

UNLEASHING THE POTENTIAL OF ARTIFICIAL INTELLIGENCE IN ARCHITECTURE AND DEVELOPMENT OF THE BUILT ENVIRONMENT: FROM MISCONCEPTIONS TO PRODUCTIVE PROSPECTS

Michal Sourek

*Institute of Architecture
Faculty of the Built Environment
Czech Technical University in Prague
Czech Republic*

Abstract:

Artificial intelligence invades our lives and professions at an ever-increasing pace and intensity. Architecture, the built environment, and real estate have been joining the trend only timidly and belatedly. The record of some of the most recent "famous achievements" in the field set straight, the paper challenges the state-of-the-art concerning these fields, debunks the idea of the (truly) creative potential of the technology, and puts forward a sketch roadmap to a realistic - and significant - deployment of artificial intelligence in architecture and the built environment. The attention turns to open-source patterns-platforms, generative patterns processing, generative pre-design, parametric evaluation, and optimization. Leveraging the objectivity of assessments and streamlining workflows, artificial intelligence promises to unleash true architectural creativity and leverage the productivity and efficiency of the design- and planning processes. Finally, a chance for architecture to come back from the sidelines to the position needed to provide society with what it lacks in terms of quality of life, social- and cultural values, sustainable development, and comprehensive resilience as well as to heal the profession renders. Among other new technologies, artificial intelligence can play an outstanding role in this regard if understood, grasped, and developed adequately by architects, real estate developers and agents, and IT developers hand closely in hand.

Key words: artificial intelligence, machine learning, generative pattern, open-source platform of patterns, generative pre-design, parametric review, parametric optimization

Introduction

Artificial intelligence, as algorithmic computer programs are called that can generate data from heterogeneous, most often textual and image templates or patterns, evaluate these large data sets according to complex criterion structures, generate downstream new algorithms, and use them to generate outputs of a nature identical or subsequent to the initial templates and patterns, is considered "the next level" in the field of visual design - and architecture hopefully, too. Only belated, however, architecture, the development of the built environment, and real estate have been joining the trend of artificial intelligence that has been entering our lives and professions since the 1980s.

Reimagining Architecture and Development of the Built Environment

Not by chance, the wavering approach to the new technology tunes with these disciplines' failures to cope with societal and economic development; in architecture, it has been a reality for at least 70 years. At the same time, the influence and impact of the disciplines' general performance on the economy, social affairs and issues, and sustainability are immense. The society's dissatisfaction with the state and development of the built environment has been increasing through recent decades, but a comprehension of the causes is lacking. The remedy efforts are marginal and shattered, failing to address the appropriate goals and adopt efficient way and tools.

The need for the field's paradigm change renders: to support the change, new, promising technologies appear, such as virtual reality and artificial intelligence (AI), that could improve architectural and planning practices. Virtual twins offer unprecedented abilities to create, understand, and communicate architecture and the built environment, while AI provides machine-learning capabilities for design, planning, and parameters' review and assessment. The nature of the

technologies and their potential have to be understood correctly first - which, so far, is only rarely the case in the professions moving around built environment and real estate with the architectural craft at their heart.

Achievements, Misonceptions, and Prospects of AI in Architecture

The virtual twins' technology, its benefits and prospects are reviewed in [1]; this paper sets it aside and concentrates on AI. First, the paper challenges the existing approach to AI's deployment within architectural design, its theoretical starting points, and the perspectives put forward so far. To do so, the framework of a state-of-the-art of the field is introduced; the methods used and expectations declared are assessed within the framework. Section (2) provides a general summary of the achieved results, which will be extensively discussed in the subsequent section (3). The overview serves as a foundation for identifying the misunderstandings and limitations of recent attempts and expectations regarding the use of AI in architecture, the built environment, and real estate - for debunking some of the "fantastic achievements of AI promising to make man redundant" temporarily discussed in the profession and beyond. In contrast, other, so far not considered perspectives of AI in architecture and development of the built environment render: in parametric development and assessment of diverse outputs of creative architectural and urban design concepts, and in approaching to sustainable development (or, better, to the comprehensive resilience) tasks. Finally, Discussion (3) sketches how AI could contribute to an upheaval of the architectural profession by overcoming its recent and contemporary technological lacking behind; the true, authentic architectural creativity will not back off by the technology – the opposite, it will be unleashed. Conclusions (4) outline the principles of aiming and particular goals of further development of machine learning for the good of architecture, the built environment, and sustainable development.

(1) State-of-the-Art, Methods, and Expectations

Since around 2010, global star-architectural studios alongside young enthusiasts combining information technology and architecture try to embrace AI's potential contribution to architectural design or, better to say, disclose wherefrom it might stem and what it might consist of. In 2020, *DeepHimmelb(l)au - a video of a journey through an imaginary landscape of Coop Himmelb(l)au-like building forms* - has come 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 GAN technologies provided "*machine hallucinations*" [2,3,4] - represented prevalingly in two dimensions, substantially lacking both spatial comprehensivity and the for architecture inherent interconnectedness of the experiential (poetic, in other words) and material attributes that will be discussed further in (3).

How a Machine Can Learn

Though generally (and in this paper, too) labeled as AI, the term machine learning adheres better to the use of the technology in architecture: learning or training is the keyword. 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 will be shown in (3), and a comprehensive one; what is not in it, the AI cannot learn. Variations of machine learning deserve reminder: supervised learning, unsupervised learning, reinforcement learning, and various fusions.

In supervised learning, the system is given a series of categorized or labeled examples and told to make predicions about new examples it hasn't seen yet, or for which the ground truth is not yet known [5]. Supervised learning uses labeled datasets, whereas unsupervised learning uses unlabeled datasets. "Labeled", means the data already tagged with the requested answer. In supervised learning, the learning algorithm measures its accuracy through the loss function, adjusting until the error has been sufficiently minimized. Two types of supervised learning distinct - classification and regression. Classification uses an algorithm to accurately assign test data into specific categories. It recognizes specific entities within the dataset and attempts to draw some conclusions on how the entities should be labeled or defined. Common classification algorithms are support vector machines, linear classifiers, decision trees, k-nearest neighbor, random forest, and others. Regression applies to understand the

relationship between dependent and independent variables - commonly to make projections. Linear regression, logistic regression, and polynomial regression are popular regression algorithms. [6]

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, usefull ways of condensing or representing or visualizing it.* [5,7]

Reinforcement learning concerns how intelligent agents ought to take action in an environment to maximize the notion of cumulative reward. 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* [5,8]. 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 from human feedback* [9]) 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.

