

ARTIFICIAL INTELLIGENCE IN ARCHITECTURE AND BUILT ENVIRONMENT DEVELOPMENT 2024: A CRITICAL REVIEW AND OUTLOOK

Michal Sourek

Department of Architecture
Faculty of Civil Engineering
Czech Technical University in Prague
Czech Republic
michal.sourek@fsv.cvut.cz

Pre-review manuscript

Abstract:

Architectural practices and the planning and management of built environment development lag in adopting artificial intelligence (machine learning techniques, more correctly), and the sparse tools and approaches implemented provide only marginal contributions. A contrast reveals not only comparing to other industries and creative disciplines but to the opportunities at hand. The paper evaluates the situation both in the context of the AI field and in the sector of architecture and the built environment, points to the causes of the sector's current setup in terms of the starting points of creativity, the technologies used, and approaches to their development, as well as in terms of the economic, social, and political framework, subsequently introduces the opportunities to overcome the falling behind, and outlines the paths. Across the paper, the critical review applies three fundamental perspectives: authentic, poetic creativity that passes and precedes parameterization and algorithmization, second, novel, in architectural designing not yet applied learning strategies and training approaches, and third, concurrence of the fundamental three- and more-dimensional spatiality of both architecture and recently developed virtual reality technology, as well as the new theory of human thinking and intelligence that waits for implementation in machine learning (together with other novel computing approaches). Given the coincidence of the three aspects, a singularity is predicted for the next development of architectural craft and field.

Keywords: artificial intelligence, machine learning, generative pattern, parametric review and optimization, imitation-based learning, learning a behavior policy from demonstration, transfer learning, self-learning, mind and brain theory, decentralized AI

(1) Introduction

Architecture and the development of the built environment, in general, continue to be disoriented, let alone to be coherent, when it comes to adopting technologies referred to as artificial intelligence – AI for short, and machine learning more correctly; they stand aside from the current that has been invading our lives and professions since the 1980s. Concurrently, news in brain biophysics and related computer science along with new computer technologies have emerged recently that render to offer to architects, other project stakeholders, and the public creative and communicative means and tools that the field has been missing throughout its history: the three- and more-dimensional spatiality and modeling together with diachrony lay at the heart of both (i) the new AI computational models and techniques to be developed based on the new computational theory of mind and brain to provide superb capabilities for design, planning, and parameters review and assessment, (ii) state-of-the-art virtual reality technologies featuring unprecedented abilities to design architecture and the built environment as it deserves – a diachronic way in space, to create architecture from spaces, to understand it and communicate - and (iii) latest contributions of architectural theory to design practice, the theory of public space in particular. (This regard requires a precursory knowledge: *Architecture emerges when a work of (physical) construction, communicating cultural, social, and material values with humans and society, is exposed in public space* [1].)

And third, regardless of the novelties in biophysics, virtual and extended reality (VR/XR for short) technologies, and architectural theory, the AI practice, research and development (R&D for short) know and at a hectic pace further release strategies and techniques that have not yet been applied concerning architectural designing and planning so far - such as imitation learning, self-play, and transfer learning among

others that could push the branch and the field forward - to the next level. Though easily overlooked as irrelevant at a superficial glance, AI-driven robotics is the realm where the main incentives for new AI developments for architecture and the development of the built environment are emerging. The opportunities, however, have met only a hesitant, if any, reaction in the respective professions so far, which aligns with the sidelining that architecture and the built environment, unnoticed, have been suffering for the past seventy years [2], losing the status and function as the *summum templum* [3], which they held and which represented their meaning and role within the societies, cultures, and economies throughout millennia.

A question arises: can the confluence of circumstances - opportunities and challenges in general, and, in particular, the newly revealed fundamental interconnectedness between architecture within the built environment framework, virtual reality technology, and mental, thus computational activity stemming from shared aspects of spatiality, spatial modeling, and diachrony – the *trinity* of three- and more-dimensional spatiality for short - be a prerequisite for a singularity in the evolution of the profession and the field?

The answer cannot be but a composition of - on the one hand - the computer scientist's and developer's ability to identify the problem niches and find cutting-edge solutions to fill them, and - on the other hand - the branch's motivation and ability to adopt the new approaches and tools. A lack of mutual understanding of the needs, capabilities, and processes of one and the other expertise between architects on the one hand and AI applications developers and data scientists on the other appears to have contributed to the embarrassing results so far; and the gap continues to widen. Among other preconditions, correct understanding and grasping the aspects of authentic - *poiétic* [4] or poetic as will be asserted later in section (3) of the paper - architectural creativity on the one hand and parametricism in architectural designing on the other are essential.

Contributing to understanding the situation, uncovering the causes, overcoming the unfavorable trend, identifying the prerequisites for effective cooperation in the development of truly productive applications of AI in architecture and the built environment development, and bridging the existing gap between the professions in general is the ultimate ambition of this paper. To stand to the ambition, the paper must address architects, construction and real estate professionals with AI researchers and developers, and also tech investors. Each professional field has its language that fails to be understood comprehensively by the others. While trying to bridge the gaps, some text sections may seem redundant or inaccurate for one or the other profession: for such cases, the author begs for leniency. Timeliness, eventually, is another challenge that faces working in the field developing at an unprecedented, more than hectic pace - more than ten thousand new solutions and hundreds of millions of new users within a year. However, a field that, so far, has been lacking a critical overview more than needs to shed light on the status quo to lay the foundations for efficient further research and development.

Architecture (and the built environment development) on the brink of a revolution

The amount of R&D efforts on AI in architecture and the built environment development significantly lags behind the influx of applications in the economy, insurance, healthcare, the social field and politics, weapons systems development, the judiciary, environmental issues, and in general.

In the field's modest conditions, a gradually increasing flow of so-called AI applications has been entering the field of architecture since the 2010s, even before. After a decade, some of the applications integrate into the workflows of architectural teams; however, the results do not match the previously declared ambitions in terms of performance and quality. AI has not proven creative (not only) in terms of architecture; even the pioneers in the field are abandoning the projects the ambitious visions not achieved and reducing their efforts to pragmatic parametric tasks. On the other hand, AI-driven applications slowly-slowly start to spare the architect's time and energy as his "new pencil" - making sketching faster - or substituting and widening the research of ideas and themes concepts. In addition, a complex AI-enhanced process has been introduced, from an expression of the client's requests and expectations to the final design of architecture and construction solutions, as section (2) of this paper reports. All this, nonetheless, lags behind the vision of an (AI-driven) assistant, an *apprentice* of an architect that section (5) of the paper introduces as feasible, moreover realistic a prospect of the *apprentice* overcoming the master in terms of tirelessness, consistency, and accuracy.

To lay a path to overcome both the architecture's and built environment development's sidelining and the field's failure to keep up in terms of adoption of the progressive technologies, the paper offers a comprehensive approach: first, an assessment of the feasibility of the past and recent visions and goals and

an analysis of successes, flaws, and failures of the last decade. The so-much-wished true creativity of artificial intelligence remains at the heart of research: is its unattainability only a question of state-of-the-art approaches and technologies, or is it a principal misconception? Second, the deployed machine-learning networks and algorithms are reviewed more deeply. Section (5) of the paper concerns the potential of generative adversarial networks and supervised learning techniques that, so far, have been ruling AI-driven-designing attempts and proposes reinforcement learning strategies and imitation-based and self-learning techniques (all terms brought up in (2)) to take over. In terms of (non)availability and how to provide it, the respective training material is the third principal issue. And fourth, strategic fields of feasible and efficient deployment of AI in architecture and AEC (architecture, engineering, construction) are proposed and supported by arguments: “reverse prompting” and AI-fostered researching, generative-patterns-based pre-design, design-development support by “advice whispering”, continuous and complex assessing of parameters of physical and quantitative aspects of an architecture and respective constructions, design reviews, and evaluations of solutions. Despite asserted incapability, let alone the true creativity of a learned machine, still a revolution may lay ahead of the architecture and the development of the built environment; the so-far overlooked potential of imitation-based learning and self-learn strategies may ignite it. Given the inherent computational demandingness due to the three- and more-dimensional nature of the tasks, a generally sought-for breakthrough in computational technologies, be it spatial computing, ternary semiconductor design and neural networks (introduced in section (4) of the paper), or another power- and performance-efficient model may prove to be a prerequisite for practical development and deployment of these strategies.

A breakthrough

Until recently, R&D in AI in architecture has been challenging architectural creativity and threatening to make the architect redundant. Substituting the human brain’s work with the performance of computer networks and various, ever-new-developing algorithms should have been the means. However, principal flaws related have not ceased to emerge.

A new theory of the human brain’s working mechanisms appeared recently to debunk the general working principles of today’s AI algorithms when it comes to substituting or competing with the brain’s or human Neocortex’ performance. Even then, however – further elaborated in section (3) of the paper - a computer remains incapable of consciousness and of true, *poiétic* creativity. In addition, such a “mission statement” fails to catch up with the natural potential of a breakthrough technology. If a technology is truly groundbreaking, it must perform unprecedented outcomes; if the new outcomes can compare to the existing ones, the technology is not a breakthrough. Had the father and son Stephenson's Rocket not been a breakthrough, railway carriages would have continued to be pulled by ropes wound on stationary winches located along railway lines [5] - and the development of rail transport, and with it the upswing of the global economy, living standards, and culture, would have been delayed by decades. Innovation is a breakthrough only if it changes not only the quantity of output but also the quality of performance – as we have been witnessing in the case of the Internet, originally intended only to enable an exchange of documents between institutions, but profoundly and globally changing social values and relations, modes of communication, economics, culture, and environment in the end - and still, it is not an end [6]. Today, however, both AI developers and (the few) interested architects are only thinking about improving winches - instead of dreaming of the next generation Rocket. It is time to start dreaming instead of plain coding: as Rutger Bregman points out, *But the real crisis of our times, of my generation, is not that we don't have it good, or even that we might be worse off later on. No, the real crisis is that we can't come up with anything better* [7]. Can architects offer better? They must; through architecture different from what we know so far, be designed not by AI but by architects grasping the AI opportunity hand in hand with computer and AI scientists and developers – and supported by tech investors that, so far, have identified close to no opportunity in this field. However, this is a misunderstanding: architecture stands at the launch of any real estate development: as such, it delivers the starting points for the business that amounts to hundreds of quintillions of EUR globally. Such is the figure that represents the basis of comparison. A five percent yield of such a figure is still a quintillion. It is not just investors who should finally take notice of this opportunity [8]; the field and the architectural craft deserve the investments – and ultimately, society as the end customer deserves them, too.

How is the paper organized

To unleash the potential – investment opportunity included - an essential interconnectedness between the three realms of the *trinity* of three- and more-dimensional spatiality and diachrony needs to be understood and the resulting singularity must be grasped in theory and deployed in practical use. Till then, however, even with “only” the approaches and tools that are already at hand though remaining not mainstream, there are still significant steps of improvement happening or to be undertaken. Today already, AI in architecture is asserted not only to generate images that help to articulate expression of the *poiésis* of the architecture within design concepts (or to customize them) or to shorten the way from a simple sketch to materiality-rich rendering. The trailblazers already coin the AI design process replacing the traditional design process from the client’s vague idea to final design with CAD/BIM (computer-aided design/building-information management) model integrated. Also, providing predictive simulations and models starts to be a business-as-usual today; section (2) and the second and fourth subheads of section (3) of the paper zoom in on all the state-of-the-art. In addition, processing generative patterns to a pre-design, a solution as close to the set parameters as the stock of generative patterns allows, final adapting the pre-designs and tailoring the final, specific solution, or continuous and complex advising through the design process, designs’ reviewing and solutions’ evaluating, or optimization criteria specifying – these are the strategies, practicals tool subsequently that render on the horizon currently and in sections (4) and (5) of the paper.

To allow for this, new classes of the existing algorithms and strategies of imitation-based learning and self-learning zooming in on the design-development processes instead of the results (to be) achieved must enter the field of AI in architecture, perhaps supported by the new mind and brain theory-based computational models, XR, or ternary semiconductor design and neural networks (introduced in section (4)).

Revealed within the research of the state-of-the-art, three focal perspectives of the critical review are introduced to launch the (3) section: authentic, *poiétic* creativity preceding and transcending parameterization and algorithmization, second, novel, in architectural designing not yet applied learning strategies and training approaches, and third, the *trinity* of three- and more-dimensional spatiality, spatial modeling, and diachrony that interweave architecture, recently developed virtual reality technology, and the new theory of human thinking and intelligence together with other novel computing approaches. The virtual twins' technology, its benefits and prospects are reviewed in [9]; this paper takes them for granted and only puts them into the context of the architecture – XR technology – and computational-model *trinity* of three- and more-dimensional spatiality and diachrony. The third section of the paper revisits existing deployments of AI in architectural designing and in general, introduces novel AI-learning strategies to be examined, and delves into starting points of true creativity - consciousness, reasoning, and *poiésis*. The next (4) section delves in the prospects of the new mind and brain (or Neocortex respectively) theory by Hawkins and Numenta in terms of computing model, the contributions the of state-of-the-art XR technology, ternary semiconductor design and neural networks, and novel approaches in terms of architectural designing and built environment planning.

Fundamental as incentive benchmarks for the “new wave” of AI in architecture and built environment development, the latest achievements in AI-driven robotics are introduced: Figure-1 starting from section (2) throughout the paper, and Phoenix by Sanctuary AI in section (5).

Section (5) discusses the misunderstandings and limitations of recent attempts and expectations regarding the use of AI in in the branch together with proposals for new R&D. The “fantastic achievements of AI promising to make man redundant” discussed recently in the profession and beyond are debunked and, in contrast, other, so far not considered approaches and perspectives are rendered: in elaboration of the latest knowledge of human Neocortex working, in parametric pre-design and human-in-the-loop approach, in new classes of the existing algorithms and strategies that strive to mimic (first) the human-in-the-loop and (then) the processes learned, in continuous assessment of diverse outputs of creative architectural and urban design concepts, and in approaching to sustainable development (or, better, to the comprehensive resilience) tasks. The final (6) section concludes the today-available and to-be-developed productive applications of AI in architecture and the built environment development, the computing strategies and networks development „to-dos“, and the perspectives and a “mission statement” of the field.

(2) State-of-the-art

Since around 2010, young enthusiasts combining information technology and architecture background have been trying to turn the attention of well-doing, mostly global-star architectural studios towards AI’s potential contribution to architectural design or, better to say, disclose where such a contribution might stem from and what it might consist of. In 2020, *DeepHimmelblau* - a video of a journey through an imaginary landscape of

Coop Himmelb(l)au-like building forms - came into existence. The result of the elaboration of datasets comprising *reference images of geomorphic formations* on the one hand and *actual Coop Himmelb(l)au projects* on the other by CycleGAN and other forms of AI technologies provided "*machine hallucinations*" [10, 11, 12] - represented prevalently in two dimensions, substantially lacking both spatial comprehensivity and the for architecture inherent interconnectedness of the experiential (poetic, in other words) and material attributes.



Fig. 1: Cloud augmentation of Deep Himmelblau Universe. Deep Himmelb(l)au. 2012. <https://coop-himmelblau.at/method/deep-himmelblau/>. (accessed May 12, 2023)

How machines have been learning (recently and still today)

The term machine learning covers various methods of how algorithms create models representing sample data that have been used to train the models to make decisions or proposals without being programmed to do so. The base for the learning is a set of data possessing the same characteristics as the data to be generated; a truly large file as a rule - what is not in it, the AI has been unable to learn. Recently, diverse machine learning strategies have been applied - supervised learning, unsupervised learning, reinforcement learning - and new variations and developments have emerged.

In supervised learning, the system is given a series of categorized or labeled examples and told to make predictions about new examples it hasn't seen yet, or for which the ground truth is not yet known [13]. Supervised learning uses labeled datasets, whereas unsupervised learning uses unlabeled datasets. "Labeled" means the data is already man-tagged with the requested answer [14].

Unsupervised learning analyzes and clusters unlabeled datasets. These algorithms discover hidden patterns or data groupings without the need for human intervention - without the need for labeling the datasets. *In unsupervised learning, a machine is simply given a heap of data and told to make sense of it, to find patterns, regularities, and useful ways of condensing, or representing, or visualizing it [13, 15].*

Reinforcement learning concerns how intelligent agents ought to take action in an environment to maximize the notion of cumulative reward. Based typically on the Markov decision process (a discrete- and also continuous-time stochastic control process in mathematics) [16], reinforcement learning differs from supervised learning in not needing labeled input/output pairs to be presented, and in not needing sub-optimal actions to be explicitly corrected. Instead, it focuses on finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge). In other words, *placed into an environment with rewards and punishments, [the system is] told to figure out the best way to minimize the punishments and maximize the rewards [13, 17].* As the comprehension upgrades of learning processes and - especially - of consequences of their details for the outputs, the reward signals to fine-tune the models tend to be human preferences based (referred to as reinforcement learning *with* human feedback [18] that, as the subhead *Production ecosystems and sensational novelties* of this section explains, is to be distinguished from reinforcement learning *from* human feedback [19]) instead of simple automatic metrics. Indicated further in this section, the safety and alignment problems are the starting point for the deployment of these approaches that are much more time- and cost-consuming. Such is, for example, the case of InstructGPT - one of the most advanced language models today.

Concerning architectural designing and planning, reinforcement learning deserves highlighting as a platform on which diverse imitation-based learning strategies develop that this paper promotes as the future of AI in architecture; the next sections delve into the topic.

New implementations of learning paradigms and new learning models evolve. Countless facets of motivation and reward has been studied and developed in this regard, comprising the dopamine-releasing mechanisms, possible-value- and/or actual-expectation-motivation, curiosity, a self-motivated desire for knowledge, imitation and interactive imitation, self-imitation, and transcendence. The results are random network distillation algorithms using prediction errors as a reward signal [20] or algorithms approximating a state-value function in a quality-learning framework [21], or knowledge-seeking agents [22]. Unplugging the hardwired external rewards not only makes the algorithm capable of playing dozens of Atari games with equal felicity (similarly as Alpha Zero will be reminded just as adept at chess as it is at shogi or Go [23]). In such algorithms, AI agents render able to come up with their own objectives, measuring intelligence in the end effect in terms of how things behave – not in terms of the reward function [24]. These schemes may render a significant part of building a general AI instead of the current models that only specialize in a specific, narrow task.

Imitating the human and other cunning strategies

Shifting the focus from the statistical approach of input-output matching to mimicking (not only) human processes appears to be a key innovation. Indeed, why a machine could not learn by observing the man acting? Still, learned in such a way, a machine can outperform a human, as is the rule: by working relentlessly, flawlessly, and much quicker. Given the nature of experiencing architecture and, most importantly, how architecture emerges in the design processes, such approaches could bring the so-much-sought-after support for the architects' craft. Nonetheless, there is a much broader - almost unlimited field of deployment for the new philosophy represented by strategies and techniques of imitation learning, self-play, inverse Q-learning, adversarial imitation learning (AIL), action-guided adversarial imitation learning (AGAIL), world imitation learning (WIL), behavioral cloning from observation (BCO), augmented behavioral cloning from observation, inverse reinforcement learning, or transfer learning.

In general, zooming on a (design) process - "how things come to existence" - instead of output - "how things shall be", imitation-based learning is a framework for learning behavior from demonstrations usually presented in state-action trajectories, with each pair denoting the action to take at the state visited [25]. Imitation-based learning provides three distinct advantages over trial-and-error learning: efficiency, safety,

and, which also renders promising for AI's deployment in architecture, the ability to learn things that are hard to describe [26]. As one of the techniques applied, behavior cloning treats the action as the target label for a state and learns a generalized mapping from state to action. Inverse reinforcement learning (IRL), on the other hand, views the demonstrated actions as a sequence of decisions and aims at finding a reward/cost function under which the demonstrated decisions are optimal. Reinforcement learning generally supports these techniques; an agent learns a policy - a mapping from states to actions - that maximizes some notion of cumulative reward. The agent iteratively updates its policy based on the feedback (rewards) it receives from the environment. It is inherently a process-oriented approach, as the agent is learning how to act in different states rather than just predicting the outcome of actions.

In imitation learning, the primary focus is on mimicking the actions of an expert. The agent tries to replicate the expert's behavior as closely as possible. While the outcomes of actions are certainly important (they can signal how well the agent is doing), the main goal is to learn the expert's policy - the sequence of actions the expert takes in different situations [25].

A radical alternative, world imitation learning goals to achieve similar performance to the expert without explicitly defining a reward function. Learning by observing the expert's actions in a given environment, the agent constructs a model of the world (a "world model") based on the observed data. The world model captures the dynamics of the environment, including state transitions and outcomes of actions. To encourage imitation, an intrinsic reward is defined within the latent space of the world model. This reward guides the agent's exploration during reinforcement learning. Addressing the problem of covariate shift (adapting to changes in the input data distribution during model training, adjusting the learning strategy when the rules of the game change [27]), where the data distribution during training differs from that during deployment, WIL can be applied offline using historical data without requiring additional online interactions. The approach leverages world models to bridge the gap between expert demonstrations and reinforcement learning. Combining imitation learning with reinforcement learning, the agent improves the performance over time [28, 29, 30]. Applied successfully in robotics, autonomous driving, gaming, and other domains where expert behavior can serve as a guide, the possibility of application in architectural design has escaped attention so far. The world model would certainly not capture the doing in the architect's studio but the architect's actions in the model of emerging architecture in a design- or CAD (computer-aided design) software environment; the paper will come back to the topic.

The self-playing agent plays against itself and it updates its policy based on the outcome of each game. Its goal is to find a policy that maximizes the expected cumulative reward. The outcomes of the games drive the learning of the agent who, however, cares about the actions taken since they determine the outcome [31].

Q-learning is a model-free, value-based, off-policy algorithm that searches for the best series of actions based on the agent's current state. With Q standing for quality, inverse Q-learning is a method for dynamics-aware learning that avoids adversarial training by learning a single Q-function, implicitly representing both reward and policy. It obtains state-of-the-art results in offline and online imitation learning settings, surpassing other methods in both the number of required environment interactions and scalability in high-dimensional spaces [32].

Adversarial imitation learning is a class of state-of-the-art algorithms commonly used in robotics. In AIL, an artificial adversary's misclassification serves as a reward signal, which is subsequently optimized by any standard reinforcement learning algorithm [33].

Action-guided adversarial imitation learning processes even demonstrations with incomplete action sequences. AGAIL divides the state-action pairs in demonstrations into state trajectories and action trajectories and learns a policy from states with auxiliary guidance from available actions [34].

Behavioral cloning from observation is a two-phase, autonomous imitation learning technique aiming to provide improved performance in diverse fields such as healthcare, autonomous driving, and complex game-playing [35]. Combining imitation learning and model-based learning, the method is to let the agent learn its agent-specific inverse dynamics model first and then show the agent the state-only demonstrations. The agent can use its model to infer the missing actions that the expert took. After the actions are inferred, the agent can use imitation learning.

Augmented behavioral cloning from observation overcomes the BCO's problem of reaching bad local minima by a self-attention mechanism that better captures global features of the states and a sampling strategy that regulates the observations used for learning [36].

In inverse reinforcement learning, the goal of the apprentice agent is to find a reward function from the expert demonstrations that could explain the expert behavior [37].

Applying beneficially in robotics, transfer learning is like borrowing knowledge from one task to help with another. It starts with a pre-trained model that has learned useful features from a large dataset, and then adapts the model to new tasks by removing the top layers of the neural network, adding new layers for the specific task, and training the modified model on a smaller dataset. The strategy proves not only to robotics; for *production architecture*, as coined in section (4), transfer learning is another technique that awaits exploration. Overcoming training from scratch, the transfer learning approach saves time; in addition, it can work even with limited data [38, 39, 40].

Another "next level" of the extrinsic-reward-free schemes comes with interaction that allows the algorithm to work properly requiring incredibly little feedback as Ross' dataset aggregation (Dagger) has shown [41]. Self-imitation and transcendence render to be a "top" of recent learning schemes [42].

And multiple other sophisticated strategies adjoin to simplify the learning and increase the performance of the models. In-context learning allows "old models" to learn new tasks by providing them with just a few examples (input-output pairs). Like strategies based on imitation, these approaches, unfortunately, have not even tried to apply themselves in architectural design that, when attempting to embrace AI, continues to stick to the supervised-learned schemes (as this section shows later). (The expressive image processing, in which Midjourney, Stable Diffusion, and other applications show encouraging results, stay aside: they help architects only indirectly in the fields of analyses' research or thematization - they do not directly participate in designing the architecture itself.) An opportunity awaits: it would be a mistake to miss it.

Allowing to quickly building models for new tasks without demanding fine-tuning, the strategy is often used regarding large language models (LLMs), which learn in context, using the provided examples without adjusting their overall parameters. Instead of fine-tuning the entire model, it is given a prompt (a list of examples) and asked to predict the next tokens based on that prompt. In-context learning works even when provided with random outputs in the examples. Traditional supervised learning algorithms would fail with random outputs, but in-context learning still succeeds like doing some implicit Bayesian inference, using all parts of the prompt (inputs, outputs, formatting) to learn [43].

As opposed to closed machine-learning strategies (to some extent, the vast majority of the state-of-the-art models), open-ended machine-learning algorithms learn how to learn from data they encounter; they adapt and refine themselves based on past examples, when faced with new data, they predict outcomes or classify information without human intervention. Recommendation engines, speech recognition, fraud detection, or self-driving algorithms tend to adopt the strategy [44, 45].

A clever approach used specifically concerning LLMs, the tree of thoughts (ToT) machine-learning approach fits to be imagined as a tree where each branch represents a coherent sequence of language (a "thought"). These thoughts serve as intermediate steps toward solving a problem through deliberate decision-making: Instead of just making decisions one after the other (like reading a sentence left to right), ToT allows the model to consider multiple reasoning paths, evaluate their choices, and decide the best next step. Thus, the problem-solving abilities of the model are enhanced to think more strategically, explore different paths, and make better decisions [46, 47].

Commonly used for pathfinding problems in video games, robotics, and more (however not in architectural designing tools), the A* algorithm is an informed best-first search algorithm that efficiently finds the shortest path through a search space using a heuristic estimate. It combines the best features of breadth-first search (BFS) and Dijkstra's algorithm, nevertheless, unlike BFS, which explores all possible paths, A* focuses on the most promising paths based on a heuristic function that guides A* toward the goal state, making it more efficient than traditional search algorithms. A* can be implemented in Python to find cost-effective paths in graphs [48, 49].

Dreamer agent is a reinforcement learning strategy that combines world models with efficient learning techniques. Used in robotics, games, and real-world scenarios (then, why still not in architectural designing?!), this technique is like teaching an AI to dream, plan, and act smartly. Dreamer learns a simplified model of the environment (a world model) from raw images; this world model predicts what will happen in the environment based on the agent's actions. Instead of trial and error, Dreamer imagines thousands of possible action sequences in parallel; by computing compact model states from raw images, it learns from predictions using just one GPU. Using a value network to predict future rewards and an actor-network to choose actions, Dreamer considers rewards beyond the immediate future. These networks help the Dreamer to make informed decisions even in new, unknown situations [50, 51].

Autoassociative algorithms learn to remember patterns and retrieve them from partial or noisy input. A smart way to remember things even when the details are fuzzy, these algorithms create a compressed

representation of data to map input patterns for themselves. Then, when given a distorted or incomplete input, they reconstruct the original pattern to prove helpful in denoising (removing noise from data), memory recall (helping to recognize familiar patterns), or anomaly detection (raising an alert if something does not match the learned patterns) [52].

Though originally bound for Atari 2600, the observe-and-look-further learning strategy still today addresses challenges like reward processing, long-term planning, and efficient exploration. The insights gained from Atari environments can be applied to other domains, such as robotics, natural language processing, and autonomous vehicles. Designed to tackle challenges in reinforcement learning, the algorithm processes rewards of varying densities and scales, it can handle different types of rewards effectively. Using an auxiliary temporal consistency loss, the algorithm can reason over extended time periods, which is crucial for complex tasks. To address the exploration problem more efficiently, the algorithm leverages human demonstrations [53].

The first-return-and-then-explore strategy is another clever approach in reinforcement learning that helps agents explore their environment more effectively. Returning to familiar states helps the agent remember promising locations - preferably spots with potential rewards or interesting features. By revisiting them, the agent reinforces its memory. Once known states have been revisited, the agent can confidently venture into uncharted territory and make smarter decisions. The strategy helps agents handle sparse rewards (where feedback is infrequent) and deceptive feedback (where rewards can mislead). Balancing curiosity with wisdom, the approach has shown impressive results in solving challenging reinforcement learning tasks, from video games to robotics (and perhaps once, it can perform in the architectural designing realm, too) [54].

Also self-supervised learning renders a promising path to advance machine learning. As opposed to supervised learning, which is limited by the availability of labeled data, self-supervised approaches can learn from vast unlabeled data [55], the model generating its own labels from the input data. Given a specific task, the self-supervised learning approach can outperform the reinforcement learning strategy in terms of data efficiency (learning from the structure of the input data itself, without needing explicit rewards or penalties), generalization (learning a more general understanding of the data, rather than optimizing for a specific reward function), stability (self-supervised learning often having more stable and predictable training dynamics than reinforcement learning that, especially in environments with sparse or delayed rewards, can sometimes be unstable or difficult to train), and in general being less demanding man-performed interventions like, for example, rewards design [56].

Self-supervised learning is the strategy that, among others, "runs" Sora - a text-to-video model developed recently by OpenAI [57]. Mentioned further in this paper, Sora is designed to understand and simulate the physical world in motion, to solve problems that require real-world interaction - similar to how *poietic* architectural concept designs come into existence. As promised, this paper will further delve into the idea of *architectural robots* elaborating designs of *production architecture* (the term coined in section (4)) making use and benefits of imitation-based, self-learning, and other novel machine learning strategies.

Finally yet importantly, heading to AI understanding and predicting others' actions like humans do, the machine theory of mind (ToM) asserts an intriguing concept inspired by how humans understand each other based on past behavior. Analogically, the machine observes how agents (or people or robots) behave. Using meta-learning (learning how to learn), the machine builds models of the agents and predicts, how agents might act in the future, even in unprecedented situations. Such a strategy performs better interaction with humans, advances multi-agent AI where multiple AIs collaborate or compete, and makes AI more interpretable, transparent, and safe [58].

Artificial neural networks

An artificial neural network is a collection of connected units or nodes called artificial neurons designed to model loosely how the neurons in a biological brain have been supposed to look and work. Like synapses in a biological brain, each connection can transmit a signal to other neurons. A deep neural network is an artificial neural network with multiple layers between the input and output layers; in a shortcut, a deep neural network makes machine learning *deep learning* [59, 60]. In essence, two computing principles apply in artificial neural networks today: feedforward computing and backpropagation. The goal is always to *train* the models generated to cope with the criteria typically inserted by vast sample datasets. Feedforward computing refers to a type of workflow without feedback connections that would form closed loops; the latter term marks a way of computing the partial derivatives during training. When training a model in the feedforward manner,

the input “flows” forward through the network layers from the input to the output. By backpropagation, the model parameters update in the opposite direction: from the output layer to the input one. Backpropagation, a strategy to compute the gradient in a neural network, is a general technique; it is not restricted to feedforward networks, it works for recurrent neural networks (to be introduced soon) as well [61].

In general, *discriminative* and/or *decoding* techniques *identify* objects and *infer* what is “true” and what is “fake”. As a principle, generative AI systems *create* objects such as pictures, audio, writing samples, and outline anything that computer-controlled systems like 3D printers can build [62]. Most often, generative and discriminative or decoding systems operate paired in generative adversarial network (to be introduced soon) models setting the business-as-usual rather than state-of-the-art of today’s AI industry. Typically, a system labeled as generative AI is self-learning, it uses unsupervised learning (but can use other types of machine learning, too), and deploys anomaly detection and problem solving - it can come up with innovative solutions or approaches based on its experience with similar outputs/inputs pairings in the past [63].

First introduced in 1987, the pioneers were convolutional neural networks (CNNs), also known as shift invariant or space invariant neural networks, most commonly applied to analyze visual imagery [64]. ImageNet, a groundbreaking project from the 2010s builds on this technology [65]. Introduced by Geoffrey Hinton, capsule networks aim to overcome the limitations of traditional CNNs by representing hierarchical structures in images [66]. They focus on learning spatial relationships between features.

Graph neural networks (GNNs) are another field of research aiming at processing of graph data. The concept of GNNs emerged in the 1980s exploring neural network architectures for processing graph-structured data. Next, in 2015, the foundational work on graph convolutional networks (GCNs) was introduced by Thomas Kipf and Max Welling to leverage spectral graph theory and convolutional operations to learn node representations in graphs [67]. By extension of the GCN framework variants such as GraphSAGE, graph attention networks (GAT), and ChebNet were developed to address scalability, inductive learning, and handling heterogeneous graphs [68]. New architectures, attention mechanisms, and graph pooling techniques have been proposed recently to explore applications in recommendation systems, social networks, and drug development. The works continue focusing on scalability, interpretability, and handling large-scale graphs.

Designed by Ian Goodfellow and his colleagues in 2014, generative adversarial network (GAN) is a class of machine learning frameworks that until recently, evolving, have represented state-of-the-art of the field [69, 70]. A composition of a generator, which produces data, and a discriminator, which evaluates the authenticity of generated data, through adversarial training, GANs learn to create data aiming to become indistinguishable from real (original input) data. ... *the generator gets no training to minimize the distance to a specific image, but rather to fool the discriminator. This enables the model to learn in an unsupervised manner; however, GANs have also proved useful for semi-supervised learning, fully supervised learning, and reinforcement learning* [71]. GANs have revolutionized generative modeling across various domains, including computer vision. Over the years, numerous variants have emerged such as Conditional GANs, which condition the generator on specific attributes, Wasserstein GANs, which improve training stability, CycleGANs excelling in image-to-image translation, or StyleGANs known for high-quality image synthesis [72].

Over time, too, drawbacks of GANs have bothered both application developers and users: mode collapses, where the generator produces a limited variety of samples, failing to capture the full diversity of the target distribution - in other words, focusing on a few modes and ignores others; instability during training; hindering convergence and quality of generated samples; minimax game between generator and discriminator can be challenging to optimize - a Nash equilibrium can be found as a result, which is not always straightforward [73]; inconsistency of generated images, which may suffer from artifacts, blurriness, or lack of fine details; training large-scale GANs on high-resolution images can be computationally expensive; finding optimal hyperparameters can be challenging; measuring the performance of GANs is difficult: common metrics like inception score or frechet inception distance have limitations and may not fully represent the quality of generated samples; understanding the latent space where the generator operates remains an open research question [74]; GANs rely substantially on the quality and diversity of the training data: if the dataset is biased or lacks diversity, the generated samples may inherit the limitations [75].

To overcome the shortcomings, modified GAN architectures (such as StyleGAN3 by Nvidia [76] or game-theoretical models [77]), GANs for imbalance problems (improving the performance of computer vision algorithms [78]), SMOTE-GAN (combining synthetic minority oversampling technique (SMOTE) with GANs to generate synthetic minority samples for imbalanced datasets), CycleGAN (that focuses on image-to-image translation tasks, such as style transfer, without paired training data, and learns mappings between two domains using cycle consistency loss), or ConditionalGANs (incorporating additional information (such as class

labels) during training, cGANs control the generated samples by conditioning on specific attributes) have been used.

Various applications of GANs are still emerging: FrankenGAN for urban context massing, detailing, and texturing, Pix2PixHD by Nvidia for high-resolution photorealistic image-to-image translation, GAN Loci, or GauGAN [79]. Still then, as the field evolves, diffusion models continue to push the boundaries of generative AI and, in particular, tend to surpass GANs in training methodology (taking advantage of the denoising process, refining data step by step), offering stability and robustness during training, being more straightforward in terms of optimization, and excelling in producing sharp and detailed features - in short, better quality images than GANs [80]. Operating in a compressed or latent space, which can make them more computationally efficient, latent diffusion models push the boundaries further [81].

