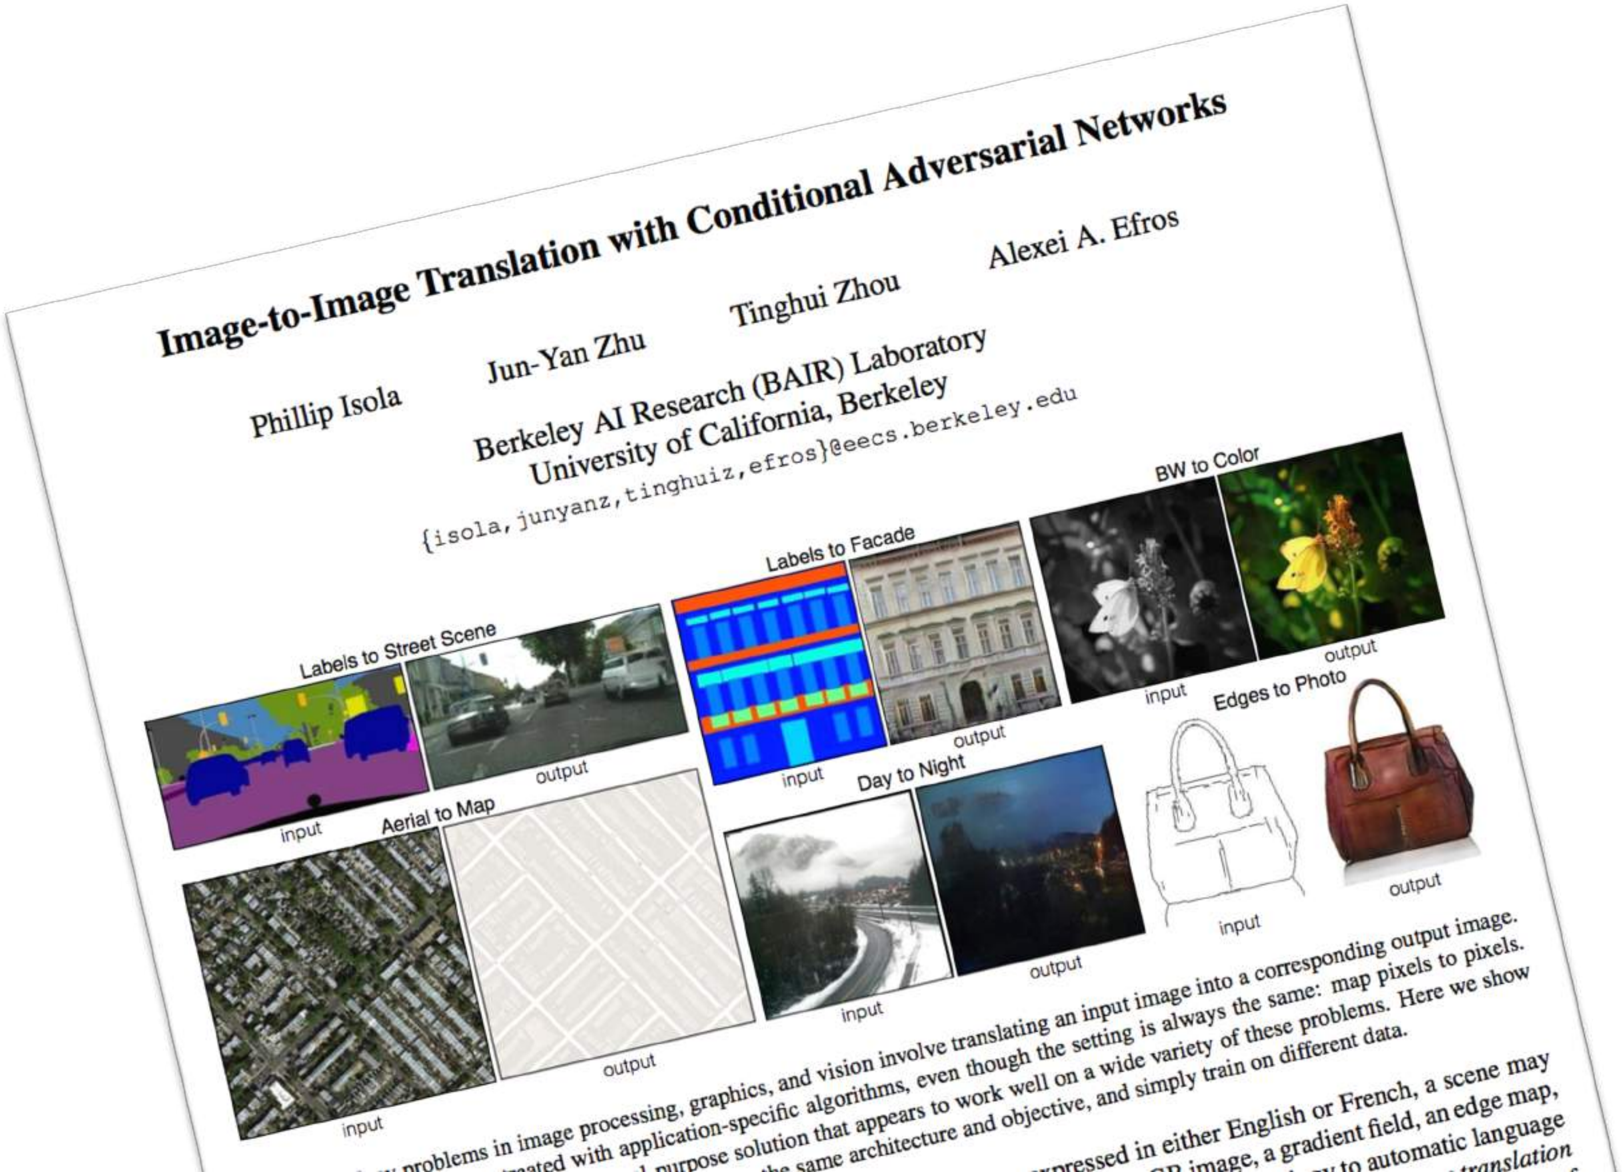

Image-to-Image Translation with Conditional Adversarial Networks

Phillip Isola Jun-Yan Zhu Tinghui Zhou Alexei A. Efros

Berkeley AI Research (BAIR) Laboratory
University of California, Berkeley
{isola,junyanz,tinghuiz,efros}@eecs.berkeley.edu

problems in image processing, graphics, and vision involve translating an input image into a corresponding output image.
treated with application-specific algorithms, even though the setting is always the same: map pixels to pixels.
purpose solution that appears to work well on a wide variety of these problems. Here we show
the same architecture and objective, and simply train on different data.

expressed in either English or French, a scene may
RGB image, a gradient field, an edge map,
to automatic language
translation