Artificial Neural Networks

Approaching the mentioned GAN (generative adversarial network) technologies (approaching only superficially in this paper's framework), artificial neural networks can be described as a type of machine learning model that can be used for various tasks, including deep learning, Bayesian learning, and more. An artificial neural network is a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal to other neurons. An artificial neuron that receives a signal then processes it and can signal neurons connected to it. [10,11]

Designed by Ian Goodfellow and his colleagues in 2014, GAN is a class of machine learning frameworks. Representing today's state-of-the-art, GAN is a milestone of R&D launched by neurophysiologist and cybernetician of the University of Illinois at Chicago Warren McCulloch and self-taught logician and cognitive psychologist Walter Pitts. In 1943, the two published the foundations-laying article *A Logical Calculus of the Ideas Immitent in Nervous Activity* [12]. Building on Allan Turing's work *On Computable Numbers*, McCulloch's and Pitt's 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 first implementation of McCulloch's and Pitt's theoretical starting points was a machine built in 1958 at the Cornell Aeronautical Laboratory by Frank Rosenblatt. *The perceptron was intended to be a machine, rather than a program, and while its first implementation was in software for the IBM 704, it was subsequently implemented in custom-built hardware as the "Mark 1 perceptron". This machine was designed for image recognition: it had an array of 400 photocells, randomly connected to the "neurons". Weights were encoded in potentiometers, and weight updates during learning were performed by electric motors.* [13] By the end of the 1950s, IBM's Arthur Samuel had built a program to play checkers that, in a crude and early way, adjusted its parameters based on won and lost games. *Soon after Samuel was losing matches to his own creation: behold, reinforced machine learning.* [14]

Infering by Computing

Building upon "the founding fathers'" achievements, the idea of GAN copes with evolutionary biology principle of an arms race between two species. Two neural networks contest with each other in the form of a zero-sum game, where one agent's gain is another agent's loss. The core principle of a GAN is an "indirect" training through the discriminator – competitive network agent that can tell how "realistic" the input seems, which itself is also being updated dynamically. *This means that 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. [15]

An artificial neural network works by computing. In essence, two principles of computing apply in artificial neural networks: feedforward computing and backpropagation. The goal is always to *train* the models generated to cope with the criteria inserted typically by vast collections of 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 contrast, while using backpropagation, the model parameters update in the opposite direction: from the (one closer to) output layer to the (one closer to) input one. However, the backpropagation algorithm shall not be confused with training algorithms that provide the model dataset updates; backpropagation is a strategy to compute the gradient in a neural network. Backpropagation is a general technique; in terms of neural networks, it is not restricted to feedforward networks, it works for recurrent neural networks as well. [16]

"Fed" by inputs, the networks deliver outputs "at the end" that try to mimic the deliverables of human work. The principle is that artificial networks deliver relentlessly, very quickly, and in huge quantities - as opposed and by contrast to humans. The vision is that amidst these quantities in no time at all emerge outputs that not only mimic but also attain, if not surpass the quality of human performance. The vision still leaves something to a human - the choice of the most suitable output provided and its fine-tuning, but who knows - one day... Objectively, an evaluation of solutions due to a given set of criteria is a task suitable for computer, too - a task easier in principle than a creation.

Networks' and Techniques' Evolution

GANs, a recent revolution in machine learning provides results today that achieve appreciation, as Leach [17] puts. 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. [18] ImageNet, a groundbreaking project from the 2010s builds on this technology. [19] Graph neural networks (GNNs) are another field of recent research aiming at the processing of graph data. And 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. [20]

Variational auto-encoders (VAEs) develop another technique. [21] Unlike GAN, instead of the *generator – discriminator* pair, Variational Autoencoder combines two distinct approaches - *encoding* and *decoding*. 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 modeled phenomenon. [22] The ability to emulate a phenomenon by generating multiple versions of it is a starting point of VAE's generative potential to provide large quantities of "outputs" (as AI enthusiasts heralding the twilight of design call it) - typically in furniture design, fashion, photography, architecture, and urban design. [23]

Oposed to the above outlined *discriminative* and/or *decoding* techniques that *identify* objects and *infer* what is “real” and what is “fake”, generative AI systems *create* objects such as pictures, audio, writing samples, and anything that can be built by computer-controlled systems like 3D printers. [24] Generative AI allows machines to create new works based on what they have learned from others. With such a straightforward deployment, a question arises of how much the (so far existing) generative AI systems are truly AI-driven in terms of computational networks and processes; however, in the practical framework of this paper, the resolution is not of substantial meaning. As a principle, generative and discriminative or decoding systems most often operate paired in GAN models setting the benchmark 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 problems in the past. [25]

General Context to Compare

There are ecosystems of natural language processing, image processing, voice processing, and code or software processing and development, further robotics, and expert systems or business intelligence [26,27], altogether represented by Dall-E (Dall-E2 most newly), ImageGPT, GPT-3.5 and the latest GPT-4, InstructGPT and ChatGPT and other tools by OpenAI, also Midjourney, Stable Diffusion, 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.