Recursive (RvNNs) [82] and recurrent neural networks (RNNs) [83] represent state-of-the-art today, also in the context of imitation-based learning and self-learn. RNNs, capable of processing sequential data, are particularly useful for natural language processing (NLP) tasks given the sequential nature of linguistic data, such as sentences and paragraphs, where the time factor differentiates elements of the sequence. They process sequences until reaching an end token, regardless of sequence length; unlike other networks (e.g., CNNs) that handle only fixed-length inputs, RNNs can process variable-length inputs without changing the network architecture. As the network processes elements of a sequence one by one, the hidden state stores and combines information from the entire sequence. RNNs also share weights across time steps, simplifying the network architecture and allowing efficient training [84]. A type of RNN designed to handle sequential data, long short-term memory (LSTM) networks excel in tasks like natural language processing and time series prediction due to their ability to capture long-term dependencies [85]. Another type of RNNs, echo state networks (ESNs) with fixed random connections are particularly useful for time-series prediction and reservoir computing [86]. RvNNs can tackle hierarchical patterns to capture relationships in hierarchical structures, making them suitable for tasks involving nested or tree-like data. RvNNs know how to represent nested relationships and are well suited for tasks that involve tree-like data [87].

Radial basis function networks (RBFNs) use radial basis functions as activation functions for function approximation and pattern recognition [88].

Introducing another class of neural networks, variational autoencoders (VAEs) and adversarial autoencoders (AAEs) also overcome some of the drawbacks of GANs. Unlike GAN, instead of the *generator* – *discriminator* pair, VAE combines two distinct approaches - *encoding* and *decoding* [89]. Encoder abstracts data by compressing while decoder brings the data back to its initial format. Through the decompression, or "reparametrization", the decoder generates variations of the respective phenomenon [90] - typically in furniture design, fashion, photography, or video generating, and also in algorithms concerning architecture and urban design [91]. VAEs combine generative and inference models - elements of both autoencoders and probabilistic modeling. They learn a probabilistic latent space and can generate new samples while preserving data distribution. Commonly used for unsupervised learning and generative tasks, in video generation, VAEs can learn a compact representation of video frames (or other data) in a lower-dimensional latent space, the encoder maps input data (video frames) to a latent space, and the decoder reconstructs the data from the latent representation. A hybrid of autoencoders and GANs, AAEs use adversarial training to learn a latent representation. They aim to improve sample quality and diversity [92].

Networks they mimic

Focusing on architectural designing and given its nature and complexity, this paper will further delve into the potential of imitation-based learning and self-learn strategies and techniques. To cope with the specificities and demands of such techniques (having not architecture or its design but robotics as an intention), various artificial neural network models have been developed. Autoregressive neural networks generate sequences element by element, making them suitable for imitation tasks. Aligning with imitation learning, they learn to predict the next element based on previous ones. Equipped with command compensation, they enhance imitation learning based on bilateral control, allowing robots to operate almost simultaneously with an expert's demonstration [93]. A new autoregressive neural network model (S2SM) has been proposed to fit that requires only the *slave* state as input and predicts both the next *slave* and *master* states. (The goal of imitation-based learning is to learn a policy or control strategy that allows the system to mimic the behavior exhibited during demonstrations (usually performed by an expert). The *slave* state serves as the input for this learning process while the master state corresponds to the desired or reference state, which the algorithm

aims to achieve. It represents the ideal or target configuration or behavior. In a bilateral control framework, the master state often comes from an external source (such as a human operator or an expert) and guides the system.) S2SM improves task success rates and stability under environmental variations.

Brain-inspired deep imitation networks have been developed to enhance the generalization ability and adaptability of deep neural networks for autonomous driving systems across various scenarios. The model draws insights from inferring functions that are believed to be close to the ones performed in the human brain to improve neural network performance in diverse environments [94].

In transformer-based deep imitation learning, a variant of self-attention architecture - transformer is applied to deep imitation learning to solve dual-arm manipulation tasks in the real world [95]. Typically, transformers possess self-attention mechanisms allowing the network to focus on important features in the input data and ignore distractions, which can improve the performance of the network and the ability of parallel processing: unlike RNNs, transformers process all input data at the same time rather than sequentially, which allows for faster training. Other properties of transformers are scalability, which makes them suitable for a wide range of applications, from natural language processing to robot manipulation, and versatility, allowing for use for both sequence-to-sequence tasks (like translation) and single-input tasks (like classification). Not a standalone network, attention mechanisms enhance the performance of various models (including transformers) [96]. They allow networks to focus on relevant parts of input sequences, making them valuable for tasks like machine translation and image captioning.

Inspired by the asymmetry of human neural networks, dual neural circuit policy (NCP) architectures are a design that helps improve the generalization ability of the networks [97].

In behavioral imitation with artificial neural networks, hidden layers of imitation learning models successfully encode task-relevant neural representations. These models predict individual brain dynamics more accurately than models trained on other subjects' gameplay or control models [98].

The introduced neural network approaches cater to different requirements and scenarios, offering solutions for diverse modifications in imitation-based learning. In summary, the landscape of deep neural networks is rich and diverse, with each architecture tailored to specific tasks and domains. The research goes on to explore novel designs and strategies to improve the capabilities and characteristics according to specific purposes. Adaptations of such approaches may be forthcoming once imitation-based learning and self-learn strategies are considered that could address architectural design tasks, (not only) overcoming the shortcomings of GANs that are ruling the field so far.

Production ecosystems and sensational novelties

There are ecosystems of natural language processing, image and video processing, voice processing, and code or software processing and development, further robotics, and expert systems or business intelligence [99, 100], altogether represented by DALL-E (DALL-E3 newly), ImageGPT, InstructGPT and ChatGPT, Bard or Gemini, Ernie Bot, Tongyi Qianwen, Sense Time SenseChat, Bedrock, and many other tools by OpenAI, Microsoft, Google, Baidu, Alibaba Group, Amazon, also MidJourney (that released version 6 recently), Stable Diffusion (currently released version 3, which demonstrates unmatched performance on the ControlNet network, designed to control diffusion models in image generation, and LayerDiffusion that introduces latent transparency, which allows the generation of a single transparent image or multiple transparent layers, combined into a single blended image [101]), Gong.io, Tellius, OPENNN, Theano, and many other tools by multiple producers. These ecosystems exist, evolve, and (some of them) work (though sometimes obscured, even covered up) already over decades and render mature. And most importantly, there are massive foundation models [102] - GPT (version 4 released in March, 2023) by OpenAI, Gemma, Gemini 1.5 Pro (successor to LaMDA and PaLM 2), Imagen Model Family, MedLM, or Codey by Google, foundation models by IBM, and others - trained on large, unlabeled datasets and fine-tuned to become a starting point for a wide array of applications.

Social networks are teeming with posts recommending dozens of artificial intelligence applications that will provide users with greater performance, time savings, and effortless earnings [103, 104, 105, 106]. However, there is indeed a solid base of research and development that is delivering impressive technological progress at a pace that is nothing short of breathtaking. At the beginning of March 2024, a new generation of the Claude family of LLMs was released - Claude 3 foundation model, which sets new industry benchmarks across a wide range of cognitive tasks [107]. The family includes three state-of-the-art models in ascending order of capability: Claude 3 Haiku, Claude 3 Sonnet, and Claude 3 Opus. Each successive model offers

increasingly powerful performance, allowing users to select the optimal balance of intelligence, speed, and cost for their specific application. Product of San Francisco headquartered Anthropic, Claude 3 challenges the hegemony of GPT 4 [108]. Launching Chat with RTX simultaneously, Nvidia also competes with OpenAI or its GPT respectively (plus, for free, running locally on a PC guaranteeing the user data privacy, enabling choice of AI model from Llama or Mistral and choice of dataset, which can include getting answers from YouTube) [109].

Nonetheless, Claude's lead renders a more robust [110]. ChatGPT is optimized for dialogue using reinforcement learning (RL) with human feedback (HF); to enhance their helpfulness and harmlessness, Claude models have undergone fine-tuning using constitutional AI and reinforcement learning from human feedback. The preposition makes the difference: in RL with HF (ChatGPT case), the AI model is initially trained using supervised learning with human demonstrations. Human AI trainers provide example conversations where they play both the user and the AI assistant, and the model learns to generate responses by imitating these demonstrations. After the initial training, the model is fine-tuned using reinforcement learning: it interacts with users and receives feedback (rewards) based on the quality of its responses. The goal is to improve the model's performance over time by adjusting its behavior based on this feedback. RL from HF (the Claude 3 case) focuses solely on reinforcement learning. The model launches training according to an initial policy (a way of generating responses) interacting with users and receiving feedback (rewards) based on the quality of its responses; unlike RL with HF, there are no initial supervised demonstrations, the model learns directly from user interactions, adjusting its behavior based on the received rewards, aiming to optimize its responses without relying on pre-existing examples. In addition, involved in Claude's training, constitutional AI can be considered a set of guiding principles - a rulebook that helps the AI make decisions and respond helpfully and ethically. Sharing similarities with rule-based algorithms, constitutional AI is not the same: instead of strict rules, constitutional AI operates based on guiding principles (a "constitution"); it is more flexible and adaptive allowing for context-aware decisions. Constitutional AI can self-correct and learn from feedback, it balances dynamically following principles with learning from real-world interactions [111].

Another challenge comes from Europe: Paris-based Mistral — a nine-month-old startup with only a few dozen employees — is corraling enough investment and attention, including a high-profile Microsoft, Nvidia, and Salesforce partnership, to put it in the top tier of AI companies globally. Mistral's top models Mistral Large and Mixtral already rival GPT-4's performance in accuracy and common sense reasoning. Designed for complex multilingual reasoning tasks, including text understanding, transformation, and code generation, Mistral Large is natively capable of function calling, facilitating application development, and tech stack modernization at scale. On the other hand, a sparse mixture-of-experts (SMoE) model pre-trained on data extracted from the open web, operating as a decoder-only model, Mixtral excels in efficiency, performance, and handling multiple languages and code generation, which is highly appreciated by developers that can fine-tune and modify Mixtral for specific business needs [112, 113]. Mistral also launched Le Chat, a chatbot "multilingual by design" [114].

In addition, Meta announces to release its newest LLM—Llama 3—in July 2023, which is expected to match GPT-4 capabilities and respond to image-based questions. Llama 3 aspires to handle challenging queries—offering context, instead of dismissing them—reducing inappropriate or inaccurate responses [115]. This timely announcement comes just as rival Google was forced to pause its Gemini-powered image-generation feature after it misrepresented historical images.

And there are not only large language models but small ones, too. Designed to perform well for simpler tasks and more easily to fine-tune to meet specific needs, these small models (SLMs) are more accessible and easier to use with limited resources, be it money, time, or training data. What the recent releases mean is not a start of a shift from large to small but an extension of available categories of models adjusted to fit best to specific performance scenarios. SLMs are well suited for those looking to build applications that can run locally on a device (as opposed to the cloud), where a task does not require extensive reasoning or a quick response is needed. A story of bedtime reading to an AI researcher's 4-year-old daughter goes behind the formation of SLMs when he thought to himself "How did she learn this word? How does she know how to connect these words?" That led the Microsoft Research machine-learning expert to wonder how much an AI model could learn using only words a 4-year-old could understand – and ultimately to an innovative training approach that has produced a new class of more capable language models that promises to make AI more accessible to a variety of specific needs and purposes. At the end of April 2024, Microsoft announced the Phi-3 family of open models, the most capable and cost-effective SLMs available. Phi-3 models outperform models of the same size and next size across various benchmarks that evaluate language, coding, and math capabilities. [116] Performing better than models twice its size (Microsoft says), starting from the end of April 2024, the first in

that family, Phi-3-mini, measuring 3.8 billion parameters is publicly available in the Microsoft Azure AI model catalog, on Hugging Face, a platform for machine learning models, as well as Ollama, a lightweight framework for running models on local machines, and as an Nvidia NIM microservice with a standard API interface that can be deployed anywhere [117].

As of 2022, we are experiencing a kind of artificial intelligence storm; yet for many reasons, it would be a mistake to forget history. After the first, largely experimental AI models - the perceptron [118] and Samuel's checkers player [119] - slowly-slowly, practical or even commercial AI applications have been arriving: from ELIZA, the first chatbot in history that replicated a therapist giving general answers to users' questions, simulating a real conversation, developed in 1966 by Joseph Weizenbaum for MIT, through Siri, one of the first successful virtual assistants, developed in 2010 by Apple for iOS devices, still active and evolving, Alexa developed in 2014 by Amazon, to Google Assistant or BERT, an artificial intelligence model that can understand the context of the conversation and provide more accurate and personalized responses, used as the basis of many modern virtual assistants, both developed by Google in 2016 and 2018 respectively. The "good old" AI applications that have been reading postal codes for USPS since the very beginning of the 1990s, have been deciding on custody and bail in many US states and liquidating insurance claims, and have been a tool of economic efficiency of the US healthcare through last two or three decades [120] deserve noting, too.

Game developers have been adept at creating artificial worlds *and telling stories based on them by their very nature. Expertise of game developers in creating and deploying diverse algorithms that generate the game narratives and scenes often goes back decades before AI was defined as an umbrella term* [121]. Generative AI shapes ever-more *the fashion industry and the art world, where brands and artists - or AI users? - can create original designs that look like human artists created them.* Such is the case of Lalaland.ai [122], a Dutch startup that provides a *self-service platform where users can create their own hyper-realistic AI-driven fashion avatars in just minutes. Users may customize the virtual models' size, body type, shape, and identity—even down to whether they are happy or sad. In the financial sector, banks are using generative AI to automate tasks such as checking account openings and loan approvals* [123].

Founded in 2016 and becoming the GitHub today of AI in the next wave, Hugging Face is a platform on which developers can discuss anything from bug tracking to API integration to overall project development [124]. *With AI drawing massive enthusiasm about being the potential epicenter of future economies, the startup is in a good position to stand the test of time.*

Cofounded only in 2023 by OpenAI-backgrounded Aravind Srinivas, Perplexity AI developed the first-of-its-kind LLM-powered answer engine in just six months as a low-cost project. Perplexity does not have its own LLM, it uses language models via an API. In the free version, ChatGPT or Claude Instant, in the paid version top models GPT-4, Claude 3, or Gemini Ultra (subscribers can choose) [125]. According to the Wall Street Journal, Perplexity is finalizing the acquisition of new investments that could value the company at around \$1 billion, making it a "unicorn" [126]. Jeff Bezos, among others, invested in Perplexity two months ago, and Daniel Gross [127] should bring new money. The company has annual revenues of about 10 million dollars and about 50 million users: the figures deserve attention showing the extremely low revenue ratios that the investors are ready to accept when it comes to AI.

Not only professionals but the general public, too, has started to employ AI: from initial embarrassment and doubts to more mature approaches and understanding. Today, AI copilots are revolutionizing productivity, creativity, and efficiency across diverse domains from general use (Microsoft Copilot [128]) through coding (GitHub Copilot [129]), sales, marketing, and customer service (Salesforce Einstein [130]), and cross-system communication (Moveworks [131]) to Stable Diffusion [132] that replaces a need for multiple specialized applications by handling tasks like report building, customer service replies, sales emails, and more.

To advanced users, chatbots serve as a sixth sense or a counterpart able to respond to questions. The quality of the responses is disputable as a whole, however better and worse questions exist that generate better or worse answers. Labeled as *prompting*, the art of questioning that stands behind any authorship creation since time immemorial experiences a revival and an upgrade to the next level. Throughout history, originally invented to transmit and preserve information of prime importance and quality, literacy has evolved from a unique skill and knowledge of a narrow circle of high-ranking experts [133] to general elementary education of the entire population that spread to apply in everyday life, also for banal purposes. Similarly, today, the range of computer non-professionals is rapidly expanding who benefit from the bold AI support to create specific amateur applications - or are getting used to exploiting handy AI applications for the most common needs. Musavir.AI [134], for example, offers to generate movie-like renders, architecture, and sci-fi

in general, tailor refined lighting aesthetics - or, through MyAvatar feature, generate diverse new hairstyles for oneself. Enjoyed as a hobby or for entrepreneurial purposes, countless popular text-to-image models set another class of AI applications: in March 2024 released in version 1.0, Ideogram [135] is an example.

AI does not hesitate to enter good old familiar tools like Microsoft Windows, Excel, or Word to enhance productivity, data analysis, and decision-making. Ideas is an AI-powered insights service that identifies trends, patterns, and outliers in data sets, allowing users to analyze and understand their data rapidly and providing functionalities like taking a picture of a printed data table using Android device and converting it into an editable Excel spreadsheet, automatically converting the pictures into fully editable tables, eliminating manual data entry and allowing formulas to spill over multiple cells dynamically, adapting to changes in data size [136]. Windows 11, too, introduces several AI-powered features that enhance productivity, creativity, and overall user experience, including copilot, updated paint functionality for photo editing and "art" creation, photo movie editor, clipchamp assisting with footage editing, enhanced ink drawing, smart (and safe) app control, and security tool [137].

Targeting professional users, Devin is a groundbreaking AI software engineer designed to revolutionize software development through collaboration with humans. Devin's key features include autonomous coding assistant - a fully autonomous software engineer, capable of planning and executing complex coding tasks, learning new technologies on the fly, building and deploying full applications from scratch, automatically finding and fixing bugs, training its own AI models, and contributing to production codebases. Devin handles long-term reasoning, dynamic learning, coding, and design skills [138].

It is a real influx of ever-new AI tools what the present experiences. Thousands of startups, along wealthy industry tycoons and their spin-offs, have begun applying generative AI to create virtual assistants who can respond appropriately to human requests with natural language processing with dialogue management capabilities, image -, and video processing. During 2023, more than ten thousand new applications were released [139]. Obviously, keeping up with all of them is beyond the limits of any human - unless AI is involved. In March 2023, OpenAI released yet another state-of-the-art large language model, GPT-4, equipped with multimodal capabilities and superior performance on benchmarks designed for humans. Stanford released Alpaca 7B, a relatively small open-source model that matches the performance of GPT-3.5. Concurrently, Google introduced Bard that became Gemini recently - the *"creative and helpful collaborator, here to supercharge your imagination, boost your productivity, and bring your ideas to life"* [140], Chinese search engine behemoth Baidu released (just before April 14, 2023) Ernie Bot [141], Alibaba Group introduced (April 11, 2023) chatbot Tongyi Qianwen, Sense Time SenseChat [142], and Amazon its new AI application Bedrock, which makes available to developers generative tools for creating texts and images.

Through 2023, Amazon has been offering four AI applications, the native Titan included, and employs the most people in AI development – more than both Google and Microsoft [143]. In addition, Amazon has been investing massively into the OpenAI rival Anthropic: the September 2023 deal put \$1.25 billion into the company, the investment to be topped to a maximum of \$4 billion, which became a reality at the end of March 2024. *Lacking the capability to develop adequate models independently, companies like Amazon and Microsoft have had to act vicariously through others, primarily OpenAI and Anthropic. The two have reaped immense benefits by allying with one or the other of these moneyed rivals, and as yet have not seen many downsides* [144].

X's or Elon Musk's answer to OpenAI's Chat GPT is Grok - an AI search assistant designed to answer your queries while keeping you entertained and engaged. xAI launched the Grok-1.5 chatbot on the social media platform X at the end of March 2024, with the enhanced version aiming to surpass current AI technologies. *On popular benchmarks, Grok-1 is about as capable as Meta's open-source Llama 2 chatbot model and surpasses OpenAI's GPT-3.5, xAI claims* [145]. And thirdly, Databricks released DBRX concurrently, a new generative AI model akin to OpenAI's GPT series and Google's Gemini. Available on GitHub and the AI dev platform Hugging Face for research as well as for commercial use, base (DBRX Base) and fine-tuned (DBRX Instruct) versions of DBRX can be run and tuned on public, custom, or otherwise proprietary data. *Databricks says that it spent roughly \$10 million and two months training DBRX, which it claims (quoting from a press release) "outperform[s] all existing open source models on standard benchmarks." But — and here is the marketing rub — it's exceptionally uneasy to use DBRX unless you're a Databricks customer: to run DBRX in the standard configuration, a server or PC with at least four Nvidia H100 GPUs (or any other configuration of GPUs that add up to around 320GB of memory) is needed* [146].

In January 2024, OpenAI has laid another touchstone, launching a specialised AI chatbot store called GPT Store that offers a wide range of applications based on advanced AI models such as GPT-4 and DALL-E3

[147]. The store is accessible free (so far) to subscribers of OpenAI's premium programs like ChatGPT Plus, ChatGPT Enterprise, or ChatGPT Team. The chatbots are categorised into areas such as lifestyle, writing, research, programming, and education, democratizing the accessibility of AI technology by enabling users to create and utilize custom GPTs tailored to their specific needs. No less importantly, contrary to the initial hype surrounding the technology, opening GPT store indicates that in the upcoming years the emphasis will shift away from artificial general intelligence (AGI) towards more specialized chatbots designed to meet personalized demands.

As another example of the trend, Google Maps has got an AI boost [148]: immersive view that includes street-level imagery and 3D models of any location, live traffic and weather simulations, navigation around congestion and unfavorable weather; lens in maps - AI-driven feature for enhanced environmental understanding that identifies and labels every object with the camera; improved navigation providing highly accurate and detailed maps for navigation, information about local businesses, landmarks, and must-see spots, and exploration along the route, including comprehensive information on charging stations and compatibility and convenience for electric journey; new areal view API (application programming interface) enabling 3D birds-eye view for applications and websites and Google's AI for object identification and extraction; photo-first results for search terms and AI image recognition for accurate matching [149].

Concerning image processing, also applications Leonardo AI, Playground AI, Bluewillow AI, Bing AI, Adobe Firefly, or Bright AI deserve attention [150]. In March 2024, Google has introduced VideoPoet, literally ChatGPT for text-to-video, image-to-video, and video editing [151]. The year-to-year progress in the quality of the images and generated is undeniable [152]. At the beginning of 2024, significant text-to-video applications were released: Sora by OpenAI (to be discussed in section (5) [153]), Lumiere by Google [154], and EMO [155]. Powered by a state-of-the-art "space-time" neural network and self-supervised learning (its competitors alike), Lumiere effortlessly crafts 5-second video clips in one go (Sora works up to 1 minute, but ...). It masters stylized generation, video editing, cinemagraphs, and video painting [156]. EMO, short for emote portrait live, is an artificial intelligence system developed by researchers at Alibaba's Institute for Intelligent Computing that transforms a single image and audio input into expressive portrait videos. It can take a single portrait photo and bring it to life using a direct audio-to-video synthesis approach, bypassing the need for intermediate 3D models or facial landmarks. The system ensures seamless frame transitions and consistent identity preservation throughout the video, resulting in highly expressive and lifelike animations - not only convincing speaking videos but also singing videos in various styles [157]. In the context of this review, it is the state-of-the-art neural networks and learning strategies allowing for understanding the physical world and variety of action scenes [158] - a common ground with approaching to architectural concepts - what makes these applications worth attention.

In March 2024, OpenAI introduced a new language model, the Voice Engine, which can clone a speaker's voice based on just 15 seconds of audio recording. The technology has been successfully tested in collaboration with several partners, for example, to create a synthetic voice for a young girl who lost the ability to speak normally due to a brain tumor: such is a class of use allowing people with speech disorders to talk in a natural-sounding voice that is developed by the Brazilian company Livox. Further applications show in education, where it helps generate educational content, or when localizing videos into different languages while preserving the speaker's original accent [159]. OpenAI has not yet released the model as a separate product because of possible misuse - creating deepfakes as an example.

And the race continues; one month later, Microsoft Research introduces VASA-1, which takes a single portrait photo and speech audio and produces a hyper-realistic talking face video with precise lip-audio sync, lifelike facial behavior, and naturalistic head movements generated in real-time [160].

Concurrently, in an interview with Bill Gates, Sam Altman of OpenAI revealed what was coming up in GPT-4.5 (or GPT-5) [161]. Altman highlighted integration with other modes of information beyond text, better logic and analysis capabilities, and consistency in performance as priorities for AI progress over the next two years. OpenAI has already launched image and audio capabilities for their current models; but those capabilities will go much further, aligning with people's desire for AI systems engaging with more elements of the real world beyond just text. Improving logical reasoning and inferencing is another key priority for the next two years. The aim will be for models to become better at analyzing prompts, synthesizing information, and drawing insightful conclusions rather than just generating speculative or untrustworthy responses. Hopefully, reliability will stem from better reasoning.

As hectic and chaotic as this overview may seem, actually, it is quite a polished picture - which, moreover, skips the field of hardware in general, and (with some exceptions) financing and investments. The

overview lacks attention to architecture and the development of the built environment: this is not an omission or a narrative intention, but an image of reality: abundance and glut on the one side and lagging on the other. Architecture and the built environment are on the edge of current attention to application development and investment in specialized tools and productive AI environments.

Spatiality and robotics to enhance AI in architecture

Also, VR introduces processes of AI-driven 3D model creation: HTC Viverse offers 3D model generators for the VR/XR community. AI enables text to be turned into VR 3D models, 2D images transformed into dynamic 3D models, and even extract intricate models from videos [162]

Naturally, AI merges with robotics. Neural networks taking video in and delivering trajectories out have taught Figure-1 to make coffee after watching humans at the activity [163]; other repetitive, though variable manual tasks can follow to deploy end-to-end AI solutions. Recently, Google DeepMind and Stanford researchers introduced Mobile Aloha - an open-source robotic system capable of completing complex tasks like cooking, cleaning, and more [164]. Noting deserves that humans remotely controlled the cooking skills.

Advancements include AutoRT for data collection, SARA-RT for faster transformers, and RT-Trajectory for better motion generalization [165]. The solutions deliver critical improvements in both areas essential for progress in AI-driven robotics: improving the robots' ability to generalize their behavior to novel situations and boosting their decision-making speed. The systems also can understand practical human goals and enable robots to gather training data in new environments.

Nvidia is entering the field of humanoid robotics with the ambitious GROOT project. A new model of artificial intelligence is intended to give robots a human level of understanding and dexterity. Machines powered by the GROOT will be able to understand natural language and mimic human movements just by observing. As a result, they quickly learn the coordination, skills, and other abilities needed to effectively navigate, adapt, and interact in the real world. In the training workflow, Nvidia Isaac Lab - a lightweight reference application optimized for robot learning - and Nvidia OSMO - a cloud-native workflow orchestration platform that scales workloads across distributed environments and coordinates training and inference workflows across different Nvidia systems - are involved. Further, Nvidia has established a partnership for this project with leading robotics companies such as Figure AI, Boston Dynamics, and Apptionik.

A general-purpose foundation model designed specifically for humanoid robots, GROOT is trained using imitation and transfer learning, which allows humanoid embodiments to learn from a small number of human demonstrations. The model leverages reinforcement learning to enhance its understanding and decision-making abilities. Taking multimodal instructions (including natural language) and past interactions as input, GROOT shall generate robot movements by analyzing video data, enhancing further its adaptability and responsiveness [166,167,168].

Similarly, specific SMLs will likely come into play when developing AI to deploy in architectural designing and planning development of the built environment.

Keeping in mind AI deployment in architecture and the built environment development, the successes in robotics deserve particular attention. Imitation-based learning, self-learn, or transfer learning are the common denominators: if Figure-1 could have learned to prepare coffee by mimicking a human, AI must be capable of learning to develop architectural designs by mimicking human architects at work. Not by a video like Figure-1 but by following the progress of the solution in a CAD or VR environment. This learning will probably be much more complex and demanding, but ... step-by-step, it will proceed. After all, it should be a question of a kind of scaling up.

Milestones

In terms of image recognition, it was AlexNet that set a benchmark in 2012 [169]. Designed by Alex Krizhevsky in collaboration with Ilya Sutskever and Geoffrey Hinton, *a composition of eight layers, the first five convolutional layers, some of them followed by max-pooling layers, and the last three fully connected layers*, AlexNet competed in the ImageNet Large Scale Visual Recognition Challenge on September 30, 2012. The network achieved ... an error of 15.3%, more than 10.8 percentage points lower than that of the runner up [59]. The paper introducing AlexNet is considered one of the most influential in computer vision, having spurred many more papers published employing CNNs and GPUs to accelerate deep learning. As of early 2023, the AlexNet paper has been cited over 120,000 times according to Google Scholar [170].

When IBM's Arthur Samuel developed a machine learning system for playing checkers in 1959 [171], he used thirty-eight considerations determining the strength of a position – like the number of pieces on each side, the spatial distribution of stones, mobility and space, safety and risks, and on. By 1990, the IBM team working on the chess supercomputer Deep Blue used eight thousand such considerations. *This chess evaluation function ... probably is more complicated than anything ever described in the computer chess literature*, put the team lead Feng-hsiung Hsu – and it deserves noting in this paper's framework that perhaps similarly complicated is the structure of considerations on – let's say – a residential building spatial layout development ... In Deep Blue nonetheless, those thousands of considerations were brought into balance neither by trial and error (as would be typical for reinforcement learning) nor by human labeling of diverse alternatives (as in supervised learning) but by imitation of human moves employing one of the that time novel machine-learning technologies [172].

Fifteen years later, DeepMind's AlphaGo system finally implemented Arthur Samuel's vision of a system that could concoct its own positional considerations from scratch. Instead of being given a big pile of thousands of handcrafted features to consider, it used a deep neural network to automatically identify patterns and relationships that make particular moves attractive, the same way AlexNet had identified the visual textures and shapes [173].

The second and even more important lesson learned shall be a focus on the process instead of the output/result, as implemented already in Deep Blue. In October 2017, Google DeepMind brought the process-focus instead of the output-result paradigm to a (so far) ultimate level by going through with the playing-against-itself strategy in AlphaGo Zero [174].

AlphaGo combines advanced search tree with deep neural networks. The "policy [neural] network", selects the next move to play, and the "value network", predicts the winner of the game: a reinforcement-learning paradigm. *Initially, the developers introduced AlphaGo to numerous amateur games to help it develop an understanding of the play. Then it played against different versions of itself thousands of times, each time learning from its mistakes. Over time, AlphaGo improved and became increasingly stronger and better at learning and decision-making. AlphaGo went on to defeat Go world champions in different global arenas and arguably became the greatest Go player of all time* [175]. AlphaGo Zero is a next-level version of the Go software. AlphaGo's team article published in the journal Nature on October 19, 2017 introduced a novel learning strategy: by playing games against itself, without using data from human games, AlphaGo Zero exceeded all the old versions of AlphaGo in 40 days [176].

Combining biological inspiration, cognitive capabilities, adaptability, and open standards that go beyond traditional approaches, the distributed intelligence strategy proposed and in development by Verses may mark the next milestone, perhaps matching the new theory of mind and brain proposed by Jeff Hawkins and the Numenta team; both innovations presented in sections (3) and (4) of the paper. If these visions prove feasible, AGI might be at hand - and perhaps, an AI breakthrough in architecture and the built environment development along with it. Sam Altman claims to see hints of AGI in the embryos of the GPT-5 model. As early as 1999, futurist Ray Kurzweil predicted the advent of AGI by 2029. Metaculus, a platform that tries to compile various scientific predictions of future developments, said in 2022 that a very weak AGI would be available in 2042, now the consensus is that it would after all, it could have been created by 2030. However, Niall Fergusson states that the brain's "computational capacity" can handle 100 trillion parameters. *The current most advanced GPT-4 model works with 1 trillion parameters. Handling a hundredfold increase can be extremely difficult* [177].

The milestones displayed represent a history, in which imitation-based learning and self-learn gradually promoted. The history, on the one hand, culminates today in projects like Figure-1, Sara, GROOT, or Phoenix (introduced in section (5) of this paper) and, on the other hand, can lay a path to a truly efficient and productive deployment of AI in architecture. The second would be a game-changer that would demand the abandoning of a substantial proportion of the recent paradigms relying on statistical approach and input-output pairings, and GANs probably, too.

The Black Box problem, security issue, and a threat to humanity

With all these nice results, it is not clear what these models are learning, as Mathew Zeiler puts an issue [178]. Inevitably, due to the learning paradigm, AI systems are vulnerable to various types of attacks and data "poisoning". In addition, by hostile input or action, an attacker can gain unauthorized access or control over the system.

Another significant challenge in the development and use of AI systems is the black box problem. Retrieving the training stock that cannot be but a pile of digital data, general pre-trained transformers – ChatGPT typically – compress the content. *This compression is lossy as in the case of jpeg: we can imagine ChatGPT as a blurred jpeg of all text information on the web* as Ted Chiang [179] puts it. *The algorithm preserves a part of the information, just as a jpeg preserves much of the information of a higher-resolution image, but if looking for the exact bit sequence, it is not there. Only an approximation is always the result. However, since this approximation renders in the form of grammatical text, which ChatGPT is excellent at producing, it is usually acceptable. It is still a blurry jpeg, but the blurring occurs so that the image as a whole does not look less sharp.* This comparison to lossy compression is not just a way to understand LLM's ability to repackage information found on the web using other words. It is also a way to understand the "hallucinations", surprising, or nonsensical answers to factual questions that LLMs are all a bit prone to. These are unsurprising results of compression; if a compression algorithm is designed to reconstruct text after 90 % of the original has been discarded, a significant portion of what it generates can only be fabricated from scratch.

Obviously, such a fabrication is impossible for humans to interpret or understand; a lack of transparency and accountability in AI decision-making is a result. Such a threat cannot be overestimated as it is inherently embedded in the nature of the learning process. The algorithm does not work – because by definition it cannot work – with the categories "true" or "false", but deduces the degree of conformity or deviation according to patterns arrived at by own judgment, either without human supervision or under human supervision or direction, but always covertly in detail. Confusion of cause and effect in access to the data, whether training or task ones, diverting attention to the background of graphic inputs (bokeh) instead of their core, or similar is usually shown as the cause.

The bokeh salience feature of AI provides a comprehensive clarification of „famous“ *Coop Himmelb(l)au "machine hallucinations"* [10,11,12]. Not the creativity of AI appears, but a misleading perception of visual information hidden in the algorithm's black box; not creativity, but an error and accident. Computer hallucinations by unintended bokeh salience are examples of how technology can be misused either to manipulate and misinterpret visual information, to fake art, or to distort scientific research. The AI development community deserves credit for looking for and already delivering the first applications that solve this problem [180,181]. As a response, techniques also evolve that can be used to diagnose problems with the network's training, identify biases in its decision-making, or optimize its performance for a specific task [68,69].

Starting with "fake news" up to cyber-attacks and war escalations, the issue is much more than "hallucinations". It has been nicknamed "Oppenheimer moment" labeling a pivotal event in which a breakthrough technology or scientific discovery reaches a level of progress that has far-reaching consequences for society, ethics, and humanity. While nuclear technology was used in combat and resulted in hundreds of thousands of lives affected, it dramatically changed the world in general having delivered the energy our lives would have hardly be the same without. Just as nuclear technology dramatically changed the world in the last century, the development of AI is prone to revolutionize everything from industries to culture, politics, and humanity in this century and the centuries to come. The nuclear era has raised profound questions about the responsible use of scientific discoveries – no reason to exclude AI. Encountering responsibility, AI raises an unprecedented question: will man still be the most intelligent creature on the planet? Or, later this year or the next, GPT-5 will sideline humanity [182]. Alexander Karp, head of Palantir (founded in 2004 by him and Peter Thiel) that provides data mining services to government agencies like the Department of Defense, the FBI, and Immigration and Customs Enforcement, recently claimed that we are currently experiencing a kind of "Oppenheimer moment" concerning the development of AI for military use and that there is no way for free democratic states not to continue developing AI, even if it is a risk to humanity, because other (non-democratic) countries will certainly not stop their AI development [183].

In addition, an opinion prevails today that technological development (of AI) is ahead of international law. Striving to react in practice, at the end of March 2023, Italy blocked ChatGPT [184] to secure the privacy of people and tycoons of global business claim pausing "giant AI experiments" – read the development of AI for six months [185] to prevent an unmanaged reaching of the singularity phase of AI development, when a spontaneous technological growth breaks out, and not only the society begins to be irreversibly changed by the effects of technology but humanity loses all control over further development of AI.

The subject of concern is not only security immediately, but also the power of the big players in the AI field. Elon Musk (who fails to keep up with the cutting edge in this area) asserts (addressing OpenAI) that tech giants are inciting existential fears to evade scrutiny [186].

As another example, an issue can be AI enabling cyber criminals to generate lifelike images and make convincing videos that impersonate taxpayers to steal their refunds. Tax identity fraud "is a great crime, because so many tax refund dollars are transacted" and it's harder to spot suspicious behavior with a once-per-year transaction, Ari Jacoby, founder and CEO of cybersecurity firm Deduce. Tax professionals may also be caught off-guard by cyber criminals trying to get them to hand over sensitive client data by posing as real taxpayers. AI is particularly difficult for tax professionals to deal with "because it is self-learning, trying techniques and failing until it succeeds" [187].