Traducción de imágenes con redes adversarias

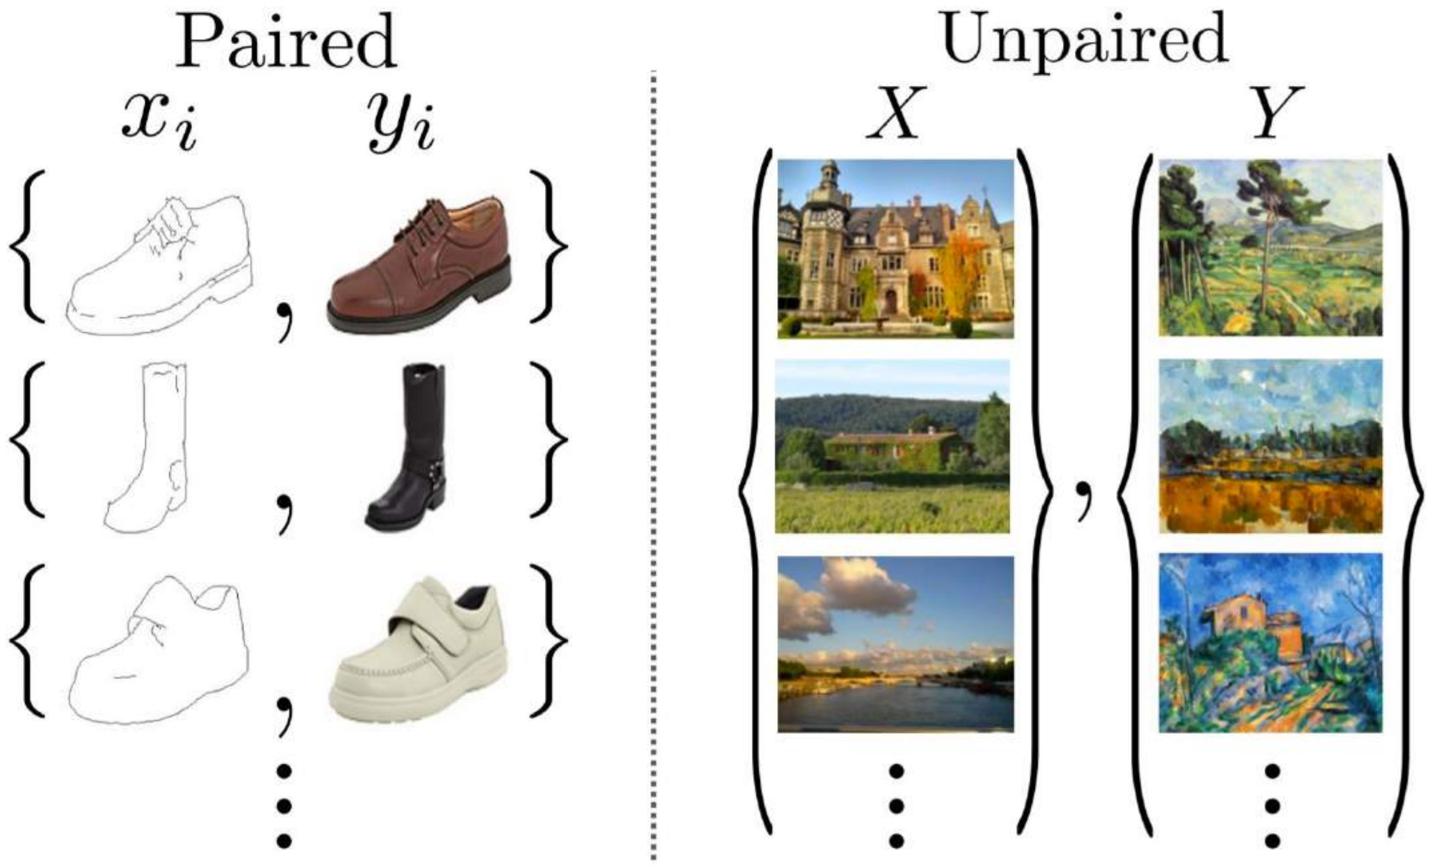

Traducción de imágenes con redes adversarias

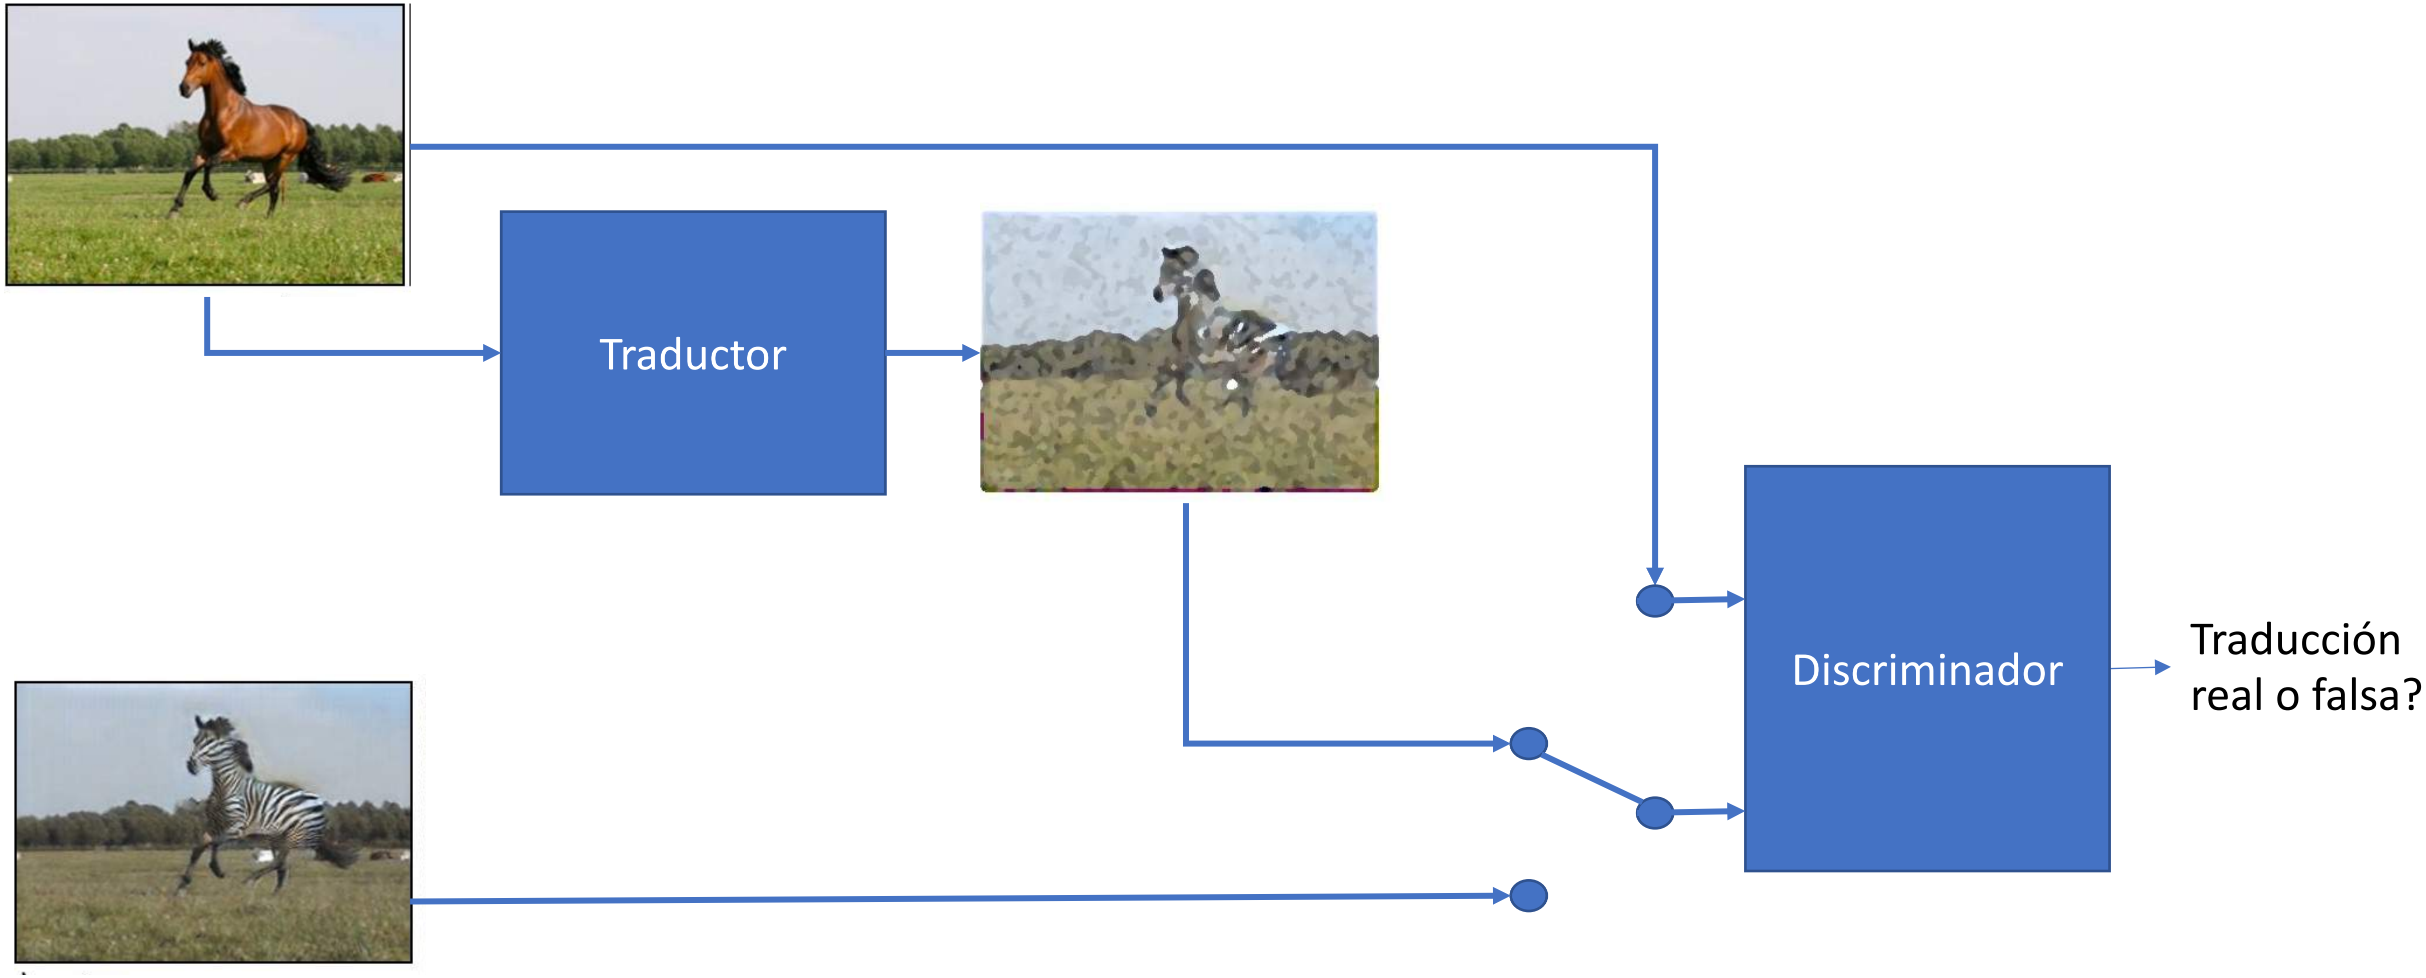

Generación de rostros realistas

thispersondoesnotexist.com

Generación de música

MuseNet

We've created MuseNet, a deep neural network that can generate 4-minute musical compositions with 10 different instruments, and can combine styles from country to Mozart to the Beatles. MuseNet was not explicitly programmed with our understanding of music, but instead discovered patterns of harmony, rhythm, and style by learning to predict the next token in hundreds of thousands of MIDI files. MuseNet uses the same general-purpose unsupervised technology as GPT-2, a large-scale transformer model trained to predict the next token in a sequence, whether audio or text.

