

An Accident in Paradise

By Paulo Gastón Flores

Thursday, October 5, 2023

During 2022, weak artificial intelligence burst onto the creative scene in a massive way. The emergence of ChatGPT and several programs that transform natural language into images enabled, respectively, the generation of coherent dialogues between humans and software, and the "creation" of high-quality images through a "prompt" (or text input). Given the brutal irruption of AI onto the public agenda, it seems necessary to first review the topic at a conceptual and general level.

The acceleration of time. The acceleration of calculation

An era of shortcuts and immediacy.

The speed of hyper-technological events overwhelms us. We are slaves to a frenetic pace. The devices and algorithms designed to streamline tasks foster a 24/7 work schedule and promote the immediacy of information transmission.

More than twenty years ago, in *Liquid Modernity*, the Polish philosopher Zygmunt Bauman described a society characterized by fluidity, constant movement, and unpredictability. It was then necessary to think of society in dynamic terms, in terms of processes, and no longer in static terms.^[1]

Liquid modernity thus contrasted with "solid" modernity. In the latter, rationalism provided physical solutions through the design of an order manifested in large factories, enormous infrastructures, and fortresses that were intended to last as long as medieval cathedrals.

In *The Art of the Motor*, the French architect and urban planner Paul Virilio argues that speed and acceleration are the driving forces behind contemporary culture, and that technology and the media have transformed our perception of space and time.^[2]

Voice assistants, *chatbots*, augmented reality, the routes selected by our cars' geolocation, social media monitoring, and medical diagnosis are just some of the technological devices and activities where AI has a decisive influence on our lives. How often do we contradict the route determined by GPS—powered by cloud-based AI—to our destination?

Our phones now come equipped with dedicated AI processors called NPUs (**Neural Processing Units**). These accelerate the complex calculations used by AI. When we use the camera, during automatic scene selection, or when adjusting white balance, the NPU is involved. AI has been silently and unnoticed in our lives.

Finally, to understand the rise of AI, we must consider that in recent years, computing power has increased exponentially. At the same time, *cloud computing and storage* have facilitated the training of sophisticated AI models based on the rental of information technologies.

- What is AI?

Let's look at IBM's definition: "Artificial intelligence harnesses computers and machines to mimic the problem-solving and decision-making capabilities of the human mind."^[3]

Reinterpreting how the brain works has been key to its development. Software receives data (either pre-prepared or collected through its own sensors, such as a camera), processes it, and responds accordingly.

AI systems are capable of adapting their behavior, analyzing the effects of previous actions, and working autonomously with minimal or no human intervention.

- Historical background

Alan Turing was a British mathematician and scientist, considered the father of AI due to his pioneering work on computing machines. Turing designed an abstract computing machine consisting of a memory and a scanner that could read and write symbols on a tape. This machine, known as the universal Turing machine, is the theoretical model for all modern computers and can simulate any computable algorithm.^[4] Turing also posed the question "Can machines think?" and proposed a method for evaluating a machine's intelligence based on its ability to mimic human behavior in conversation. This method, known as the Turing test, is one of the most widely used criteria for measuring the progress of AI.^[5]

Marvin Minsky was an American scientist and mathematician, and is also considered one of the fathers of AI for his pioneering work in computing, neural networks, robotics, and cognitive science. Minsky coined the term "artificial intelligence" in 1956, along with other colleagues, at a landmark conference that marked the beginning of this discipline. He designed and built several machines and robots capable of learning, reasoning, and solving problems, such as Snarc, the first neural network simulator; the Advanced Artificial Arm, which could execute fourteen different movements; and the Constructor, a robot that used visual feedback for guidance.