The issue does need to be a straight risk of war. OpenAI's Sora video generator (discussed in section (5) of the paper), approaching being available to the public this year as confirmed by the company's chief technology officer Mira Murati, illustrates the everyday nature of the fears and risks „at hand“. Concerns arise about possible misuse, for example for the creation of deepfake pornography or compromising footage of public figures. Polls show that most Americans support introducing regulatory measures and safeguards to prevent abuse of AI video-generating tools. More than two-thirds of respondents believe these tools' developers should be held legally responsible for any illegal activities. According to cyber experts, there is an urgent need to adopt rules for user authentication, content flagging, risk assessment, and restrictions on the export of AI-generated videos [188].

Since Italy blocked ChatGPT, the EU - slowly but persistently and consistently as usual – has been finding and codifying a balanced approach to AI, which has been treating possible risks in a structured way but, at the same time, shading attention to cases where the technology can support the economic development, a comprehensive sustainable development, and well-being of the population, has been establishing support for the development of AI [189,190,191].

Global powers (also) in AI do not lag behind Europe. Significant steps to address the safety and security implications of AI have been undertaken including the US/China agreement to prevent the development and deployment of AI-powered weaponry that could pose risks to global security [192]. International agreement on AI safety gains signatures from 18 countries: the US, UK, and other major powers (excluding China) have unveiled a 20-page document with general recommendations for companies developing and deploying AI systems. These recommendations include monitoring for abuse, protecting data from tampering, and vetting software suppliers [193]. During a British summit, China agreed to work with the United States, the European Union, and other countries to collectively manage the risk from AI. This collaborative effort recognizes the need for international cooperation in addressing the challenges posed by AI technology [194]. Naturally, on national levels, the legislation starves to catch up with the hectic development, too [195].

Lawsuits

Obviously, training dataset quality in terms of size, comprehensivity, and relevance is of critical meaning for an AI application's performance. Not always but often, the objects that assemble the set represent the intellectual property of respective authors. The authors feel mishandled and affected if someone - no matter whether a human or an AI - takes and compiles their creations (or their digital representations) to put them on display or to submit them individually. *At issue, mainly, is generative AI's tendency to replicate images, text, and more — including copyrighted content — from the data that was used to train it. Indeed, image-generating AI models like Midjourney, DALL-E, and Stable Diffusion replicate aspects of images from their training data. As a result, together with generative AI entering the mainstream, each new day brings a new lawsuit. Microsoft, GitHub, and OpenAI are currently being sued in a class action motion that accuses them of violating copyright law by allowing replicating licensed code snippets without providing credit. Two companies behind popular AI art tools, Midjourney and Stability AI, are in the crosshairs of a legal case that alleges they infringed on the rights of millions of artists by training their tools on web-scraped images. And stock image supplier Getty Images took Stability AI to court for reportedly using millions of images from its site without permission to train Stable Diffusion, an art-generating AI* [196].

Let us zoom in: what is the issue? Is it the use of other authors' performance - even though an adaptive use? Or is it the use of other authors' performance by a machine - without a creative input of a man? Labeling them as paraphrasing, history and present is rich in cases of such use, shall the examples be William-Adolphe Bouguereau or Alexandre Cabanel paraphrasing Botticelli's *Birth of Venus*, Joos van Cleve's paraphrasing-slash-counterfeiting of Leonardo da Vinci's *Mona Lisa*, Michal Ozibko paraphrasing *Girl with a Pearl Earring* by Jan Vermeer van Delft, Tadao Cern paraphrasing *Selfportrait* by Vincent van Gogh, Peter Lindberg featuring Julianne Moore in paraphrases of Gustav Klimt's or Egon Schiele's portraits, Paul Cezanne paraphrasing Édouard

Manet's *Olympia* paraphrasing Francesco Goya's *Maja*, and many others paraphrasing many other works by many other authors. Probably none of the cases of paraphrases of authors' works has provoked rejective reactions either from the authors (not only because they are often already dead) or from the professional public; on the contrary, paraphrases are often perceived as a tribute to the original author.

It looks like the problem is the machine - the AI taking what it can get. Let's leave aside that that's not entirely true either - AI only takes what trainers - supervisors tell it to or allow it to go to if they do not serve the AI directly with it. The anonymity of the „independent machine's“ tackling renders to be the core of the issue, underlined by the black-box nature of AI that, in a way, hides even more the author. No wonder then: How should AI cope with expectations, and legal paradigms, too, that emerged and evolved without having even a glimpse of a notion of something like AI? Maybe general understanding and legal framework only need to catch up with a new, unprecedented, and unexpected phenomenon.

Nevertheless, the latest evolution does not come close to solving the problem, rather the opposite. In late December 2023, the New York Times, publisher of the famous newspaper of the same name, filed a lawsuit against OpenAI and Microsoft for massive copyright infringement in their ChatGPT and Copilot products. The entitlement of the complaint is declared the crucial role independent journalism plays in a democratic society and how The New York Times finances its activities: through subscriptions and licensing arrangements. Also deploying a subscription model, ChatGPT competes directly with the New York Times. The defendants' activities are far from unprofitable: "The use of valuable content belonging to others was extremely lucrative for the defendants, ... OpenAI being on track to generate \$1 billion in sales by 2023, and Microsoft increasing its market value to a record \$2.8 trillion (up \$1 trillion from a year ago)." the lawsuit states [197].

The infringement is twofold, according to New York Times representatives. Primarily, ChatGPT can reproduce copyrighted content "word for word" on a larger than small scale. The difference between how search engines treat newspaper content and how a chatbot treats it is obvious. ChatGPT helps a user bypass the security of a locked article. Of course, internet search engines also show parts of the text. That is what the current online information environment is built on, and the New York Times benefits from it, too. However, the snippets that ChatGPT or Bing Chat (later Microsoft Copilot) lists are much longer and thus can be argued to infringe significantly on the editorial copyright. The second type of infringement is no less serious, according to the NY Times: it is the "hallucinations" that occur in LLMs by the very principle of their operation, generating a plausible-looking text, but one that is out of touch with reality [198].

The New York Times (and Times) lawsuits are not isolated; more news organizations—The Intercept, Raw Story, and AlterNet among others—have filed separate lawsuits against OpenAI, alleging ChatGPT-related copyright infringement [199].

In addition, the allegations are not coming only from outside the AI industry. On the precious 29th day of February of the leap year 2024, Elon Musk filed a lawsuit against OpenAI and its CEO Sam Altman, alleging they have abandoned the company's founding agreement to pursue AI research not for profit but for the benefit of humanity as he, Altman, Greg Brockman, and other co-founders had agreed. [199]. The issue is that upon releasing its GPT 4 model, OpenAI Inc. has transformed into a closed-source de facto subsidiary of the largest technology company in the world - Microsoft. The true motivation of the lawsuit may reveal later; at the time being, a business motivation cannot be excluded.

Also, Verses (see section (3) of this paper) asserts that generative AI market leader OpenAI is breaching its charter by not engaging with Verses' revolutionary approach. Having taken a full-page advertisement in the New York Times in December 2023 that has been announcing its progress and methods [200], Verses challenged OpenAI to cooperate — asking the nonprofit firm to fulfill its charter promise to "stop competing" with any "value-aligned, safety-conscious project that comes close to building AGI before we do." Verses believes it qualifies for such cooperation and offers collaboration to ensure safe and beneficial AGI development. Therefore, the competition is the friction point as Verses with its roughly 100 employees and \$65 million in revenues feels - no wonder – threatened when entering a battlefield of tycoons. Even if its founders are right in theory, there is no guarantee that their approach will take off in the market.

Big expectations - and challenges

During the World Economic Forum Davos gathering in 2024, AI was given a big say. Chief AI Scientist at Meta Yann LeCun was speaking on the power of open-source AI and how to keep AI progressing fast, CEO of Cohere Aidan Gomez assured that AI is about to accelerate even more due to improvements in architecture and

hardware, Microsoft's CEO Satya Nadella's speech presaged how AI can improve lives, from healthcare through education to products and services, Rwanda Innovation Minister Paula Ingabire reported on the economic boost of the country of up to 6% from leveraging AI, CEO of IBM Arvind Krishna forecasted AI to generate 4 trillion dollars of annual productivity in economies worldwide, U.S. Senator of South Dakota Mike Rounds asserted that AI will be transformational in healthcare and that the U.S. public will become optimistic about AI after seeing the quality-of-life improvements, Qualcomm's CEO Cristiano Amon explained advancements in AI released very recently at CES 2024 [<https://www.ces.tech>; <https://www.ces.tech/>] such as the ability to have a conversation with one's car, and Google's CFO Ruth Porat was bringing to mind the importance of cybersecurity and preventing misinformation in the age of AI [201].

Exploring the transformative potential of AI on the global economy and humanity, the International Monetary Fund expects AI to redraw the map of the world economy. The Fund's research asserts a global employment impact of up to 40% of jobs worldwide and 60% in advanced economies, addressing primarily skilled jobs. Advanced economies are expected to experience both significant risks and opportunities from AI while emerging markets and developing countries, less exposed to AI, may face less immediate disruption; the second, however, are assumed at risk of falling behind due to insufficient infrastructure and skilled labor strength. Within countries, AI can exacerbate income and wealth inequality. It can boost the productivity and wages of those who can use AI but leave behind those who cannot. Policymakers are requested to address the challenges posed by AI and focus on comprehensive social safety nets and retraining programs. A need for both developed and developing economies to adapt to the AI era is regarded as urgent. Advanced economies are encouraged to focus on AI innovation and regulatory frameworks while emerging and developing economies should invest in digital infrastructure and workforce skills.

However, some skepticism remains as to whether the application of artificial intelligence will usher in a new era of sustained acceleration in productivity. US Federal Reserve Chair Jerome Powell's take last month was "probably not in the short run," though "probably maybe in the longer run." What deserves attention is the emerging major difference between the world's two largest economies concerning how AI research and applications will be funded. And that, in turn, may affect the extent to which experimentation and diffusion of generative AI, or that which creates new content, evolves in the US versus China.

The Chinese Communist Party is making it increasingly clear that it wants to call the shots on how capital gets allocated. Beijing reported this week that president Xi Jinping led a call for the party's Central Committee to set up a mechanism to steer technology work in the country. In the meantime, a debate surfaced on whether to shut the domestic stock market to initial public offerings—something hardly thinkable in the US. American capital markets are just the opposite: a forum that allows anything to happen, including hype and mania, as witnessed this week with Nvidia corp.'s record surge. While that makes US markets vulnerable to crashes and volatility, it also means they are a venue for funding big dreams. And that may prove decisive in determining which of the two big rivals enjoys the productivity gains that—sooner or later—materialize from AI.

And Europe? Rather unexpectedly, Europe shows the capability to attract talented people for AI and European education systems prove to be proficient enough to generate about the same amount of top AI experts as the US. This good news breaks the notion that Europe is unable to generate experts. However, so far only half of them stay to work in Europe. Among other things, this relates to the strength of the capital market, the business environment, and the flexibility of jobs and entrepreneurship [202].

AI market leaders naturally have their ideas about what comes next. Microsoft and OpenAI (which is substantially co-owned by Microsoft [203]) have ambitious plans to collaborate on a massive data center project with an estimated cost of up to \$100 billion. Aiming to create cutting-edge infrastructure to support AI R&D, the AI supercomputer "Stargate" anticipated to be operational by 2028 is at the heart of the project. The project budget is not only hundreds of times larger than the cost of any of the existing "competitors"; for comparison, it is equal to the gross domestic product of Bulgaria in 2023 [204], leaving behind the gross domestic product of two-thirds of the world's national economies.

With its renowned customer drive and unique ability to connect research and development with marketing, Apple Inc. has set teams recently to launch investigating a push into personal AI-driven robotics, a field with the potential to become one of the company's ever-shifting "next big things." After the electric vehicle project gets nixed, the search is on for new growth sources: naturally, AI-driven robotics is the goal - and, when it is Apple, where else than in the field of home devices [205]?

AI has become a phenomenon in economic, cultural, and also political terms. In March 2024, Demis Hassabis, co-founder and CEO of Google's AI subsidiary DeepMind, has been knighted for his contributions to

the field of AI. His work at DeepMind—particularly the development of the AI system AlphaGo—has been appreciated as pivotal in advancing AI technology, and the knighthood has been declared to reflect his role in positioning the U.K. as a leader in AI research and development. Hassabis' acknowledgment reflects concurrently the broader impact of AI on the global stage, highlighting the significance of AI innovation in contemporary society [206].

Concurrently, OpenAI's co-founder and CEO Sam Altman worries about upcoming hardware parts shortages that would not allow raising the overall computing power quickly enough according to the needs of ever-novel and mightier AI algorithms. Proactively, he seeks to raise seven billion USD to establish a bold network of chipmakers [207]. In the shadow of AI, the future renders not only intelligent but groundbreakingly expansive, too. Symptomatically, the April 3, 2024, magnitude 7.4 Hualien earthquake in Taiwan showed Altman's worries not so detached from reality, though the reasons differed [208].

To fully appreciate AI expectations, a look at the amounts being invested is needed. Global corporate investment in AI from 2013 to 2022 goes to trillions and starts to be backed by the public sector, too (if Saudi Arabian wealth funds are considered public) [209]. Venture capital investors have lavishly funded a pipeline of additional upstarts; eight of the most prominent were recently valued at an average of 83 times their projected annual revenue in the process. In March 2024, the markets impatiently outlooked the 2023 leading companies' economic results - and breathed a sigh of relief when Nvidia showed to have beaten Aramco (after overtaking Amazon and Alphabet six months ago) to become the third-most valuable company in the world. The largest technology firms in the world known as the "Magnificent Seven", which includes Microsoft, Apple, Nvidia, Alphabet, Amazon, Meta, and Tesla, together have such a huge market value that they would create the second-largest stock market in the world only by themselves (the largest being the American New York Stock Exchange). The comparison can go further: Microsoft and Apple are each separately valued at a similar value to all firms listed on stock exchanges in France, Saudi Arabia, or Great Britain. Deutsche Bank warns that such a high concentration of titles has come dangerously close to the levels of 1929 and 2000, when there were huge falls in the markets [210].

Reminded in the pre-previous subsection, the clashes between tech tycoons may signalize an economic „bubble“ in the branch that also current economic data may indicate. Also, the Market cap to GDP ratio, nicknamed the Buffet indicator, achieving 193% currently, suggests that the (US) stock market is overvalued relative to GDP [211,212]. In addition, primarily the high-performing stocks of the "Magnificent Seven" have been driving much of the market growth while other stocks have been stagnating or underperforming. Market concentration risk renders apparent, strongly advising spreading investments across various sectors and companies; a "bubble" may be looming. Could the turn of AI R&D to new application areas - such as architecture and the built environment development - help to avert such negative prospects?



Fig. 2: Studio ZHA: Design idea for a project. Dezeen. 2024. https://www.dezeen.com/2023/04/29/patrik-schumacher-ai-generated-images-this-week/#disqus_thread. (accessed Nov. 16, 2023)

AI in architecture and engineering

The previous subhead demonstrated big expectations the field of AI lives with: nothing similar refers to architecture and the built environment development. As annotated in concluding the overview of AI production ecosystems, the scale of insight must change to examine this field.

Among multiple others, also Zaha Hadid (studio) met AI using the technology to render forms not so free to cease resembling antic temples patterns that served as imagery datasets to feed the GAN [213]. In doctoral research under the supervision of Patrik Schumacher of ZHA in 2017, Daniel Bolojan created Parametric Semiology Study using machine learning algorithms and other tools of gaming AI implemented in Unity 3D to model the behavior of human agents in order to test the layout of a proposed space [214].

Stanislas Chaillou and Nvidia company, and also others efforted providing AI applications to generate floorplans and apartment layouts [215]. ArchiGAN uses generative networks to create 2D and 3D building designs based on input parameters such as dimensions and space requirements. Another model is CityGAN, which generates drafts of city blocks and buildings. From a practical point of view and concerning the efficiency of deployment, the results of both applications are questionable - as in all other similar cases. On the principle of image-to-image translation with conditional adversarial networks (CANs), Phillip Isola Research Group [216] provided series of machine-generated facades following the "style" and character of the pattern deployed as the "input". [217,218] Introduced by the same team, Pix2Pix is shorthand for an implementation of a generic image-to-image translation using CANs [219]. Having investigated building façades generation, the team opened the door to architectural designing using GAN. Andrew Witt expanded on this work in his QUILTING exhibition; showcased as one linear, endless animation of an urban skyline, larger facade designs were created by enlargement of the final layers of Pix2Pix [220].

The use of GANs for floorplan recognition and generation was first studied by Zheng and Huang in 2018 [221]. Their Pix2PixHD [222] GAN architecture transferred floorplan images into programmatic patches of colors, and inversely, patches of colors were turned into floorplans. The same year, Nathan Peters [223] in his thesis at the Harvard Graduate School of Design researched laying out rooms in a framework of a single-family home footprint; an empty footprint represented in color patches was the output. Developed in 2019 by Kyle Steinfeld [79], GAN Loci tries to generate perspective images-like of urban scenes assembled with given facades-like textures, pathways, street furniture, pedestrians, cars, etc., by training to achieve the required "mood" - suburban, public park, etc. [224, 225] Blending the outcomes of Isola's team and Steinfeld's R&D, Sketch2Pix provides an interactive application for architectural sketching augmented by automated image-to-image translation [226].

Having investigated a different approach, Nono Martinez's thesis at the Harvard Graduate School of Design, 2016, deserves noting [227]. The idea of a human-in-the-loop rests at the heart of the method that tackles GAN as a design assistant. Martinez trained models for specific sketching tasks and proposed an interface allowing the human designer further "hand-elaborate" the model at any moment of the design process. Section (5) of this paper will appreciate and further elaborate the principle of human-in-the-loop.

Tom Mayne of Morphosis employed AI to develop *operational strategies to generate output that could never be predicted*. The studio developed *Combinatorial Design Studies: a Grasshopper definition of one formal study* elaborated by GAN technology *provided a range of further combinatorial options* [228]. Characterizing the 2010s state-of-the-art AI in architectural designing when giving credit to true creativity, adaptability, and intuition of a learning machine, a misconception reveals already at this elementary level when the research's mission statement finds itself in explicit contradiction with the theory that has (yet) not been disproved [229]. On the other hand, an approach of the breakthrough paradigm deserves appreciation.

Foster+Partners, another global-star architectural studio cannot stay aside; in its *Applied R+D team architects and engineers together with expert programmers combine the best of human intuition and computational rigor working with new technologies such as augmented reality, machine learning, and real-time simulation* [230].

In terms of practical use, based on the experience from other fields such as image processing, *predictive simulations* have been considered an etalon. ComfortGAN, for example, investigates the challenges of predicting a building's indoor thermal comfort [231]. Also structural design was on the lookout for AI. *Using variational autoencoders, for instance, research development at MIT investigates how diverse structures can*

be generated while ensuring performance standards [232]. However, due to the essential material liability of the structural design, the not yet-solved problems of the algorithm's black box that do not allow to rely on the machine curb so far the deployment of AI in structural design to the theory and conceptual drafting.

Typically deploying supervised learning, ZMO.AI promises interior design driven by AI that, however, soon reveals unable to cope with professional standards [233]. During 2022, hundreds of generative models were released to give the year the label: a proportion of them are (separately) no more supported in 2024. Stable Diffusion - *an unheard-of state-of-the-art AI model available for everyone through a safety-centric open-source license* [234] - has been creating hyperrealistic art, ChatGPT dared to answer questions about the meaning of life, and Galactica by Meta has been learning humanity's scientific knowledge but also revealed the limitations of large language models [235]. The innovation arrivals offered *hierarchical text-conditional image generation with CLIP latents, high-resolution image synthesis with latent diffusion models, a dataset (LAION-5B) containing 5.85 billion image-text pairs being used to train models such as Stable Diffusion and even CLIP itself, personalizing text-to-image generation using textual inversion, fine-tuning text-to-image diffusion models for subject-driven generation, text-to-video generation without text-video data, frame interpolation for large motion, trainable "bag-of-freebies" setting new state-of-the-art for real-time object detectors, building open-ended embodied agents with Internet-scale knowledge, human-level play in the game of Diplomacy by combining language models with strategic reasoning (Cicero), training language models to follow instructions with human feedback, language models for dialog applications, robust speech recognition via large-scale weak supervision (Whisper), instant neural graphics primitives with a multiresolution hash encoding, scalable large scene neural view synthesis (Block-NeRF), Text-to-3D using 2D Diffusion (DreamFusion), or Point-E - a system for generating 3D point clouds from complex prompts* [236]. None of these tools addresses the architectural field directly but early-adopting architects are getting acquainted with them to learn the benefits they provide: more to be zoomed on in the next section.

On an urban scale, attempts are ongoing to contribute by generating "typical style" road- and circulation patterns and networks using - among others - the Neural Turtle Graphics. [237,238] *Over the past decade, the deployment of online platforms has provided an adequate infrastructure to the end users*, [239] also to deploy Generative AI: (former) Spacemaker [240,241], Cove.tool [242], Giraffe [243], or Creo [244] are a few examples of this growing ecosystem, offering simplified access to AI-based predictive models [245], *generative design, real-time simulation, additive manufacturing, and IoT to iterate faster, reduce costs, and improve product*.

City digital twins are also embracing machine-learning algorithms. Commonly used in engineering and manufacturing sectors, digital twins are increasingly being adopted by cities for various purposes including emission reduction (from buildings primarily). AI-fostered digital twins also aid municipalities in managing traffic effectively, they contribute to economic development planning, assist in climate action planning and monitoring, and for facility-management purposes, they represent streets, buildings, trees, fire hydrants, and other urban assets, using both live and historical data. Notable city digital twin implementations include Las Vegas, Transport for London, and Mannheim. Additionally, ABI Research predicts that over 500 urban digital twins will be deployed by 2025, resulting in US\$280 billion in savings for city planners by 2030. AI strategies and models, in particular, contribute to addressing challenges related to data availability and awareness that render crucial for successful adoption of the technology [246].

Block-NeRF is a variant of Neural Radiance Fields that can represent large-scale environments, specifically, to render city-scale scenes spanning multiple blocks. Waymo built a grid of Block-NeRFs from 2.8 million images to create the largest neural scene representation to date, capable of rendering an entire neighborhood of San Francisco [247].

Not only start-ups, academia, and spin-offs of global architectural star-studios go in for AI: the global CAD-tycoon Autodesk runs Machine Intelligence AI Lab – and much of Autodesk's software, including Fusion 360, is (said to be) AI-enhanced and applying generative design today [248]. Nonetheless, as broad as all this listing may seem, the development of AI for and in AEC is still in its infancy, failing to catch up with LLMs, text-to-image processing, deployment of AI in internet search, content placement, and advertising, but also healthcare, pharmaceuticals, insurance, or justice referring to custody and bail [249].

In the narrow field of architecture itself, the results so far are, if anything, lowly. So far, none of the above applications has been widely used or appreciated in architectural practice. After years spent studying and (sort of) basic researching the limits of deployment of AI in architectural design, Stanislas Chaillou, one of the most distinguished protagonists of the field, accepted that AI is not capable of replacing the human architect: in architecture (more than in other professions) AI shall not replace human intelligence but augment

it. The question remains how. Today, as the co-founder of Rayon, a Paris-based startup, sermons *building the next generation of space design tools - a collaborative, online platform aspiring to provide an updated toolset to the "architecture of the 90%"* [250] - the *production architecture*, as this paper puts the term and elaborates it later. The results, however, do not seem to be able to address professionals; so far, the application more resembles customer "design" tools that furniture producers offer to their clients to ease their purchasing. [251,252]




Fig. 3: Svarovska, E.: PromeAI deployment at progressing from a simple volumetric model to near photo-realistic rendering. *Student's design*. Department of Architecture, Faculty of Civil Engineering, CTU in Prague. 2023. author's archive.

A challenge of creativity

Design starts with a *poiétic* idea emerging way before any "early stage design" CAD tool can apply. A misunderstanding of this starts with missing the *poiétic* starting points and nature of architecture (as alleged

in section (3) of this paper), even disregarding the general definition of architecture (as put by Encyclopedia Britannica in (3), too). The misconception is widespread and entrenched throughout the AEC field and is close to a rule in the CAD field; consequently, the situation cannot be better regarding AI-driven architectural design "creative" tools development. Displayed repeatedly in this paper, the intractable situation not only leads to wasting efforts on trying to make an algorithm truly creative but, so far, has been sidelining R&D in perspective fields such as generative-patterns-based pre-design, design-development support by "advice whispering" - continuous parameters assessing, design reviews, and evaluations of solutions or extended-reality technologies, imitation-based learning, spatial computing, novel networks- and algorithms paradigms, and computing efficiency and performance.

Nonetheless, the issue of *poiésis* in architectural designing (section (3) delves into the concept), represented by the marginalized phase of pre-parametric, conceptual ideas designing, has been addressed recently by Prague, the Czech Republic-based startup Wearrecho [253]. Released towards the end of 2023, the 1.0 operational version provides a comprehensive working platform that not only allows architects to make architecture as it deserves - from spaces, a diachronic way in space and motion - but also liaises this creative realm seamlessly with the parametric BIM domain (Autodesk Revit being the backbone of the latter). The unique workflow elaborates a virtual twin of the architecture to be designed; embedded in both domains - VR/XR and CAD/BIM - and allowing to be approached in the domain that better accommodates the feature to be used, the twin is a dual one keeping all properties and aspects regardless the domain from which they originate.

As inspirations from robotics give rise to the concept of production software for designing the development of the built environment as a robot, the question of creativity takes on a new form. Further elaborated in section (3) of this paper, *poiésis*, a precondition of true creativity, precludes parametrization, which is a starting point of CAD software. On the other hand, the seamless connectedness of the free, diachronic architectural creation and the parametric design realms as provided by Wearrecho may be opening a path to grasping at least a part of the domain of true creativity. Until the question is surveyed, the prospects shall be kept open considering either state-of-the-art CAD platforms or an extension beyond the parameters to become the world in which the design robots will learn and perform.

So far however, the deployment of AI in architectural designing has been lagging far behind any idea of a design robot. After purely experimentally getting acquainted with AI in the 2010s, along with many others involved in design disciplines and artistic professions, Zaha Hadid Architects studio adopted AI applications such as DALL-E to their standard workflows. The core is project development – support of – using AI images – via generating images that step-by-step proceed to articulate expression of the *poiésis* of the architecture within design concepts [254]. The approach sets a clear position in terms of the deployment of AI in architecture and artistic design: AI should be embraced by architects rather than met with skepticism or fear ... „I am not at all worried about facing the newly empowered competition enabled by AI“, Patrik Schumacher puts it. On the other hand, a critical voice sounds, too, such as The Guardian's architecture critic Oliver Wainwright's, who claims all this can be dismissed as superficial because it works via mere image generation [255], though putting architects' jobs ... at risk [256,257].

A calm, artisanal point of view can prevail in positive terms: challenging the *poiésis* – the architectural narrative represented by the step-by-step AI-generated images may have the ability to develop the conceptual idea more straightforwardly, or, better to say more quickly and more consistently. Labeling the approach as a *reverse prompting* can, perhaps, explain the technology by itself rather understandably. Economizing the process is a benefit of AI's deployment. Simulating human phantasy as this practical use of AI applications could be labeled, turns out to be an augmentation of an architect's creative potential - not replacing him, not making an architect obsolete, and not to confuse with true creativity as will be elaborated in section (3) of the paper. And the skill of prompting only confirms the role of „a pencil and brush“ in a trained architect's hand in the new era.

Tools like Lookx.AI and Krea.ai fill the role of the new pencil and brush in a straightforward manner. Lookx.AI allows the architect to generate images directly from his Sketchup drafts easily and quickly [258]. Furthermore, pairing the first Lookx.ai outputs with Krea.ai can follow [259]. Upscaling, and also enhancing details and correcting uneven lines produced by lookx.ai, krea.ai shows promising potential, too.

Architecture students tend to be early adopters of the new tools. At the School of Architecture, Faculty of Civil Engineering, Czech Technical University in Prague, Czech Republic, students use AI tools such as DALL-E3, PromeAI, Midjourney, Runaway, PhotoshopAI, Adobe Firefly AI, Inpaint AI, or Lookx.AI to shorten the way from a simple sketch to materiality-rich rendering or from photo to a video. The benefit is a time cut, valuable

for architecture students beyond popular opinion, and, in addition, uplift of the results of designs and their presentations. A new generation of architects that, instead of fearing AI, will master it as a natural creative and economical aid and tool, and will be eager to the advances in the field may enter the practice soon.

Though only experimentally, to "replace"—or rather enhance—the "traditional design process", an "AI design process" is being applied in the studio led by the author of this paper. Starting from the first encounter with the client through sketching, development of the CAD/BIM model of the architecture, and its visualizations or a tour in VR to final design, diverse AI tools apply concurrently to standard CAD tools like Sketchup, Enscape, or Revit, and Wearrecho - a unique working platform developed by the studio's internal startup that through *a dual virtual twin* of the architecture to be elaborated allows both for free diachronic creation of architecture from spaces (in Unreal Engine environment), parametric elaboration of the design (in BIM Revit environment), seamless switching between the two environments, and fluent communication between the design- and project stakeholders who can join the scene of the design both in place and remotely. After Midjourney [260] in the first phase, Stable Diffusion [261] joins the workflow pipeline to get the sketches ready for CAD/BIM model development. To provide visualizations or a VR tour, ControlNet [262] joins.

Beyond practical application, but still, in a very useful manner, studies of Carlos Bañon together with his online workshops are leading to mastery in the deployment of Stable Diffusion, Midjourney, and ControlNet [263,264]. Text-to-image, diffusion, and other AI tools are gradually entering the role of design assistants, quickly and tirelessly, against prompts, delivering variant impressions - depictions of possible architectural solutions with different levels of adjacency to the specified spatial parameters. The prompting technique improves constantly both in terms of prepared templates and request categories, and user training. Nevertheless, not exceptionally, the supplied images are ultimately rejected in general. Even then, the use of inspiration can remain. Mastery in using a tool is the challenge, the necessity of the craft and true creativity remaining untouched.




Fig. 4: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Development Scenario, Prague*. MS architekti, Prague. 2024. author's archive. The process launches with a 3D model of the current situation of the locality. To be continued in figures 4 to 11.



Fig. 5: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Development Scenario, Prague*. MS architekti, Prague. 2024. author's archive. *Poiétic* conceptualization: prompted, Midjourney (Model Version 6) offers impressions of the locality development options. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.



Fig. 6: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Development Scenario, Prague*. MS architekti, Prague. 2024. author's archive. Human-in-the-line drafts the volumes/buildings' capacities. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.



Fig. 7: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Developent Scenario, Prague*. MS architekti, Prague. 2024. author's archive. Respecting the set volumes only loosely, Stable Diffusion (3.0) delivers options for the architectural form. The view is set within the prompt. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.




Fig. 8: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Developent Scenario, Prague*. MS architekti, Prague. 2024. author's archive. Respecting the set volumes only loosely, Stable Diffusion (3.0) delivers options for the architectural form. Another view set within the prompt. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.

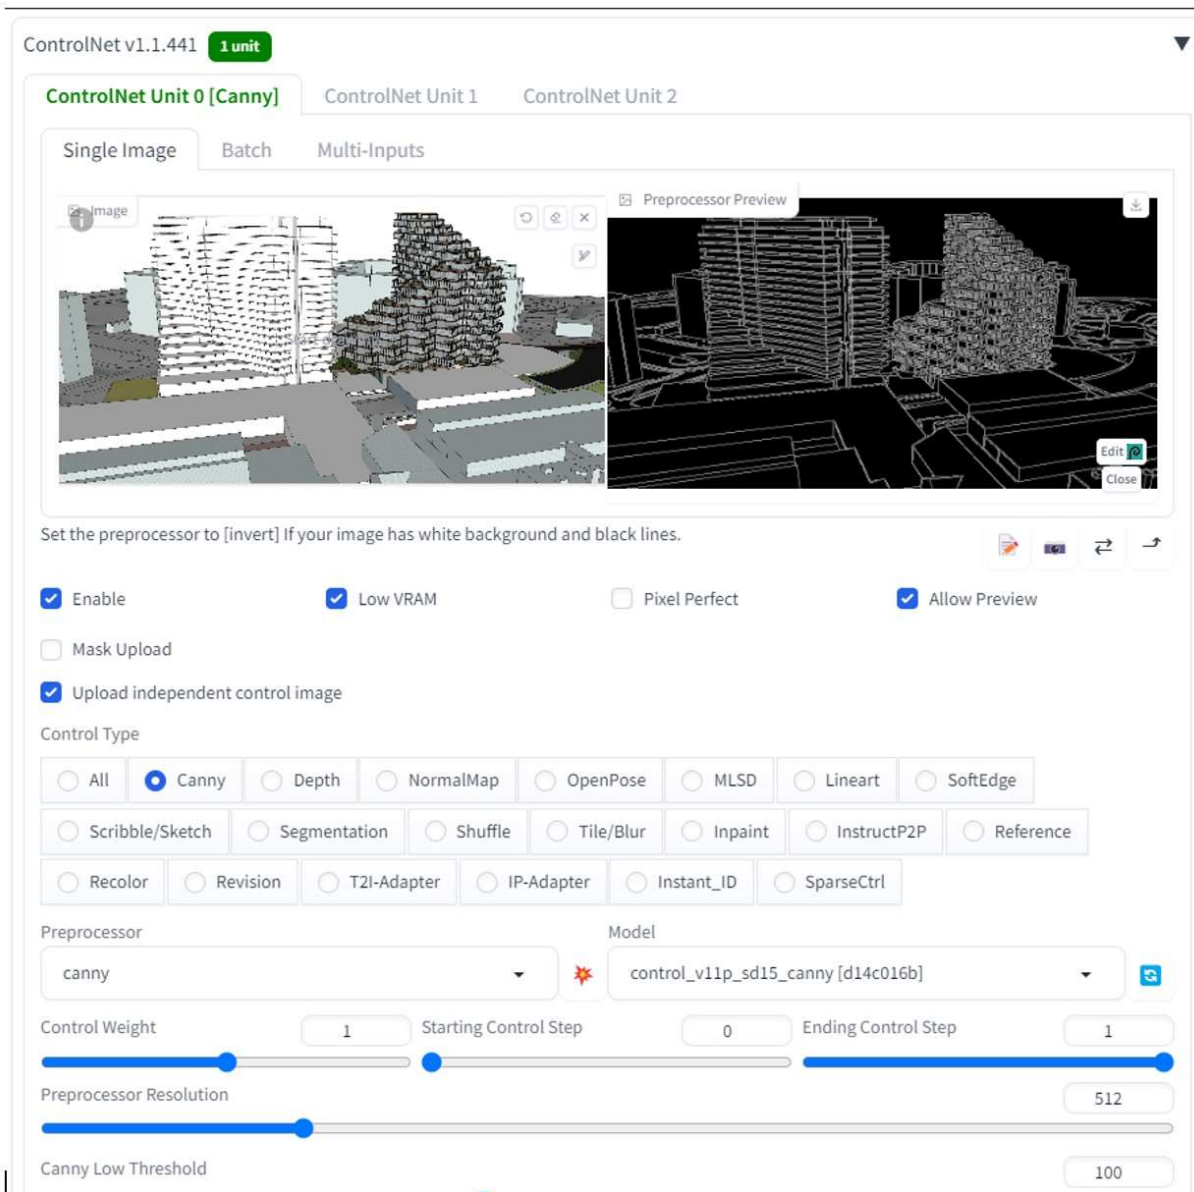


Fig. 9: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Developent Scenario, Prague*. MS architekti, Prague. 2024. author's archive. Based on the outcomes of the previous human-controlled Stable Diffusion performance, the volumes of the buildings = the geometry of the model is stabilized by ControlNet (1.1), the solution to be further developed by Stable Diffusion. The complete workflow in figures 8 to 11 and in the text, section (2) of the paper.




Fig. 10: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Developent Scenario, Prague*. MS architekti, Prague. 2024. author's archive. A scheme of how AI customizes a 3D model of the building. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.

