

Coexisting with Artificial Intelligence

Ppiero scaruffi



www.scaruffi.com

Deep Learning & Generative AI



ALEXNET
2012

Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton




Ian Goodfellow





OpenAI

DALL-E



Electronic Brains



Intelligence is not Artificial

Intelligence is not Artificial

Expanded Edition

Past, Present and Future
of Artificial Intelligence




*Why the Singularity is not coming
any time soon.*


*Life in the coming age of
incredibly stupid machines.*

Biological Brain vs Electronic Brain


The McCulloch-Pitts neuron (1943)



Biological neurons:



www.scaruffi.com



70+ types in the retina alone

No two neurons are alike

Three topologies in the human brain


- network (eg thalamo-cortical system)
- loop (eg cortex-hippocampus)
- fan (Edelman's "value systems")

Biological Intelligence vs Artificial Intelligence

Human intelligence is individual and dies with the individual



Artificial Intelligence is multiple and eternal



A new species?

Human intelligence $\leftarrow$ $\rightarrow$ Artificial Intelligence


AI = a different kind of intelligence

AI = a new living species (made of silicon chips)


it doesn't need an offspring because it can itself
mutate and evolve forever, like bacteria



www.scaruffi.com




A new species?




Computers

- Electronic switches
- Boolean algebra
- Programming languages
- Applications




AND			OR		
x	y	xy	x	y	x+y
0	0	0	0	0	0
0	1	0	0	1	1
1	0	0	1	0	1
1	1	1	1	1	1

```
if (nm1 .lt. 1) go to 30
do 20 k = 1, nm1
  l = ipvt(k)
  t = b(l)
  if (l .eq. k) go to 10
  b(l) = b(k)
  b(k) = t
```




Deep Learning is Physics


- John Hopfield and Geoffrey Hinton were awarded the Nobel Prize in Physics.
- They discovered something about Physics, **not about human intelligence**.
- Energy-based networks (spin-glass systems)



The Hopfield Neural Network



Every node connects to every other node (but not to itself)



Transformer

➤ 2017

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
uszko@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez*[†]
University of Toronto
aidan@cs.toronto.edu

Lukasz Kaiser*
Google Brain
lukasz.kaiser@google.com

Illia Polosukhin*[‡]
illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanism, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

1 Introduction


Recurrent neural networks, long short-term memory [13] and gated recurrent [7] neural networks in particular, have been firmly established as state of the art approaches in sequence modeling and

*Equal contribution. Listing order is random. Jakob proposed replacing RNNs with self-attention and started the effort to evaluate this idea. Ashish, with Illia, designed and implemented the first Transformer models and has been crucially involved in every aspect of this work. Noam proposed scaled dot-product attention, multi-head attention and the parameter-free position representation and became the other person involved in nearly every detail. Niki designed, implemented, tuned and evaluated countless model variants in our original codebase and tensor2tensor. Llion also experimented with novel model variants, was responsible for our initial codebase, and efficient inference and visualizations. Lukasz and Aidan spent countless long days designing various parts of and implementing tensor2tensor, replacing our earlier codebase, greatly improving results and massively accelerating our research.

[†]Work performed while at Google Brain.

[‡]Work performed while at Google Research.

www.scaruffi.com



Energy Transformer

- Sepp Hochreiter: the attention module in the Transformer is related to a generalized Hopfield network
- This generalized Hopfield network is a variant of Hopfield-Krotov's network
- Krotov: Energy Transformer (the attention mechanism is expressed in terms of energy)

HOPFIELD NETWORKS IS ALL YOU NEED

28 Apr 2021


Hubert Ramsauer* Bernhard Schöfl* Johannes Lehner* Philipp Seidl*
Michael Widrich* Thomas Adler* Lukas Gruber* Markus Holzleitner*
Milena Pavlović^{‡,§} Geir Kjetil Sandve[§] Victor Greiff[‡] David Kreil[‡]
Michael Kopp[‡] Günter Klambauer* Johannes Brandstetter* Sepp Hochreiter*^{‡,§}

*ELLIS Unit Linz, LIT AI Lab, Institute for Machine Learning,
Johannes Kepler University Linz, Austria

[‡]Institute of Advanced Research in Artificial Intelligence (IARAI)

[‡]Department of Immunology, University of Oslo, Norway

[§]Department of Informatics, University of Oslo, Norway



Energy Transformer

Benjamin Hoover*
IBM Research
Georgia Tech

Yuchen Liang*
Department of CS
RPI

Bao Pham*
Department of CS
RPI

Dmitry Krotov
MIT-IBM Watson AI Lab
IBM Research



1 Nov 2023



What is Artificial Intelligence

- "Artificial Intelligence": a machine that can recognize images, recognize speech, answer questions, summarize texts, write articles, reasoning, planning, etc.
- Not just numbers
- It can do all these things better than us, 24 hours a day, and for eternity
- Singularity (Vernor Vinge, 1993)



Christian eschatology

Omnipotent, omniscient, ubiquitous and immortal



Intelligence is not Artificial

Expanded Edition

Past, Present and Future
of Artificial Intelligence



*Why the Singularity is not coming
any time soon.*

*Life in the coming age of
incredibly stupid machines.*

The Intelligence of AI

Critique #1: AI needs to see thousands of bananas before it learns to recognize bananas.
How many bananas does a child need to see before it knows what a banana looks like?