Minsky also founded and directed the MIT Artificial Intelligence Laboratory, where he trained several generations of AI scientists and developers, including Ray Kurzweil, Gerald Sussman, and Patrick Winston.^[6]

Ray Kurzweil ultimately developed several innovative AI-related products, such as music synthesizers, speech recognition systems, book scanners, and reading systems for the blind. He proposed the idea of the technological singularity, which is the moment when AI will surpass human intelligence and merge with it, creating a new species of hybrid beings.^[7]

This singularity is defined by a rapid increase in intelligence, unlike other technologies such as writing: "The Singularity will allow us to transcend the limitations of our biological bodies and brains."

In a short period of time, these scientists contributed significant discoveries and solutions in this field of science. The availability of computing power in the second half of the 20th century limited the development of AI until well into the 21st century, where the growth of this same power and the renewed interest of the scientific and business communities exponentially increased its evolution.

- Basic System Components

Artificial Neurons and Neural Networks

A biological neuron is an electrically excitable, interconnected nerve cell within the brain that processes and transmits information through electrical and chemical signals. Neurons are connected to each other to form a neural network.

What, then, is an artificial neuron?

An artificial neuron is a mathematical function conceived from a model of biological neurons. They are the basic units of an artificial neural network. The first artificial neuron was developed by Frank Rosenblatt and was called a "perceptron".^[8]




Diagram of an artificial neuron

Neural Networks

We could then define a neural network as a set of artificial neuron units connected to each other to transmit signals. The input information passes through the neural network (where it undergoes various operations) producing certain output values (results).




Diagram of an Artificial Neural Network (ANN) with three interconnected layers. Source: Adapted from Isasi and Galván (2004).

There are many types of neural networks (convolutional, adversarial, etc.) and different ways to classify them, but the most successful so far are the so-called transformers. They are very good at natural language processing, understanding context better than other types of networks. OpenAI's recent ChatGPT (Chat Generative Pre-Trained Transformer) application is the best example of this.

Machine Learning Machine learning is a branch of artificial intelligence that allows machines to learn without being explicitly programmed to do so. It is an essential skill for creating systems capable of identifying patterns in data to make predictions. This technology is present in countless applications, such as Netflix and Spotify recommendations, Gmail's smart replies, and the speech capabilities of Siri and Alexa.

In short, machine learning is about pattern recognition, and it can transform a sample of data into a computer program capable of drawing inferences from new datasets for which it has not been previously trained. [\[9\]](#)

Among the main types of machine learning, we can mention:

- Supervised Learning

In supervised learning, training information (datasets) is used. For example, to determine whether an image shows a cat or a dog, a model is trained with thousands of images, and we tell it which ones are of dogs and which ones are of cats (labeling). After many examples, given a new image (without labels), the model will be able to determine

whether it is a cat or a dog. This problem is called classification. What distinguishes supervised learning is that it is trained by providing many examples and, from that, can generalize to new cases.

- Unsupervised Learning

Unlike supervised learning, here we don't have a true value or label. Unsupervised learning models aim to understand and abstract patterns from information directly. For example, an Airbnb model will receive the locations of the houses and determine how to group them. This problem is called clustering. Although it may sound complex, it's very similar to how we think with new information. For example, how do we define what constellations are? From observing the stars, we can deduce certain patterns.

- Reinforcement Learning

In this technique, models learn from experience. In a self-driving car, when it makes a bad decision, it is "punished." Through rewards and punishments, its AI learns to perform the task in the best way. This is a trial-and-error technique to optimize its performance. This is one of the most promising techniques because it does not require large amounts of data.^[10] Unintentionally, every time we play with a new application, for example, replacing faces in a film sequence, we are helping to train an AI, contributing to data labeling and verifying the quality of the results.

It is clear that the databases of images, texts, scientific formulas, etc., used to train artificial intelligences are colossal. These systems appropriate all of human culture created and then process it at extraordinary speeds.

- The Appropriation of Human Culture

In our generalist training as architects and professionals in the construction and design industry, the apprehension of architectural culture, history, and technique helps us create and develop our forms, functional organizations, and spatial qualities.