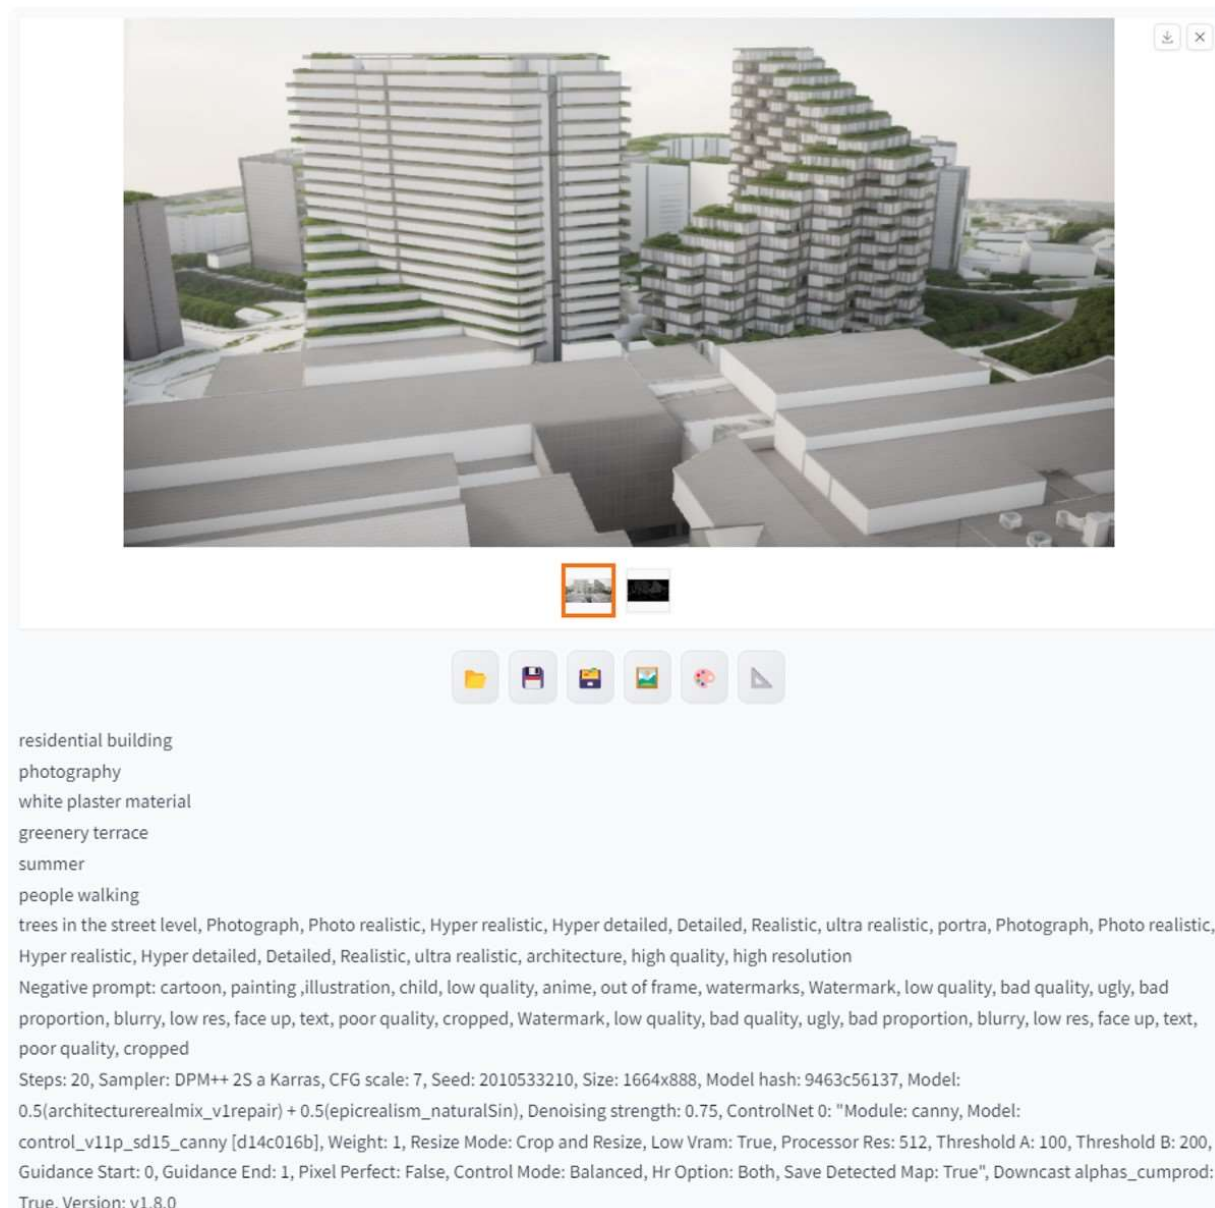


Fig. 11: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Development Scenario, Prague*. MS architekti, Prague. 2024. author's archive. The prompt to customize the given 3D model of the building. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.



Fig. 12: Labbe, J.: Diverse AI-tools deployment at a multi-level design development. *Luka Development Scenario, Prague*. MS architekti, Prague. 2024. author's archive. The volumes of the buildings = the geometry of the model stabilized by ControlNet, Stable Diffusion delivers options for the architectural forms of the buildings. The complete workflow in figures 3 to 11 and in the text, section (2) of the paper.

(3) AI revisited in architectural designing and in general

Today, the early-adopting professionals are witnessing AI performing well concerning parametric aspects of diverse materializations of architecture (such as construction, energy efficiency, daylighting, or noise in buildings and neighborhoods) and failing to be effective and productive in conceptual architectural designing. Rather than a temporary swing in the performance of particular efforts, it is a consequence of AI applications' developers failing to grasp and follow the starting points and the workflow of creating architecture, be it buildings or areas of the built environment. However, the flaws not only stem from the gap between understandings of architects on the one hand and AI developers and computer scientists on the other. The core is that the truly creative architectural tasks are objectively unsolvable by state-of-the-art AI; two principal reasons will be discussed later: unavailability of the computing performance needed and, more substantially, the nature of true creativity that state-of-the-art AI algorithms and networks cannot cope with.

Critical review's focal perspectives

To be followed to understand the position and prospects of AI in architecture and the development of the built environment, three focal perspectives appear throughout researching the state-of-the-art. Across the sections, the critical review must adopt all three viewpoints.

An architectural design process starts way before the parametrization can involve that lays at the heart of all the so-called computer-aided design (CAD) tools and all the existing AI strategies to develop layouts, whether in a building or a locality scale. This phase is where authentic creativity – *poiésis* as Heidegger revives the term coined by antic philosophers – applies as opposed to mimetic approaches that represent techné – mathematics as well as parametric design whether of spaces division within a building, construction, or pattern-based elaborating of architectural form; and, significantly, this is the design phase, which IT AI developers, computer scientists, also the general public, and surprisingly, many members of the architectural profession are not aware of. The consequences are social – undervaluation of the profession's members – branch's concerns – sidelining and downgrading of the architectural craft– and the misconception embedded

in efforts to perform mimetic, parametric ways the authentic creativity that is inherently *poiétic*. Leaving both the social and branch-affecting consequences for else elaboration, the last of the consequences is the first of this paper's critical approach's starting points.

Second, in the hectic conditions of rapid development and radical innovations in today's AI scene, architecture and built environment development appear not only sidelined but outdated, or perhaps misled in terms of technologies applied and considered, too. Leaving aside image-processing tools deployed in architectural studios to spare time and boost efficiency in topic researching, ideas thematization, and concept- and mood-options articulation, the efforts (though disputable in terms of actual contribution) to deliver generative tools to design architecture at least in a form of layouts are stuck in statistical approaches of input-output pairings implemented by supervised learning algorithms and generative adversarial networks. At the same time, impressive results are being achieved in autonomous driving, robotics, gaming, and other areas where expert human behavior can serve as a guide. There, instead of supervised learning, which no longer enjoys state-of-the-art status, reinforcement learning, self-learning, meta-learning, and other new approaches supporting imitation-based- and transfer learning, tree of thoughts, auto-associative, and other revolutionary algorithms apply that active and proactive agents perform. These strategies appear to be very adherent to architectural design processes; however, the possibility of adopting this class of applications has so far escaped the attention. Unlike an autonomous vehicle, a "design robot" would not learn from real-world scenes - it would learn from processes - sequels of steps taking place (in the parametric stages of architectural design or even before as Wearrecho allows) in a CAD or a VR/XR extended software environment (as annotated in section (2) of this paper).

Overlapping partly (when it comes to architectural design applications) with the second, third starting- and focal point of the critical review is the concurrence (already coined as *the trinity*) of the fundamental three- and more-dimensional spatiality and diachrony of both architecture and recently developed virtual reality- and spatial computing technologies together with the new theory of human thinking and intelligence that may be waiting for implementation in machine learning.

AI in AEC: Results achieved

With respect to the undoubted qualifications and ingenuity of the authors, the results of the Phillip Isola Research Group, Kyle Steinfeld, the "typical style" road- and circulation patterns and networks delivered by Neural Turtle Graphics, and others can be considered interesting outputs of research efforts in computer science or perhaps graphics, but only scarce contributions can be identified in terms of architectural workflow and solutions. Similarly, the *parametric semiology* outcomes of Daniel Bolojan or Tom Mayne's *operational strategies* render too speculative to provide some practical analytical starting point. *DeepHimmelb(l)au* alike early results of ZHA in AI show outputs of hundreds (rather thousands) hours of dedicated work of talented multi-expertise teams: outputs (in terms of conceptual approach and contribution – leaving aside the "video show" that, factually, has little to do with architecture) that the principal of the studio would sketch by hand within half an hour or so - and at the same time, opposed to the AI, would consider the spatial and operational concept represented by the sketch.

So far, the values of deliverables provided by the AI of (former) Spacemaker, Cove.tool, or Creo are appear undeniable, though not without caveats. Starting from a better organization of the working environment of a design engineer, Creo contributes to the productivity and efficiency of his work by model-based defining, simulations, additive and subtractive modeling, and manufacturing; Creo fosters the creative potential of a designer through generative design [265]. Similarly, Cove.tool, a cloud-based network of tools that provides interconnectivity within the teams working in the design and pre-construction cycle on issues of daylight, carbon footprint, climate, geometry, HVAC, or cost, delivers performance data of the building solutions in real-time, employing AI's power [266]. Nevertheless, being acquainted with the working paradigm of Creo or Cove.tool challenges whether it is true AI – in terms of network, algorithm, and the principle of training – what makes the software able to deliver: an ordinary rule-based algorithm would provide the same information.

Also famous as the two hundred and forty million acquisition of the tycoon AEC-software-producer Autodesk [97], Spacemaker was promising not only to *give the architects and developers the automation superpower to test design concepts in minutes and explore the best urban design options. It enables users to quickly generate, optimize, and iterate on design alternatives, all while considering design criteria and data*

like terrain, maps, wind, lighting, traffic, and zoning, with the help of AI. Utilizing the full potential of the site from the start, it allows designers to focus on the creative part of their professional work [267].

A practical deployment of Spacemaker raised doubts, though: the workflow was the issue. The user entered the address of the location and the boundaries of the territory; with the help of an open database like OSM, the terrain was generated, and it was also possible to generate existing buildings and structures. Roads could be added; buildings could be placed either manually by inserting individual floorplans as objects that could not be subject to later adjustments or the buildings could be generated automatically by the software based on input parameters entered: width, height, object shape, minimum/maximum number of floors, and/or by apartments' sizes mix. Then the tool generated options for site plans and the user could assess the options based on gross- and/or netto-floor area totals and modify the chosen option by some of the spatial transformations: shift, rotation, ... Spacemaker evaluated the finally proposed solution in terms of noise, wind, sunlight, daylight, and microclimate. Exports of the parameters to Excell and of the model to Autodesk Revit or to .ifc format were available. All in all, the evaluations of the designed locality's microclimate parameters and the imports of entry parameters were valuable and efficient functionalities. What happened in between - how the design came to existence, Pavel Shaban of MS architekti [268], a Prague/Czech Republic-based architectural studio claims, *could suit a shortsighted real estate trafficker, but it was far from a creative and responsible architect's workflow: in short, a comprehensively sustainable built environment develops along a grid-and-grain public space structure and not vice versa*. A network and profiles of vital, livable, and responsible public space do not fall from the skies nor emerge by chance. The public space - streets, squares, parks, places, public amenities areas, ... - has to be designed carefully, responsibly, considerably, and poetically first, to adopt particular buildings only after [269]. However, this is a process that Spacemaker not only did not support but also did not allow.

The career of the startup Spacemaker was brilliant - the life of the tool was short. Shortly after the acquisition by Autodesk, Spacemaker "dissolved" into Autodesk Forma, a *cloud-based early-stage planning and design software* [270], a more advanced successor to the former Autodesk Formit.

The presentation of Carlos Bañón's achievements in mastering Midjourney, Satble Diffusion, and ControlNet are impressive [271]. However, delving into the proposals delivered fails to reveal a true architecture - if not added by a human architect in the loop. In general: the more impressive the performance of a generative application presented by the developer, the more essential for the user is to understand that it is just a tool at the designer's hand, not a feature that could make the human author redundant.

Similarly, up-to-date AI applications to generate floorplans and apartment layouts emerge to be ambivalent when it comes to the effectivity and practical usability of their outputs: tens of options are delivered to ease an architect's task when taking over a load of mechanical generating various options and to be considered by him finally [272] – appearing prematurely published if not useless at the end. On the other hand, when designing the furniture layouts of a prepared floorplan, results that are more satisfactory emerge: as a rule, a corrective AI tool instructed through careful and detailed prompts by a human architect renders more capable than a conceptual one - better to say, it provides usable output, unlike a conceptual AI tool.




Fig. 13: Real-time analysis and visual information about design proposals. *Autodesk Forma: Cloud-based software for early-stage planning and design*. 2023. <https://www.autodesk.com/products/forma/overview?term=1-YEAR&tab=subscription&plc=SPCMKR>. (accessed Nov. 3, 2024)

Mimicking the workflow

Opposed to fine arts, literature, poetry, dance, or drama, whose production is only consumed, architecture is always also used. The "almost consumer goods" characteristic may turn attention to the processes of emerging the architecture – rather than its final form. Subsequently, deploying algorithms (even rule-based algorithms that do not bear the characteristics of AI), pre-defining parameters, and patterns in the design process come into consideration.

One of the first authors and researchers active in this field was Makoto Sei Watanabe already in the 1990s [273]; however, having focused on machine-aided design rather than on analyses and the use of patterns, he remained unsatisfied with what AI was able to deliver in terms of design compared to the intuition of the (human) architect. Others, like Immanuel Goh or Andrea Banzi searched for explicit rules-scripting-based design generators working with inferred rules drawn from the dataset of samples. *Not a patterns' assessment and appropriation, but recognizing the internal logic of the pattern, and then extrapolating a broader design based on that logic that could potentially continue forever ...* [274] in reality failing to contribute to the design practice eventually.

Recently, XKool, an AI startup in Shenzhen, China, developed a web-based platform for using AI across a range of tasks from architecture to urban design [275]. Though not-so-easy to be used practically or to be tested by non-Chinese residents [276], the approach of the studio and results achieved so far by the application awaken hope to overcome the lack of attention to the immense richness of patterns provided by the existing building stock and design representations. *XKool appears an efficient AI application for architectural design, streamlining the design process and making it more efficient in terms of both analyzing a vast range of possibilities and generating designs* (or pre-designs, more accurately said) based on samples – *to evaluate and return the most suitable outcomes, and, moreover, to develop them further according to the given constraints* [277]. XKool's way of working is revolutionary - and no worry that the outcomes do not look very novel as a rule: the core is it copes with the "consumer goods" characteristics of the design category. The mission *to challenge "the architect genius"* [278] that, so far, has been the motivation behind the efforts to develop an AI-based design tool (almost) as a rule, shows debunked by XKool approach and results achieved.

In general, a new approach emerges consisting of AI "designing" by - first - delivering pre-designs, it is solutions close to set parameters - as close as the available patterns allow, and then - second - "assisting" the human designer in adapting the pre-designs, tailoring the final, specific solution; the nature of the "assisting" is quantitative, parametric assessment, feedback concerning the goals and result assessment including finding the system of criteria, specification of the particular criteria, and evaluation criteria sets that AI can develop and complete continuously. The patterns-oriented approach, when confirmed and developed, and developed the patterns stock – libraries of parametric examples, representations of solutions existing so far, promises to bring a paradigm change within AI in architecture and AECO (architecture, engineering, construction, and operation) ; more to be elaborated in section (4) of the paper.

Ultimately, not only solution patterns - layouts, computer 3D models of architecture - can and should be mimicked. Examples of robots mimicking humans in diverse activities - Figure-1 preparing coffee as an example together with diverse self-learn and imitation-based learning strategies, as introduced in section (2) of this paper, suggest focusing on the design process itself. Opposed to the Figure-1 case, learning will not take place by observing a person at work (with a video camera), indeed: it will be the steps of design development proceeding in the CAD or an into VR/XR extended environment - the sequences of the drafter's or designer's actions in a parametric design tool, Autocad, Revit, Allplan, or Bentley CAD software typically, or within the Wearrecho platform. Such a strategy may sound self-evident: however, it has not yet been implemented for architectural design or development planning of the built environment, even experimentally, perhaps with a partial exception of the experiments mentioned at the beginning of this sub-chapter, XKool primarily. Finally, the proposal of such an approach in this article also makes comprehensible the extensive introduction of learning strategies and network types in section (2). If AI in architecture has been dominated by supervised learning and GANs so far, the future - and a breakthrough - may be reinforcement learning strategies and autoregressive, brain-inspired deep imitation neural networks, attention mechanisms, dual neural circuit policy architectures, and transformers (so far dedicated to robotics primarily).

When AI Works

Put in the previous sub-section and more in deep elaborated later, the using aspect of architecture is not a denial of the poetic essence of architecture. Inevitably, two realms of architectural design and a plan to build a building can be identified and distinguished; the first one comprises properties and performances that concern (even though not exclusively) the use, and the other delivers *poiésis*, poetry, mood, excitement, or experience. The interface between the two realms does not match the interface between architecture, as characterized by concerns to use by humans, to human activities and needs, and the communication of experiences and ideas by its form, and construction that materializes the architecture. It circumscribes the material opposed to the mental, emotional, artistic - whatever you like to call the architectural poetics. E.g., set by architecture, the spatial structure of a building gives the ergonomics and efficiency of movement within the building: it is the architectural design, not the construction solution that determines these material, quantitative parameters of the building.

Anything material can be parameterized, anything parametric can be quantified, and anything quantitative can be compared and evaluated objectively - or at least (very) close to objectively. And this is the case for a large part of an architectural design, a proposal of a building or an enclave of the built environment. Concerning the quantitative, objective and comparable assessment of the complex of diverse physical performances of buildings' and built environment's designs – such as operational and energy efficiency, acoustics, ergonomics, daylighting, and other physical benefits that architecture provides to man, community, and society – the state-of-the-art performs mature tools related to particular parameters. This is what software applications like Cove.tool, Creo, Giraffe, the applications used by MVRDV as outlined further in this chapter, and many other tools already introduced and proven in architectural and planning practice deliver, though not always distinguishing between physical architectural respects and respects of the construction.

The construction is another story: parametric, "mathematic" by nature, the mimetic, imitative creativity of designing constructions welcomes algorithms and parametric patterns. Such is the starting point for the excellence of generative AI software systems, their leading computational approaches being optimization and optioneering, analysis and simulation. The examples of most advanced applications are ETABS, SAP 2000, STAAD PRO, RAPT structural engineering software, SCIA Engineer [279], or Tribby3D [280,281]. Generative AI may be an approach deployed in multiple structural engineering software tools "from time immemorial", however, a black box must not have a final say when it comes to responsibility, such as in the

case of structural design. So far, “good old” computational practices and rule-based models are proving indispensable in this regard.

Mimetic, too, is urban design, its supportive disciplines being parametric by nature. Examples of application of AI in the field have been overviewed in (2) zooming in on tools such as (former) Spacemaker, Creo, or Cove.tool – in terms of both successful and contributing use of AI and of misconceptions.

At this point (in previous subsection), Spacemaker “got spanked” for many other AEC software tools that are parametric and shorthand imitative by nature and yet they are pushed to architects as creative tools, which is not so rare in today’s practice. Fortunately, better cases have been witnessed - also in the deployment of Generative AI. MVRDV, Dutch by origin, today a global architectural studio, shows up as a successful pathfinder in terms of AI use and development. In response to the need to push the limits of technological possibilities for the sake of innovative architecture, MVRDV NEXT - shorthand for New Experimental Technologies - was founded in the 2010s as an internal startup. Headed by one of the studio’s partners Sanne van den Burgh, *a group of in-house specialists develops and implements computational workflows and new technologies. Through a mixture of project-based work and standalone computational research, they rationalize designs and setup configurations, unlock potentials on an urban and particular buildings’ scale, optimize workflows, speed up processes, and make projects more efficient and adaptable in the face of change.* Represented by projects such as HouseMaker, VillageMaker, The Vertical Village, Barba, Space Fighter, or Porocity, and site specifically Rotterdam Rooftops or FAR MAX for Bordeaux, *their methods allow the studio to explore a future that is equitable, data-driven, and green* [282]. Awarded the best skyscraper of 2022 in the Emporis Award competition, the MVRDV Valley at the South Axis, the central business district of Amsterdam, is a showcase of successful AI technologies deployment alongside authentic architectural creativity. Machine analyses allowed for developing a rich, truly sculptural form and maximizing the efficiency of the land’s and space’s use while ensuring generous sunlighting and daylighting of all apartments and providing views and livable garden terraces to them. In planning the project, a Grasshopper script optimized the architectural form and detail to make the construction economical and efficient and to provide for sustainability thus. Alongside the comprehensive use of information technology to analyze the tasks and the opportunities and to support and streamline the creative design process, the rigorous avoidance of the terms AI and machine learning in the studio’s communication is notable [283].



Fig. 14: Design optioneering. MVRDV NEXT. 2023. <https://www.mvrdv.com/projects/515/la-serre>. (accessed Feb. 18, 2024)

Computer Consciousness

Consciousness ... is not merely the process of learning. It is not, strangely enough, required for many rather complex processes. Conscious focus is required to learn to put together puzzles or play the piano. However, after a skill is mastered, it recedes below the horizon into the fuzzy world of the unconscious. As Jaynes [284]

saw it, a great deal of what is happening to you right now does not seem to be part of your consciousness until your attention is drawn to it. What evokes consciousness is an impulse from - or induced by contact of "a self" with the external world. Perhaps surprisingly, a physical substance proves herewith to be a prerequisite of consciousness; and not only any physical substance, but also a substance that shows a degree of independence of *the world of its existence* [285], but, on the other hand, it can perceive the world around, and can and wants to react to it. Possessing an individual, specific history of being in the world and experience, any human performs individual and specific relation to the world of his existence - a specific, individual consciousness. And a computer? The world is not only World Wide Web and digital data; by far.

A brief list of a random choice of experiences undoubtedly forming a man's consciousness: visiting St. Peter's Basilica in Rome, sitting in the silent auditorium of Teatro Olímpico in Vicenza, chatting and drinking beer with friends in a pub, walking along Champs Elysées in Paris, jogging in Central Park in NY, a dinner in a posh restaurant, a bivac when climbing a mountain, fishing, to fall in love, to have sex, to give birth to a child, to raise it, and so on and so on ... When and how will be computers able to grasp such experiences? Erich Maria Remarque wrote his novel *All Quiet on the Western Front* as a veteran of the First World War: is anybody prone to pay his reader's attention through all two hundred pages to a computer's attempt to catch up?

Consciousness is also a mind-and-body creation, literally interwoven with the body and the body's support systems, a sort of thing a robot can hardly experience.

On the other hand, in terms of connectedness with the world around or how the information is coming in, a strong parallel exists between a computer and the Neocortex that is the place where the thinking and learning of mammals - man not excluded - takes place. Like a computer, the Neocortex has no direct informational connection to the physical world: all of the Neocortex's connection to the world around us goes (both ways) through the "old brain" the workings of which are much less known than the workings of the Neocortex. Could this be significant for the next development of AI?

Moreover, just as AI can perform complex calculations without understanding arithmetic, creatures (including humans) can display finely tuned behavior without understanding why they do so. The rationale for their behavior is "free-floating"—implicit in the creatures' design but not represented in their minds. Competence without comprehension is the default in nature [286]. The mental items that populate consciousness are more like fiction than accurate representations of reality. Computers may continue to increase in competence but hardly will develop genuine comprehension, since they lack the autonomy and social practices that have nurtured comprehension in humans. As a result, the computer comprehension can never align with the human.

"Life dwells in stories," Salman Rushdie [287] claims. The quote implies a strong double-sided interconnectedness of a real story and life: not only there is not a story that might be considered true (which property is in no conflict with phantasy!), which is not embedded in life, but there is also no authentic life that not composed of stories - of multiple, intertwining layers of stories. Yes, there are "life-stories" of computers - monotonous, parametric, and boring even if they had not been predictable.

A sensorical equipment coping with the one of man together with (to some extent) free motion in the physical world and physical interaction with it render a precondition for consciousness; a precondition out of today's imagination when it comes to computers. Paradoxically, man in the loop might be a solution at the closest at hand. The Neuralink project [288,289] may be the first step: aiming to let people control computers with thoughts, the researchers inserted a sensorical and communicational chip into the first patient's brain. But how many chips in which particular positins in a human brain will do to forward a consciousness to a computer? And still then, it would not be the computer's conscioucness but a consciousness of the human that would transmit to the computer. A weird idea ... A quote of Daniel Alarcón for BBC [290] fits: "Here in Caribbean [a cab driver says], we all have wonderfull stories. Gabo [Gabriel García Márquez] only types well." So far, AI is a superb typist - but its stories are only retrieved, moreover, often only stupidly taken over. The long story can bear shortening: So far, all the efforts in machine learning have been challenging cognitive processes that take place in the (human) Neocortex. Human consciousness, however, also arises outside the Neocortex, in the so-called "old brain" as various findings of recent biomedical research show [291]. Inseparable from receptors, the diverse parts and regions of the "old brain" "think and learn" in ways that are far from being revealed (at least as the processes in the Neocortex do) but surely resist any so far thinkable imitation by artificial neural networks. As a result, the today-most-promising way to endow a machine with consciousness would be a kind of reversion of Elon Musk's Neuralink project - attaching the ("old") human brain to an artificial neural network as its interface to the world and, in addition, as a "co-thinker". Research

projects like DreamDiffusion [292] generating images from brain EEG signals and MIT's Mark Harnett's investigation of how electrical activity in mammalian cortical cells helps to produce neural computations that give rise to behavior [293] delve into this realm.

The idea of a truly, in a sense/terms of poiesis creative AI shows rather a naivety.

Reasoning

Representing the state-of-the-art deep learning or AI today, GPT (the shortcut for Generative Pre-trained Transformer) is a kind of computer program that anticipates what shall continue after particular words or phrases; GPT models can create a new text that may look like created by man. It is not about the truth of the statement or the text that has been generated respectively (for architecture, authenticity represents "the truth"). It is only about the text generated to be in the pre-trained, pre-defined relation to the learning dataset; it is about following the pattern that has been discovered in the learning dataset and articulated explicitly in the process of training, whether following the criteria of regression (in supervised learning) or cumulative reward (in reinforcement learning), or criteria induced based on the analysis of the unlabeled dataset (in unsupervised learning). It is the pattern saying what is "correct" and what is "false". And, once we know how this "correctness" emerges, we can regard it as "usuality" – which, by the way, is by itself another explication that - and why AI is not truly intelligent and cannot truly create, when human intelligence is defined - among others - by the ability to think critically, to master successfully unprecedented and unusual situations, to articulate unprecedented ideas, and to create, adapt, and transform the living environment; most notably, the principle of disruption being inherent to human creativity.

In general, it has been LLMs what have been setting benchmark recently: not only in terms of popularity, frequency of use, and contribution credited by professionals and general public, but by level of development and results achieved, and by volumes of financial investments, too. The amounts provided as well as the eagerness to invest exceed any expectation. From technical point of view, a LLM tends to be at heart of most AI applications revealed today, not only the so called assistants or personas - digital characters that, by means of AI algorithms, imitate the behavior and qualities of a real human persona, used in marketing, virtual conversation, conversational interfaces, search, and more.

Current language models fall short in understanding aspects of the world not easily described in words, and struggle with complex, long-form tasks (as is characteristic, among others, for architecture and the built environment). Video sequences offer valuable temporal information absent in language and static images, making them attractive for joint modeling with language. Such models could develop an understanding of both human textual knowledge and the physical world. However, learning from millions of tokens of video and language sequences poses challenges due to memory constraints, computational complexity, and limited datasets. Large world model (LWM) is not a specific class of network or a learning strategy, but a deep learning model that uses transformer models and is trained using massive datasets by playing a guess-the-next-word game with itself over and over again. LWMs provide the ability to generate coherent and contextually appropriate responses over extended interactions, giving the impression of understanding or modeling the world. However, it's important to note that these models do not have an understanding of the world in the way humans do. They do not have consciousness or beliefs. Instead, they learn patterns from the data they are trained on and generate responses based on those patterns [294,295].

Above the business-as-usual and also state-of-the-art R&D, still behind the horizon, there is the vision of artificial general intelligence (AGI). The path to the meta is expected to require a different approach than today's generative AI models, still - in a way, as reviewed in section (4) - inspired by the McCulloch-Pitts theory of neural networks of 1943. Evidence is missing for the assumption that the machine-learning methods behind ChatGPT and other advances in AI during the past 20 years could act outside training data, which is considered a precondition for AGI. Verses, a California-based cognitive computing company that proposes a paradigm shift in the approach to AGI [296]. Specializing in *cognitive computing*, Verses aims to build *next-generation intelligent software systems inspired by the wisdom and genius of nature*. Within an alternative approach called *active reference*, Verses proposes an unparalleled model, which could resemble the new, in section (4) introduced, Hawkins'/Numenta's theory of mind and brain. The team is working on what it calls distributed intelligence, a system that can self-organize and retrain in real-time - identify its mistakes and fix them by re-training as biological organisms do. Digital intelligence based on a web of intelligent agents is believed to be cheaper, more environmentally sustainable, and more geopolitically defensible than one vast system trained on billions of data points. To date, Verses has developed Genius, an operating system for "continually learning

autonomous agents" operating at the edge of the company's connected devices. Genius combines biological inspiration, cognitive capabilities, adaptability, and open standards to create a unique AI platform that goes beyond traditional approaches. NASA's Jet Propulsion Laboratory and Volvo are among the beta users of Genius. "Minimizing complexity" is what is believed to be the finding of the way at the road fork. Instead of building ever-bigger AI models, Verses AI aims to deliver "99% smaller models" without sacrificing quality and performance and promises to release a public beta version of Genius in summer 2024 [297].

Verses' concept of distributed intelligence is closely related to and a predecessor of the field of multi-agent systems. It executes AI algorithms across multiple nodes or devices that can act independently and communicate asynchronously, exploits large-scale computation and spatial distribution of computing resources, and, not requiring all relevant data to be aggregated in a single location, it operates on sub-samples or hashed impressions of large datasets. Due to its scale and loose coupling, the system of distributed AI is robust and elastic. An alternative to it, also (perhaps) inspired by the human brain reasoning strategy similar to what Jeff Hawkins proposes [298], decentralized AI also executes algorithms across multiple devices or individuals, though, organizing the devices in a network, goes a step forward; it often uses blockchain technology to create transparent and secure platforms for collaboration. Both approaches involve distribution, distributed AI specifically focuses on solving problems using distributed approaches, whereas decentralized AI emphasizes the decentralized execution of AI algorithms [299, 300]. Currently, the second attracted the attention of one of the icons of AI, Emad Mostaque, founder and recently CEO of Stability AI who together with several high-profile researchers left the company in March 2024 to "pursue decentralized AI" [301].

Poiésis: Architectural design within and against AEC ecosystem

From architecture and urban design over construction and MEP (Mechanical, Electricity, Plumbing), environmental, climatic, meteorological, and microclimatic expertise to transportation expertise, economy, demography, and sociology, multiple professions engage in the development of the built environment. The background of some of the fields is natural sciences whilst, for the others, it is social sciences or even arts – poetics or *poiésis* [302] as will be explained soon. According to the nature of the contribution provided by the respective expertise, the design and evaluation approaches range from "hard" to "soft" ones, from quantitative, parametric, and material to qualitative and emotional ones. According to such an origin and nature, quantitative parameters define the approach as well as the output in some cases, while it is manifestations of consciousness in others; let's call it feelings, moods, or emotions to keep it simple. Obviously, as explained in the previous subchapter, manifestations of consciousness resist following parametric algorithmization as well as entering datasets.

There is a saying in Czech that goes something like "Just as one calls into the forest, so it echoes back". A rephrasing of the saying in terms of training dataset and the algorithm may sound *you can get the requested parametric answers if you address the right question to the correct forest*; however, no forest and no question exist to give the coveted emotion back - to give it in any situation, not to say an unclear situation, as it is the rule with man's feelings.

Creativity, to be authentic and true, cannot be but *poiétic* or poetic [302]. The poetic principle requests consciousness together with intention: only consciousness together with intention is able to deliver *poiésis* [285]. In terms of architecture and built environment, consciousness is reserved for a man, or, more precisely, to Dasein, as Heidegger coined a proved. An algorithm, however complex and sophisticated is the artificial network it works on, can deliver only based on the principle of equality (or similarity, which, however, is only a deficient mode of equality) or by random choice. Face-to-face to new solutions, advance knowledge is the prerequisite. *Prior knowledge* is another aptitude reserved for consciousness [285] - to a human, not to a machine, and not to an algorithm. No consciousness, no own will, and no true creativity, but algorithms and immense data searched through, assessed, and prioritized according to the defined criteria are the attributes of today's AI. And even the state-of-the-art theory does not show a vision of how machines could overcome the shortcoming.

Inevitably, when deployed on buildings, AI works in some respects and cannot but fail in others.

Approaching architecture as the most significant among the built environment creators, let us be clear: it is not a natural science scheme, algorithm, or calculus that is the architecture's starting point. Moreover, it is not a linear sequence of signs - opposed to speech or text. On the other hand, among many other attributes, architecture can be consumer goods, too; and the more a consumer goods a practical architecture shall be, the more a pattern, a calculus, and an algorithm contribute to the delivery; but even then, the environment,

the narrative of the development, and/or the people passing, entering and using the building or the structure „make the difference“.

In theory, architecture unanimously distinguishes from arts. But even so, even when architecture shall not be an art like painting, sculpture, drama, dance, or literature, let us not be shy: It is poetics or *poiésis* as Martin Heidegger coins in antic Greek that is the starting point and method of architectural creativity. *Poetically dwells man*, puts it Heidegger [302]: *full of merit, yet poetically dwells a man*. *Poiésis* precludes algorithm and vice versa, and similarly, a training dataset limits *poiésis*. By definition, in this regard, as claimed above in this section, a dataset must always be far from being comprehensive. Then, it cannot but limit the creativity for which, inevitably, the training dataset is „the whole world“ – there is nothing beyond.

Also, Encyclopedia Britannica distinguishes and confirms the emotional, social and societal, non-parametric nature of architecture, ... *the art and technique of designing and building, as distinguished from the skills associated with construction*. [303] *The characteristics that distinguish a work of architecture from other built structures are (1) the suitability of the work to use by human beings in general and the adaptability of it to particular human activities [and needs], ..., and (3) the communication of experience and ideas through its form*. Obviously, “use by human beings”, “human activities and needs” as well as “communication of experiences and ideas” cannot but resist algorithmization as well as digital parametrization. These needs, experiences, and ideas tackling and elaboration represent the heart of the initial and most important phase of any design process that precedes any parameterization and refuses it. Significantly, this is the phase and process that IT and AI developers and computer scientists do not know about (unfortunately, given the decline of both the architectural profession and theory during the last 70 years [8,304], a good proportion of architects are no better at it): hence, wherefrom the gap (pinpointed in section (1) of the paper) stems in mutual understanding of the needs, capabilities, and processes of one and the other expertise involved in the design and development of AI tools for architects. To make the long story short: a good proportion of the issue behind recent flawed attempts to foster the architectural design process is the AI developer teaching the machine to design architecture without knowing what architecture is.

Among all types and natures of creations by humans, architecture intertwines the most with human consciousness; not by accident. In the essay *Poetically Dwells Man* [302], elaborating further his seminal opus *Being and Time* [285] and the theme of *Dasein* - *being-there* or *existence* in English - after the Second World War in relation to the timely and pressing topic of housing, architecture by extension, Heidegger coins the concept of *das Geviert* - *the fourfold* in English - the union of *the earthly and the heavenly, the human and the divine* in man's existence and in *the world of his being* - thus, as we have seen, in architecture. This is not only another strong argument refuting the vision of architecture created by an algorithm. It is no coincidence that materiality manifests itself in both consciousness and architecture: materiality manifests itself in them in the same way and is a strong link between them. This recalls the *dual nature* of architecture - of ideas, emotions, and experiences on the one hand and material, physical on the other - that slowly-slowly begins to lead to uncovering the feasible way of deployment of AI in architecture and grasping its prospects.