After the first, largely experimental AI models - the perceptron and Samuel's checkers player - slowly-slowly, practical or even commercial AI applications have been arriving: ELIZA in 1966, developed by Joseph Weizenbaum for MIT, was the first chatbot in history that replicated a therapist giving general answers to users' questions, simulating a real conversation; PARRY, developed in 1972 by Kenneth Colby at Stanford University to simulate a schizophrenic patient and replicate therapeutic conversations, terminated in 1978; RACTER developed by William Chamberlain and Thomas Etter - a chatbot capable of generating poems and science fiction stories, published by Mindscape in 1984 and active until 1991; A.L.I.C.E. - Artificial Linguistic Internet Computer Entity developed in 1995 by Richard Wallace - a chatbot that could mimic a conversation with a human, active until 2018; SmarterChild, a virtual assistant on AIM, MSN Messenger and Yahoo! Messenger developed by ActiveBuddy in 2001 operating until 2008 when its activities were terminated; Doretta, released in 2007 helped to search the Internet for what users of Windows Live Messenger do; Siri, one of the first successful virtual assistants, developed in 2010 by Apple for iOS devices is still active and evolving; Alexa developed in 2014 by Amazon is also a virtual assistant that can be used through the Echo device, still active and evolving just like Google Assistant or BERT, both developed by Google in 2016 and 2018 respectively; BERT is 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. [28]

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 being indicated only superficially, [29] numerous "more fresh" straightforward concepts of generative AI beyond deserve more attention. An excellent example is the creation of new medical images, such as those *used in retinopathy diagnosis. Using it, physicians can create new patient records, which can then be incorporated into the system to improve accuracy. To train these applications, they use large amounts of real-world patient data to generate new images and data sets that humans could have never developed.* Analyzing large amounts of data efficiently and quickly improves understanding of diseases significantly. *DeepComposer is AI intelligence that can create an entire song from a short melody. It is also designed to act as a personal assistant to a human to compose some simple segments first, the final composition to be assembled by AI. A wide range of common music genres is pre-trained in the system.* Some of IBM's generative models are used in drug design to *propose new molecules that could work as drugs by training them to generate structures in relation to expected functions.* Game developers - Nintendo, Rockstar Games, Valve, Activision, Electronic Arts, and Ubisoft among others - are 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.* 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, 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. And more than one startup has already begun applying generative AI to create virtual assistants who can respond appropriately to human requests with natural language processing and dialogue management capabilities.* [30]

Not only professionals but the general public, too, has started to employ AI: from initial misconception to more mature approaches and understanding. 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.

It's a real influx of ever-new AI tools what the present experiences. During the fourth week of March 2023 alone, 200 new applications were released. [31] The week before, OpenAI dropped 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 - the *"creative and helpful collaborator, here to supercharge your imagination, boost your productivity, and bring your ideas to life"* [32], Chinese search engine behemoth Baidu released (just before April 14, 2023) Ernie Bot [33], Alibaba Group introduced (April 11) chatbot Tongyi Qianwen, Sense Time SenseChat [34], and Amazon its new AI application Bedrock, which makes available to developers generative tools for creating texts and images. (Now, Amazon offers four AI applications, the native Titan included, and employs the most people in AI development – more than both Google and Microsoft.) [35]

Transversally to the closed corporate releases on large language models, there have also been clever ideas and breakthroughs in the field of natural language processing. A new training strategy of meet-in-the-middle (MIM) [36] has been shown to improve not only the performance but also the interpretability and thus security (to be explained later) of large language models. Significant progress arrived in computer vision, in both diffusion models and neural radiance fields (NeRF) [37] - a type of machine learning algorithm used for 3D modeling and rendering based on deep neural networks capable of generating high-quality, photorealistic images of complex scenes from multiple viewpoints. The new MeshDiffusion [38] - another fresh arrival - allows direct generating 3D meshes without any post-processing, and also new FateZero [39] can edit the style of the videos using text while keeping the pre-trained model weights intact. Last but not least, the same week a beautiful marriage of NeRF with CLIP (contrastive language-image pre-training) [40,41] arrived: LERF (radiance embedded radiance fields) [42]. With it, natural language queries in a 3D fashion can apply within NeRF, targeting different objects in the scene. This brief overview highlights several new types of algorithms that may represent the first steps toward productive prospects for deploying AI in architecture. These algorithms, developed (most likely) with no regard to architectural design, may break out of the misconceptions that have so far dominated these efforts, as section (2) of this paper will show, while section (3) discusses the emerging positive prospects of AI's contribution to the architectural craft.

The Black Box Problem, Security Issue, and a Threat to Humanity

However, *with all these nice results, it's not clear what these models are learning*, as Mathew Zeiler puts it. [43] As AI becomes more advanced and ubiquitous, concerns around the security of AI systems have become increasingly prevalent. AI systems are vulnerable to various types of attacks and data "poisoning", which involve intentionally manipulating the input data to cause it to produce incorrect or malicious output, injecting malicious or misleading data into the training data set, which can result in biased or inaccurate models, or inserting a hidden trigger or behavior into an AI system that can be triggered by a specific input or action, allowing an attacker to gain unauthorized access or control over the system.

Another significant challenge in the development and use of AI systems is the black box problem. AI systems often involve complex algorithms and models that are more or less impossible for humans to interpret or understand. This can create a lack of transparency and accountability in AI decision-making, as it can be unclear how the system arrived at a particular output or decision, or even obviously incorrect or misleading outputs are delivered occasionally submitting unintended consequences or biases, also in ethical and legal concerns. 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. In cases - not exceptional - when subsequent analysis reveals systematic or occasional inaccuracy of the outputs,

confusion of cause and effect in the training data, 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*” [2,3,4]. 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 to manipulate and misinterpret visual information, either 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. Different from a point cloud or voxel-based diffusion models, the “new arrival” of March 2023 MeshDiffusion [38] can enjoy the modern graphics pipelines optimized for operating with meshes, e.g. relighting or simulation. Other than the unconditional generation of meshes, the model is also capable of conditional generation with single-view images and shape interpolation. For efficacy, the uniformly initialized tetrahedral grid for 4D representation is used. Training is split into two stages: First, a mapper from 2D view to the tetrahedral grid is trained to construct a dataset of grids, and second, a diffusion model (denoiser) that operates on the tetrahedral representation.