We build upon the work of others: we interpret historical buildings; we understand our environments; we listen to and interpret a commission. Some may say they design on a blank sheet of paper, but they surely do so using all the culture incorporated into their thinking with every movement of a pencil, mouse, or three-dimensional visualization. The cloud of visual noise begins to take shape as we imagine and advance in our project. The mind is not blank.

We generally consider that good design is a consequence of understanding general culture, interpreting the historical moment, but also applying academic knowledge. Our projects are the organization of knowledge.

The Children of the Sand Crawler

Let's take as an example one of the cult objects in the architectural world of recent decades: the Sand Crawler from Star Wars. We could say that through an operation of morphological analogy,^[11] several contemporary works "could" have a certain similarity to the sand crawler, half vehicle, half building.



Sandcrawler. Star Wars Episode IV: A New Hope, 1977



From left to right: 1) Casa da Música in Porto, Portugal. Rem Koolhaas. 2005, 2) Single-family home. Moon Hoon Studio. Yongin-si, South Korea. 2013, 3) "Sandcrawler" building, headquarters of Lucasfilm Singapore, by Andrew Bromberg of Aedas. 2014.

Certain projects transport us to other worlds, not strictly architectural ones. Hybridization, deformation, readaptation, and reinterpretation of forms, functions, and spaces are constants in the design process. Balanced, intentional, systemic, and/or hierarchical composition is one of the mechanisms for producing new forms. We start from pre-existing conditions, from often unconscious apprehensions.

- State of the art

State of the art Natural language processors.

Pre-trained generative transformers (GPTs) are a type of large language model (LLM). The first GPT was presented in 2018 by OpenAI. GPT models are artificial neural networks based on a neural network architecture called *Transformer*. They are pre-trained on large datasets of unlabeled text and are capable of generating content similar to what a human could generate.

Generative artificial intelligence

Generative AI, is a type of artificial intelligence (AI) system capable of generating text, images, or other media in response to commands. Generative AI models learn the patterns and structure of their input training data and then generate new data that has similar characteristics. The ChatGPT is the basis for many of the most advanced generative AI applications.

- An example exercise

Dall-E is an artificial intelligence system that can create realistic images and art from a natural language description.

Let's try some text inputs that will give the reader a general idea of how the system works:

In art First

Let's recall some famous works from the history of painting: Saturn Devouring His Son by Francisco Goya, the Mona Lisa by Leonardo da Vinci, and Autumn by Giuseppe Arcimboldo, which incorporates fruits, plants, and flowers to recreate the human face.



From left to right: 1) *Saturn Devouring His Son*. Francisco Goya. 1820-1823, 2) *The Mona Lisa (La Gioconda)*. Leonardo Da Vinci. 1503-1509, 3) *Autumn*. Giuseppe Arcimboldo. 1573.

Using Dall-E, let's type (that is, enter a prompt): "Saturn Devouring His Son and the Mona Lisa Looking at the Camera. Portrait. Oil Painting. Style of Giuseppe Arcimboldo. High Detail." The system then returns a series of very suggestive images:



From left to right: Three results from the previous query

Disturbing, isn't it?

Note the near-fusion of the lips and bulging eyes in image 1, note the child inside a "ring" and their anguished expression in image 2, note the figure glancing sideways at a child inside a double "ring" in image 3.

In all of them, the child is surrounded by a ring—Saturn's ring?

The larger figures have symmetrical hairstyles, like the Mona Lisa?

The images received were possibly not what was expected. An act of digital creativity? A hierarchical balance of pre-existing elements?

Let's look at the production of architectural images (I ask the reader to understand the simplicity and the stylistic request for practical reasons):

A prompt that reads: "A 600-square-meter house in a gated community in Greater Buenos Aires. Minimalist style. High-definition image."



"A 600-square-meter house in a gated community in Greater Buenos Aires. Minimalist style. High-definition image."