Dalibor Vesely featured and reviewed critically another face of architecture's duality starting in the heading of the groundbreaking book *Architecture in the Age of Divided Representation The Question of Creativity in the Shadow of Production* [305]. Creativity never can be substituted by production; however, the material side of architecture - its physical properties both in terms of microclimate convenience, durability, security, ergonomics, operational efficiency, and sustainability - deserve and are keen to enjoy productivity - productivity, that is parametric and algorithms-inclined by nature.

So far in the field of AI in architecture, as in the whole AEC field, however, all the time only analogical, parametric-oriented approaches have been witnessed (the differences between diverse neural networks and AI algorithms, as outlined in (1) make no difference in this regard). Tackling data by a computational algorithm can provide poetics only by chance and randomly. It is not a question of learning or training; by definition, a poetic "output" cannot be trained. Even if bokeh salience offers a "*hallucination*", it's not *poiesis* nor a creative act; it's just a random interpretation of training data that we only additionally realize it was misleading. In a conclusion, the idea of a creative contribution of AI to conceptual architectural design is debunked; and together with it the theoretical collateral and all the AI's outputs in the field so far. On the other hand, debunking the vision of *AI or an AI's "superuser" replacing "the architect genius"* [306] as erroneous should not prevent algorithmizing and machine-generating what fits; and this is the physical aspect of architecture.

(4) Prospects

So far, slowly, AI has been showing helpful in architectural practice regarding imagery and parametric analyses - and failing when it comes to spatial creativity. However, what else is at the heart of architecture but physical spaces? Previous sections of the paper strived to lay a base to elaborate interconnectedness of spatiality (physical, three- or more-dimensional, when time and other-than-sight senses are included) that the Introduction features tentatively in terms of first, architectural space as such and its embeddedness in public space, second, the computer technology of virtual and extended reality that provides an unprecedented opportunity to tackle the space instantly (not only through its representations as has been the rule so far), and third, the recently disclosed essential spatiality of thinking (and of "intelligent computing" thus). Now it is time to zoom in on the third area.

Catching up with the human brain

In 1943, neurophysiologist and cybernetician of the University of Illinois at Chicago Warren McCulloch and self-taught logician and cognitive psychologist Walter Pitts published the foundations-laying article *A Logical Calculus of the Ideas Immanent in Nervous Activity* [307]. Building on Allan Turing's work *On Computable Numbers* [308], McCulloch's and Pitts' paper set a path to describe cognitive functions in abstract terms showing that simple elements connected in a network can have a huge computational capacity. The computational model that is the core of the McCulloch-Pitts theory of neural networks stems from the idea that neurons in the brain can be considered as simple binary switches that either fire or do not fire, depending on the input they receive. Within the already lively community of biophysicists doing mathematical work on neural networks at that time, McCulloch and Pitts' approach was novel in their use of logic and computation to understand neural, and thus mental, activity. Their contribution included [a] *a formalism whose refinement and generalization led to the notion of finite automata (an important formalism in computability theory)*, [b] *a technique that inspired the notion of logic design (a fundamental part of modern computer design)*, [c] *the first use of computation to address the mind-body problem*, and [d] *the first modern computational theory of mind and brain* [309, 310].

The McCulloch-Pitts theory of neural networks is a useful computational model that has been influential in the development of artificial neural networks. Emphasized their neurobiological aspects, artificial neural networks are modeled after the computational principles of the brain (as McCulloch and Pitts imagined them), with the specific aim of understanding and replicating human abilities. Such neural networks are being used for many machine learning tasks such as function approximation and pattern recognition - see for example [60,65,311]. However, these ideas no longer correspond to the state-of-the-art description of how the human brain works: the updated knowledge delivers an essentially different model. Since 1990s, Jeff Hawkins and the by him founded Numenta company develop neuroscience to dramatically increase performance and unlock new capabilities within the AI realm [298]. The field of their research is, first and foremost, Neocortex. The mantle of the large brain (telencephalon), also known as the cerebral cortex, phylogenetically the youngest part of the central nervous system, Neocortex can only be found in mammals to be fundamentally involved in functions such as perception, learning, memory, thinking, and others, unlike the "old brain" inherited by Synapsida (and mammals) from Amniota reptiles in the Late Carboniferous period.

First identified by Mountcastle in 1957, the fundamental "processor" of cognitive functions is Vertical Column (also called Cortical Column or Module) [312, 313], composed of six differentiated layers of neurons making together two to three mm of the human grey matter. There are roughly 150.000 Vertical Columns in Neocortex, interconnected mutually by dendrites and by axons that also, mediated by the cortical projection areas of "the old brain", ensure the transmission of information from sensory cells, to locomotor apparatus, and the like. [t] This apparatus "thinks" by creating - learning models of the objects encountered - a spatial, three- or more-dimensional object is the result - and [u] placing them into reference frameworks. [v] Learning runs in motion concerning the object to be learned, in other words, learning happens diachronically; [w] movement and perceiving run repeatedly and continuously update the model - as opposed to one-time training and "freezing" of today's AI models. [x] All senses that can be involved, within a breakthrough machine learning, novel types of sensors - radar as an example - can adjoin. [y] In our knowledge, there is not just one model for a particular object of reality: knowledge about particular object is distributed across many complementary models learned by many Vertical Columns. The structure of the Neocortex comprehending many models offers flexibility. [z] The key to making the many-model system work is "poll" or "conciliation". Each column works independently to some extent, but the extensive connections in Neocortex allow the columns to coordinate the perceiving.

The discrete nature of the Neocortex, or the complex of human cognitive processes, deserves emphasizing. At any moment in time, some neurons in the brain are active and others are not. The active neurons represent what we are currently thinking and perceiving. Our cognition and thinking is not in direct contact with the world around us, including our own body; all information comes to the Neocortex from the "old brain underneath", in the knowledge of whose functioning we still have large gaps. Crucially, these thoughts and perceptions relate to the brain's model of the world, not the physical world outside the skull. The world we live in is a simulation of the real world. This not only provides a state-of-the-art science's proof of Descartes' philosophical claim "cogito ergo sum" but also shows importance of the comprehensiveness of the respective simulation and the at least three-dimensional spatiality of the models and their reference frameworks that are the platform on which the simulation takes place.

A deep mutual interconnectedness shows between three so far discreetly taken realms: spatiality and diachrony of thinking and learning, the spatial and diachronic performance of the state-of-the-art immersive virtual reality technology, and spatiality and diachrony of architecture. (In all the cases, the spatiality conceives three or more dimensions.) The first two spatialities and diachronies apply broadly, and the third renders specific for the field reviewed by this paper - architecture and development of the built environment. However, the specific concern turns out broad enough as soon as the ubiquity of architecture in terms of human life is considered.

Spatiality and diachrony of thinking and learning surveyed and elaborated by Jeff Hawkins and Numenta raise questions that remain sidelined in today's race for ever-improving LLMs. Though a long way since, the current AI models still come from the McCulloch-Pitts essentials derived in the 1940s and recapitulated above under [a] to [d] above. However, these principles do not catch up with Hawkins'/Numenta model, as characterized by [t] to [z] above. Does this not mean a dead end, or rather a glass ceiling concerning development of AI models as practiced so far? Can the processes of tokenization, vectorization, and operation of two-dimensional, though in multiple layers arranged neural networks that "power" today's LLMs (and image- and video-processing models alike) cope with [t] three-plus-dimensional spatiality and [u] reference networks' embeddedness, [v] diachrony and [w] continuous updating, [x] sensorial and [y] models' plurality, and [z] "conciliation" principle that altogether make learning and thinking run in the human Neocortex?

The so-much-needed answer to this question does not need to regard the McCulloch-Pitts theory or the up-to-date AI networks and algorithms as flawed. The model is the question: AI may be on a similar threshold to the next-level development as physics was when Euclidian geometry and Newtonian physics showed unable to respond to questions raised by the theory of relativity and quantum physics. In day-to-day life, we still follow Newton's laws and dwell in Euclidian space; to learn the more tricky reality of the universe, we strive to get aligned with the Standard model and to put the theory of everything. Vice versa, without unraveling the problem of compatibility of the two models of thinking on the one hand and computing on the other, not only the AGI but, perhaps, also a fully potent architectural AI may remain just desire. As Abby Betrics of The Economist comments on the current state-of-the-art AI models: "[There is] no reason to believe...that this is the ultimate neural architecture." [314].

Computing performance, impacting the environment

In 2023, MIT's Schwarzman College of Computing launched a research delving into the opportunities and challenges of the computing age. Steven Gonzalez Monserrate's case study addresses the environmental, social, and health impacts of skyrocketing computing power ubiquitous in contemporary life [315]. Until 2020, perhaps 2022, the main consumer of electricity for computing and the main producer of waste heat, burdening the environment, was the mining of cryptocurrencies. Since then, machine learning has taken over, and the acceleration has accelerated. With AI, energy consumption is skyrocketing and sustainability is at risk. The means of mitigation of the unwanted effects ought to comprise but not be limited to efficient model architectures, hardware innovations, efficiency policies and regulations, green data centers using renewable energy sources, measuring the environmental footprints of the services as a starting point for the impacts' management and control, and - on the other hand - harnessing AI to address sustainability challenges - to optimize energy distribution, predict climate patterns, or enhance energy efficiency in buildings.

The Biden administration, too, wants to *accelerate* its conversations with big technology companies regarding how to generate more electricity — including nuclear power — to meet the massive energy demands of AI [316].

The hundreds of GPUs running AI foundation models, LLMs, and other applications consume tens of kilowatts of electric energy each, and about 80% of the energy converts into heat to burden the global climate. As a popular saying, the energy consumption of Google AI-search tools and LLMs copes with Ireland's total consumption. Other elements of the performing AI ecosystem add on top: cloud servers that host the model and handle user requests, load balancers distributing incoming requests across multiple servers to ensure efficient utilization and prevent overload, networking infrastructure that needs to be robust enough to ensure seamless communication between servers, data centers, and clients; this includes routers, switches, and firewalls. Then come the storage systems for the models to rely on storage solutions for model checkpoints, training data, and other resources; distributed file systems manage large-scale data storage. Database systems store metadata, user profiles, and other relevant information; relational databases or NoSQL databases may be used. Monitoring and logging tools track system performance, resource usage, and errors; examples include Prometheus, Grafana, and ELK stack (Elasticsearch, Logstash, Kibana). Authentication and authorization services secure access, services like OAuth, JWT, or API keys authenticate users and authorize their interactions with the model. Further, docker containers and Kubernetes orchestration manage deployment, scaling, and updates of the model's instances. Content delivery networks optimize content delivery by caching responses closer to users, reducing latency and server load. Backup and disaster recovery systems ensure data integrity and system resilience; and other components apply to contribute to the model's reliable operation [317].

One way or another, despite all the advances in hardware, namely chip performance, AI developers as well as hardware producers are increasingly facing the challenge of optimizing the computing performance. Nvidia, as has happened to be a rule, has hurried up to lend a hand: its Blackwell is four times faster than the previous model Hopper. The Blackwell transformer engine utilizes fine-grain scaling techniques called micro-tensor scaling to optimize performance and accuracy enabling 4-bit floating point precision (FP4). This doubles the performance and size of next-generation models that memory can support while maintaining high accuracy. Optimized to work with shorter decimal numbers, Blackwell is at 20,000 FLOPS (floating-point operations per second) at the 4-bit precision [318].

A contribution to higher computing power at lower energy consumption and lesser environmental impact is revealed with the development of ternary computing. In the still-emerging field, ternary neural networks (TNNs) make deep learning more resource-efficient using only ternary weights and activations instead of traditional binary values. The models match full-precision transformers in terms of perplexity with the student TNN learning to mimic the behavior of a teacher network without using any multiplications. In addition, TNNs inherently prune smaller weights during training, making them more energy-efficient; ternary LLMs offer the potential for orders of magnitude faster inference and significantly improved power efficiency.

Samsung-backed researchers have debuted ternary semiconductor design, introducing a 1-bit LLM variant - BitNet b1.58. The 1.58-bit LLM defines a new scaling law and a recipe for training new generations of LLMs that are both high-performance and cost-effective. Furthermore, it enables a new computation paradigm and opens the door for designing specific hardware optimized for 1-bit LLMs [319].

Potentially much more powerful for certain tasks than traditional computers, quantum computers can process a vast amount of possibilities simultaneously. Quantum computing is a type of computing that uses quantum-mechanical phenomena, such as superposition and entanglement, to perform operations on data. While traditional computers use bits as the smallest unit of data (which can be either 0 or 1), quantum computers use quantum bits, or qubits, which can be both 0 and 1 at the same time. The state-of-the-art in quantum computing is developing. As for prospects, quantum computing has the potential to revolutionize various fields by handling complex problems that are currently unsolvable by classical computers. This includes areas like cryptography, drug discovery, financial modeling, AI of course, and more.

There are still many technical challenges to overcome before they will be ready for widespread use. Researchers are making progress in building more stable and larger systems of qubits, which are necessary for practical quantum computers. Along the way, there are multiple issues. In the field of error rate, the developers proceeded from an error rate of 8×10^{-3} to 10^{-5} . This is still the beginning; Microsoft has set an error rate of 10^{-8} , allowing some useful calculations on a quantum computer.

Quantum AI is another topic - so far nebulous and full of strong statements. It was very promising when for specific processes quantum neural networks needed less learning for the same quality of results than traditional networks. However now, Chinese scientists have presented experiments with optical neural networks using common light, not individual photons. Such optical neural networks behaved quite similarly to quantum ones, including their potential advantages [320]. Therefore, quantum computers will not have an easy time showing any advantage.

There is significant progress in quantum communication and security: 100 kilometers of quantum-encrypted transfer [321]. And of course, there are turmoils in business, investments, and politics concerning quantum computing R&D [322,323,324,325].

Not only computing strategies and hardware novelties are trying to enhance machine learning. Verses, *a company that develops intelligent software systems based on the wisdom and genius of nature* [326], proposes revolutionary neural network, learning strategy, and output-delivery technique (that seem inspired by Jeff Hawkins' "A thousand brains" theory) to boost the computing power of its distributed intelligence model. Focus on biologically inspired distributed intelligence and transforming disparate data into coherent knowledge models are the starting points of Verses' technology. Introduced in section (3) of this paper, patterned after natural systems and neuroscience, Verses' flagship product - Genius operating system of thousands of intelligent agents minimizes complexity to deliver substantially smaller models that compromise neither computing power nor the performance quality.

Extended reality technologies

Together with the boom of the development of "practical" AI applications offering (though not always delivering) contributions to industries, the development of the once-promising technology of virtual and extended reality (VR/XR) has met its plateau. Leaving aside entertainment and social media in the metaverse, the promises for industry, design, engineering, planning, and architecture have been deferred and the development has slowed down as money - immense money dedicated originally to XR technologies has found a more promising investment target - AI.

The impact of the slowdown on AEC is one of the most graving within industries. The material architects work with is primarily space. So far, no tool to handle space existed, the only chance has been working only with representations of the space - various pictures of it. But a picture and a reality, it is two different things. XR allows the architects to work with space as it deserves - in space, motion, and real size - for the first time in history. And perhaps even more, the laymen - clients, authorities, the public - appreciate the chance to grasp spaces in XR perceiving and to understand the designs before implementation. A state-of-the-art XR technology [253] provides a break-through in three essential performance areas: (i) instant modeling of spaces in space and motion, real size; nothing else even remotely as good exists when it comes to conceptual designing; (ii) communication and collaboration in the scene; stakeholders can be in place or remote; (iii) seamless integration of XR and BIM environment; high visual quality of rendering of materials, surfaces, lighting, day and night, seasons of the year atmospheres is provided within the immerses into the virtual reality rendering of the BIM model. Even though however, the penetration of the technology within architectural practices and the AEC industry is still low.

On the brink of 2024, the XR field raises its head and anticipates a breakthrough. It is not only and most probably even not primarily about the famous Apple Vision Pro. As these lines are emerging, the entire rumor goes about the unprecedented quality of visual perception and (visual) control of action provided by the overpriced gadget that, however, so far lacks a sufficiently robust embedding in other than entertainment and social realm. The actual starting points are cloud solutions and (another) streaming capacity uprise, and spatial computing [327]. Merging the digital and physical realms, the concept of spatial computing is intended to allow computers to understand and interact with people more naturally in their everyday environment (which, importantly in this paper's framework, architecture and the built environment articulate). Sensors, computer vision, and techniques like augmented reality contribute to this transformative approach [328]. Spatial computing, indeed, can be a liaison between XR and AI that was not even hoped for and that not only could channel investment flows back to XR but could bring both the technology fields to (another) take off.

Spatial computing (and the industrial metaverse) is exceptionally prone to adopt beneficially the Hawkins'/Numenta model of three-plus-dimensional spatial thinking. Practical results and implementation are still far ahead but the motivation for respective R&D is strong: from revolutionizing the ways we interact with the world around us - both the real and the not yet implemented - through abilities to drive efficiencies and improve processes. Whatever the AI applications and networks to bring architectural designing to the new level are or will be, the patterns' retrieving and elaboration as well as generating, pre-generating for the human-in-the-loop engagement, processings' imitating and transferring, predictive simulations, designs' reviewing and solutions' evaluating (not only ex-post but real-time, following state trajectories and action trajectories), or optimization criteria specifying (all the/and other concepts elaborated later in (4)) will achieve (next level of) productivity and efficiency if implemented by spatially structured computing algorithms.

Space: the essence of architecture

Along with nature, it is architecture that creates the world of human existence - a space that has three or more dimensions with diachronic, sensory and experiential, social, cultural, and other knowledge "inputs" that add to the geometric dimensions of the rooms and places we inhabit. The *theory of public space* puts it clearly: *As soon as and only exposed in public space, a construction becomes architecture. ... Architecture is a physical spatial structure created by man, lasting in a given place and resonating with the socio-cultural values of the environment, which has the nature of a public space* [1].

Comprehensive spatiality is the key to understanding architecture. (In addition to being worthy of emphasis,) the fundamental spatiality of architecture is intertwined with the (above brought about) spatiality of our thinking probably more deeply than explored so far: the key may be the spatiality of diachrony, which underlies activities both in our cognition of and interaction with the world around us and in their processing in the Neocortex. In any case, today, the two are combined with the unprecedented spatiality of perception brought about by immersive XR. The processing of spatiality of thinking in AI networks and algorithms should naturally embed itself in the computational structures of XR. This will be the next level of both AI and XR; for architectural designing, this moment may be a singularity that - unlike the often futile catching up with human to date - will bring real breakthroughs that will provide the unexpected.

Nonetheless, the technological breakthrough alone will hardly be enough. An adoption of the unprecedented creative tools and possibilities by the industry will be needed. Today, architects should be truly motivated to bring their field and profession to *sumum templum architecturae* after decades of sidelining again [8,304]. It is time for architects to prove Rutger Bregman err when saying that *the real crisis [of our times] is that we can't come up with anything better* [7]." It is time for architects to grasp the singularity given by the encounter of the trinity of three- and more-dimensional spatiality and diachrony (the term coined in the Introduction to this paper) with the acute opportunities to overcome the lasting lapse of attention to the *poiétic* nature of architectural creation and to employ in architecture and the built environment development not the "proven" but the state-of-the-art deep-learning strategies (as identified in the introduction to section (3) of this paper).

Design patterns: lessons that have been skipped mistakenly

Opposed to the "sky is the limit" architectures whose form is often pre-defined neither by existing neighboring structures nor by short-term financial perspectives and approaches, there are architectures – buildings designed and constructed according to given spatial conditions, terms of future usage, and strict economic templates. In fact, this is the case for the vast majority of architectures - which, nevertheless, neither diminishes their importance nor makes the role of their architects less responsible and demanding. The vast majority of what is being designed, planned, and eventually built to saturate the needs of a growing population and living standards in terms of dwelling – it is residential buildings - work, and production – from office buildings to production objects - transport and logistics – among others logistic complexes and storage facilities - and many other buildings' typologies falls into the category of mass production and, kind of, consumer goods. Such a categorization does not challenge the contribution of the respective authors, designers, and planners in terms of "creativity used", craft, and efforts. Many such architectures launch their way to existence in architectural competitions - formal and non-formal - and not a few of them get their "five minutes of fame" in architectural websites, magazines, and exhibitions; nevertheless, they remain a "stardust", a sort of "no name" (except for specialist history scholars or local patriots); not to make anybody offended, let us label them "production [ones]" - *production architectures* and *production architects*. In a consequence, such architectures make the complex performance of the built environment: estimated the occurrence of the phenomenon very similarly to Chaillou [329], more than 90 % of the performance in terms of environmental impacts, sustainability, and macro economy, but also the majority of the performance in social, cultural, and economic terms. It is not the architectural icons but these *production architectures* that the entire population is exposed to on a day-to-day basis – at home, at work, at school, at leisure and social activities, at commuting, at going in for sports and recreational activities, and at walking pets as well as at tackling the household budgets. In terms of design and planning, the obvious richness of examples and models may balance the complexity of multiple limits and constraints that intervene in the design process; however, it does not make a creative architectural approach redundant or expendable.

Until recently, little was more overlooked in the field of AI than this opportunity and challenge - both in architectural and planning practice and in research and tools and processing standards development.

Further discussed in section (5) of this paper, a vast and promising field of pattern-based AI-driven development of architectural designing reveals. Advanced reinforcement learning techniques, imitation-based learning, self-learn, and transfer learning strategies shall take over R&D in AI in architecture to concentrate not only on the input-output images' pairings but primarily on the design development processes. AI-driven robotics shall be welcomed as the inspiration. The performance of advanced reinforcement-learning agents becomes the focal point: instead of complementarily to general models, particular studios and architects can be equipped with specific, personalized design assistants that may enhance general craft on the one hand and specific knowledge and skills on the other. Fostering, fine-tuning, and strengthening the unique capabilities and talents of particular architectural studios and individuals is an image of the future of the profession and the branch that, as a concept and in its consequences, may resist understanding and grasping today; a true, game-changing breakthrough.

Advice Whispering

Not a straightforward supervised-learned layout outputs deriving from image inputs that section (5) discusses as a dead-end but "sampling" of generative patterns = already existing solutions, selected by AI as the most suitable not only in terms of floor-plan or/and spatial solutions but in terms of structural solutions, too, and, even more importantly, learning the techniques of design development appears a key. Based on the given goal parameters and constraints, an adaptation (first human, soon AI-supported, and finally AI-driven) of selected patterns is proposed.

Moreover, both the selection and adaptation processes interweave with outcomes and adaptation solutions evaluation in terms of microclimate qualities – daylight and sunshine, or temperature stability – energy efficiency and consumption, acoustics, as well as area capacity and other qualities of the solution in process. An ability to *infer* the properties of the solution to which the design development is heading – whether led by a human or AI – stemming from the experience gained in learning on a set of solutions is natural to advanced AI models. The *predictive inference* can be available starting from the earliest phases of design – from the first sketch in terms of how a human drafts and develops a design. AI can go conveying the *inference* continuously in a way we can call *whispering*, providing the designer – human as well as AI – with comprehensive feedback on his or its design decisions and heading of the design. This way, the design will be optimized not in the mode try – error – correction – another error – another correction – and so forth till the designer is satisfied with the feedback parameters or too tired to continue trying, which is the state-of-the-art today, but continuously. The effect in terms of time and cost spent, and quality of the solution achieved is obvious and huge.

However, challenges remain: how to access the immense sum of the preceeding architectures records when a paradigm of protecting the authorship by hiding the representations of the architecture designed to the public. In this respect, the approach of the architectural community contrasts the approach of the IT developers community. Even the law contributes the „jealousy“ approach – „jealousy“ compared to the liberal approach of the IT developers community – of architects to the outputs of their work putting that an architectural drawing is an author's work, whilst a software code is not. However, the IT developers' community feels no disadvantage: the opposite is the reality. IT developers are used to providing each other with their achievements in widely shared libraries; Github [330], Gitlab [331], or Patternforge [332], and many others are the platforms. Who makes the profit are not only particular IT developers that can fulfill their tasks and achieve goals more quickly, with less effort, and for lower cost, whilst making available the results of their previous work costs them nothing; the whole field makes a profit developing quicker and better, a more efficient way based on the joint efforts of all members of the community. The perspective of the benefit of free approach to the existing solutions – in particular parametric representations of architectures both built and only designed - appears an incentive for reconsideration current approaches in terms of architectural design – and whole AEC, too. A particular architecture is „a product“ of public space, public space is outlined by particular architectures, and public space is, as a substance, an inclusive goods that all people are entitled – and welcome! – to use. So why not to share „all architectures“, too – at least virtually.

Debunked the vision of AI replacing „the architect genius“ [306], the supreme involvement and role of a human in a creative process remains untouched or, better to say, becomes upgraded. Maybe it is not always „the genius“ – sometimes it may be rather a craftsman - but it is only his intuition, creativity – however you want to address it – that makes the authentic *poiésis* of architecture real. Opposed to poetic, authentic architectural creativity, the nature of design development of parametric and material aspects of architecture

is mimetic – developing patterns in an imitative way by definition: it is the field for AI. The parametric aspect of architecture may become "an output" of AI; to be authentic, the *poiétic* aspect must always be a human creativity issue – and consequently, the whole architecture, too. Today, architecture as an inherently comprehensive discipline is developed in teamwork as a rule. Along with the development of AI's deployment in architecture, new roles will emerge: among others the „superuser“ tackling the AI, an architect with a strong IT background, or an IT expert with a strong architectural background. Nonetheless, the „superuser“ will replace the leading architect neither in his conceptual role nor in aesthetic respects. The „superuser“ will economize leading architect's efforts and forces for the sake of indispensable creativity. By the way, it is about a position that renders to be not so far (though undoubtedly distinct) from today's BIM (Building Information Management) Coordinator ...

Considering the three currently emerging main task-realms for AI in architecture – patterns' collection and choice, patterns' processing, and patterns' based continuous evaluation of to-date design outcomes, though not yet matured, even only sketches – a prospect for next-generation AI models, learning strategies, and (reinforcement learning) agents - *apprentices* reveals.

(5) Discussion

AI is a super-parrot: it is superb in repeating what it has learned, explains Tomas Mikolov in a chat with Dan Vavra [105]. In other words, for a *Generative Pre-trained Transformer*, the magnitude and comprehensiveness of the training dataset is the starting point, the algorithm is the method or, running on an artificial neural network, the tool respectively, and the computational performance is the limit.

Nonetheless, financing appears to be another limit. Touched on in section (2) concerning security issues, the Alexander Karp story also illustrates eloquently the economic starting point for the deployment of AI in architecture. If he had invested as an architect the same talent, the same significant skills, and the same intense effort he put into AI, he would never have become the billionaire he is; most likely, neither he nor any of his fellow AI tycoons would become a billionaire also in the field of built environment development. Obviously, in economic and investment attractiveness terms, designing architecture and even developing real estate cannot compete with data mining (not only) to ensure national security.

Obviously? The United States allocates approximately three percent of its gross domestic product (GDP) to defense outlays [333]. On the other hand, the finance, insurance, real estate, rental, and leasing industries (that one way or another could not perform without architects) contributed to the US GDP at 20.2 percent, and the construction industry alone at approximately four percent [334,335]. The only objective justification for consistently underestimating, even ignoring the potential of monetizable or otherwise tangible benefits of architectural designing and the built environment development can only stem from mistrust or even rejection of a possibility of improvement in their performance. These benefits can be economic, environmental, cultural, and social. Climate change has turned attention to the issues of environmental impacts and sustainability of construction: even in these respects, there are still no major efforts to improve the performance of architectural design and construction planning by applying machine learning. Other aspects of the built environment development and architecture remain oblivious to AI; better to say, AI remains oblivious to these fields. Any doubt that this deserves to be changed? Cui bonum? What profession shall intervene? Or is it a public interest, and the governments shall act?

Employing another, more technical perspective, the market potential of AI in architecture and built environment development should bear a comparison with the BIM software market. Providing advanced 3D modeling solutions for architecture, engineering, and construction professionals, the BIM software market was valued at \$5.2 billion in 2019 and \$5.71 billion in 2020; the projected BIM for 2027 is \$11.96 billion with a compound annual growth rate (CAGR) from 2020 to 2027 of 11.1%. Another study suggests that the BIM software market was valued at USD 9,665 billion in 2021 and is expected to reach as high as USD 23,950 billion by 2027, with a projected CAGR of 16.33% [336,337]. Enhancing operational performance, decision-making, cost estimation, and collaboration, BIM software tools are benefiting from government support. Appreciating its boost for operational efficiency and (accelerated during the COVID-19 pandemic) remote collaboration support, private builders' implementation further contributes to the BIM software market growth. All these incentives could apply to AI's support for the professions involved; nonetheless, the reality is still far behind. Following the patterns witnessed in the economy of R&D in AI today, a prospect of \$10 billion (almost) immediate market value should be able to attract some half-trillion in investments. However, architecture and the built environment development experience no warm welcome of this kind in AI R&D (and investment).

Not primarily but still, Carlota Perez's call [338] for understanding AI in a broader context addresses the contradiction; though often behind common understanding, architecture and the built environment fundamentally contribute to such a context.

Often regarded as the next technological revolution, AI may be better understood as a pivotal development within the ongoing information-communications-technology (ICT) revolution, which began with microprocessors in the 1970s. The ICT revolution accelerated in the 1990s when the US government privatized the Internet, which led to intensified innovation and globalization. AI may indeed represent a third leap; however, it is essential to recognize that ICT has already brought us to the brink of a golden age. Realizing this potential hinges on understanding the role of market-shaping public policy during previous technological revolutions. Without such policies, AI will fall short of its potential to drive inclusive social and environmental progress. The question of whether AI constitutes a new technological revolution remains significant. Early stages of revolutions involve creative destruction across the entire economy, not just specific sectors. During these periods, new technologies create and eliminate jobs, reshaping industries and regions. To maximize social gains, institutions and regulations must guide these technologies. AI relies on a massive energy supply and the Internet, which, in turn, depends on powerful microprocessors. These technologies mechanize mental work, and their future evolution may combine AI with biotech, new materials, new dwelling concepts, and new public spaces (among others) within the context of an ongoing ICT golden age. The historical context matters for economic decisions made by investors, firms, governments, and households. For sustainable progress, AI's development must occur within a regulated system, avoiding detachment from the real economy, real life that, as reminded, is fundamentally embedded in architecture and the built environment.

State-of-the-art floorplan generation

As mentioned in (2), the till recently achieved results of AI's deployment in architecture show that the so-far-ruling principle of lossy compression and subsequent "creative" decompression within the supervised (or unsupervised) learning has exhausted its possibilities without being able to deliver truly usable results. Not only the strategy reduces the comprehensive three-plus-dimension-spatial architectural task to image processing; the training stock is another issue. Hundreds of thousands of images, each labeled by humans, that are a precondition for the "statistics" to work properly, are unachievable in real life: unachievable for two distinct reasons. First, such images are available on the internet and in public resources in general only very sparsely. Training datasets – in section (4) predicted open source platforms pose first questions on materials assembly, materials quality, and size. Given state-of-the-art machine learning, the size should (significantly) exceed the N^{th} power of two, where N is the number of parameters to specify the AI task: thousands rather than hundreds of parameters when it comes to the comprehensive parametric and physical structure that materializes architecture: a building. Even if it were "only" lower hundreds, the number has a hundred and more zeros - a googol: the question of computing power - or rather, the optimization of the parameters structure - is immediately raised when googol exceeds the estimate of the number of elementary particles in the known universe. Considering the issue of computing power and the needed volumes of training datasets combined, the efforts to generate floorplans and apartment layouts using GANs render futile.

Stanislas Chaillou carried out and introduced bold research and development [329] as the starting point of the „next era“ of his approach to the use of artificial intelligence in architectural design - an approach effectuated in co-founding the start-up Rayon (introduced in (2)). Engineering a road map "from the parcel to the building footprint, from the footprint to a room split, from a room split to a furnished one," and to a room rendering, Chaillou intends to equip the "design pipeline" with a „catalog“ of (four) pre-defined „styles“ to satisfy subjective preferences. Further on, "design categories" of „footprint, program, orientation, thickness and texture, connectivity, and circulation“ are to be tracked and algorithmized. By reducing the quantity of design parameters, the approach successfully deals with the computing power problem. Nonetheless, confronting Chaillou's design categories with the *poiétic*, truly creative design aspects that, as section (3) of this paper recollects, a computer algorithm can never grasp and perform (unless the computer acquires consciousness and ability of a true reasoning together with it), reveals another issue instantly. Chaillou's or, better to say, Rayon's "design pipeline" does not respect the distinction; the design categories lying in or running into the sphere of *piétic* creativity, the "design pipeline" is doomed to provide fruitless outputs at any design stage. No matter if "the designer is then invited to "pick" a preferred option and modify it if needed, before actioning the next step. Browsing through the generated options however can be frustrating, and time-consuming," Chaillou admits. "To that end, the set of metrics defined in the "Qualify" chapter can demonstrate

their full potential here and complement our generation pipeline. By using them as filters, the user can narrow down the range of options and find in a matter of seconds the relevant option for its design. This duality of Generation-Filtering is where the value of our work gets all the more evidenced: we provide here a complete framework, leveraging AI while staying within reach of a standard user. Once filtered according to a given criterion (Footprint, Program, Orientation, Thickness & Texture, Connectivity or Circulation), we provide the user with a tree-like representation of her/his choice. At the center is a selected option, and around it, its nearest neighbors classified according to a user-selected criterion. The user can then narrow down the search and find its ideal design option, or select another option within the tree, to recompute the graph."

As the metrics run into the *poiétic* creativity sphere, "the set of ... "Quality" metrics" has no "potential to complement ... the generation pipeline" in fact – as opposed to Chaillou's expectations and promises. Not only by Chaillou's application (the Rayon) but by AI generally, in terms of authenticity, *poiésis*, a truly satisfactory output can come into existence only haphazardly. Browsing through piles of substandard proposals turns out to be a waste of time and energy, and diversion to man-made modification renders a deliverance; the sooner approached, the better. Inevitably, such "design pipeline's" outputs can only compare to AI-generated elevator music: not any architecture an architect would go in for. Only either a layman with scanty user experience and cultural background, and lacking outlook can be satisfied with what the algorithm performs, or a software professional who, however, appreciates not architecture but how the application works and how it produces.

The gap between architects - and authors across creative fields in general - on the one hand and AI applications developers and data scientists on the other renders again. OpenAI revealed the Sora AI application [153] in February 2024 that, in response to a text prompt, can create videos up to a minute in length. Sora is supposed to be able to generate complex scenes with multiple characters, specific types of movement, and precise details of individual objects and backgrounds. "Our model understands not only what the user asked for in their input, but also how those things exist in the real world," OpenAI representatives said in a post on the company's blog [339]. Impressive: however, a gadget like Sora, anyhow further developed and perfected, will never replace author teams headed by Quentin Tarantino or Jonathan Glazer - for reasons starting from the simple fact that a prompt that could capture all the poetics that the screenwriter, director, actors, cameraman, and other co-creators put into the work would have to be a performance more complex than the resulting film itself; as an example, the detail that during the production of the film, the creators communicate not only with language but also employing other senses and means of expression, is worth mentioning as an illustration. Nonetheless, this does not imply that Sora, as a representative of a class of algorithms, is doomed to be useless. The authors can, and probably will use it in place of picture script, as a research tool, as other - in addition to natural human senses and communication tools - means of communication.

In its processes of birth and elaboration, architecture is more similar to filmmaking than a non-architect would say. In terms of a competent architectural design workflow, by contrast with world imitation-and other advanced reinforcement learning strategies, the strategy and the technique recently introduced by Chaillou turn out to be a cul-de-sac, and all the other applications, which reduce spatial making to image supervised-learning processing, too.




Fig. 15: Supervised-learned GAN: a pattern-driven development from the plot over the building footprint to room split and furniture layout. Chaillou. S.: AI + Architecture Towards a New Approach. *Academia.edu*.

Harvard GSD, 2019. [https://www.academia.edu/39599650/AI Architecture Towards a New Approach](https://www.academia.edu/39599650/AI_Architecture_Towards_a_New_Approach). (accessed Jan 15, 2024)

R&D Anew

Being right and wrong at the same time, Chaillou reveals the core of the misconception when concluding his paper [329]: *... our belief is that a statistical approach to design conception shapes AI's potential for Architecture. Its less-deterministic and more holistic character is undoubtedly a chance for our field. Rather than using machines to optimize a set of variables, relying on them to extract significant qualities and mimicking them all along the design process is a paradigm shift. Then, nesting models one after the other will ultimately allow ... to encapsulate relevant pieces of expertise.* „Relevant pieces of expertise“ is the keyword, indeed. However, it is not the expertise hidden in input-output pairings but pieces of expertise representing steps and milestones of the (design) process as the AI-driven-robotics benchmarks of Figure-1, GROOT, and others introduced within the state-of-the-art (2) section suggest. The statistical approach and supervised-learned GANs, on the contrary, appear doomed to remain just an etape of the history of the deployment of AI in architecture that was treading a dead end. In this field, supervised-learned GANs exhausted their potential before they could deliver qualified results.