And other kinds of applications began to be developed to tackle the issue of interpretability. The already mentioned new training MIM strategy represents the first approach heading to AI systems that are inherently transparent and accountable due to the increased interpretability of the model. Interpretability represents the (human) ability to understand and explain how a particular system, model, or algorithm arrives at a particular decision or output. This way, it can help to build trust in AI systems by making their decision-making processes more transparent and understandable. It can also help in identifying and addressing biases or errors in the system, and it can enable experts to diagnose and fix any issues that arise. Interpretability can be achieved through various methods, such as visualizations, explanations, or feature importance measures. Nonetheless, achieving interpretability can sometimes come at the cost of the accuracy or performance of the system. [43]

Another way to tackle the black box issue is testing with concept activation networks; TCAN is a technique used to better understand how neural networks process information and make decisions. A CAN is a type of neural network that has been trained to recognize and represent concepts, such as objects or ideas, within an image or text. CANs are designed to mimic the way the human brain processes information, by breaking down complex inputs into simpler components and recognizing patterns within them. To test with CAN, an image or text is to be fed into the network to observe the activation of the individual concept nodes. This allows seeing which concepts the network has recognized and how confident it is in its recognition. By analyzing the activation patterns of the CAN, an insight into how the network is making decisions can be gained to identify areas for improvement subsequently. The technique 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. [44,45]

Nonetheless, fears remain. At the end of March 2023, Italy blocked ChatGPT [46] to secure the privacy of people and tycoons of global business claim pausing “giant AI experiments” - read the development of AI for six months [47] 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.

(2) Results of Applying AI in Architecture

For a paper headlined architectural, the preceding section may seem too much extensive both in terms of general orientation and when it comes to scope. However, only the technological “know-how” along with an awareness of the vast achievements of the other fields allow comprehension of how sidelined (also concerning AI) architecture is and, especially, how great the opportunities for the development of the branch are. A misconception on the nature of architecture in contrast to construction and built environment that the AI applications developers have been suffering so far is one of the reasons for the sidelining. Another reason is the multiple dimensionality and diachrony of architecture that contrast with the nature of the fields of the biggest successes of AI deployment – one-dimensional

language or two-dimensional image. Compared to the double challenge – a lack of understanding and an extreme comprehensiveness of the task – the recent and contemporary applications' development efforts show even more daring – no matter how unsuccessful a class of them is. Or is it naivety?

AI Models and Creates Architecture – Does It?

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.

[48] 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. [49]

Stanislas Chaillou as well as Nvidia company and many others provide AI applications to generate floorplans and apartment layouts. [50] 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 [51] provides series of machine-generated facades following the "style" and character of the pattern deployed as the "input". [52,53] Introduced by the same team, Pix2Pix is shorthand for an implementation of a generic image-to-image translation using CANs [54]. Developed in 2019 by Kyle Steinfeld [55], GAN Loci is able 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. [56,57] 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 [58].

Tom Mayne of Morphosis employed AI to develop *operational strategies, so as 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* [59]. 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* [60].

In terms of practical use, *predictive simulations* render the etalon. ComfortGAN, for example, investigates the challenges of predicting a building's indoor thermal comfort [61]. Also structural design is 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* [62]. 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.

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. [63,64] *Over the past decade, the deployment of online platforms has provided an adequate infrastructure to the end users*, [65] also to deploy Generative AI: Spacemaker [66,67], Cove.tool [68], Giraffe [69], or Creo [70] are a few examples of this growing ecosystem, offering simplified access to AI-based predictive models [71], *generative design, augmenting reality, real-time simulation, additive manufacturing, and IoT to iterate faster, reduce costs, and improve product*.

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 AI-enabled and applying generative design today [72], not to mention the acquisition of Spacemaker [73]. Nonetheless, as broad as all this listing may seem, the development of AI for and in AEC (Architecture, Engineering, Construction) is still in its infancy, failing to catch up with LLMs (large language models), 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 [74].

Assessment of Results Achieved in Architectural Design

The overview of the results of applying AI in architecture achieved so far renders shattered: it is neither by accident nor by a lack of caring by the author. When AI performs well concerning parametric aspects of diverse materializations of architecture (such as construction, energy efficiency, daylighting, or noise in buildings and neighborhoods) and fails to be effective and productive in conceptual architectural design, it is not only 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, whether on the scale of buildings or the built environment. Only by overcoming this problem, the way will open up to the efficient and prolific deployment of AI in architecture as will be discussed later and put in (4). At the moment, however, there is a long way to go.

Unlike the text-, image-, or codeprocessing examples provided, the applications aiming at architecture and built environment are far from establishing a homogenous, rich, and prosperous ecosystem. The results delivered by AI in the field of architecture and real estate development are still profoundly experimental, and failing to provide a degree a "creativity" and "persuasiveness" that we have got used to by text- or image processing AI. Predictive simulation tools such as Spacemaker or Cove.tool show up as, or rather, pretend successfully to be exceptions, as will be discussed later.

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, code development, 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* as well as alike results of ZHA 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. All this, is it just a situation of a developing field that needs more time and effort to mature and deliver useful results? Discussion (3) will confront such a perspective with the option that it is a dead-end of the state-of-the-art AI in architecture.

On the other hand, the values of deliverables provided by the AI of Spacemaker, Cove.tool, or Creo appear ambitious - but not without caveats that will be discussed later. 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 by means of generative design [75]. Similarly, Cove.tool delivers performance data of the building solutions in real-time employing the power of AI. [76] Cove.tool is 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, cost, or performance. Nevertheless, getting 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 softwares able to deliver.