Now we want it a little bigger:



"A 1200-square-meter house in a gated community in Greater Buenos Aires. Minimalist style. High-definition image."

Very large and not very minimalist, don't you think?



"A 40-square-meter refuge in a Mendoza vineyard, minimalist style, with the Andes Mountains as a backdrop. Sunset. High-definition image."

The "Outpainting" function also allows the DALL-E model to extend an image beyond its original borders, adding entirely new visual elements while maintaining the original style of the artwork, but allowing for modifications to its context. It simply uses a description provided by the user.

In this case, the model "created" an environment, but this environment was requested. The clearer the request, the clearer the model.



Girl with a Pearl Earring, Johannes Vermeer. 1665. In this environment reconstruction done with Con Dall_E, using "Outpainting", the algorithm recreates the environment according to the style of the painting.

Other 2D Image Generators

- Midjourney: "Midjourney is an independent research lab exploring new ways of thinking and expanding the imaginative powers of humankind. We are a small, self-funded team focused on design, human infrastructure, and artificial intelligence."

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

- Stablediffusion:

“Stable Diffusion is a text-to-image latent diffusion model capable of generating photorealistic images from any text input. It cultivates autonomous freedom to produce incredible images and allows billions of people to create stunning art in seconds.”

<https://stablediffusionweb.com/#features>

- Dreamstudio

<https://dreamstudio.ai/generate>

Three-Dimensional Model Generators

Natural language models with three-dimensional output are being developed simultaneously. Below are some links that demonstrate their development:

<https://journals.sagepub.com/doi/epub/10.1177/1478077118800982>

<https://research.nvidia.com/labs/toronto-ai/LION/>

<https://www.intel.com/content/www/us/en/newsroom/news/intel-introduces-3d-generative-ai-model.html>

- Generative urban design and artificial intelligence

Today, several public and private entities worldwide are dedicated to incorporating AI into disciplines related to urban and environmental development. **As previously mentioned, managing and processing large amounts of information is no longer an impediment to solving complex systems.**

Today, possessing the right data on each field of knowledge—geographic characteristics, climate, population characteristics, waste management, the economy, etc.—is the fuel for analysis and decision-making in large cities.

Urban sciences (such as urban economics), after analyzing this data, must respond to the needs of physical and environmental planning. The resulting urban design, based on prior planning, must provide the conceptual and design framework for how to build the city holistically. Its transportation and mobility network, land use, housing, and economic activities must be intelligently interconnected to achieve not only sustainability but also energy efficiency, public and private economic investment, urban aesthetics, and democratic citizen participation.

In all cases, data science is fundamental. To make large-scale decisions, this data must be organized and prioritized by expert teams to subsequently study the interaction of different factors.

If cities are extremely complex systems, AI would seem to be the right tool to help manage and administer these vast amounts of information.

So-called smart cities use artificial intelligence to collect public information in order to then efficiently manage resources. These databases are used to predict behaviors that can be beneficial.

Below are two initiatives, clear examples of the current state of affairs. One was carried out by an internationally renowned architecture firm, and the other by a group of experts based in Paris.

Scout, by KPF

Scout is an interactive 3D web platform for visual design and data exploration. It is developed by KPF Urban Interface, a research and design arm within Kohn Pedersen Fox Architects. “Scout is a shared web platform that helps our global firm gain rapid, data-driven insights, present them to clients, and engage with the community. Through Scout, designers and collaborators can easily explore and compare thousands of options, make more informed decisions,

and enjoy the creative freedom of visualizing results in real time. By automating certain design components, Scout frees up focus for deeper development, innovation, and craftsmanship. It is software-independent, easily accessible, and brings together complex design scenarios with a single click. Whether for two minutes or two hours, different types of users can coexist and gain quick insights or deep insights”.