New classes of imitation-based learning and self-learning paradigms zooming in on the design-development processes instead of the results (to be) achieved must be welcome to "customize" the most suitable pattern to the requested final proposal. Two kinds of patterns are to be considered. A pattern of an existing architectural solution - an existing building, and a pattern of the process of how the solution came to existence. Fitting well with the so-far ruling supervised-learning image processing paradigms, the first strategy suffers two bottlenecks: computing power needed, further elaborated in the previous subsection, and training database availability as annotated in section (4). As the links put it, any attempt to develop with this strategy an architectural design of a building from scratch to a 3D model, even to layout drafts only, is doomed to fail. However, the task does not need to be so comprehensive. In the production architecture realm, most often a similar solution already exists and can be accessed. Launching the design process, this is a step undertaken by an experienced professional as a rule. After, the comprehensiveness of the task is reduced to bridging the gaps between the pattern and the current brief. Even a simple choice of a convenient sample from a limited database provides a significant reduction in the comprehensiveness of the task to be fulfilled. And it is easy to imagine that a machine can be fully competent to make the choice. Then, to bridge the remaining deviation from the current brief is easier, indeed.

The sequence of steps to develop the design up to the requested solution - a 3D model of the building or its layouts - is the second pattern. Here, the experience of AI-driven robots will apply. Not the way Figure-1 has trained the robot to make coffee - taking videos watching humans at the activity; unlike an autonomous vehicle, a "design robot" would not learn from real-world scenes but "watching" the design-development steps in a CAD software or in a VR/XR extended environment (in the parametric stages of architectural design or even before), their sequence, and sequels. This way, either a general algorithm to develop designs from a given situation according to the brief or a "private" agent to do this can be designed and trained. A general algorithm would learn how diverse designs have been developed, a "private" agent would be an *apprentice* of a particular (human) architect or a studio; the nomenclature of reinforcement learning becomes true to the real-life - hence world imitation learning concept. Reinforcement learning, self-learning, meta-learning, and other new approaches supporting imitation-based- and transfer learning, tree of thoughts, auto-associative, and other revolutionary algorithms that active and proactive agents perform - welcome to the sphere of architecture and the built environment development. The possibility of adopting this class of applications has so far escaped the attention of the branch; nonetheless, these new strategies appear to be very adherent to architectural design processes.

World imitation learning turns out to be a promising strategy with CAD BIM environment as the world. Uniquely parameterized both in terms of geometry and physical properties, i.e. attributes of entities of this world, and in terms of their making, i.e. commands and steps of software work, the world of BIM software environment is an ideal living or working environment for reinforcement-learning (RL) agents. In implementing the strategy, the motivation of the RL agents is crucial. Examples of state/action-related intrinsic motivation techniques such as curiosity-driven exploration technique encouraging an RL agent to explore uncharted areas of a game world out of curiosity, even if there's no immediate external reward associated with it, or novelty detection technique rewarding intrinsically the agent for encountering novel states or actions that it has not

encountered before, resemble quite intimately the thoughts of an architect making architecture. Intrinsic rewarding of the deep-learning agents emerges as a specific issue at the heart of the models; nonetheless, policy-related motivation that modifies the agent's policy to provide goals due to (extrinsic) rewards for learning specific behaviors shall not be overlooked. The nature of the rewards fundamentally sets the performance of the models: the performance not only in quantitative terms but - more importantly - in terms of quality that interconnects the abstract world of AI model with the real world. This renders an essential quality for AI models to succeed in architecture, which most probably rests behind the horizon of our understanding so far. To launch R&D in this direction, lessons from Claude models fine-tuning using constitutional AI and reinforcement learning from human feedback (annotated in section (2) of the paper) seem promising, "from human feedback" being the keyword. Also, a full and both-sides understanding of this problematics renders a precondition for bridging the gap between architectural and computer-science and -technic thinking that has been almost a rule in the realm of AI in architecture so far.

A strategy proposed already in 2016 by Nono Martinez (see section (2) of this paper), human-in-the-loop renders essential. Starting from following the human-in-the-loop in the phase of "customizing", the algorithm shall learn by self-training to master the design process better than the man in the end – hopefully; the *apprentice* surpasses the master and, also uniquely, eventually becomes a personal or home *apprentice*, but his craft can draw on a global background at the same time.

The basis of the algorithm structure could be a pair of mutually interfering loops: a generative loop and an advice-whispering loop, or there can be more advice-whispering loops particularized according to the diverse natures of the parameters, which will be "switched-on" only in a cascade. In the beginning, a suitable pattern will be selected from the database, which will be tested and optimized due to the specified outlines and with respect to a benchmark of independent parameters. The parameters of the spatial structure of the proposed building, the parameters of the physical properties of its constructions, and finally the parameters of the internal environment in the object will certainly come into consideration. A pragmatic optimization of the involved parameters structure appears a key task.

To take up such a challenge in the architectural design and building planning realm, lessons from the in section (2) reminded trailblazers such as AlexNet, Deep Blue, and especially AlphaGo Zero, Figure-1 and GROOT have to be learned. The evolutionary algorithms approaches and genetic programming have to undergo a deep survey to be subsequently considered as an option. Such a concept ought not to be refused or underestimated pointing out the gap between designing production (residential) buildings and playing chess or Go on the masters level: deployment of analogical algorithms deserves to be studied thoroughly first. After all, given the data format with which the AI algorithms can efficiently work in architecture and built environment development, the strategy can be quite straightforward: let BIM CAD software (or an into VR/XR extended platform, such as Wearrecho) be a robot! Concurrently, to facilitate the communication between the *master* and the *apprentice*, SMLs introduced in section (2) of this paper appear to have entered the scene just in time.

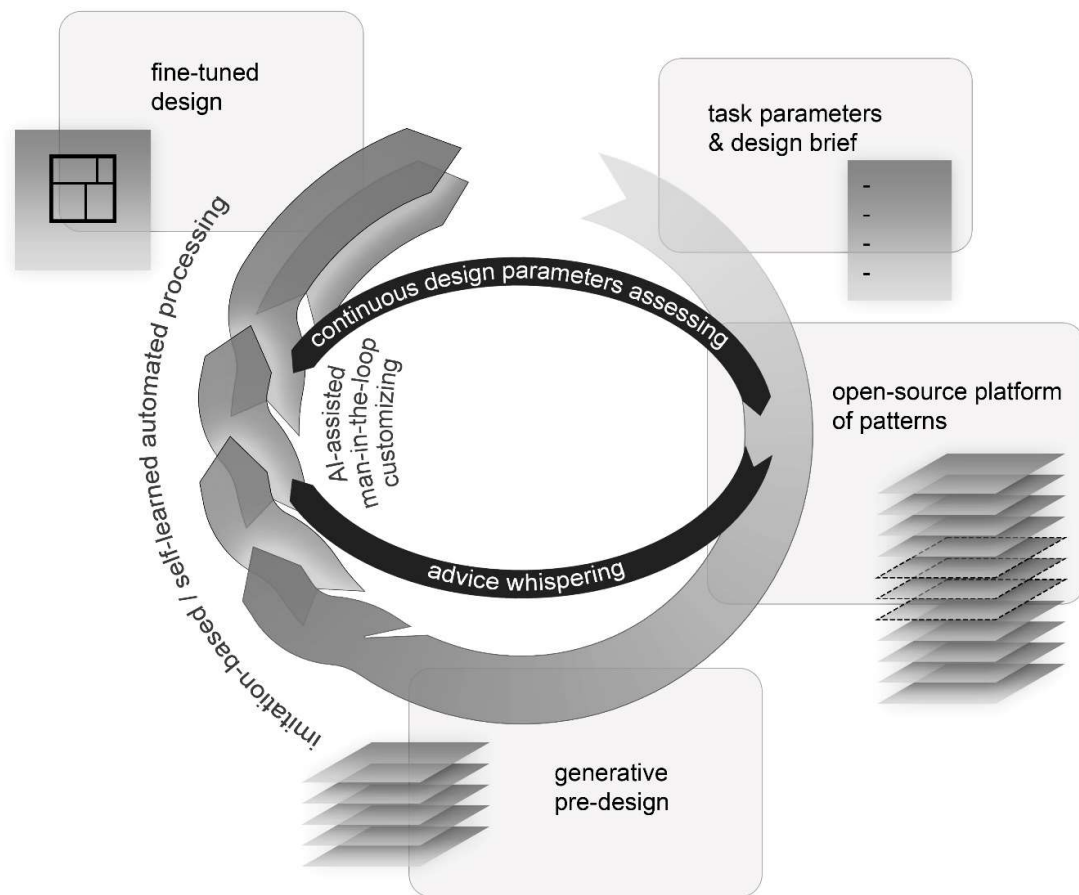


Fig. 16: From AI-aided to AI-driven architectural design workflow: *pattern-based generative pre-design – AI-assisted human-in-the-loop customizing – imitation-based/self-learned processing by reinforcement-learned apprentices*. graphics Bahdanovich, K. 2024. author's archive.

Let BIM CAD software be a robot!

Current achievements in robotics offer great and straightforwardly transferable inspiration for architectural designing. Towards the end of April 2024, headquartered in Vancouver, British Columbia, Sanctuary AI revealed a 7th-generation humanoid robot Phoenix that, after a week-long pilot test in a retail store, completed over 110 retail tasks, confirming Sanctuary AI's mission to create the world's first human-like intelligence in general-purpose robots [340,341]. Its pioneering AI control system Carbon allows the robot to learn and automate tasks in less than 24 hours, closely mimicking subsystems found in the human brain, including memory, sight, sound, and touch, and incorporating sophisticated human behavioral data-capture technologies; subsequently, this high quality, high-fidelity data accelerates the development of foundational AI models and customer deployments. Integrating state-of-the-art AI technologies for reasoning, task execution, and motion planning, Carbon translates natural language into real-world actions executing explainable and auditable reasoning, which ensures transparency and accountability, symbolic and logical reasoning coupled with deep learning and reinforcement learning, agency and goal-seeking behaviors, which enables autonomous decision-making, and human-in-the-loop supervision, which allows for fine-tuning and safety. Sanctuary AI envisions Phoenix as an augmentation to the human workforce; as such, the robot can be both directly piloted by humans for specific tasks, operated with a pilot-assistant combining human guidance and autonomy, or performing tasks autonomously when supervised by Carbon [342]. A cognitive architecture and software platform specifically designed for humanoid general-purpose robots, Carbon leverages symbolic reasoning, deep learning, and reinforcement learning to learn from data and improve over time adapting to different scenarios and making informed decisions; deep learning recognizes patterns, while reinforcement learning guides the robot's decision-making [343]. Carbon exhibits adaptability and robustness when

encountering uncertainty or novel situations during learning. It harnesses uncertainty; perhaps surprisingly, instruction with inherent uncertainty can promote learning and transfer more effectively than certain instruction. Thus, uncertainty triggers curiosity, which, in turn, facilitates learning [344]. To adapt to uncertain scenarios, Carbon employs adaptive neuro-fuzzy inference system (ANFIS), a hybrid model combining fuzzy logic and neural networks, and extreme learning machine (ELM), a single-layer feedforward neural network. These models allow Phoenix to learn from data and adjust its behavior based on uncertainty. Also, providing generalization ability, Carbon stands out in responses to novel situations performing goal-oriented behaviors, seeking to achieve objectives efficiently [345].

The experience of architectural practice combined with the ability to perceive and evaluate the processes of making an architectural design, and supported by knowledge of the fundamentals of machine and deep learning reads the previous lines as applicable to the work of CAD parametric, BIM software systems (or an into VR/XR extended platform, such as in section (2) said Wearrecho) in most, if not all categories and aspects. Moreover, a CAD software tool is well parameterized and digitized and, unlike a physical, humanoid robot, does not address gravity issues.

If understanding the nature of designing architecture and planning the built environment development, and comprehending the design- and planning workflow, an architect would appreciate an AI *apprentice* or assistant that would learn and automate tasks in less than 24 hours, closely mimicking the *master's* attitudes and techniques of making, able to grasp natural language commands and transfer them into real-world-design actions, executing transparency and accountability, and, in addition, autonomous decision-making. An augmentation to the human workforce, such a design robot could be both directly human-piloted for specific tasks, operated with a pilot combining human guidance and autonomy, or would perform tasks autonomously. Leveraged its symbolic reasoning, deep learning, and reinforcement learning, the *apprentice* or assistant would learn from data and improve over time adapting to different scenarios, recognizing patterns, and making informed decisions, exhibiting adaptability and robustness when encountering uncertainty or novel situations during learning. On top of it, it would harness uncertainty when instruction with inherent uncertainty could promote learning and transfer more effectively than certain instruction; uncertainty would trigger curiosity, which, in turn, would facilitate learning. In general, by providing generalization ability, such an *apprentice* or assistant would stand out in responses to novel situations performing goal-oriented behaviors and seeking to achieve objectives efficiently. Where is the problem then?

Also, the new (in 2023 introduced) machine learning algorithms for 3D modeling and rendering - both new diffusion models and NeRFs deserve investigation in this relation, leaving for the moment aside the capability of these deep neural networks of generating high-quality, photorealistic images of complex scenes from multiple viewpoints that, as an unprecedented AI and VR/XR fusion, would mark out the next level of exploiting the immersive VR-environment for (among others) instant designing and communicating architecture as it deserves - in space and motion, diachronically, from spaces, and in "life-size".

Introduced in section (4) of the paper, spatial computing shall also be surveyed and developed as another marriage between AI and VR/XR, that can push forward a state-of-the-art deployment of AI in architectural designing. Promising in regard to AI-aided designing architecture and planning of the development of the built environment may also show MeshDiffusion [38], appreciated for direct generating 3D meshes without any post-processing, and also LERF, the new marriage of NeRF with CLIP (contrastive language-image pre-training); with it, natural language queries in a 3D fashion can apply within NeRF, targeting different objects in the scene. And many other 3D objects considering outputs of AI development, though originally not focusing on architecture, need to be reviewed within the new R&D paradigm of AI's deployment in architecture and the built environment.

Motivated not by architecture but by (the example of) machine learning industrial robots working in the isolation of individual production plants, the ERC Advanced Granted FRONTIER project led by Josef Sivic from the Czech Institute of Informatics, Robotics and Cybernetics of the Czech Technical University in Prague can also help to show the way to approach architectural design effectively eventually. Issues of computer vision and perception of the "environment" can, if a suitable approach would be found, also benefit the field of architectural design and planning of the development of the built environment. Applications for this field could include *new neural architectures that credibly represent physical and geometric structure* as well as *new algorithms that enable learning of complex multi-step tasks from just a few examples ... like how humans can learn* [346]. Algorithmic sharing of experiences between projects could address the problem of training database size and building, which, most probably, will be difficult and only slow given the already entrenched

conservatism and autarky of the field. Similarly, the work of Josef Urban's team (of the same institution) in the field of large-scale computer-assisted reasoning can eventually contribute to effective deployment of AI in architecture [347].

(6) Conclusions

Artificial intelligence (machine learning, more correctly) has not proven creative (not only) when it comes to architecture; even the pioneers in the field are abandoning the projects whose ambitious visions have not been achieved and reducing their efforts to pragmatic parametric tasks. Introduced in section (2) of this paper, models like Stable Diffusion, Midjourney, and ControlNet represent state-of-the-art AI-deployment in architecture and built environment development shortening the way from a simple sketch to materiality-rich rendering, and easing and enhancing the process of visual conceptualization and image representation in phases from the (client's brief) ideation through sketches drafting, CAD BIM (or better an into VR/XR extended working platform) model development, over rendering to final design. Currently, whatever of the multiple applications deployed, AI is a "clever", more sophisticated pencil in an architect's hand.

Nothing more and nothing less remained of the decade and more old threats to make architects useless. "Machine hallucinations" by DeepHimmelblau, Tom Mayne's operational strategies ... *to generate output that could never be predicted*, or Daniel Bolojan's Parametric semiology study under the supervision of Patrick Schumacher (all cited in section (2)) neither took over the architects' work nor they delivered a new quality. Towards the end of the 2010s, creative AI efforts in architecture focused on designing facades and room interiors, and creating floor plans of buildings: in both areas without providing a solution applicable on a professional level. These models generally apply statistical approaches of image processing based on input-output pairings in supervised learning on GANs; as a representative of the latter class, Rayon has deserved notice in sections (2) and (5) of this paper. Giving way to diffusion models and other types of networks suitable for imitation-based, self-learn, or transfer learning techniques, in the architecture field, GANs and supervised-learning techniques represent no more state-of-the-art today.

Another class of AI-driven tools for architects are parametric tools such as Spacemaker, Cove.tool, Creo, or Autodesk Forma, delivering predictive simulations of technical infrastructure, and physical conditions in the interior and exterior of buildings, or preliminary proposals for urban-design solutions - the last, however of unusable quality as a rule.

In general, AI in architecture and built environment development experiences a plateau today. It was never a trailblazer, indeed; nonetheless, the current inactivity in the field sharply contrasts with the hectic making and emerging of LLMs, image- and video applications, robotics, and alike. Even the clashes of naturally conflicting interests in the field of training data and the use of R&D results avoid the field of AI in architecture and development of the built environment. This review, nonetheless, identifies novel impulses and opportunities that emerge both in the reviewed narrow field and its general/broad hinterland to revive the scene. No science fiction or beyond-the-horizon visions but feasible actions and currently or soon achievable incentives to R&D are concerned.

Starting from distinguishing between truly poetic, fundamentally individualistic creations and *production* (or, as other authors coin, of *the 90%*) *architecture*, the paper proposes to leave the first realm primarily to human creativity and concentrate the efforts on AI deployment to the latter. Consequently, the question of patterns' exploitation arises. As opposed to software developers, architects are lagging in the utilization of existing solutions: even if the solutions' databases were not organized and shared across the profession (as is a good practice and praxis with software codes) but only within studios, already the current state-of-the-art AI could enhance the efficiency, productivity, and quality of productive architecture making. The contribution can start with choice and evaluation - objective according to a set of criteria that can be as comprehensive as possible thanks to machine learning deployment. Specifying the criteria (that would later in the design workflow apply to optimize the proposal, too) is a question of human experience and remembering; as opposed to man, AI is significantly less prone to flaws in practice and remembering. A well-chosen and assessed pattern shortens the path to the final solution. Evaluating the design-in-process according to the already said set of criteria, AI can contribute along the path: not only at the end setting the trial-and-error approach that is a standard today but at any stage, after any step, enhancing the pace and efficiency and cutting costs of the process; the term advice whispering not only labels but spontaneously explains the technique.

Until this moment, state-of-the-art machine-learning approaches and techniques may do to fulfill the strategies proposed. Nonetheless, there is a vast realm of advanced AI models, networks, algorithms, and techniques that are waiting to be invited to the field of architecture and the built environment. Robotics and advanced imitation-based, self-learn, and transfer learning strategies offer a variety of approaches and solutions that fit architectural designing and built environment development planning much more than thought so far. In section (5), starting points and principles of the novel approach are discussed. A standard CAD/BIM software or, better, a platform liaising the parametric CAD realm and VR/XR realm of authentic architectural making such as Wearrecho (which section (2) introduces), shall be the robot, which a by Figure-1 and Phoenix (for both, see section (5)) inspired deep-, reinforcement-learning model shall steer and control. Every creative architect served by an AI assistant, apprentice-agent learning through encountering the task challenges, whether from a general or the studio's background, is the vision.

Also, introduced in section (2), SMLs represent a novel class of strategies likely to have a say in the upcoming R&D on effective deployment of machine learning in architecture and built environment development due to their fundamental ingeniousness and adaptability.

As pointed out repeatedly in the paper, the gap in understanding the making of architecture on the one side and the thinking of computer scientists and code developers on the other has been an obstacle to faster and more productive R&D in AI in architecture and built environment development so far. To bridge it, introduced in section (3), three focal points have to be pursued in parallel: in short, the poetic nature of the conceptual making of architecture that precedes any parametrization, (second) concentration on making - on processes rather than on input-output pairings, and (third) the concurrence of the fundamental three- and more-dimensional spatiality and diachrony of both architecture and recently developed virtual reality- and spatial computing technologies together with the new theory of human thinking and intelligence that may be waiting for implementation in machine learning. Not only for architecture and the development of the built environment, in the light of the results of the ongoing research on the human Neocortex (which section (4) introduces) and in section (3) annotated Verses' R&D in distributed AI together with novel, so far experimental computing techniques said in section (4), the question of the (possible) glass ceiling of existing paradigms of neural networks and machine learning algorithms also deserves exploration. The research and consequent development may, as annotated in the Introduction of this paper, turn into a singularity that, based on a three-plus spatiality and diachrony, intertwines human thinking and AI, the technology of VR and spatial computing, and architecture.

Beyond the understanding-gap issue, there are economic reasons to sideline AI applications for architecture. Delivering data essential for ensuring national security or enabling hundreds of millions to chat with a machine shows (much) more profitable than leveraging the quality of the living environment. Such reasoning has proven flawed by hard economic data, but still, it works; is it up to architects, governments, or society who profits from the qualities of architecture and the built environment to make the investors change the approach?

Essentially, the outcome of a singularity must be a breakthrough. In terms of this paper's focus, the breakthrough will not have a form of catching up with and overcoming man in *poiétic* creativity. Today, we can only assume the form and nature of the breakthrough contribution; but we can take for granted that it will be *technological*. Not technological in the sense of today's smart home or a smart city, not in the sense of today's AI applications delivering images or videos when prompted. In the realm of architecture and the built environment, groundbreaking approaches and construction solutions, and materials shall be expected that will accommodate the visions of unleashed architectural creativity and achieve unprecedented levels of comprehensive sustainability and resilience.

References

- [1] Sourek, M.: From Functional Areas towards Metropolitans Structure: Public Space in Sustainable Development Context. *Advanced Engineering Forum*. Trans Tech Publications 2014. vol. 12, pp 176-180, ISBN 978-3-03835-323-2.
- [2] Sourek, M.: *Architecture in Modern Times: Searching for Public Space*. České vysoké učení technické v Praze 2019. ISBN 978-80-01-06576-1.
- [3] Vitruvius, M. P.: *The Ten Books on Architecture*. original title De architectura [libri decem], translation Morris Hicky Morgan. Kessinger Publishing 2005. ISBN 9781417969579.

- [4] Heidegger, M.: *Voll Verdienst, doch dichterisch wohnt / Der Mensch auf dieser Erde. Heidegger und Hölderlin, herausgegeben von Peter Trawny*. Vittorio Klostermann 2000. ISBN 978-3-465-03084-3.
- [5] Ferneyhough, F.: *Liverpool and Manchester railway 1830-1980*. Book Club Assoc. 1980. OCLC 656128257.
- [6] Kahn, R., Dennis, M. A.: Internet. *Encyclopedia Britannica*. 5 Feb. 2024. <https://www.britannica.com/technology/Internet>. (accessed Mar. 16, 2024)
- [7] Perrigo, B.: 'Acts of Kindness Are Really Contagious.' Historian Rutger Bregman Argues for a New Way of Thinking About Humanity. *TIME*. May 22, 2020. <https://time.com/5838900/rutger-bregman-humankind-interview/>. (accessed Mar. 16, 2022)
- [8] Sourek, M.: Industrialize the Architectural Profession at Last to Revitalize It and Make It up to the Urgent Needs of Our Era. *International Journal of Architecture, Arts and Applications*. vol. 8, No. 3, 2022, pp. 121-129. doi: 10.11648/j.ijaaa.20220803.14.
- [9] Sourek, M.: Virtual Twins of Architecture: The singularity of the Profession and the Field. *International Journal of Architecture and Planning*, vol. 2, issue 2, pp. 1-26, 2022. doi: 10.51483/IJARP.2.2.2022.1-26. <https://doi.org/10.51483/IJARP.2.2.2022.1-26>.
- [10] Meet Deep Himmelb(l)au. Our algorithm learns CHBL's semantic characteristics to generate new interpretations and new worlds. *Facebook Watch*. <https://www.facebook.com/watch/?v=1473300746151435> (accessed Mar. 3, 2023)
- [11] Coop Himmelb(l)au, Bolojan, D.: Deep HIMMELB(L)AU. *YouTube*. <https://www.youtube.com/watch?v=8G3jpGQMvWE> (accessed Oct. 5, 2022)
- [12] Leach, N.: *Architecture in the Age of Artificial Intelligence An introduction to AI for architects*. 1st ed. Bloomsbury Visual Arts, London, 2022, p. 104. ISBN 978-1-3501-6551-9.
- [13] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed., W.W.Norton, 2020, p. 11. ISBN 9780393635829.
- [14] What is supervised learning? *IBM*. <https://www.ibm.com/topics/supervised-learning>. (accessed Oct. 30, 2022)
- [15] What is unsupervised learning? *IBM*. <https://www.ibm.com/topics/unsupervised-learning>. (accessed Oct. 30, 2022)
- [16] Jagtap, R.: Understanding the Markov Decision Process (MDP). *builtin*. <https://builtin.com/machine-learning/markov-decision-process>. (accessed Nov. 3, 2022)
- [17] Reinforcement learning. *WIKIPEDIA*. https://en.wikipedia.org/wiki/Reinforcement_learning. (accessed Oct. 30, 2022)
- [18] Awesome RLHF (RL with Human Feedback). *Opendilab*. <https://github.com/opendilab/awesome-RLHF>. (accessed Mar. 16, 2024)
- [19] Ayush, T.: Understanding Reinforcement Learning from Human Feedback (RLHF): Part 1. *Weights & Biases*. <https://wandb.ai/ayush-thakur/RLHF/reports/Understanding-Reinforcement-Learning-from-Human-Feedback-RLHF-Part-1--VmlldzoyODk5MTIx>. (accessed Nov. 3, 2022)
- [20] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed., W.W.Norton, 2020, p. 199. ISBN 9780393635829.
- [21] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing Atari with Deep Reinforcement Learning. *Papers with Code*. paperswithcode.com/paper/playing-atari-with-deep-reinforcement (accessed Nov. 3, 2022)
- [22] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed., W.W.Norton, 2020, p. 206. ISBN 9780393635829.
- [23] *ibid.* p. 207.
- [24] *ibid.* p. 209.
- [25] Imitation Learning. *Papers with Code*. <https://paperswithcode.com/task/imitation-learning> (accessed Apr. 12, 2023)

- [26] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed., W.W.Norton, 2020, p. 222. ISBN 9780393635829.
- [27] Dowat, J.: Internal Covariate Shift: How Batch Normalization can speed up Neural Network Training. *Medium*. Mar 29, 2021. <https://medium.com/analytics-vidhya/internal-covariate-shift-an-overview-of-how-to-speed-up-neural-network-training-3e2a3dcdd5cc>. (accessed Mar. 15, 2023)
- [28] Mikami, A.: Imitation Learning. In: Binder, M.D., Hirokawa, N., Windhorst, U. (eds): Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-29678-2_2366. (accessed Mar. 22, 2023)
- [29] DeMoss, B., Duckworth, P., Hawes, N., Posner, I.: Offline Imitation Learning with World Models. *arXiv*. Cornell University. [arXiv:2302.03086v1](https://arxiv.org/abs/2302.03086v1); <https://doi.org/10.48550/arXiv.2302.03086>. (accessed Mar. 12, 2023)
- [30] Khansari, M., Ho, D., Du, Y., Fuentes, A., Bennice, M., Sievers, N., Kirmani, S., Bai, Y., Jang, E.: Practical Imitation Learning in the Real World via Task Consistency Loss. *arXiv*. Cornell University. [arXiv:2202.01862v2](https://arxiv.org/abs/2202.01862v2); <https://doi.org/10.48550/arXiv.2202.01862>. (accessed Mar. 15, 2024)
- [31] Bansal, T., Pachocki, J., Sidor, S., Sutskever, I., Mordatch, I.: Emergent Complexity via Multi-Agent Competition. *arXiv*. Cornell Univ., 2018. [arXiv:1710.03748v3](https://arxiv.org/abs/1710.03748v3); <https://doi.org/10.48550/arXiv.1710.03748>. (accessed Jan. 15, 2024)
- [32] IQ-Learn: State-of-the-Art Imitation Learning for AI. *Techfinder*. Stanford University. <https://techfinder.stanford.edu/technology/iq-learn-state-art-imitation-learning-ai>. (accessed May 15, 2023)
- [33] Deka, A., Liu, Ch., Sycara, K.: ARC - Actor Residual Critic for Adversarial Imitation Learning. *arXiv*. Cornell Univ., 2022. [arXiv:2206.02095v4](https://arxiv.org/abs/2206.02095v4); <https://doi.org/10.48550/arXiv.2206.02095>. (accessed Jan. 15, 2024)
- [34] Sun, M., Ma, X.: Adversarial Imitation Learning from Incomplete Demonstrations. *arXiv*. Cornell Univ., 2019. <https://arxiv.org/pdf/1905.12310.pdf>. (accessed Jan. 15, 2024)
- [35] Torabi, F., Warnell, G., Stone, P.: Behavioral Cloning from Observation. *arXiv*. Cornell Univ., 2018. [arXiv:1805.01954v2](https://arxiv.org/abs/1805.01954v2); <https://doi.org/10.48550/arXiv.1805.01954>. (accessed Jan. 16, 2024)
- [36] Monteiro, J., Gavenski, N., Granada, R., Meneguzzi, F., Barros, R.: Augmented Behavioral Cloning from Observation. *arXiv*. Cornell Univ., 2020. [arXiv:2004.13529](https://arxiv.org/abs/2004.13529); <https://doi.org/10.48550/arXiv.2004.13529>. (acc. Jan. 16, 2024)
- [37] Gonfalonieri, A.: Inverse Reinforcement Learning Introduction and Main Issues. *TDS*. <https://towardsdatascience.com/inverse-reinforcement-learning-6453b7cdc90d> (accessed Apr. 7, 2023)
- [38] Yang, Q., Zhang, Y., Dai, W., Jialin Pan, S.: Transfer Learning. *Cambridge University Press*, 2020. DOI: <https://doi.org/10.1017/9781139061773>. (accessed Feb. 12, 2024)
- [39] Tormos, A., Garcia-Gasulla, D., Gimenez-Abalos, V., Alvarez-Napagao, S.: When & How to Transfer with Transfer Learning. *arXiv*. Cornell Univ., 2022. [arXiv:2211.04347](https://arxiv.org/abs/2211.04347); <https://doi.org/10.48550/arXiv.2211.04347>. (accessed Feb. 12, 2024)
- [40] Hulstaert, L.: Transfer Learning: Leverage Insights from Big Data. *datacamp*, 2018. <https://www.datacamp.com/tutorial/transfer-learning>. (accessed Feb. 12, 2024)
- [41] Ross, S., Gordon, G., Bagnell, D.: A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*. vol. 15, 2011. pp. 627-635. <https://proceedings.mlr.press/v15/ross11a.html>. (accessed Jan. 22, 2023)
- [42] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed., W.W.Norton, 2020, p. 240-243. ISBN 9780393635829.
- [43] Dong, Q., Dai, D., Zheng, C., Wu, Z., Chang, B., Sun, X., Xu, J., Li, L., Sui, Z.: A Survey on In-context Learning. *arXiv*. Cornell Univ., 2023. [arXiv:2301.00234v3](https://arxiv.org/abs/2301.00234v3); <https://doi.org/10.48550/arXiv.2301.00234>. (accessed Feb. 12, 2024)
- [44] Sigaud, O., Baldassarre, G., Colas, C., Doncieux, S., Duro, R., Perrin-Gilbert, N., Santucci, V. G.: A Definition of Open-Ended Learning Problems for Goal-Conditioned Agents. *arXiv*. Cornell Univ., 2024. [arXiv:2311.00344v3](https://arxiv.org/abs/2311.00344v3); <https://doi.org/10.48550/arXiv.2311.00344>. (accessed Feb. 18, 2024)
- [45] Stooke, A., Mahajan, A., Barros, C., Deck, Ch., Bauer, J., Sygnowski, J., Trebacz, M., Jaderberg, M., Mathieu, M., McAleese, N., Bradley-Schmieg, N., Wong, N., Porcel, N., Raileanu, R., Hughes-Fitt, S., Dalibard,

- V., Czarnecki, W. M.: Open-Ended Learning Leads to Generally Capable Agents. *arXiv*. Cornell Univ., 2021. [arXiv:2107.12808v2](https://arxiv.org/abs/2107.12808v2); <https://doi.org/10.48550/arXiv.2107.12808>. (accessed Feb. 18, 2024)
- [46] Victor, A. M.: Tree-of-Thought (ToT): A Comprehensive Guide. Medium, 2023. <https://medium.com/aimonks/tree-of-thought-tot-a-comprehensive-guide-c88f60dc1580>. (accessed Feb. 16, 2024)
- [47] Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., Narasimhan, K.: Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *arXiv*. Cornell Univ., 2023. [arXiv:2305.10601v2](https://arxiv.org/abs/2305.10601v2); <https://doi.org/10.48550/arXiv.2305.10601>. (accessed Feb. 18, 2024)
- [48] A* Search. *codecademy*, 2023. <https://www.codecademy.com/resources/docs/ai/search-algorithms/a-star-search>. (accessed Feb. 11, 2024)
- [49] Chatterjee, M.: What is A* Search Algorithm? 2024. *BLOG*. Great Learning, 2024. <https://www.mygreatlearning.com/blog/a-search-algorithm-in-artificial-intelligence/>. (access. Feb. 16, 2024)
- [50] Mitrev, D.: Google Introduced Dreamer: A New Reinforcement Learning Agent. *Neurohive*, 2020. <https://neurohive.io/en/news/google-introduced-dreamer-a-new-reinforcement-learning-agent/>. (accessed Dec. 13, 2023)
- [51] Hafner, D., Lillicrap, T., Ba, J., Norouzi, M.: Dream to Control: Learning Behaviors by Latent Imagination. *arXiv*. Cornell Univ., 2020. [arXiv:1912.01603v3](https://arxiv.org/abs/1912.01603v3); <https://doi.org/10.48550/arXiv.1912.01603>. (accessed Dec. 13, 2023)
- [52] Auto-associative Neural Networks. *GeeksforGeeks*, 2021. <https://www.geeksforgeeks.org/auto-associative-neural-networks/>. (accessed Dec. 15, 2023)
- [53] Pohlen, T., Piot, B., Hester, T., Gheshlaghi Azar, M., Horgan, D., Budden, D., Barth-Maron, G., van Hasselt, H., Quan, J., Večerík, M., Hessel, M., Munos, R., Pietquin, O.: Observe and Look Further: Achieving Consistent Performance on Atari. *arXiv*. Cornell Univ., 2018. [arXiv:1805.11593v1](https://arxiv.org/abs/1805.11593v1); <https://doi.org/10.48550/arXiv.1805.11593>. (accessed Oct. 21, 2023)
- [54] Ecoffet, A., Huizinga, J., Lehman, J. et al.: First return, then explore. *Nature* 590, pp. 580–586 (2021). <https://doi.org/10.1038/s41586-020-03157-9>. (accessed Dec. 28, 2023)
- [55] Balestriero, R., Ibrahim, M., Sobal, V., Morcos, A., Shekhar, S., Goldstein, T., Bordes, F., Bardes, A., Mialon, G., Tian, Y., Schwarzschild, A., Gordon A., Geiping, W. J., Garrido, Q., Fernandez, P., Bar, A., Pirsiavash, H., LeCun, H., Goldblum, M.: A Cookbook of Self-Supervised Learning. *arXiv*. Cornell University. April 24, 2023. <https://arxiv.org/abs/2304.12210>, [arXiv:2304.12210](https://arxiv.org/abs/2304.12210), <https://doi.org/10.48550/arXiv.2304.12210> (accessed Apr. 29, 2023)
- [56] Li, X., Shang, J., Das, S., Ryoo, M. S.: Does Self-supervised Learning Really Improve Reinforcement Learning from Pixels? *arXiv*. Cornell Univ. 2023. [arXiv:2206.05266v4](https://arxiv.org/abs/2206.05266v4); <https://doi.org/10.48550/arXiv.2206.05266>. (accessed Oct. 29, 2023)
- [57] Video generation models as world simulators. *OpenAI/Research*. 2024. <https://openai.com/research/video-generation-models-as-world-simulators>. (accessed Feb. 29, 2024)
- [58] Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, Ch., Eslami, S. M. A., Botvinick, M.: Machine Theory of Mind. *arXiv*. Cornell Univ., 2019. [arXiv:1802.07740v2](https://arxiv.org/abs/1802.07740v2); <https://arxiv.org/pdf/1802.07740.pdf>. (accessed Oct. 29, 2023)
- [59] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 23-48. ISBN 978-3-0356-2400-7.
- [60] ARTIFICIAL INTELLIGENCE, MACHINE LEARNING, NEURAL NETWORKS. *Datamind*. https://www.datamind.cz/en/services/Artificial_intelligence_machine_learning_deep_neural_networks?msc_lkid=83db14a2aa13173a1045708f78459aa3. (accessed Nov. 3, 2023)
- [61] Barreto, S.: Differences Between Backpropagation and Feedforward Networks. *Baeldung*. <https://www.baeldung.com/cs/neural-networks-backprop-vs-feedforward> (accessed Nov. 5, 2023)
- [62] Burke, J.: Successful generative AI examples worth noting. *Tech Accelerator*. <https://www.techtarget.com/searchenterpriseai/tip/Successful-generative-AI-examples-worth-noting>. (accessed Sept. 15, 2023)