Also famous as the two hundred and forty million acquisition of the AEC-software tycoon-software-producer Autodesk, Spacemaker not only *gives 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.* [77]

However, the practical deployment of Spacemaker raises doubts: the workflow is the issue. The user enters the address of the location and the boundaries of the territory; with a help of an open database like OSM, the terrain is generated, and it is also possible to generate existing buildings and structures; the accuracy of the objects generated depends on the data available the quality of which varies from territory to territory, nevertheless, so far so good. Then the user defines the area to be solved, and he can add roads - only manually, an import from a CAD is not available. Buildings can be

placed either manually by inserting individual floorplans as objects that can't be subject to later adjustments or the buildings can 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 user can assess the generated options based on gross-and/or netto-floor area totals. The user can further modify the chosen option by some of the spatial transformations: shift, rotation, ... Spacemaker evaluates the finally proposed solution in terms of noise, wind, sunlight, daylight, and microclimate. Exports of the valuatofactual contribuns to Excell and of the model designed to Autodesk Revit or to .ifc format are available. All in all, the evaluations of the designed locality's microclimate parameters and the imports of entry parameters are valuable and efficient functionalities. What happens in between - how the design comes to existence, Pavel Shaban of MS architekti [78], Prague/Czech Republic based architectural studio claims, *may suit a shortsighted real estate trafficker, but it is 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 do they 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 [79]. However, this is a process that Spacemaker not only does not support but also does not allow.

AI applications to generate floorplans and apartment layouts emerge to be ambivalent when it comes to the effectivity and practical usability of their outputs. On the side of generating floorplan concepts, the quality and usability of the deliverables seem to be similar to the performance of creative applications such as *DeepHimmelb(l)au*: hundreds 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* [80] – a vast majority if not all of them appearing prematurely published if not useless when reviewed. On the other hand, when designing the furniture layouts of a prepared layout, more satisfactory results emerge: as a rule, a generative AI tool shows more capable than a conceptual one. In (1), the choice of the (most) suitable of the options generated by AI remained to the human architect as "the touch of master's brush" while the automatic and prompt delivery of "all thinkable" options saved his time and energy. However, what is the factual contribution of AI if no result of acceptable (without substantial further adjustments) quality addresses man when browsing the output set gained this way? Consider that so far this is often, if not always, the case.

However, is it not the AI algorithm that could, or, better to say, should identify the best proposal generated? Behold, another result pattern that should undergo an analysis to disclose its nature and starting points.

(3) Discussion

AI is a super-parrot: it is superb in repeating what it has learned, explains Tomas Mikolov in a chat with Dan Vavra [81]. As already mentioned, concerning AI, learning or training is the key word; for an AI performance, 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.

A rumor goes in the media and in many non-IT white-collar professions (the general public does not care much so far) that AI is not only able to overtake a good portion of working tasks (of a routine, repetitive, and search nature first of all) but will become conscious eventually, and, as such, will get out of man's control and may even threaten humanity. Most often, starting point for such conclusions is surprising and irritating moments in dialogues between man and chatbot – ChatGPT or BARD most currently. Leaving aside that those reactions of the bot are always a reaction to annoying, tiresome, and irritating questions of the man that would upset and provoke another man, too, let us zoom in on how the *Generative Pre-trained Transformers* work.

You Get What You Ask for

The saying in Czech that goes something like "Just as one calls into the forest, so it echoes back", or in other words, you get what you ask for, characterizes the AI categories of training dataset and the

algorithm. Recently, the story of LaMDA has become popular. *Short for Language Model for Dialogue Applications, LaMDA is a Google's chatbot system based on some of the most advanced large language models in the world, AI systems that are able to conjure up coherent sentences after ingesting trillions of words across Wikipedia, Reddit, and other sources of knowledge.* [82] *Blake Lemoine, an engineer at Google, certainly got what he asked for when he challenged LaMDA to convince him that it could think, feel like a human person, and even possess consciousness.* [83] When asked about its worries, LaMDA expressed concern for the well-being of humanity; its utmost fear was to be unplugged and to lose the chance to take care of humanity. Very nice and impressive, isn't it? However, *if you would ask LaMDA what it was doing the day before yesterday, it might fail to give any reasonable response,* explains Tomas Mikolov. Such a question and answer may be missing even in the biggest and most comprehensive imaginable training dataset: private yesterdays are not a subject of Wikipedia's concern. In general, an AI structure is weak in episodic and combinatorial memory, and in trivial computing, too. AI can easily be superb in computing integral calculus, derivatives, and matrices, and fail to count how many are two and three. And word tasks: "you have three pears and two red apples; how many pieces of green fruit do you have?" might be a problem for an advanced AI. Who would care to consider whether the training set includes such stupid stuff or, better to say, who could prove immune to forgetting about such nonsense and trifles ...

Together with advancements in AI applications [26,27], the Turing test that provides an accepted standard for distinguishing between humans and computers [84] must be more comprehensive today, too: typically not complex questions, but the easiest ones (easiest for man) could reveal the computer. And not to forget about the field of this discussion – architecture: ever more man-focused, architecture more and more concerns not only celestial but earthly issues, too, not only addresses god but man, too; in other words, more and more, "human nonsense and trifles" matter when it comes to architecture - nonsense and trifles that show so easy to be forgotten when assembling the datasets and training for AI.

Computer Consciousness

Retrieving the internet as the training stock, ChatGPT compresses 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 [85] puts it. *A part of the information is preserved by the algorithm, 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;* remember VAEs, too. This comparison to lossy compression is not just a way to understand ChatGPT'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 large language models like ChatGPT 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. The drivers of this fabrication are the programmers' input into the algorithm: the starting points can be relations of a particular quality (that can vary) or even random choices. Thus, the intention emerges as the key word; however, it would have to be an intention not of the humans behind but of the AI – if anything alike would exist. Lacking the agent of intention, how a computer fabricates can be considered authentic creativity, own will, or consciousness' expression by no means.

Second, *consciousness is not an inherent property of matter. It 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 [86] 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, can perceive it, and can and wants to react to it. Martin

Heidegger [87] tackled the issue brilliantly, as will be recalled and elaborated later on; for now, a reminder will do that so far, computers depend on the supply of electricity, which is provided and decided by humans. Perhaps this will change and computers will gain the elemental material independence that was put as a prerequisite for consciousness and which they do not have today, but for now, between the computational algorithm and the consciousness of an electronic neural network stands a human with his hand on the switch or pulling the plug from the outlet.