<https://ui.kpf.com/smarter-city>

<https://scout.build/>

URBAN AI

“For years, ‘Smart Cities’ have been the ultimate form of urban progress. This concept can be defined as the use of information technologies to improve the quality of life for citizens while simultaneously reducing operating costs for cities and urban stakeholders. The ongoing digital transformation, the abundance of urban data, and the emergence of artificial intelligence technologies have supported this vision as cities around the world have become “smarter.”

However, this paradigm alone is incomplete. Cities are more than just an optimized platform that offers a frictionless user experience. Civic engagement cannot be reduced to the transmission of data, and human beings are not machines...”

“We call for the urbanization of artificial intelligence.”

<https://urbanai.fr/>

- Some Speculations

Let's let our imaginations run wild and speculate about the potential impact of AI on our profession. Some of this software is already being developed...

What if:

-1. There was a mobile application, 2. That, through sequential text input (or speech), 3. Allowed us to "design" houses, 4. For a modest monthly subscription, 5. Or for a fixed amount per "project".

What if:

-1. There was an application, 2. That, through sequential text input (or speech), 3. Allowed us to "design" horizontal properties, 4. Adapting the proposal to the urban code (or Argentine or international urban codes; the limit would be determined by the database) for each lot, 5. Calculating costs simultaneously (since it would be linked to an updated database), 6. Creating the construction documentation in BIM instantly, 7. Quoting the project instantly. *Municipal approval would be automatic because the corresponding building code would be applied without error.

What if:

-1. An application existed, 2. That, through sequential text input (or speech), 3. Allowed the "design" of hospitals (using a database of all major hospitals worldwide), 4. Generating technical documentation immediately (almost instantaneously due to the computing power currently available in the cloud), 5. Constructing the building using 3D printing? [\[12\]](#)

And so on, infinitely, with more complex systems: airplanes, airports, cities. Complexity would not be a problem.

Many of these proposals already exist:

<https://www.sensetime.com/en/product-business?categoryId=1077&gioNav=1>

<https://www.archdaily.cl/cl/937274/inteligencia-artificial-crea-planos-en-planta-y-estilos-generativos-con-machine-learning-en-harvard>

<https://ovacen.com/programa-ia-planos-distribucion-vivienda-edificio/>

Just three years ago, some reports claimed that design-related professions (engineering, architecture, etc.) would be the last to decline. With the emergence of generative graphics applications based on Transformers, this assessment has changed. A recent paper by OpenAI—the creators of ChatGPT—indicates that not only creative professions will be affected, but also lawyers, mathematicians, managers, accountants and auditors, financial analysts, and others, who will be among the hardest hit. Biologists, photographers, and architects will be one step below. ^[13]

- Reflections on our creative paradise

It is clear that the emergence of Generative Artificial Intelligence bursts forth as something unexpected and disruptive, like an accident in our creative paradise, untouchable and exclusive—until now—to our minds.

Perhaps due to the similarity in the ways our brain incorporates and processes information and how AI uses neural networks, generative applications are having a profound impact on our activities. That world that takes us so long to build appears threatened by something that appropriates human knowledge and modifies and interprets it on demand. “On demand,” like streaming.

Databases are fed by the collection, hierarchization, and processing of billions of works, images, and texts produced by humanity (just as we nourish our knowledge by appropriating human culture). **So, what happens to copyright?**

This constantly evolving digital cocktail shaker increasingly optimizes results, replacing (not just complementing) the work carried out by many disciplines.

In “The End of Work” ^[14] (1995), Jeremy Rifkin addresses—among other topics—the replacement of work activities by the emergence of new technologies. He argues that we are entering a new phase of human history, characterized by what already seems to be a permanent and inevitable decline of what we have understood as work until now. He observes that the number of underemployed or unemployed people is growing at a dizzying pace. Even some of those newly entering the labor market are becoming victims of the technological revolution.

Today more than ever, the sustainability of work is at stake.