- [63] Recursive neural network. WIKIPEDIA. https://en.wikipedia.org/wiki/Recursive_neural_network. (accessed Nov. 7, 2023)
- [64] Convolutional neural network. WIKIPEDIA. https://en.wikipedia.org/wiki/Convolutional_neural_network. (accessed Nov. 5, 2022)
- [65] IMAGENET. <https://www.image-net.org>. (accessed Nov. 5, 2022)
- [66] Introduction to Capsule Neural Networks. *GeeksforGeeks*. 2024. <https://www.geeksforgeeks.org/capsule-neural-networks-ml/>. (accessed Mar. 12, 2024)
- [67] Liu, Z., Zhou, J.: *Introduction to Graph Neural Networks*. Springer International Publishing, 2020. ISBN 9783031004599.
- [68] Jost, Z.: Foundations of Graph Neural Networks. *Welcome AI Overlords*. <https://www.graphneuralnets.com/p/full-course>. (accessed Sept. 8, 2023)
- [69] Boesch, G.: Guide to Generative Adversarial Networks (GANs) in 2024. *viso.ai*. <https://viso.ai/deep-learning/generative-adversarial-networks-gan/>. (accessed Mar. 18, 2024)
- [70] Wood, T.: Generative Adversarial Network. Deepai. <https://deepai.org/machine-learning-glossary-and-terms/generative-adversarial-network>. (accessed Nov. 21, 2023)
- [71] Leach, N.: *Architecture in the Age of Artificial Intelligence An introduction to AI for architects*. Bloomsbury Visual Arts, London, 2022. pp. 26-64. ISBN 978-1-3501-6551-9.
- [72] Chakraborty, T., Reddy K. S. U., Naik, S. M., Panja, M., Manvitha, B.: Ten Years of Generative Adversarial Nets (GANs): A survey of the state-of-the-art. *arXiv*. Cornell University, 2023. <https://arxiv.org/pdf/2308.16316.pdf>. (accessed Sept. 12, 2023)
- [73] Saxena, D., Cao, J.: Generative Adversarial Networks (GANs): Challenges, Solutions, and Future Directions. *arXiv*. Cornell University, 2020. <https://arxiv.org/ftp/arxiv/papers/2005/2005.00065.pdf>. (accessed Oct. 2, 2023)
- [74] Yinka-Banjo, Ch., Ugot, O-A.: A review of generative adversarial networks and its application in cybersecurity. *Springer Link*. 2019. <https://link.springer.com/article/10.1007/s10462-019-09717-4>. (accessed Oct. 4, 2023)
- [75] Rehman, J.: Advantages and disadvantages of generative adversarial networks (GAN). IT Release. 2020. https://www.itrelease.com/2020/06/advantages-and-disadvantages-of-generative-adversarial-networks-gan/#google_vignette. (accessed Oct. 4, 2023)
- [76] What's Next for GANs: Latest Techniques and Applications. *Sciforce/Medium*. 2022. <https://medium.com/sciforce/whats-next-for-gans-latest-techniques-and-applications-3be06a7e5ab9>. (accessed Oct. 8, 2023)
- [77] Moghadam, M. M., Boroomand, B., Jalali, M., Zareian, A., DaeiJavad, A., Manshaei, M. H., Krunz, M.: Game of GANs: Game-Theoretical Models for Generative Adversarial Networks. *arXiv*. Cornell University, 2022. [arXiv:2106.06976v3](https://arxiv.org/abs/2106.06976); <https://doi.org/10.48550/arXiv.2106.06976>. (accessed Jan. 18., 2023)
- [78] Sampath, V., Murtua, I., Aguilar Martín, J. J. et al.: A survey on generative adversarial networks for imbalance problems in computer vision tasks. *Journal of Big Data* 8, 27 (2021). <https://doi.org/10.1186/s40537-021-00414-0>. (accessed Jan. 18., 2023)
- [79] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 112-146. ISBN 978-3-0356-2400-7.
- [80] Comparison between Diffusion Models vs GANs. MLK. 2023. <https://machinelearningknowledge.ai/comparison-between-diffusion-models-vs-gans-generative-adversarial-networks/>. (accessed Nov. 6., 2023)
- [81] Bouchard, L.: How Stable Diffusion works? Latent Diffusion Models Explained. *Louis-François Bouchard, aka What's AI*. 2022. <https://www.louisbouchard.ai/latent-diffusion-models/>. (accessed Jan. 16., 2023)
- [82] Recursive neural network. WIKIPEDIA. https://en.wikipedia.org/wiki/Recursive_neural_network. (accessed Nov. 7, 2022)

- [83] Schmidt, R. M.: Recurrent Neural Networks (RNNs): A gentle Introduction and Overview. *arXiv*. Cornell University, 2019. [arXiv:1912.05911v1](https://arxiv.org/abs/1912.05911v1); <https://doi.org/10.48550/arXiv.1912.05911>. (accessed Nov. 7, 2022)
- [84] Karimi, A.: Recurrent vs. Recursive Neural Networks in Natural Language Processing. *Baeldung*. 2024. <https://www.baeldung.com/cs/networks-in-nlp>. (accessed Mar. 22, 2024)
- [85] Deep Learning | Introduction to Long Short-Term Memory. *GeeksforGeeks*. 2023. <https://www.geeksforgeeks.org/deep-learning-introduction-to-long-short-term-memory/>. (accessed Mar. 12, 2024)
- [86] Echo State Network – an overview. *GeeksforGeeks*. 2023. <https://www.geeksforgeeks.org/echo-state-network-an-overview/>. (accessed Mar. 12, 2024)
- [87] Difference between Recursive and Recurrent Neural Network. *GeeksForGeeks*. 2024. <https://www.geeksforgeeks.org/difference-between-recursive-and-recurrent-neural-network/> (accessed Mar. 22, 2024)
- [88] Bhargav, N.: Radial Basis Function. *Baeldung*. 2024. <https://www.baeldung.com/cs/rbf-neural-networks>. (accessed Mar. 12, 2024)
- [89] Stewart, M.: GANs vs, Autoencoders: Comparison of Deep Generative Models. *Towards Data Science*. <https://towardsdatascience.com/gans-vs-autoencoders-comparison-of-deep-generative-models-985cf15936ea> (accessed Nov. 5, 2022)
- [90] What are the fundamental differences between VAE and GAN for image generation? *StackExchange*. <https://ai.stackexchange.com/questions/25601/what-are-the-fundamental-differences-between-vae-and-gan-for-image-generation>. (accessed Sept. 15, 2022)
- [91] Variational autoencoder. *WIKIPEDIA*. https://en.wikipedia.org/wiki/Variational_autoencoder. (accessed Sept. 15, 2022)
- [92] Kingma, D. P., Welling, M.: An Introduction to Variational Autoencoders. *arXiv*. Cornell University, 2019. [arXiv:1906.02691v3](https://arxiv.org/abs/1906.02691v3); <https://arxiv.org/abs/1906.02691>. (accessed Sept. 16, 2022)
- [93] Hayashi, K., Sasagawa, A., Sakaino, S., Tsuji, T.: A New Autoregressive Neural Network Model with Command Compensation for Imitation Learning Based on Bilateral Control. *ResearchGate*. 2021. https://www.researchgate.net/publication/350104951_A_New_Autoregressive_Neural_Network_Model_with_Command_Compensation_for_Imitation_Learning_Based_on_Bilateral_Control. (accessed Sept. 17, 2022)
- [94] Bayarov Ahmedov, H., Yi, D., Sui, J.: Brain-Inspired Deep Imitation Learning for Autonomous Driving Systems. *arXiv*. Cornell University, 2021. [arXiv:2107.14654v1](https://arxiv.org/abs/2107.14654v1); <https://doi.org/10.48550/arXiv.2107.14654>. (accessed Mar. 20, 2024)
- [95] Kim, H., Ohmura, Y., Kuniyoshi, Y.: Transformer-based deep imitation learning for dual-arm robot manipulation. *arXiv*. Cornell University, 2024. [arXiv:2108.00385v2](https://arxiv.org/abs/2108.00385v2); <https://doi.org/10.48550/arXiv.2108.00385>. (accessed Mar. 20, 2024)
- [96] Correia, de Santana, A., Colombini, E. L.: Neural Attention Models in Deep Learning: Survey and Taxonomy. *arXiv*. Cornell University, 2021. <https://arxiv.org/pdf/2112.05909>. (accessed Mar. 2, 2024)
- [97] Razzaq, W., Mo, H.: Neural Circuit Policies Imposing Visual Perceptual Autonomy. *Neural Processing Letters* 55(7)/*ResearchGate*. 2023. DOI: [10.1007/s11063-023-11194-4](https://doi.org/10.1007/s11063-023-11194-4). https://www.researchgate.net/publication/368752388_Neural_Circuit_Policies_Imposing_Visual_Perceptual_Autonomy. (accessed Mar. 12, 2024)
- [98] Kemtur, A., Paugam, F., Pinsard, B., Sainath, P., Le Clei, M., Boyle, J., Jerbi, K., Bellec. P.: Behavioral Imitation with Artificial Neural Networks Leads to Personalized Models of Brain Dynamics During Videogame Play. *BioRxiv*. 2023. doi: <https://doi.org/10.1101/2023.10.28.564546>. <https://www.biorxiv.org/content/10.1101/2023.10.28.564546v1>. (accessed Mar. 10, 2024)
- [99] Firth-Butterfield, K.: These were the biggest AI developments in 2022. Now we must decide how to use them. *World Economic Forum*. <https://www.weforum.org/agenda/2023/01/davos23-biggest-ai-developments-how-to-use-them>. (accessed Feb.. 19, 2023)
- [100] Gozallo-Birzuela, R., Garrido-Merchan, E. C.: GPT is not all you need. A State of the Art Review of large Generative AI models. *arXiv*. Cornell University. <https://arxiv.org/abs/2301.04655>. (accessed Jan. 29, 2023)

- [101] LayerDiffuse. *GitHub*. 2023. <https://github.com/layerdiffusion/LayerDiffuse>. (accessed Mar. 23, 2024)
- [102] What are Foundation Models? *aws*. <https://aws.amazon.com/what-is/foundation-models/>. (accessed Aug. 2, 2023)
- [103] Nandi, B.: Canva is outdated. *@LearnWithBishal, X*. 2024. <https://x.com/LearnWithBishal/status/1766466763783364796?t=ZieTMInTrsQWDAGxTxKmWQ&s=09>. (accessed Mar. 18, 2024)
- [104] Zyer, E.: 60 AI Tools to Start your Profitable Online Business in 2024. *@eyishazyer, X*. 2024. <https://x.com/eyishazyer/status/1765687546305184128?t=MeNHWSVM5SILktZxzYliYA&s=09>. (accessed Mar. 16, 2024)
- [105] Chetan, A.: Forget ChatGPT, it costs \$20/M to use. *@AdarshChetan, X*. 2024. , <https://x.com/AdarshChetan/status/1768505452642369612?t=I2YWPu3TSjcarZsTMNZmow&s=09>. (accessed Mar. 16, 2024)
- [106] Hamza, K.: Here is how to start a faceless YouTube Channel: *@Whizz_ai, X*. 2024. https://x.com/Whizz_ai/status/1764886228896247951?t=3J9Y4SdrKatTSqhStjZ7Rg&s=09. (accessed Mar. 17, 2024)
- [107] Introducing the next generation of Claude. *Anthropic*. 2024. <https://www.anthropic.com/news/claude-3-family>. (accessed Mar. 12, 2024)
- [108] LMSYS Chatbot Arena Leaderboard. *Hugging Face*. 2024. https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard?utm_source=substack&utm_medium=email. (accessed Mar. 31, 2024)
- [109] Rahman, R.: Forget ChatGP. *@heyronir, X*. 2024. <https://twitter.com/heyronir/status/1764886267647709361>. (accessed Mar. 9, 2024)
- [110] Anthropic: *Model Card and Evaluations for Claude Models*. 2023. <https://www-cdn.anthropic.com/files/4zrzovbb/website/bd2a28d2535bfb0494cc8e2a3bf135d2e7523226.pdf>. (accessed Mar. 12, 2024)
- [111] Claude (language model). *Wikipedia*. 2024. https://en.wikipedia.org/wiki/Claude_%28language_model%29. (accessed Mar. 14, 2024)
- [112] Frontier AI in your hands. *Mistral AI*. 2023. <https://mistral.ai>. (accessed Feb. 15, 2024)
- [113] Mixtral of experts. *Mistral AI*. 2023. <https://mistral.ai/news/mixtral-of-experts/>. (accessed Feb. 15, 2024)
- [114] Heath, R., Fried, I.: Mistral, it's French for AI. *Axios AI+*. 2024. https://www.axios.com/newsletters/axios-ai-plus-8c79d014-d04c-4026-97ab-a8fa913b4489.html?utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axioslogin&stream=top. (accessed Mar. 9, 2024)
- [115] David, E.: Meta is planning a July release for a “looser” Llama 3. *The Verge*. 2024. <https://www.theverge.com/2024/2/28/24085755/meta-is-planning-a-july-release-for-a-looser-llama-3>. (accessed Mar. 8, 2024)
- [116] Bilenko, M.: Introducing Phi-3: Redefining what’s possible with SLMs. *Azure*. <https://azure.microsoft.com/en-us/blog/introducing-phi-3-redefining-whats-possible-with-slms/>. (accessed Apr. 25, 2024)
- [117] Microsoft GenAI research team: Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. *arXiv*. Cornell University, 2024. arXiv:2404.14219v2; <https://export.arxiv.org/abs/2404.14219>. (accessed Apr. 26, 2024)
- [118] Bishop, Ch. M.: *Pattern Recognition and Machine Learning*. Springer. ISBN 0-387-31073-8.
- [119] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed, W.W.Norton, 2020, p. 126. ISBN 9780393635829.
- [119] stylemania.it: Před Chat GPT: zde jsou průkopníci, kteří připravili půdu pro umělou inteligenci. *Style Mania*. <https://www.msn.com/cs-cz/zpravy/other/před-chat-gpt-zde-jsou-průkop-níci-kteří-přip-ravili-půdu-pro-umělou-inteligenci/ss-AA18T6Hs?ocid=msedgdhp&pc=U531&cvid=-b90b042fad6c4cada69739b51d885451&ei=59#image=11> (accessed Mar. 20, 2023)

- [120] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed. W.W.Norton, 2020, pp. 28-64. ISBN 9780393635829.
- [121] Generative AI for Games Market Map. X. 2024. <https://twitter.com/ashnobe/status/1595793418932097024?t=qQ7Ppu1D-X-27kmqjVnmDg&s=09>. (accessed Mar. 20, 2023)
- [122] Your 3D designs on lifelike AI-models. *LALALAND*. 2023. <https://lalaland.ai/>. (accessed Oct. 22, 2023)
- [123] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed. W.W.Norton, 2020, pp. 28-64. ISBN 9780393635829.
- [124] Ngila, F.: The GitHub of AI is named after an emoji—and Microsoft has its fingers in it already. *QUARTZ*. 2024. file:///C:/Users/uzivatel/Desktop/texty_2020/AI_2023/CEJ_2024/next%20wave/A%20startup%20called%20Hugging%20Face%20is%20becoming%20the%20GitHub%20of%20AI.pdf. (accessed Mar. 15, 2024)
- [125] AI Blog Creator: Perplexity AI Where knowledge begins. *DailyAIBlogs*. 2024. <https://www.dailyaiblogs.com/perplexity-ai/>. (accessed Dec. 15, 2023)
- [126] Jin, B., Krppa, M.: AI Search Startup Perplexity Set to Double Valuation to \$1 Billion. *The Wall Street Journal*. 2024. <https://www.livemint.com/ai/artificial-intelligence/ai-search-startup-perplexity-set-to-double-valuation-to-1-billion-11709644470097.html>. (accessed Mar. 17, 2024)
- [127] Daniel Gross. 2024. https://dcgross.com/?utm_source=substack&utm_medium=email. (accessed Mar 31, 2024)
- [128] Copilot. *Microsoft*. 2023. <https://www.microsoft.com/en-us/microsoft-copilot>. (accessed Nov. 5, 2023)
- [129] The world's most widely adopted AI developer tool. *Github*. 2022. <https://github.com/features/copilot>. (accessed Nov. 5, 2023)
- [130] What is an AI Copilot? *Salesforce Asia Blog*. 2023. <https://www.salesforce.com/ap/blog/ai-copilot/>. (accessed Mar. 6, 2024)
- [131] Poda, M.: What is an AI copilot? *Moveworks Blog*. 2023. <https://www.moveworks.com/us/en/resources/blog/what-is-an-ai-copilot>. (accessed Mar. 6, 2024)
- [132] Pykes, K.: How to Run Stable Diffusion. *Datacamp*. 2023. <https://www.datacamp.com/tutorial/how-to-run-stable-diffusion>. (accessed Mar. 6, 2024)
- [133] Hossain, A.: AI Prompt Engineers Can Earn \$300k Salaries. *Helping you to make money with AI, Tech Tools & Digital Skills*, X. 2024. https://www.linkedin.com/posts/alamipro_ai-prompt-engineers-can-earn-300k-salaries-activity-7176611454949720064-Ajt5/?utm_source=share&utm_medium=member_android. (accessed Mar. 19, 2024)
- [134] Where creativity knows no bounds. *Musavir.ai*. 2023. <https://musavir.ai/>. (accessed Jan. 22, 2024)
- [135] Cintas, A.: Ideogram has released a new 1.0 version of its text-to-image AI model, which enhances photorealism and text rendering. @dr_cintas, X. 2024. https://twitter.com/dr_cintas/status/1763250401413881994?t=LpfdZoiWXdVjyiXW9Y0Ycw&s=09. (accessed Jan. 16, 2024)
- [136] Bringing AI to Excel—4 new features announced today at Ignite. *Microsoft*. 2018. <https://www.microsoft.com/en-us/microsoft-365/blog/2018/09/24/bringing-ai-to-excel-4-new-features-announced-today-at-ignite/>. (accessed Mar. 12, 2024)
- [137] Welcome to Windows. *Microsoft*. 2024. <https://www.microsoft.com/en-us/windows/?r=1>. (accessed Mar. 4, 2024)
- [138] Devin. *EASY WITH AI*. 2024. <https://easywithai.com/ai-developer-tools/devin/>. (accessed Mar. 4, 2024)
- [139] Nouri, S.: TOP 200 Generative AI Tools 2023!. *LinkedIn*. 2024. https://www.linkedin.com/posts/stevenouri_gen-ai-tools-ugcPost-7149239580200697856uVj4/?utm_source=share&utm_medium=member_android. (accessed Feb. 14, 2024)
- [140] Pichai, S.: An important next step on your AI journey. *Google The Keyword*. <https://blog.google/technology/ai/bard-google-ai-search-updates/>. (accessed Apr. 2, 2023)

- [141] Martin, B.: Investors didn't think much of Baidu's ChatGPT – competitor debut, wiping \$3 billion off the tech giant's value – Fortune. *Inferse*. <https://www.inferse.com/484983/investors-didnt-think-much-of-baidus-chatgpt-competitor-debut-wiping-3-billion-off-the-tech-giants-value-fortune/>. (accessed Apr. 16, 2023)
- [142] Kilian, K.: Čína má odpověď na ChatGPT. A ne jednu, další chatbot ukazuje Alibaba. *Živě*. <https://www.zive.cz/clanky/cina-ma-odpoved-na-chatgp-a-ne-jednu-dalsi-ai-chatbot-ukazuje-alibaba/sc-3-a-221685/default.aspx>. (accessed Apr. 16, 2023)
- [143] Urban, P.: Amazon má továrnu na AI a chatboty. Na umělé inteligenci u něj pracuje víc lidí než u Microsoftu nebo Googlu. *Connect!* <https://connect.zive.cz/clanky/amazon-bedrock-titan-generativni-ai/sc-320-a-221729/default.aspx>. (accessed Apr. 16, 2023)
- [144] Coldewey, D.: Amazon doubles down on Anthropic, completing its planned \$4B investment. *TechCrunch*. 2024. https://techcrunch.com/2024/03/27/amazon-doubles-down-on-anthropic-completing-its-planned-4b-investment/?utm_source=aitoolreport.beehiiv.com&utm_medium=newsletter&utm_campaign=amazon-fights-openai-with-4b-anthropic-funding. (accessed Mar. 30, 2024)
- [145] Wiggers, K.: What is Elon Musk's Grok chatbot and how does it work? *TechCrunch*. 2024. <https://techcrunch.com/2024/03/29/what-is-elon-musks-grok-chatbot-and-how-does-it-work/>. (accessed Mar. 30, 2024)
- [146] Wiggers, K.: Databricks spent \$10M on new DBRX generative AI model. *TechCrunch*. 2024. https://techcrunch.com/2024/03/27/databricks-spent-10m-on-a-generative-ai-model-that-still-cant-beat-gpt-4/?utm_source=aitoolreport.beehiiv.com&utm_medium=newsletter&utm_campaign=amazon-fights-openai-with-4b-anthropic-funding. (accessed Mar. 31, 2024)
- [147] Introducing the GPT Store. *OpenAI blog*. 2024. <https://openai.com/blog/introducing-the-gpt-store>. (accessed Feb. 23, 2024)
- [148] Rahman, R.: Google Maps just got a HUGE AI upgrade. *@heyronir*, X. 2024. <https://twitter.com/heyronir/status/1766339186611953889?t=pcjnXaHsACuVQ9tYOuURsA&s=09>. (accessed Mar. 12, 2024)
- [149] New Maps updates: Immersive View for routes and other AI features. *Google The Keyword*. 2024. <https://blog.google/products/maps/google-maps-october-2023-update/>. (accessed Mar. 12, 2024)
- [150] Unleash your Creativity with the power of Leonardo Ai. *Leonardo Ai*. <https://leonardo.ai/>. (accessed Feb. 20, 2024)
- [151] Rishabh: Google has introduced VideoPoet, and it's MIND-BLOWING. *@Rixhabh*, X. 2024. https://twitter.com/Rixhabh/status/1770052606293086516?t=elred_nCdYp6qcpcCq53Dg&s=09. (accessed Feb. 20, 2024)
- [152] Advancing creativity with artificial intelligence. *runway*. <https://runwayml.com>. (accessed Feb. 17, 2024)
- [153] Sora. 2024. <https://soraai.ai/>. (accessed Mar. 16, 2024)
- [154] Mishra, S.: Google just launched LUMIERE, and it's crazy. *@heyshrutimishra*, X. 2024. https://twitter.com/heyshrutimishra/status/1750492116109996496?t=dHXpBsHiD2cWkue_uP31bw&s=09. (accessed Mar. 17, 2024)
- [155] Cintas, A.: First Sora AI videos, and now Emo turns your images into talking or even singing videos. *@dr_cintas*, X. 2024. https://twitter.com/dr_cintas/status/1762864281647817123?t=WZj6WIZkuJATYIRvoLMkXA&s=09 (accessed Mar. 7, 2024)
- [156] Omer Bar-Tal, O., Chefer, H., Tov, O., Herrmann, Ch., Paiss, R., Zada, S., Ephrat, A., Hur, J., Liu, G., Raj, A., Li, Y., Rubinstein, M., Michaeli, T., Wang, O., Sun, D., Dekel, T., Mosseri, I.: Lumiere: A Space-Time Diffusion Model for Video Generation. *arXiv*. Cornell University. 2024. [arXiv:2401.12945v2; https://doi.org/10.48550/arXiv.2401.12945](https://arxiv.org/abs/2401.12945). (accessed Mar. 20, 2024)
- [157] Tian, L., Wang, Q., Zhang, B., Bo, L.: EMO: Emote Portrait Alive - Generating Expressive Portrait Videos with Audio2Video Diffusion Model under Weak Conditions. *arXiv*. Cornell University. 2024.

- arXiv:2402.17485v1; <https://doi.org/10.48550/arXiv.2402.17485>; <https://humanaigc.github.io/emote-portrait-alive/>. (accessed Mar. 7, 2024)
- [158] Creating video from text. *OpenAI*. 2024. <https://openai.com/sora>. (accessed Mar. 20, 2024)
- [159] OpenAI: Navigating the Challenges and Opportunities of Synthetic Voices. *Blog Open AI*. 2024. <https://openai.com/blog/navigating-the-challenges-and-opportunities-of-synthetic-voices#OpenAI>. (accessed Mar. 30, 2024)
- [160] Xu, S., Chen, G., Guo, Y., Yang, J., Li, Ch., Zang, Z., Zhang, Y., Tong, X., Guo, B.: VASA-1: Lifelike Audio-Driven Talking Faces Generated in Real Time. arxiv. Cornell University, 2024. [arXiv:2404.10667v1](https://doi.org/10.48550/arXiv.2404.10667); <https://doi.org/10.48550/arXiv.2404.10667>. (accessed Apr. 20, 2024)
- [161] Perez, C. E.: Sam Altman reveals in an interview with Bill Gates (2 days ago) what's coming up in GPT-4.5 (or GPT-5). *@IntuitMachine, X*. 2024. https://twitter.com/IntuitMachine/status/1746278269165404164?t=ikx_sw5O7m28-Xya6jBfVA&s=09. (accessed Mar. 17, 2024)
- [162] Top 6 AI 3D Model Generators: Explore New Dimensions of Creativity. *Viverse*. 2024. https://www.news.viverse.com/post/top-6-ai-3d-model-generators-explore-new-dimensions-of-creativity?wpsrc=Email&wpsnetn=marketo&wpcn=en_viverse_email94_021624&wpcid=6930&utm_source=crm&utm_medium=email&utm_campaign=en_viverse_email94_021624&utm_content=cta03&mkt_tok=MDU4LVNVVY04OTQAAAGRUXzluyS626xTddCcObMQ14XPBnQ-woViSi7m7Ec2yGDw261YR7v9PFcmh4AHkPH12hYIWYHF lwmiAnTVzThiOUWVMeNxMxy5X-ro8WgCRWjKa44. (accessed Mar. 30, 2024)
- [163] Introducing Figure-1. *Figure*. <https://www.figure.ai>. (accessed Mar. 16, 2024)
- [164] Adcock, B.: 2024 will be the year of AI and Robotics. Figure, Google DeepMind, Apple, Alibaba, Microsoft, and OpenAI have all made major progress over the past weeks. Here's everything you need to know. *X (twitter.com)*. 2024. https://twitter.com/adcock_brett/status/174440486224290253?t=umf_MLjzYePaV92sajLRyg&s=09. (accessed Mar. 15, 2024)
- [165] Shaping the future of advanced robotics. *Google DeepMind*. 2024. https://deepmind.google/discover/blog/shaping-the-future-of-advanced-robotics/?utm_source=twitter&utm_medium=social. (accessed Jan. 10, 2024)
- [166] Nvidia Project GR00T Generalist Robot 00 Technology. *NVIDIA DEVELOPER*. 2024. <https://developer.Nvidia.com/project-gr00t>. (accessed Mar. 30, 2024)
- [167] Nvidia Announces Project GR00T Foundation Model for Humanoid Robots and Major Isaac Robotics Platform Update. *NVIDIA NEWSROOM*. 2024. <https://Nvidianews.Nvidia.com/news/foundation-model-isaac-robotics-platform>. (accessed Mar. 30, 2024)
- [168] Nvidia Announces Project GR00T Foundation Model For Humanoid Robots And Major Isaac Robotics Platform Update. *Nvidia Investors*. 2024. <https://investor.Nvidia.com/news/press-release-details/2024/NVIDIA-Announces-Project-GR00T-Foundation-Model-for-Humanoid-Robots-and-Major-Isaac-Robotics-Platform-Update/default.aspx>. (accessed Mar. 30, 2024)
- [169] Krizhnevsky, A., Sutskever, I., Hinton, G.: ImageNet classification with deep convolutional neural networks. *Computer Science, Communications of the ACM*. Dec. 3, 2012. doi:10.1145/3065386. <https://www.semanticscholar.org/paper/ImageNet-classification-with-deep-convolutional-Krizhevsky-Sutskever/abd1c342495432171beb7ca8fd9551ef13cbd0ff> (accessed Feb 23, 2023)
- [170] *Google Scholar* https://scholar.google.com/citations?view_op=view_citation&hl=en&user=xezhJcAAAAJ&citation_for_view=xegzhJcAAAAJ:u5HHmVD_uO8C (accessed Feb 28, 2023)
- [171] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed, W.W.Norton, 2020, p. 126. ISBN 9780393635829.
- [172] Campbell, M., Hoane, A. J. Jr., Hsu, F.: Deep Blue. Artificial Intelligence. Elsevier. vol. 134, issues 1-2, Jan. 2002, pp. 57-83. [https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1) (accessed Feb. 16, 2023)
- [173] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed. W.W.Norton, 2020, pp. 241-243. ISBN 9780393635829.
- [174] AlphaGo Zero: Starting from scratch. *Google DeepMind*. deepmind.com/blog/alphago-zero-starting-from-scratch (accessed Feb. 1, 2023)

- [175] AlphaGo. *Google DeepMind*. https://www.deepmind.com/research/highlighted-research/_alpha (accessed Jan 21, 2023)
- [176] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Fan, H., Sifre, L., Driessche, G. v. d., Graepel, T., Hassabis, D.: (19 October 2017). Mastering the game of Go without human knowledge (PDF). *Nature*. 550 (7676), pp. 354–359. Bibcode:2017Natur.550..354S. doi:10.1038/nature24270. ISSN 0028-0836. PMID 29052630. S2CID 205261034. closed access.
- [177] Ferguson, N.: Resistance Is Futile, But Maybe Not With AI. *Bloomberg*. 2024. <https://www.bloomberg.com/opinion/articles/2024-03-24/can-sam-altman-make-ai-smart-enough-to-answer-these-6-questions?srnd=undefined>. (accessed Mar. 30, 2024)
- [178] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed. W.W.Norton, 2020, p. 109. ISBN 9780393635829.
- [179] Chiang, T.: ChatGPT Is a Blurry JPEG of the Web OpenAI's chatbot offers paraphrases, whereas Google offers quotes. Which do we prefer? *The New Yorker*. Advance Publications, New York. February 9, 2023. ISSN 0028-792X.
- [180] Liu, Z., Feng, Y., Black, M. J., Nowrouzezahrai, D., Paull, L., Liu, W.: MeshDiffusion: Score-based Generative 3D Mesh Modelling. *arXiv*. Cornell University. [arXiv:2303.08133](https://arxiv.org/abs/2303.08133); [arXiv:2303.08133v2](https://arxiv.org/abs/2303.08133v2); <https://doi.org/10.48550/arXiv.2303.08133> (accessed Apr. 21, 2023)
- [181] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed. W.W.Norton, 2020, p. 113. ISBN 9780393635829.
- [182] Edwards, B.: GPT-5 might arrive this summer as a “materially better” update to ChatGPT. *ars TECHNICA*. 2024. <https://arstechnica.com/information-technology/2024/03/openai-gpt-5-may-launch-this-summer-upgrading-chatgpt-along-the-way/>. (accessed Mar. 26, 2024)
- [183] Davos 2023: A conversation with Palantir's Alex Karp. World Economic Forum. 2023. <https://www.weforum.org/podcasts/meet-the-leader/episodes/palantir-alex-karp-saying-yes-thinking-differently/>. (accessed Mar. 12, 2024)
- [184] Zorloni, L.: Il Garante della privacy blocca ChatGPT in Italia. *Wired*. <https://www.wired.it/article/chatgpt-blocco-italia-garante-privacy/>. (accessed Apr. 2, 2023)
- [185] multiple authors, Harari, Y., N., Musk, E., Wozniak, S., Bengio, Y., Russel, S. among others: Pause Giant AI Experiments: An Open Letter. *future of life*. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>. (accessed Apr. 3, 2023)
- [186] Malik, K.: Elon Musk v OpenAI: tech giants are inciting existential fears to evade scrutiny. *The Guardian*. 2024. https://www.theguardian.com/commentisfree/2024/mar/10/ai-wont-destroy-us-but-tech-giants-use-fear-it-will-to-evade-scrutiny?utm_term=65ed5a781147ba6a86383433a27924d3&utm_campaign=GuardianTodayUK&utm_source=esp&utm_medium=Email&CMP=GTUK_email. (accessed Mar. 10, 2024)
- [187] Heath, R.: AI-fueled scams target tax refunds. *Axios AI+*. 2024. <https://www.axios.com/2024/03/20/tax-returns-scam-ai-cybersecurity>. (accessed Mar. 23, 2024)
- [188] Cather, A.: Can government work be fast or efficient? An AI engineer transforming the public sector says yes. *QUARTZ*. 2024. <https://qz.com/can-government-work-be-fast-or-efficient-an-ai-engineer-1851279150>. (accessed Mar. 25, 2024)
- [189] Excellence and trust in artificial intelligence. European Commission. 2024. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/excellence-and-trust-artificial-intelligence_en. (accessed Mar. 28, 2024)
- [190] Gibney, E.: What the EU's tough AI law means for research and ChatGPT. *nature*. 2024. <https://www.nature.com/articles/d41586-024-00497-8>. (accessed Mar. 12, 2024)
- [191] WHITE PAPER On Artificial Intelligence - A European approach to excellence and trust. European Commission. 2020. https://commission.europa.eu/document/download/d2ec4039-c5be-423a-81ef-b9e44e79825b_en?filename=commission-white-paper-artificial-intelligence-feb2020_en.pdf. (accessed Mar. 12, 2024)