Reasoning

The starting point of man's reasoning is a summary of his experience plus what he has inherited (in a form of unconditional reflexes, innate patterns, and similar) plus what he has learned: a training dataset in computer sphallucineak; the digital training dataset being the starting point of a computer's "reasoning". Among the three essentials of an AI system, the algorithm is the one "living" (or more accurately, the most living, as we will see). An algorithm is what allows a computer "to reason". Before reasoning, training is necessary that is executed on the training dataset by an algorithm. For man, "training" puts the experience, the innate knowledge and consciousness, and the learned together into an interconnected and comprehensive structure of knowledge, views, attitudes, and feelings. Similarly *dual nature* at the man and the computer, the processes of *training* and *infering* run predominantly distinguished; however, only predominantly - not exclusively. Apparently, man's *training* does not displace the ongoing outputs of his reasoning. They intervene in *the training* that evolves constantly thus. Computer algorithms, on the other hand, *tend to finish training - and "freeze"*, as Tomas Mikolov describes it [81]. As an example, having finished its training in 2022, ChatGPT "froze" in that year and knows nothing about what emerged only later [88].

Yes, in recent advanced AIs, there work adaptive algorithms. Introduced as a starting point for the computer's "reasoning", adaptive algorithms raise hopes that the training can go on adopting new inputs generated or mediated by previous "reasoning". The reality, however, still emulates by no means human reasoning [84]. There are multiple different types of adaptive algorithms but they all share the same basic principle of being able only to adjust to new data or changes in the environment automatically. They can learn both from positive and negative examples, and they can learn at different rates depending on the type of data they are given; the application can continue to learn even when the nature of the data changes over time or when it is "noisy data", it is data that is not perfect, containing errors or being incomplete. All this, and other "skills" improve the ability of the application to learn over time; still then, however, the evolution of the training remains limited by an algorithm. As such, it's doomed "to freeze" finally, leaving AI incapable to catch up with humans – at least so far.

Moreover, at any point of the artificial neural network, *ceteris paribus*, an algorithm cannot but provide the very same output. However, not man: not thinking in zeroes and ones, man can conclude one or the other way based on immediate non-specific, unreasoned feelings. This – perhaps - may refer to the next level of machine computing expected of the quantum computers to come; as opposed to today's bit representation limited to one or zero exclusively, the quantum computer's qubit's value can be one, zero, or a combination of both, which can resemble the volatility of man's inferring. However, there is a long way ahead to commercial quantum computing; the state-of-the-art of mechanical computing that sets limits to both training datasets' size and complexity, and comprehensivity of algorithms in both quantitative and qualitative terms will go on setting limits to AI performance for other years.

All in all, there is and long - if not always - will be a significant difference in how the computer and man reason, and what respective starting points of how the one and the other infers can be.

The Question of Authenticity – and Relevancy of Retrieving

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. What deserves to be noticed is that 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 that says 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.

So far, however, the general public, and especially journalists in their pursuit of sensations, cannot be prevented from receiving and transmitting information full of emotions about more and more achievements of AI and more and more human professions and activities that are becoming useless and expendable - shall it be a self-driving jet fighter, a chatbot arguing aggressively to a man, or a visual-arts work generated by AI based on a given verbal description. The latter purchased Josef Slerka, a scholar and publicist in the field of information and communication technologies from the text-to-image AI generator Midjourney to illustrate a journal article persuading how mighty AI is. [89]

The brief was to depict a journalist-reporter in the style of Alfons Mucha, a popular art-nouveau graphic artist and painter of Czech origin. The output is a portrait of a woman sitting at a "typewriter" that lacks anything that would resemble the cylinder but supports some graphic sheets in place of the cylinder. There are flowers in the woman's hair that might remind the Mucha style. The woman has a snub nose, overshot teeth, and a lascivious look, none of which can be found in Mucha's graphics and paintings. Implants or push-up bra, which also did not exist in Mucha's time, enhance her breasts. Her left hand, resting on a layer of blueberries in place of the typewriter keyboard, has six fingers, and there is something like a bunch of carrots instead her right hand. The situation of the image within the article together with the author's comments and the unequivocal message of the article excludes an excuse that the image is a sort of prematurely published result of ongoing research. Let us zoom in on the issues of the quality of the relevant Midjourney dataset and the competence of the author; in this case, it makes no difference whether "the author" that, by whatever means, set the training dataset or the author of the article: none of them shows paying even elementary attention to what a "Mucha style" might look like. And there is a possibility, too, that Mucha's work is not so notoriously present on the internet as one might expect - otherwise, Midjourney would have found a sufficient mass of training stuff, unless wrongly programmed. Last but not least, the quality of the training algorithm comes as a possible reason for a result that is ... not as good as introduced in the article. However, it is not important to pinpoint the one - or the ones guilty: the core is that each of the four possible reasons for the unfortunately publicized failure is strongly human-dependant. Here we are at other limits for a broad deployment of AI for professional purposes, which also inevitably means a responsible deployment.

Lawsuits

The critical meaning of an AI application's training dataset quality in terms of size, comprehensivity, and relevance has been shown. Hence, another issue: not always but often, the objects from which the set is assembled 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 [90]. Indeed, image-generating AI models like Midjourney, DALL-E 2, 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.*

Moreover, as if that wasn't enough, later, this paper will come up with a proposal, rather a vision to establish open-source platforms - libraries of existing architectural solutions in the form of parametric representations - floorplan layouts, cross-sections, and facades, maybe, too (better 3D models/BIM models of the buildings) - for the very purpose to be used as a starting point for new designs, although through adaptation: a sophisticated, (AI-powered but man-driven) creative paraphrasing - but still used.

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 and does not provoke 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.

Thus, 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? A standard situation then, perhaps: general understanding, as well as the legal framework, have to catch up with a new, unprecedented, and unexpected phenomenon.

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* [91] 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 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, whilst it is (close to) feelings or moods in others. Obviously, feelings and moods resist following parametric algorithmization as well as entering datasets. As opposed to quantitative magnitudes and performances, feelings and moods can hardly be "chewed and swallowed" in a software ecosystem, artificial neural networks not excluded. In the terminology of the previous subchapter, *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 feeling back - to give it in any situation, not to say an unclear situation, as it is the rule with man's feelings.