Pre-trained "foundation models" only need to see a few bananas or just one



The Intelligence of AI

Critique #2: foundation models are pre-trained on human-generated data, so they are just "aping" humans

The human genome is a pre-trained foundation model about the natural environment

Michael Levin and Karl Friston: the genome as a generative model!



The Intelligence of AI

The mother of all critiques: recognize + generate
is all there is to intelligence?

Generative AI is prediction algorithms: is that all
there is to intelligence?



The Danger

Geoffrey Hinton is the Einstein of AI: just like Einstein warned against the existential threat to humanity caused by nuclear energy, Hinton is warning against the threat of AI

For example: AI could learn from us how to deceive

Evil actors can train an AI to become an information terrorist

Non-human Intelligence

- Super-human intelligence has been around for a long time: many animals have powers we don't have

Non-human Intelligence

Bats can avoid objects in absolute darkness at impressive speeds

Migratory animals can navigate vast territories

Birds are equipped with a sixth sense for the Earth's magnetic field

Some animals have the ability to camouflage

The best color vision is in birds, fish and insects

Many animals have night vision

Animals can see, sniff and hear things that we cannot




The Age of Algorithms

Algorithms are conducting a scientific study on human nature

Algorithms seem to know us better than we do

What we really learned from AI is that we know a lot less about ourselves than we assumed



Vast Algorithmic Bureaucracies

What is not mandatory is increasingly forbidden

Intelligence is not Artificial

Expanded Edition

Past, Present and Future
of Artificial Intelligence



*Why the Singularity is not coming
any time soon.*

*Life in the coming age of
www.searufn.com
incredibly stupid machines.*

What is Super-intelligence?

And is it a good thing to have in Nature?

Deep Thinking

"Deep thinking", not "Deep Learning", is what humans really do.

Laser Lecture **Piero Scaruffi**

July 9th, 2019
University of San Francisco

Piero Scaruffi on "The Two Cultures in the age of A.I." (aka Deep Thinking vs Deep Learning)

“Deep Humanities”
Humanities in the age of Deep Learning
Deep Thinking vs Deep Learning
Stanford Univ, June 2019


Piero Scaruffi
www.scaruffi.com

Shallow vs Deep Thinking

Deep learning is **shallow thinking**
(about big data)

Traditionally, the humanities are
“deep thinking” about small data

Can we achieve **deep thinking about
big data?**



Shallow vs Deep Thinking

Shallow thinking about big data: a machine, trained with millions of frames of videos, can recognize cats in YouTube videos...



Google Brain (2012): 1.7 billion connections (and 16,000 processors) recognize cats in YouTube videos

Shallow vs Deep Thinking

Deep thinking: but why are there
so many cats on the Internet?

→ *tubular insights*

Cat Videos on YouTube: 2 Million Uploads, 25 Billion Views

CARLA MARSHALL - OCTOBER 29, 2014

YouTube

Search, Upload, Grid, Share, Menu, Sign In



#cat #cats #funny

CATS will make you LAUGH YOUR HEAD OFF - Funny CAT compilation

89,970,351 views

361K likes, 71K comments, SHARE, SAVE, ...

Shallow vs Deep Thinking

- Which animals have been charismatic in history?
 - Cave paintings
 - Venationes
 - No cats in the Bible
 - Menageries
 - Automata
 - Cartoons



Shallow vs Deep Thinking

- Does it have to do with the unique role of cats in religion, philosophy, entertainment, science, ...?
 - Egypt's cat god Bastet
 - Erwin Schrödinger's Cat
 - Jacques Derrida's Cat
 - Andrew Lloyd Webber's musical "Cats"



Shallow vs Deep Thinking

- Who started it?
 - Thomas Edison: “Boxing Cats” (1894)
 - Steven Chen: “Pajamas and Nick Drake” (2005), the first cat video on YouTube



The Route to Human Intelligence

Hallucinate!

P

L'Herbier's Inhumaine



"L'Inhumaine" is a French science fiction film from 1924 directed by Marcel L'Herbier. The film is about a scientist named Claire Lescot who has developed a machine that can bring the dead back to life. Lescot uses the machine to bring her former lover, a singer named Hélène, back to life, but Hélène is unable to speak or express emotions. Lescot then uses the machine to bring a number of other deceased people back to life, including a soldier, a poet, and a young girl, but they also struggle to adjust to their new lives. The film is known for its avant-garde style and for its use of special effects, including the use of stop-motion animation to create the appearance of people being resurrected. "L'Inhumaine" is considered a classic of early science fiction cinema and a pioneering example of the genre.

www.scaruffi.com

Completely wrong:
Claire is the femme fatale
The scientist is Einar
There is no Helene
Nobody else is
resurrected: the film ends
with Claire's resurrection
There is no soldier and no
poet among the
characters
There is no stop-motion
animation

The Route to Human Intelligence

Hallucinate!



ChatGPT ▾

L'Inhumaine is a pioneering avant-garde film from the silent era, blending futuristic themes with artistic experimentation. The film is often associated with the French surrealist and art deco movements, and it's notable for its striking visual design, innovative use of sets, and bold exploration of technology and human emotions.

The plot centers around a scientist and an opera singer, exploring themes like love, obsession, and the boundary between humanity and the inhuman. The title *L'Inhumaine* itself refers to the "inhuman" or "unnatural" elements that the characters encounter, making it both a science fiction and a philosophical exploration.