In “The Three Laws of Robotics” ^[15], Isaac Asimov proposes a kind of control in the management of robotics:

First Law:

A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Second Law:

A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.

Third Law

A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Just as Asimov proposes this operational framework, should we regulate and manage the use of AI? Could we do so even if we wanted to? Are algorithms under the control of the majority or do they belong to a very select technological elite? Does responsible and ethical AI exist? [\[16\]](#)

The development of technologies linked to **AI represents an exponentially growing business for corporations**. It is also an invaluable asset for governments.

The current technological race between the major powers for extreme ultraviolet lithography (a technology used for the production of high-tech chips) seems to be heading in a direction that is not prudent. The need to have the technology (computing power) to train intelligent systems has become a top global strategic objective. Everyone wants to be the first to develop the best system and claim the throne of the state of the art in AI.

Faced with the emergence of this disruptive technology that directly impacts creative and organizational activity, **what do we do? Do we sit and wait? Do we let others speak for us? Do we let the IT/AI industry design for us? Or do we actively intervene?**

And how? The debate remains open.

The development of AI may become a loss for some. For a select few, it will represent an opportunity, a tool for decision-making. For most, it is currently an obscure science that breeds uncertainty.

While we are currently in the midst of developing weak artificial intelligence, where AI is limited to being applicable to a specific type of problem, we are on the path to developing strong or general artificial intelligence. This AI will be able to perform any activity carried out by any person. Or all of them together. When this singularity is reached, we will surely be talking about self-aware and self-preserving systems.

We are plunging into a fleeting and uncertain era.

- [1] Bauman, Zygmunt. *Liquid Modernity* (2000). Polity Press y Blackwell Publishers Ltd.
- [2] Virilio, Paul (1994). *El arte del motor. Aceleración y realidad virtual*. Manantial.
- [3] ¿Qué es la inteligencia artificial (IA)? [Recuperado 5-10-2023] [International Business Machines \(IBM\)](#)
- [4] *Alan Turing padre de la inteligencia artificial*. [Recuperado 5-10-2023] [Fundación Aquae](#).
- [5] *Alan Turing y la Inteligencia Artificial*. [Recuperado 5-10-2023] International Business Machines ([IBM](#))
- [6] Marvin Minsky: 'La inteligencia artificial nos recuerda que no es una gran cosa ser una persona' [Entrevista de Héctor D'Amico] (8 de febrero de 2016). [La Nación](#)
- [7] Kurzweil: la IA será inteligente en 2029 y nos fusionaremos con ella en 2045. [Observatorio de Inteligencia Artificial](#)
- [8] Ramírez, Fran. *Historia de la IA: Frank Rosenblatt y el Mark I Perceptrón, el primer ordenador fabricado específicamente para crear redes neuronales en 1957* (19 de julio de 2018). [Telefónica Tech](#)
- [9] 'Machine learning': ¿qué es y cómo funciona? (8 de noviembre de 2019) [BBVA](#)
- [10] ¿Qué es el aprendizaje supervisado? [International Business Machines \(IBM\)](#)
- [11] Analogías formales, excéntricas y morfológicas (1990). Taller CBC Bózzoli.
- [12] Hospital del futuro. Oma y Buro Happold proyectan el distrito de salud Al Daayan en Doha. [Metalocus](#)
- [13] Tyna Eloundou, Sam Manning, Pamela Mishkin y Daniel Rock (agosto 2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.
- [14] Rifkin, Jeremy (1995). *The End of Work. The Decline of the Global Labor Force and the Dawn of the Post-Market Era* [El fin del trabajo. El declive de la fuerza del trabajo global y el nacimiento de la era posmercado]. G. P. Putnam's Sons.
- [15] Asimov, Isaac (1942). *Círculo Vicioso [Runaround]*. Ed. Street & Smith. En esta publicación aparecen las tres leyes de la robótica por primera vez.
- [16] Ética de la IA. [International Business Machines \(IBM\)](#)