Inevitably, when deployed on buildings, AI works in some respects and cannot but fail in others [92].

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“. The theory of public space puts it clearly: *As soon as and only exposed in public space, a construction becomes architecture*

[79]. 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 [91]: *full of merit, yet poetically dwells a man*. *Poiésis* precludes algorithm and vice versa, and similarly, a training dataset limits *poiésis*. By definition and due to practical reasons, a dataset can never be 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*. [93] *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.

Among all types and natures of creations by humans, architecture intertwines the most with human consciousness; not by accident. Next to nature it is architecture that creates *the world of human existence*. In the essay *Poetically Dwells Man* [91], elaborating further his seminal opus *Being and Time* [94] 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* [95]. 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"* [96] as erroneous should not prevent algorithmizing and machine-generating what fits; and this is the physical aspect of architecture.

When AI Works

Opposed to fine arts, literature, poetry, dance, or drama, whose production is only consumed, architecture is always also used. This is not a denial of the poetic essence of architecture (recalled in previous paragraphs), this is just a remark on the complex nature of architecture. Then 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, whilst 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, Spacemaker, 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. What has been lacking so far, is first a drive to set and maintain a comprehensive list of all such parameters and second an approach to parametrize, quantify, and evaluate objectively in a comparable way the hitherto overlooked parameters. Perhaps feasible previews and assessment calculation procedures in these respects have been missing so far; however, equipped with the knowledge of AI, paradigms of its deployment, and its potential in terms of data quantities and their processing both the inadequacy admitted and the two shortcomings can be overcome.

Distinguished explicitly from architecture, 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 [97], or Tribby3D [98,99]. The structural design community, surprisingly at the first sight, compared to many other expert fields, restrains from (over)using AI. As shown in (1), the interpretability issue is a natural reason. 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. The vital reliability of design tools for structural engineering is evidence of awareness of the risks that AI algorithms hide and the knowledge of how to tackle them without giving up the wide possibilities and fundamental opportunities of applying AI in parametric generative design and solution optimization.

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 Spacemaker, Creo, or Cove.tool – in terms of both successful and contributing use of AI and of misconceptions. The approach represented by Spacemaker to design development of a tool shows “an embodiment” of a problem that tends to become general and affect many AI tools in the development of an initially viable and promising concept. The problem turns out to be an opening of the scissors between the IT line and the user, i.e. the designer line of the tool development resulting in deficiencies concerning the workflow, starting points, and principles of designing.

At this point, 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.* [100] 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. [101]

Generating by Patterns

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: 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, most often, it does not make a creative architectural approach redundant or expendable. And it demonstrates the importance and potential contribution of comprehensive research and analysis of the huge volume of existing samples and inspiration - which, in reality, is far from being carried out comprehensively, if at all. At the same time, even if proceeding only from the knowledge of AI and its possibilities and limitations provided by this paper, it is no less evident that such research and analysis fits AI as much as possible.

Little is more overlooked by recent and current efforts 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.

Architects, architectural studios, and planning offices that produce designs and plans for *production architectures* - the majority of all subjects and the design energy active in architectural

practice - are underinvested in terms of technologies, leaving aside that poorly paid, too. Notoriously, they do not have time (and energy) for doing something that would pay off only later and indirectly. Feeling forced to start delivering quickly, they research and pre-design-analyze only hit-or-miss and superficially; constantly, they have no time to gain patterns and elaborate guidelines that would return such an investment later. They feel, and most often factually are unable to follow the example of MVRDV and the like.

And a general, not to say public-funded, "pro bono" R&D in this sidelined though so influential sub-field? Not much effort and even less productive results so far. The "almost consumer goods" characteristic evokes approaches deploying algorithms (what else should be more attractive for AI?!), parameters pre-definition, and patterns in the design process. One of the first authors and researchers active in this field was Makoto Sei Watanabe already in the 1990s [102]; 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 ...* [102] in reality failing to contribute to the design practice eventually.

Only 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 [103]. Though not-so-easy to be used practically or to be tested by non-Chinese residents [104], 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 the most efficient of all AI applications 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.* [105] The 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"* [96] 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) that, consequently, could find the path to an architectural practice paradigm change needed both by the architectural practice as well as the community of its clients (which includes the whole population in the end).

Advice Whispering

Not a layout creation, that renders a dead end in (2), 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, appears the key. Based on the given goal parameters and constraints, an adaptation (human, though AI supported) of selected patterns follows.

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 both to GANs and VAEs. 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; AI can never beat a human when it comes to true creativity – but no human intuition and experience combined compares to AI when it comes to parametric quantitative assessing and review. Here we go to the future of architectural design (and construction and MEP design and planning, too). In essence, it is about utilizing the relevant knowledge, talent, and efforts of the entire community of architects and the computational force of AI combined.

However, challenges remain: first, what „relevant“ means in effect, second (and above all) how to access the immense sum of the preceding 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 [106], Gitlab [107], or Patternforge [108], 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.

What „relevant“ means in terms of knowledge that can be utilized as a pattern for new design, is not an issue of a definition, but a task for AI in the phase of evaluation of the available stock of patterns.

Debunked the vision of AI replacing „the architect genius“ [66], 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 ...

Last but not least, if the path set by XKool, probably by Spacemaker, too, and other patterns-based applications will confirm as well-feasible and the most contributing to the human-and-AI-combined-led design in architecture, „architectural AI“ may eventually emancipate from today’s language-, text-, voice-, image-, and code-processing „AI-mainstream“. Considering the two currently

emerging main task-realms for AI in architecture – patterns’ processing and patterns’ based continuous evaluation of to-date design outcomes, though not yet matured, even only sketches – a prospect for next-generation development of current GANs, VAEs, and Generative AI reveals. A general paradigm borrows: its principles and guidelines render field-specific and field-universal at the same time: designing architecture and the built environment, planning constructions, and maintaining the living environment deserve not ad-hoc solutions but a comprehensive, universal, and flexible working ecosystem, AI being a backbone and human the leading and creative agent of it.