Clasificación de patologías en imágenes médicas

Clasificación de patologías en imágenes de rayos X

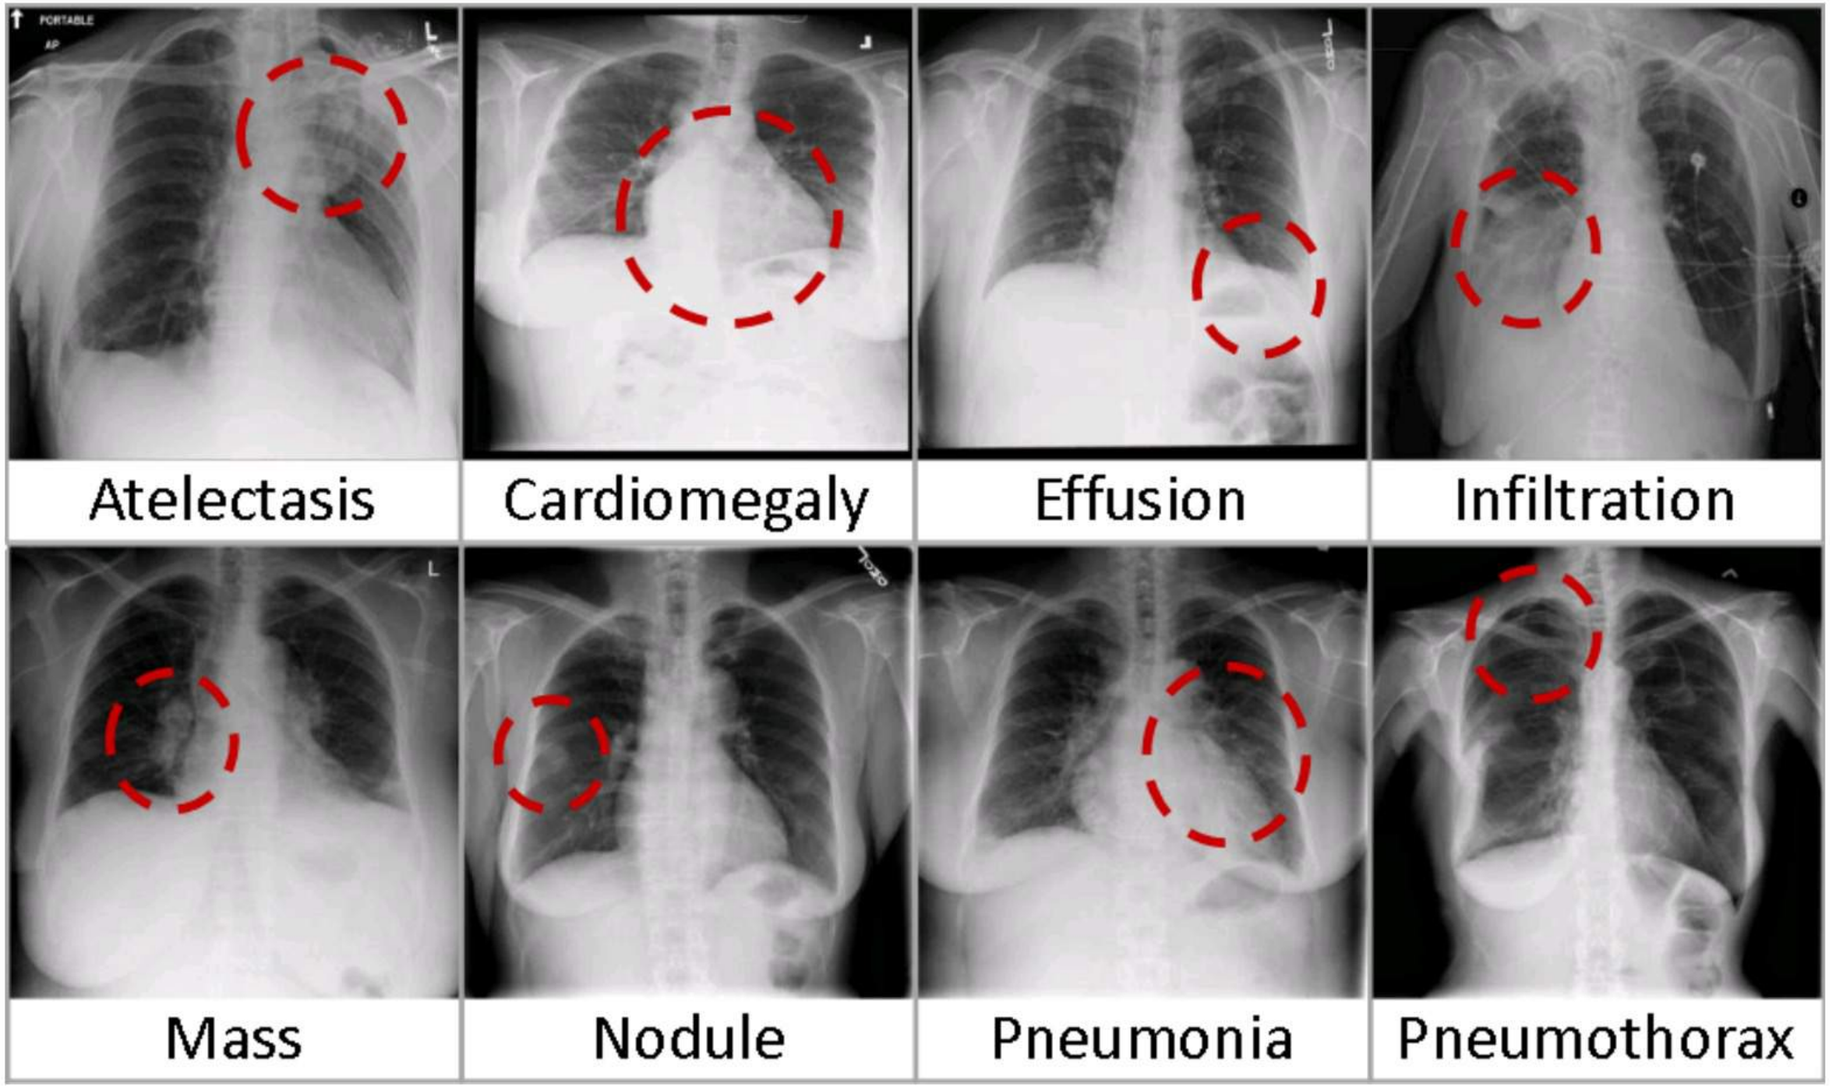

Clasificación de patologías en imágenes de rayos X con localización automática

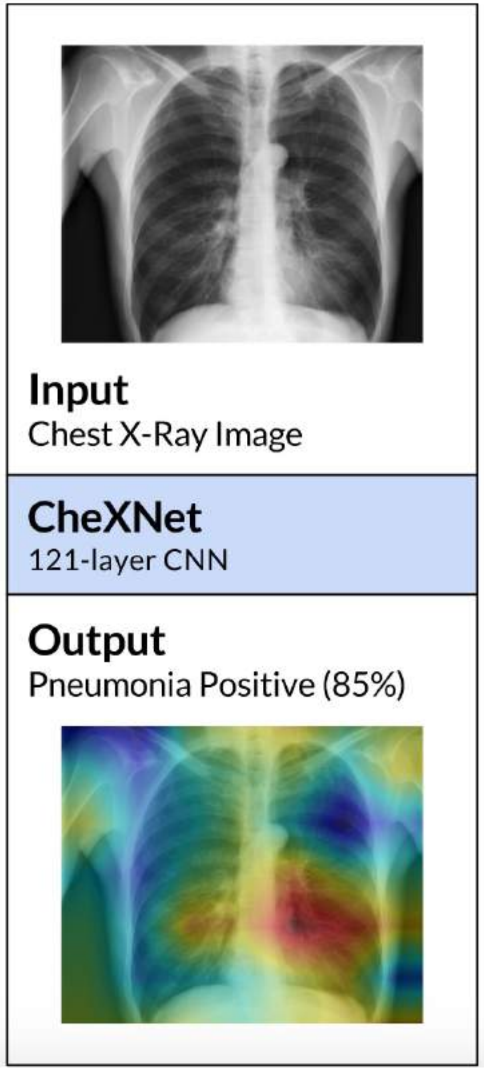

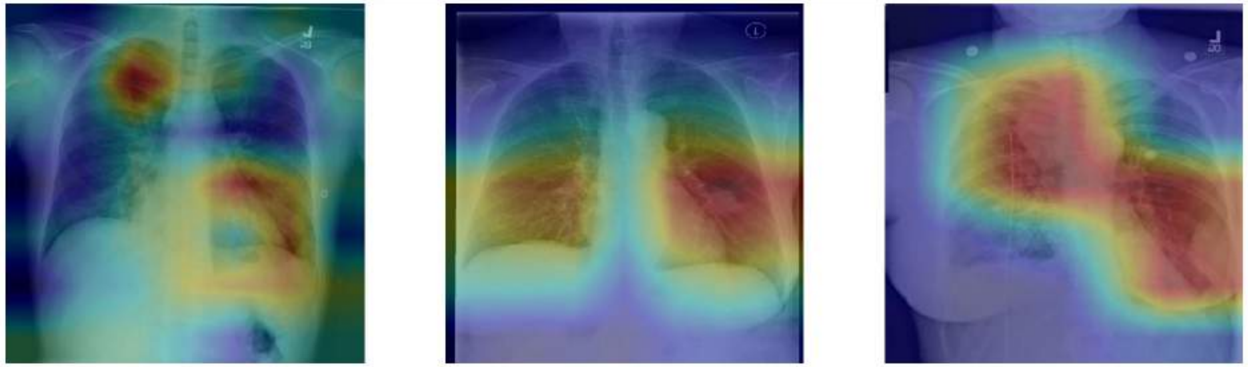

(a) Patient with multifocal community acquired pneumonia. The model correctly detects the airspace disease in the left lower and right upper lobes to arrive at the pneumonia diagnosis.

(b) Patient with a left lung nodule. The model identifies the left lower lobe lung nodule and correctly classifies the pathology.

(c) Patient with primary lung malignancy and two large masses, one in the left lower lobe and one in the right upper lobe adjacent to the mediastinum. The model correctly identifies both masses in the X-ray.

(d) Patient with a right-sided pneumothroax and chest tube. The model detects the abnormal lung to correctly predict the presence of pneumothorax (collapsed lung).

(e) Patient with a large right pleural effusion (fluid in the pleural space). The model correctly labels the effusion and focuses on the right lower chest.

(f) Patient with congestive heart failure and cardiomegaly (enlarged heart). The model correctly identifies the enlarged cardiac silhouette.

Video resumen

Video

https://www.youtube.com/watch?v=1G0e-mR9a4k&ab_channel=NVIDIA

Clase 1

Introducción al aprendizaje profundo

Enzo Ferrante

eferrante@sinc.unl.edu.ar

@enzoferrante