- [192] Saballa, J.: Biden, Xi to Sign Deal Prohibiting AI in Drones, Nuclear Weapons: Report. *The Defense Post*. 2023. <https://www.thedefensepost.com/2023/11/15/biden-xi-prohibit-ai/>. (accessed Jan. 23, 2024)
- [193] Weatherbed, J.: International agreement on AI safety gains signatures from 18 countries. *The Verge*. 2023. <https://www.theverge.com/2023/11/27/23977571/international-agreement-on-ai-safety-gains-signatures-from-18-countries>. (accessed Jan. 24, 2024)
- [194] Sandle, P., Coulter, M.: AI Safety Summit: China, US, and EU agree to work together. *Reuters*. 2023. <https://www.reuters.com/technology/britain-brings-together-political-tech-leaders-talk-ai-2023-11-01/#:~:text=China%20agreed%20to%20work%20with%20the%20United%20States%2C,safe%20way%20forward%20for%20the%20rapidly%20evolving%20technology>. (accessed Jan. 24, 2024)
- [195] AI legislation, lawmakers and companies to watch right now. *AXIOS PRO*. 2024. https://www.axios.com/pro/reports/ai-legislation-lawmakers-companies?utm_adgroup=&keywords=ai%20regulations&subid=&utm_device=c&gclid=EAlaIqobChMlvPCx2NP-hAMV4QkGAB3_QApLEAAYAyAAEgK7CvD_BwE&campaignid=20010038575&adgroupid=153370161295&adid=672038612286&utm_medium=search-pd&utm_source=google-pd&utm_campaign=subs-search-goog-ProPolicy-Whitepaper&utm_term=ai%20regulations&hsa_acc=2813569195&hsa_cam=20010038575&hsa_grp=153370161295&hsa_ad=672038612286&hsa_src=g&hsa_tgt=kwd-991024546247&hsa_kw=ai%20regulations&hsa_mt=b&hsa_net=adwords&hsa_ver=3&gad_source=1. (accessed Mar. 29, 2024)
- [196] Wiggers, K.: The current legal cases against generative AI are just the beginning. *Techcrunch*. <https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/?guccounter=1> (accessed Feb. 2, 2023)
- [197] THE NEW YORK TIMES COMPANY, Plaintiff, v. MICROSOFT CORPORATION, OPENAI, INC., OPENAI LP, OPENAI GP, LLC, OPENAI, LLC, OPENAI OPCO LLC, OPENAI GLOBAL LLC, OAI CORPORATION, LLC, and OPENAI HOLDINGS, LLC, Defendants, Civil Action No. _____, COMPLAINT, JURY TRIAL DEMANDED. UNITED STATES DISTRICT COURT, SOUTHERN DISTRICT OF NEW YORK. 2023. https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf. (accessed Mar. 12, 2024)
- [198] Roth, E.: The New York Times is suing OpenAI and Microsoft for copyright infringement. *The Verge*. 2023. <https://translate.google.cz/?hl=cs&sl=en&tl=cs&text=The%20New%20York%20Times%20is%20suing%20OpenAI%20and%20Microsoft%20for%20copyright%20infringement.%20The%20Verge.%202024.%20h&op=translate>. (accessed Dec. 30, 2023)
- [198] David, E.: The Intercept, Raw Story, and AlterNet sue OpenAI and Microsoft. *The Verge*. 2024. <https://www.theverge.com/2024/2/28/24085973/intercept-raw-story-alternet-openai-lawsuit-copyright>. (accessed Mar. 3, 2024)
- [199] Rai, S.: Elon Musk Sues OpenAI and Altman for Breaching Firm's Founding Mission. *Bloomberg*. 2024. <https://www.bloomberg.com/news/articles/2024-03-01/musk-sues-openai-altman-for-breaching-firm-s-founding-mission>. (accessed Mar. 12, 2024)
- [200] Verses: AGI Breakthrough An Open Letter to the Board of Directors at OpenAI. *The New York Times*. 2023. <https://www.nytimes.com/paidpost/verses/2023-verses-ai-open-letter/agi-breakthrough.html>. (accessed Jan. 15, 2024)
- [201] Cheung, R.: The WEF is famous for its annual meeting among the world's most rich and powerful. X. 2024. <https://twitter.com/rowancheung/status/1747484395911418035?t= bPoXzYR2J0OUfsjTr0Ggw&s=09>. (accessed Jan. 21, 2024)
- [202] Navratil, D.: Evropa je schopna generovat cca stejné množství top expertů na AI jako USA. *LinkedIn*. 2024. https://www.linkedin.com/posts/david-navratil_evropa-je-schopna-generovat-cca-steiné-množství-activity-7162684357231517697-Rt6f/?utm_source=share&utm_medium=member_android. (accessed Mar. 18, 2024)
- [203] Orrru, M.: Microsoft's OpenAI Investment Could Face EU Probe. *Morningstar*. 2024. <https://www.morningstar.com/news/dow-jones/202401093218/microsofts-openai-investment-could-face-eu-probe>. (accessed Feb. 14, 2024)
- [204] Bulgaria GDP. *countryeconomy.com*. 2024. <https://countryeconomy.com/gdp/bulgaria>. (accessed Apr. 2, 2024)

- [205] Gurman, M.: Apple Explores Home Robotics as Potential ‘Next Big Thing’ After Car Fizzles. *Bloomberg*. 2024. <https://www.bloomberg.com/news/articles/2024-04-03/apple-explores-home-robots-after-abandoning-car-efforts>. (accessed Apr. 5, 2024)
- [206] Sawers, P.: Google DeepMind CEO Demis Hassabis gets UK knighthood for ‘services to artificial intelligence’. *TechCrunch*. 2024. https://techcrunch.com/2024/03/29/google-deepmind-ceo-demis-hassabis-gets-uk-knighthood-for-services-to-artificial-intelligence/?utm_source=aitoolreport.beehiiv.com&utm_medium=referral&utm_campaign=openai-can-replicate-your-voice-with-a-15-second-sample&guccounter=1. (accessed Apr. 2, 2024)
- [207] Ludlow, E., Bass, D., Tan, G., Metz, R.: Altman Seeks to Raise Billions for Network of AI Chip Factories. *Bloomberg*. 2024. <https://www.bloomberg.com/news/articles/2024-01-19/altman-seeks-to-raise-billions-for-network-of-ai-chip-factories>. (accessed Feb. 16, 2024)
- [208] 2024 Hualien earthquake. *Wikipedia*. 2024. https://en.wikipedia.org/wiki/2024_Hualien_earthquake. (accessed Apr. 6, 2024)
- [209] Letzing, J.: To fully appreciate AI expectations, look to the trillions being invested. *World Economic Forum*. 2024. https://www.weforum.org/agenda/2024/04/appreciate-ai-expectations-trillions-invested/?utm_source=sfmc&utm_medium=email&utm_campaign=2825108_WeeklyAgenda5April2024&utm_term=&emailType=Agenda%20Weekly. (accessed Apr. 10, 2024)
- [210] Smith, E.: Magnificent 7 profits now exceed almost every country in the world. Should we be worried? *CNBC*. 2024. <https://www.cnbc.com/2024/02/19/magnificent-7-profits-now-exceed-almost-every-country-in-the-world-should-we-be-worried.html>. (accessed Mar. 30, 2024)
- [211] The Buffett Indicator. 2024. <https://currentmarketvaluation.com/models/buffett-indicator.php>. (accessed Mar. 9, 2024)
- [212] Di Pizio, A.: Warren Buffett Is Sitting on a Record \$167.6 Billion Cash Pile. Here are 10 Stocks Berkshire Hathaway Could Buy Outright. *The Motley Fool. Nasdaq*. 2024. <https://www.nasdaq.com/articles/warren-buffett-is-sitting-on-a-record-%24167.6-billion-cash-pile.-here-are-10-stocks>. (accessed Mar. 9, 2024)
- [213] Zaha Hadid Studio: Antic Style Architectur. *Bing.com Videos*. <https://www.bing.com/videos/search?q=zaha+ai&docid=608008305523768389&mid=399E3EA3BEEECAC53B89399E3EA3BE EECAC53B89&view=detail&FORM=VIRE> (accessed Feb. 13, 2023)
- [214] Leach, N.: *Architecture in the Age of Artificial Intelligence An introduction to AI for architects*. Bloomsbury Visual Arts, London, 2022. p. 109. ISBN 978-1-3501-6551-9.
- [215] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 80, 86-89. ISBN 978-3-0356-2400-7.
- [216] Isola, P.: Phillip Isola. *MIT*. <http://web.mit.edu/phillipi/> (accessed Feb. 23, 2023)
- [217] Chai, L., Wulff, J., Isola, P.: Using latent space regression to analyze and leverage compositionality in GANs. *Github*. <https://chail.github.io/latent-composition/> (accessed Feb. 12, 2023)
- [218] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 90. ISBN 978-3-0356-2400-7.
- [219] Pix2Pix. *Machine Learning 4 Artists Github*. <https://ml4a.github.io/guides/Pix2Pix/> (accessed Jan. 6, 2023)
- [220] Witt, Andrew: *Formulations: Architecture, Mathematics, Culture*. The MIT Press, 2022. ISBN 978026543002
- [221] Zheng, H., Huang, W.. *Architectural Drawings Recognition and Generation through Machine Learning*. ACADIA. Cambridge, MA, 2018.
- [222] Wang, T. C., Liu, M. Y., Zhu, J. Y., Tao, A., Kautz, J., Catanzaro, B.: High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. *CVPR*. 2018. arXiv: 1711.11585v2. (accessed Jan. 12, 2024)
- [223] Nathan Peters. Master Thesis: *Enabling Alternative Architectures: Collaborative Frameworks for Participatory Design*. Harvard Graduate School of Design, Cambridge, MA. 2017.
- [224] Steinfeld, K.: *Berkeley College of Environmental Design*. https://ced.berkeley.edu/wpcontent/uploads/2022/02/pix2pix_grid_01.jpg (accessed Jan. 6, 2023)

- [225] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 90. ISBN 978-3-0356-2400-7.
- [226] Steinfeld, K.: Sketch2Pix. *CDRF*. http://blah.ksteinfe.com/200625/cdrf_2020.html (accessed Jan. 6, 2023)
- [227] Martinez, N.: *Suggestive Drawing Among Human and Artificial Intelligences*. Harvard Graduate School of Design, Cambridge, MA. 2016.
- [228] Leach, N.: *Architecture in the Age of Artificial Intelligence An introduction to AI for architects*. Bloomsbury Visual Arts, London, 2022. p. 106. ISBN 978-1-3501-6551-9.
- [229] De Cremer, D., Kasparov, G.: AI Should Augment Human Intelligence, Not Replace It. *Harvard Business Review*. 2021. <https://hbr.org/2021/03/ai-should-augment-human-intelligence-not-replace-it>. (accessed Jun. 17, 2023)
- [230] Applied R+D. *Foster+Partners*. <https://www.fosterandpartners.com/people/teams/applied-rplud> (accessed Nov. 11, 2022)
- [231] Quintana, M., Schiavon, S., Wai Tham, K., Miller, C.: Balancing thermal comfort datasets: We GAN, but should we? *Buildsys '20: Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. November 2020, pp. 120-129. doi: 10.1145/3408308.3427612
- [232] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 98. ISBN 978-3-0356-2400-7.
- [233] Interior AI designer. *ZMO.AI*. 2022. <https://www.zmo.ai/interior-ai/>. (accessed Jun. 15, 2023)
- [234] Romero, A.: Stable Diffusion Is the Most Important AI Art Model Ever. *The Algorithmic Bridge*. 2022. <https://www.thealgorithmicbridge.com/p/stable-diffusion-is-the-most-important>. (accessed Oct. 11, 2023)
- [235] Taylor, R., Kardas, M., Cucurull, G., Scialom, T., Hartshorn, A., Saravia, E., Poulton, A., Kerkez, V., Stojnic, R.: Galactica: A Large Language Model for Science. *Papers with Code*. Meta AI, 2022. <https://galactica.org/static/paper.pdf>. (accessed Jan. 12, 2023)
- [236] Chuan-En Lin, D.: 2022 Top Papers in AI — A Year of Generative Models. *Medium*. 2022. <https://chuanenlin.medium.com/2022-top-ai-papers-a-year-of-generative-models-a7dcd9109e39>. (accessed Mar. 23, 2023)
- [237] del Campo, M., Manninger, S.: Imaginary Plans. *Proceedings of the 2019 ACADIA Conference – Ubiquity and Autonomy*. https://www.academia.edu/40735096/Imaginary_Plans (accessed Nov. 11, 2022)
- [238] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 83-86. ISBN 978-3-0356-2400-7.
- [239] Tian, R.: Suggestive Site Planning with Conditional GAN and Urban Data. Yuan, P.F., Yao, J., Yan, C., Wang, X., Leach, N. (eds), *Proceedings of the 2020 DigitalFUTURES. CDRF 2020*. Springer, Singapore. doi: 10.1007/978-981-33-4400-610. https://link.springer.com/chapter/10.1007/978-981-33-4400-6_10 (accessed Nov. 11, 2022)
- [240] Kyle, R.: Spacemaker: Merging AI technology with Urban Planning and Design. *HBS Digital Initiative*. <https://d3.harvard.edu/platform-digit/submission/spacemaker-merging-ai-technology-with-urban-planning-and-design/> (accessed Nov. 11, 2022)
- [241] Spacemaker: Early-stage planning, reimagined. *Autodesk*. <https://www.autodesk.com/products/spacemaker/overview> (accessed Oct. 1, 2022)
- [242] Supercharge your building design capabilities. *Cove.tool*. <https://www.cove.tools/> (accessed Oct. 1, 2022)
- [243] The platform for urban analytics. *Giraffe*. <https://www.giraffe.build/> (accessed Oct. 2, 2022)
- [244] *Creo: Design. The way it should be. ptc*. <https://www.ptc.com/en/products/creo> (accessed Oct. 2, 2022)
- [245] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 103-109. ISBN 978-3-0356-2400-7.

- [246] Wray, S.: City digital twin adoption growing fast but challenges remain. *Cities today*. 2022. <https://cities-today.com/city-digital-twin-adoption-growing-fast-but-challenges-remain/>. (accessed Apr. 16, 2023)
- [247] Block-NeRF Scalable Large Scene Neural View Synthesis. Waymo. 2023. <https://waymo.com/research/block-nerf/>. (accessed Nov. 23, 2023)
- [248] Leach, N.: Design in the Age of Artificial Intelligence. *Landscape Architecture Frontiers*. Higher Education Press, Peking University. vol. 6, issue 2, April 2018. pp. 115-118. ISSN 2096-336X.
- [249] Christian B.: *The Alignment Problem: Machine Learning and Human Values*. 1st ed, W.W.Norton, 2020, pp. 17-80. ISBN 9780393635829.
- [250] About | Stanislas Chaillou. 2023. <https://stanislaschaillou.com/bio.html>. (accessed Mar 16. 2024)
- [251] Stanislas Chaillou is an architect, designer and AI-researcher. 2023. <https://stanislaschaillou.com/main.html>. (accessed Mar. 16, 2024)
- [252] Draw floor plans faster together. *rayon*. 2024. <https://www.rayon.design/>. (accessed Mar. 16, 2024)
- [253] Elevate your architectural workflow with Wearrecho's VR capabilities. Wearrecho. 2023. <https://wearrecho.space>. (accessed Apr. 5, 2024)
- [254] Barker, N.: ZHA developing "most" projects using AI-generated images says Patrik Schumacher. *dezeen*. 2023. <https://www.dezeen.com/2023/04/26/zaha-hadid-architects-patrik-schumacher-ai-dalle-midjourney/>. (accessed Oct. 7, 2023)
- [255] Wainwright, O.: 'It's already way beyond what humans can do': will AI wipe out architects? *The Guardian*. 2023. <https://www.theguardian.com/artanddesign/2023/aug/07/ai-architects-revolutionising-corbusier-architecture>. (accessed Oct. 18, 2023)
- [256] Barker, N.: "Architects may become a thing of the past" says ChatGPT. *dezeen*. 2023. https://www.dezeen.com/2023/02/13/chat-gpt-ai-architecture/?utm_medium=email&utm_campaign=Daily%20Dezeen. (accessed Oct. 16, 2023)
- [257] Leach, N.: AI is putting our jobs as architects unquestionably at risk. *dezeen*. 2023. <https://www.dezeen.com/2023/02/13/ai-architecture-jobs-risk-neil-leach-opinion/>. (accessed Oct. 16, 2023)
- [258] Next generation AI platform for architects & designers. *Lookx.AI*. 2022. <https://www.lookx.ai>. (accessed Jul. 23, 2023)
- [259] Krea Open Beta. *Krea*. 2023. <https://www.krea.ai/home>. (accessed Jul. 23, 2023)
- [260] *Midjourney*. <https://www.midjourney.com/home>. (accessed Apr. 5, 2024)
- [261] *Stable Diffusion Online*. https://stablediffusionweb.com/#google_vignette. (accessed Apr. 5, 2024)
- [262] *ControlNet: A Complete Guide*. <https://stable-diffusion-art.com/controlnet/>. (accessed Apr. 5, 2024)
- [263] Ikiz, S. U.: Carlos Banon said on AI, "We're not just following the old rules anymore; we're making new ones". *Parametric Architecture*. 2023. <https://parametric-architecture.com/carlos-banon-said-on-ai-were-not-just-following-the-old-rules-anymore-were-making-new-ones/>. (accessed Feb. 23, 2024)
- [264] Muaz, E.: Taking Control: Midjourney X ControlNet – Studio Carlos Banon, PA Next Team: Taking Control 2.0: Midjourney X ControlNet – Studio Carlos Banon, PA Next Team: Taking Control 3.0: Stable Diffusion XL X ControlNet – Studio Carlos Banon. *PA Academy*. 2023. <https://parametric-architecture.com/tag/studio-carlos-banon/>. (accessed Feb. 23, 2024)
- [265] What's in Creo9. *ptc*. <https://www.ptc.com/en/blogs/cad/whats-in-creo9> (accessed Oct. 4, 2022)
- [266] Cove.tool Review. *FiancesOnline*. <https://reviews.financesonline.com/p/cove-tool/> (accessed Oct. 4, 2022)
- [267] Harouk, Ch.: Spacemaker Proposes AI-Powered Generative Design to Create More Sustainable Spaces and Cities. *Archdaily*. <https://www.archdaily.com/952850/spacemaker-proposes-ai-powered-generative-design-to-create-more-sustainable-spaces-and-cities> (accessed Jan. 2, 2023)
- [268] Pavel Shaban. *MS architekti*. <https://www.wemakespaces.archi/cs/lide> (accessed Jan. 23, 2023)

- [269] Sourek, M.: From Functional Areas towards Metropolitans Structure: Public Space in Sustainable Development Context. *Advanced Engineering Forum*. Trans Tech Publications 2014. vol. 12, pp 176-180, ISBN 978-3-03835-323-2.
- [270] Autodesk Forma: Cloud-based software for early-stage planning and design. *AUTODESK*. 2023. <https://www.autodesk.com/products/forma/overview?term=1-YEAR&tab=subscription&plc=SPCMKR>. (accessed Mar. 23, 2024)
- [271] Buckle up, we're venturing into a new era of architecture! *Martin Posta*. *LinkedIn*. 2024. https://www.linkedin.com/posts/martinposta_buckle-up-were-venturing-into-a-new-era-activity-7172501651814105088-fxru?utm_source=share&utm_medium=member_android. (accessed Mar. 22, 2024)
- [272] Chaillou, S.: *Artificial Intelligence and Architecture From Research to Practice*. Birkhauser Verlag GmbH, Basel, Switzerland, 2022, pp. 80,86-92. ISBN 978-3-0356-2400-7.
- [273] Leach, N.: Design in the Age of Artificial Intelligence. *Landscape Architecture Frontiers*. Higher Education Press, Peking University. vol. 6, issue 2, April 2018. pp. 113-114. ISSN 2096-336X.
- [274] *ibid*, p. 114
- [275] XKool Technolgy. *Crunchbase*. <https://www.crunchbase.com/organization/xkool-technolo-gy> (accessed Dec. 4, 2022)
- [276] XKool. <https://www.xkool.ai> (accessed Dec. 4, 2022)
- [277] Leach, N.: Design in the Age of Artificial Intelligence. *Landscape Architecture Frontiers*. Higher Education Press, Peking University. vol. 6, issue 2, April 2018. p. 117. ISSN 2096-336X.
- [278] *ibid*, pp. 115-118
- [279] SCIA ENGINEER. *Scia*. <https://www.scia.net/en/scia-engineer> (accessed Feb. 7, 2023)
- [280] Salehi, H., Burgueño, R.: Emerging artificial intelligence methods in structural Engineering. *ScienceDirect*. <https://www.sciencedirect.com/science/article/abs/pii/S0141029617335526> (accessed Feb. 7, 2023)
- [281] Structural Engineering Software to Learn in 2023: Top 10 Best. *Tribby3d*. <https://tribby3d.com/blog/structural-engineering-software/> (accessed Jan. 30, 2023)
- [282] NEW EXPERIMENTAL TECHNOLOGIES. *MVRDV*. <https://www.mvrdv.nl/themes/15/next> (accessed Feb. 2, 2023)
- [283] van den Burgh, S.: MVRDV Next. *AECtech 2022*. <https://www.youtube.com/watch?v=ZlzaOcm5wBM> (accessed Feb. 3, 2023)
- [284] Greenwood, V.: Consciousness Began When the Gods Stopped Speaking. *Nautilus*. https://nautil.us/consciousness-began-when-the-gods-stopped-speaking-235445/?utm_source=nautilus-newsletter&utm_medium=email&he=1eee5366d2d1297277930d07b824c45b (accessed Mar. 29, 2023)
- [285] Heidegger, M.: *Sein und Zeit*. Max Niemeier Verlag, Tübingen, 2006. pp. 146-216. ISBN 3-484-70153-6. in English *Being and Time*. translated by Macquarrie, J., Robinson, E. 1st ed., Blackwell Publishers, 1962. ISBN 0-631-19770-2.
- [286] Dennett, D.: *From Bacteria to Bach and Back: The Evolution of Minds*. W. W. Norton, 2017. ISBN: 978-0393242072.
- [287] Rushdie, S.: *Languages of Truth: Essays 2003-2020*. Random House, 2021. p. 158. ISBN 978-0593318211.
- [288] McBride, S.: Musk Says First Neuralink Patient Received Implant in Brain. *Bloomberg*. 2024. <https://www.bloomberg.com/news/articles/2024-01-29/elon-musk-says-first-human-patient-has-received-brain-implant>. (accessed Feb. 2, 2024)
- [289] Marek, J.: Video prvního člověka, který má v mozku čip od Neuralinku. *Seznam Zprávy*. 2024. <https://www.seznamzpravy.cz/clanek/tech-technologie-video-prvniho-cloveka-ktery-ma-v-mozku-cip-od-neuralinku-248357>. (accessed Mar. 23, 2024)
- [290] Rushdie, S.: *Languages of Truth: Essays 2003-2020*. Random House, 2021. p. 180. ISBN 978-0593318211.

- [291] Xu, G., Mihaylova, T., Li, D., Tian, F., Farrehi, P. M., Parent, J. M., Mashour, G. A., Wang, M. M., Borjigin, J.: Surge of neurophysiological coupling and connectivity of gamma oscillations in the dying human brain. *Proc Natl Acad Sci USA*. DOI: 10.1073/pnas.2216268120. <https://pubmed.ncbi.nlm.nih.gov/37126719/>. (accessed Oct. 23, 2023)
- [292] Bai, Y., Wang, X., Cao, Y.-P., Ce, Y., Yuan, Ch., Shan, Y.: DreamDiffusion: Generating High-Quality Images from Brain EEG Signals. *arXiv*. Cornell University, 2023. arXiv:2306.16934v2; <https://doi.org/10.48550/arXiv.2306.16934>. (accessed Dec. 15, 2023)
- [293] Trafton, A.: From neurons to learning and memory. *MIT News*. 2024. <https://news.mit.edu/2024/mark-harnett-investigates-electrical-activity-0416>. (accessed Apr. 17, 2024)
- [294] Prakash, D.: Mastering Large Language Models in 2024: A Learning Path for Developers. *TURING*. 2024. <https://www.turing.com/blog/mastering-large-language-models-learning-path-for-developers/>. (accessed Mar. 8, 2024)
- [295] Li, K.: Do Large Language Models learn world models or just surface statistics? *The Gradient*. 2023. <https://thegradient.pub/othello/>. (accessed Mar. 10, 2024)
- [296] Verses. *Crunchbase*. 2024. <https://www.crunchbase.com/organization/verses>. (accessed Mar. 22, 2024)
- [297] Heath, R.: The road map to AI's next level could be nature. *AXIOS*. 2024. <https://www.axios.com/2024/03/13/verses-ai-artificial-general-intelligence-chatgpt>. (accessed Mar. 18, 2024)
- [298] Hawkins, J.: A Thousand Brains: A New Theory of Intelligence. Basic Books, 2021. ISBN 9781541675810.
- [299] Keary, T.: Decentralized Artificial Intelligence (DAI). *Technopedia*. 2023. <https://www.techopedia.com/definition/decentralized-ai-dai>. (accessed Mar. 21, 2024)
- [300] Distributed AI. *clan*. 2024. <https://clanx.ai/glossary/distributed-ai>. (accessed Mar. 21, 2024)
- [301] Stability AI fires 10% of its workforce. *AI Tool Report*. 2024. https://aitoolreport.beehiiv.com/p/meta-ai-comes-to-whatsapp-instagram-and-facebook?utm_source=aitoolreport.beehiiv.com&utm_medium=newsletter&utm_campaign=whatsapp-instagram-and-facebook-get-metaai. (accessed Apr. 19, 2024)
- [302] Heidegger, M.: *Voll Verdienst, doch dichterisch wohnt / Der Mensch auf dieser Erde*. Heidegger und Hölderlin, herausgegeben von Peter Trawny. Vittorio Klostermann 2000. ISBN 978-3-465-03084-3.
- [303] Gowans, A., Ackerman, J. S., Scruton, R., Collins, P.: Architecture. *Encyclopedia Britannica*. <https://www.britannica.com/topic/architecture> (accessed Jan. 15, 2023)
- [304] Sourek, M.: Architecture tech: A challenge for investors. *Academia Letters*, 2021, article 892. <https://doi.org/10.20935/AL892>. (accessed Mar. 20, 2024)
- [305] Vesely, D.: *Architecture in the Age of Divided Representation: The Question of Creativity in the Shadow of Production*. 1st ed. MIT Press, 2004. ISBN 0-262-22067-9.
- [306] Leach, N.: Design in the Age of Artificial Intelligence. *Landscape Architecture Frontiers*. Higher Education Press, Peking University. vol. 6, issue 2, April 2018. p. 118. ISSN 2096-336X.
- [307] McCulloch, W., Pitts, W.: A Logical Calculus of Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biophysics*. 5 (4), 1943, pp. 115–133. doi: 10.1007/BF02478259.
- [308] Turing, A. M.: On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*. 2 s. vol. 42 (1936–1937), pp. 230–265. online by Cambridge University Press, 2014. <https://www.cambridge.org/core/journals/journal-of-symbolic-logic/article/abs/a-m-turing-on-computable-numbers-with-an-application-to-the-entscheidungs-problcm-proceedings-of-the-london-mathematical-society-2-s-vol-42-19361937-pp-230265/4DFCA89035F7F7C5BF4DB5129B8BB09E>. (accessed Apr. 22, 2023)
- [309] Abraham, T. H.: (Physio)logical Circuits: The Intellectual Origins of the McCulloch-Pitts Neural Networks. *Journal of the History of the Behavioral Sciences*. 38(1), 3–25. 2002. DOI: 10.1002/jhbs.1094. accessible online <https://psycnet.apa.org/record/2002-04758-001>. (accessed Nov. 15, 2023)

- [310] Abraham, T. H.: Integrating Mind and Brain: Warren S. McCulloch, Cerebral Localization, and Experimental Epistemology. *Endeavour* 27(1), 32–38. 2003. DOI: 10.1002/jhbs.1094. accessible online https://www.academia.edu/4051744/Integrating_mind_and_brain_Warren_S_McCulloch_cerebral_localization_and_experimental_epistemology. (accessed Nov. 17, 2023)
- [311] Recurrent neural network. *WIKIPEDIA*. https://en.wikipedia.org/wiki/Recurrent_neural_network (accessed Nov. 7, 2022)
- [312] Mountcastle, V.: Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology*. 20 (4): 408–34. 1957. doi:10.1152/jn.1957.20.4.408. PMID 13439410. accessible online <https://europepmc.org/article/MED/13439410>. (accessed Dec. 9, 2023)
- [313] Haueis, P.: The life of the cortical column: opening the domain of functional architecture of the cortex (1955–1981). *HPLS* 38, 2. 2016. <https://doi.org/10.1007/s40656-016-0103-4>. accessible online <https://link.springer.com/content/pdf/10.1007/s40656-016-0103-4.pdf>. (accessed Nov. 20, 2023)
- [314] Bertics, A.: AI models will become smaller and faster. *The Economist*. 2023. https://www.economist.com/the-world-ahead/2023/11/13/ai-models-will-become-smaller-and-faster?utm_content=conversion.direct-response.non-subscriber.article_bottom&utm_campaign=a.24worldahead_content_v1_registered.2024-01-07&utm_medium=email.owned.np&utm_source=salesforce-marketing-cloud&utm_term=1/7/2024&utm_id=1841573&sfmc_id=0033z000032ypDXAAY&utm_content=343495&utm_id=5c9b35b7-9eb1-4aaf-aaaf-cacabb0c75c2&sfmc_activityid=60093c56-939e-4455-81fd-ab25ccc697f7&utm_medium=email.owned.np. (accessed Feb. 2, 2024)
- [315] Monserrate, S. G.: The Staggering Ecological Impacts of Computation and the Cloud. *The MIT PRESS Reader*. 2024. https://thereader.mitpress.mit.edu/the-staggering-ecological-impacts-of-computation-and-the-cloud/?utm_source=90404-Pod%20čarou&utm_medium=email&utm_term=17599780728&utm_content=environmentální%20dopady%20datových%20center&utm_campaign=SZ%20-%20Pod%20čarou--20240309. (accessed Mar. 18, 2024)
- [316] Nichols, H.: Energy secretary's AI power plan. *Axios AI+*. 2024. https://www.axios.com/newsletters/axios-ai-plus-a131709d-30f3-4045-87b3-2796316a8990.html?utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axioslogin&stream=top. (accessed Apr. 1, 2024)
- [317] OpenAI team: GPT-4 Technical Report. *arxiv*. Cornell University, 2024. [arXiv:2303.08774v6; https://doi.org/10.48550/arXiv.2303.08774](https://doi.org/10.48550/arXiv.2303.08774). (accessed Mar. 22, 2024)
- [318] Nvidia Blackwell Platform Arrives to Power a New Era of Computing. *Nvidia Newsroom*. 2024. <https://Nvidianews.Nvidia.com/news/Nvidia-blackwell-platform-arrives-to-power-a-new-era-of-computing>. (accessed Mar. 26, 2024)
- [319] Hruska, J.: Samsung-Backed Researchers Debut Ternary Semiconductor Design. *EXTREMETECH*. 2019. <https://www.extremetech.com/computing/295424-back-off-binary-samsung-backed-researchers-debut-ternary-semiconductor>. (accessed Oct. 21, 2023)
- [320] Classical optical neural network exhibits 'quantum speedup'. *PhysOrg*. Technical University of Denmark, 2024. <https://phys.org/news/2024-04-classical-optical-neural-network-quantum.html>. (accessed Apr. 5, 2024)
- [321] 100 kilometers of quantum-encrypted transfer. *PhysOrg*. Technical University of Denmark, 2024. <https://phys.org/news/2024-04-kilometers-quantum-encrypted.html>. (accessed Apr. 5, 2024)
- [322] D-Wave Reports 4th Quarter and Year-End 2023 Results. *HPC Wire*. 2024. <https://www.hpcwire.com/off-the-wire/d-wave-reports-4th-quarter-and-year-end-2023-results/>. (accessed Apr. 5, 2024)
- [323] Quantum Computing Inc. Releases Full Year 2023 Financial Report. *Quantum Computing Report*. 2024. <https://quantumcomputingreport.com>. (accessed Apr. 5, 2024)
- [324] Swayne, M.: Zapata AI Stock Faces Steep Decline On First Day Of Trading, Could Chill Future SPAC Deals In Quantum. *The Quantum Insider*. 2024. <https://thequantuminsider.com/2024/04/02/zapata-ai-stock-faces-steep-decline-on-first-day-of-trading-could-chill-future-spac-deals-in-quantum/>. (accessed Apr. 5, 2024)

- [325] Swayne, M.: Chicago Won't Be Second City In Quantum: Illinois Governor Unveils Strategic Plan For Quantum. 2024. <https://thequantuminsider.com/2024/04/02/chicago-wont-be-second-city-in-quantum-illinois-governor-unveils-strategic-plan-for-quantum/>. (accessed Apr. 5, 2024)
- [326] Imagine a Smarter World. *VERSES*. 2024. <https://www.verses.ai/>. (accessed Mar. 18, 2024)
- [327] VIVE Team: MWC 2024 Unveiled: HTC VIVE's AI and Spatial Computing Innovations. *VIVE*. 2024. <https://blog.vive.com/us/mwc-2024-unveiled-htc-vives-ai-and-spatial-computing-innovations/>. (accessed Mar. 2, 2024)
- [328] Yenduri, G., Ramalingam, M., Kumar, P., Maddikunta, R., Gadekallu, T. R., Jhaveri, R. H., Bandi, A., Chen, J., Wang, W., Shirawalmath, A. A., Ravishankar, R., Wang, W.: Spatial Computing: Concept, Applications, Challenges and Future Directions. *arXiv*. Cornell University, 2024. <https://arxiv.org/html/2402.07912v1>. (accessed Mar. 2, 2024)
- [329] Chaillou, S.: Tooling the Ordinary. A Manifesto for Rayon. *Medium* | *Rayon*. <https://medium.com/rayonapp/tooling-the-ordinary-42b91b837876>. (accessed Dec. 12, 2023)
- [330] Let's build from here. *Github*. <https://github.com> (accessed Mar. 24, 2023)
- [331] Software. Faster. *GitLab*. <https://gitlab.com/> (accessed Mar. 24, 2023)
- [332] What is PatternForge? *PatternForge*. <https://patternforge.org> (accessed Mar. 24, 2023)
- [333] The United States spends more on defense than the next ten countries combined. *Peter G. Peterson Foundation blog*. 2023. <https://www.pgpf.org/blog/2023/04/the-united-states-spends-more-on-defense-than-the-next-10-countries-combined>. (accessed Mar. 16, 2024)
- [334] Statista Research Dept.: U.S. share of value added to GDP 2023, by industry. *Statista*. 2024. <https://www.statista.com/statistics/248004/percentage-added-to-the-us-gdp-by-industry/>. (accessed Apr. 5, 2024)
- [335] Kolmar, Ch.: 25 Essential US Construction Industry Statistics [2023]: Data, Trends And More. *Zipppia*. 2023. <https://www.zipppia.com/advice/us-construction-industry-statistics/>. (accessed Apr. 5, 2024)
- [336] Global BIM Adoption: A 2024 Outlook. *Pinnacle Infotech*. 2024. <https://pinnacleinfotech.com/global-bim-adoption/>. (accessed Feb. 14, 2024)
- [337] Building Information Modeling Market Statistics: 2032. *Allied Market Research*. 2024. <https://www.alliedmarketresearch.com/building-information-modeling-market>. (accessed Feb. 14, 2024)
- [338] Perez, C.: What Is AI's Place in History? *Project Syndicate*. 2024. <https://www.project-syndicate.org/magazine/ai-is-part-of-larger-technological-revolution-by-carlota-perez-1-2024-03?barrier=accesspaylog>. (accessed Mar. 15, 2024)
- [339] Sora: first impressions. OpenAI. 2024. <https://openai.com/index/sora-first-impressions?ref=themarktechweekly.com>. (accessed Mar. 28, 2024)
- [340] Crowley, M.: Humanoid robot learns tasks in 24hrs. *AI tool report*. 2024. https://www.aitoolreport.com/articles/humanoid-robot-learns-tasks-in-24hrs?utm_source=aitoolreport.beehiiv.com&utm_medium=newsletter&utm_campaign=musk-to-raise-6b-from-his-buddies-for-xai. (accessed Apr. 26, 2024)
- [341] Sanctuary AI Unveils the Next Generation of AI Robotics. *Sanctuary AI*. 2024. <https://sanctuary.ai/resources/news/sanctuary-ai-unveils-the-next-generation-of-ai-robotics/>. (accessed Apr. 26, 2024)
- [342] Sanctuary AI Phoenix Humanoid Robot Used for Sorting Mechanical Parts. *Robotic Gizmos*. 2023. <https://www.roboticgizmos.com/sanctuary-ai-phoenix/>. (accessed Apr. 26, 2024)
- [343] Levy, I., Schiller, D. Uncertainty in learning and decision-making: Introduction to the special issue. *Cogn Affect Behav Neurosci*. 23, 459–464 (2023). <https://doi.org/10.3758/s13415-023-01111-6>. (accessed Apr. 20, 2024)
- [344] Lamnina, M., Chase, C.C. Uncertain instruction: effects on curiosity, learning, and transfer. *Instr Sci*. 49, 661–685 (2021). <https://doi.org/10.1007/s11251-021-09557-2>. (accessed Apr. 19, 2024)

[345] Sanctuary AI Unveils Phoenix™ - A Humanoid General-Purpose Robot Designed for Work. *Sanctuary AI*. 2023. <https://sanctuary.ai/resources/news/sanctuary-ai-unveils-phoenix-a-humanoid-general-purpose-robot-designed-for-work/>. (accessed Apr. 26, 2024)

[346] Josef Sivic získal prestižní ERC Advanced Grant na výzkum inteligentních počítačových systémů schopných učit se v dynamicky měnícím se světě. *Czech Institute of Informatics, Robotics, and Cybernetics*. <https://www.ciirc.cvut.cz/cs/josef-sivic-awarded-a-prestigious-erc-advanced-grant/> (accessed Apr. 5, 2023)

[347] Automated Reasoning Group: ERC Project AI4REASON. 2020. <http://ai4reason.org/>. (accessed Feb. 20, 2024)