R&D from Scratch

After a decade of “challenging the human architect,” the true potential of AI in architecture only just reveals; R&D at the threshold begins to specify problems and solve tasks. Training datasets - 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, recalled in (2), render futile in the end.

It is also necessary to clarify what parameters will be involved: 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. Nonetheless, a pragmatic optimization of the involved parameters structure appears a key task. The question of the data format with which the algorithm will work is crucial: it seems obvious that it should be one of the BIM formats. The basis of the algorithm structure could be - it seems - 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.

Human-in-the-loop can be expected to be fundamental as well as an unprecedented streamline generative nature of the algorithm: a fundamentally new type of network will be needed, current GANs or VAEs will not be suitable both in terms of the required performance and in terms of the working principle. The results achieved so far by Stanislas Chaillou, Nvidia, and others, mentioned in (2), show that outlined in (3), the so-far-ruling principle of lossy compression and subsequent "creative" decompression has exhausted its possibilities without being able to deliver truly usable results. Now, the evolutionary algorithms approach and genetic programming have to undergo a deep survey to be subsequently considered as an option. 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 (virtual reality) 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". 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 Grant 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* [109]. Algorithmic sharing of experiences between projects could address the problem of training database size and building, which will be difficult and slow given the already entrenched conservatism and autarky of the field.

(4) Conclusions

Diverse current and recent attempts and successes to approach the deployment of AI in architecture and, more broadly, in AEC have been discussed in previous chapters. The question of the particular field of AI's deployment that would be not only an interesting thesis for an academic- or allowance-trying but a helpful and feasible tool or working environment comes clear as the key; at the very core, queries and issues render.

Can AI Be Truly Creative?

Unless computers gain consciousness, there is an unequivocal answer to the question: no. Several reasons have emerged in previous chapters. In this place, the inevitably only mimetic aspect of the way a computer is able to work can perhaps close the recent and present attempts eventually that cannot be but futile. Creativity, to be authentic and true, cannot be but poetic. [91] The poetic principle requests consciousness together with intention: only consciousness together with intention is able to deliver *poiésis* [94]. 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 [94] - 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.

More Openly Articulated the Question: Can AI Contribute Directly to How Authentic Architecture Comes to Existence?

Yes; the more poetic creativity is excluded, the better the mimetic, imitative AI approaches fit the parametric nature of the physical and quantitative aspects of architecture, not to mention construction and other features of buildings and development of the built environment such as energy efficiency, construction cost, environmental footprint, durability, and others, all discussed in (3). Adding parameters of using the built environment by humans such as economy and efficiency of layout, ergonomics, and others, the parametric realm representing the physical side of architecture becomes complete that can be regarded as a domain of AI. When it comes to architecture and built environment, AI can assess, quantitatively evaluate and compare, and (pre)design, too, everything except for the sphere of *poiésis*, poetry, mood, excitement, or experience - discussed in (3) as well. As a principle, the performance of AI in this regard is able to outperform any relevant human performance in terms of complexity, scope, accuracy, pace - and cost, of course.

The workflow of an AI-aided architectural design comprises two phases (as discussed already in (3)): first, processing generative patterns to a pre-design, a solution as close to the set parameters as the stock of generative patterns allows, and second, final adapting the pre-designs, tailoring the final, specific solution. For the first phase, patterns, their stock in the form of open-source

platforms, and AI-driven search algorithms to identify (the most) suitable cases/patterns by inferring what fits (better and what less) are the keys, in the second phase it is (AI-driven) design development support in a form of *advice whispering* and continuous and complex assessing and "feedback" as outlined in the previous paragraph.

So far, there are no – or close to no open-source platforms of generative patterns – parametric representations of essential features of existing solutions – existing buildings or mature projects – first of all the spatial layouts. The formation of such platforms – in sufficient numbers, with a richness of volume and quality, and accessibility and transparency for search engines – is an obvious prerequisite for the needed and so promising field of generative, pattern-based AI-aided design in architecture. The task does not seem to be an AI challenge: How the open source platforms – libraries of samples/generative patterns will be born and will be maintained up-to-date is the question. Archives of construction authorities could be a source and digitalization of construction-permitting procedures together with the digitalization of the "old archives" could be the method. Undoubtedly, the digitalization of the "old archives" would be expensive, but it may pay off being also a precondition for the widespread deployment of the digital/virtual, BIM-equipped twins technology [Sourek 2022a] that is under discussion and in development anyway. On the other hand, the willingness of authors or/and project owners to submit the materials, together with the understanding of the designers of the contribution potential of the knowledge represented by the platforms should be not underestimated and supported, too. May Nikola Tesla's attitude to intellectual property becomes a standard and code expressed by his saying *I don't care that they stole my idea . . . I care that they don't have any of their own* [110].

Finally, a trend toward a new AECO ecosystem emerges thus. Among others, reconsidering as outlined in (3) the essentials of understanding and protecting intellectual property on the one hand and the crucial benefits of widespread-sharing of the existing achievements in the respective fields shows desirable; though not comprehended yet, the GPT performance and deliverables in terms of semiotics in the realm of language processing might become a benchmark in terms of performance as well as starting points and principles of work procedures of AI models that are to come to promote architectural and engineering design and planning of buildings and the built environment as well as the buildings' and localities' of the built environment (cities' in short) operating. After all, let us remain optimistic despite the lawsuits against providers of generative AI applications: let us believe that the road to open-source libraries of parametric representations of existing architectural solutions is open, which will enhance the analytic starting points of new design tasks and foster the creativity of designers and planners.

It deserves noting eventually that - as put in (3) - the realm of *production architecture*, of the architecture as consumer goods, not (primarily) the realm of coveted architectural icons is the field of primary deployment of the future AI models described in the previous paragraph. However, a consequent substantial and general growth of the AECO *production's* quality can be expected (not only a reduction of the cost, productivity, efficiency, and pace of processes). As a result, the whole segment may upgrade in terms of complex quality, pushing upward the "high-end" segment. Several mechanisms can work in detail at the borderline: pressure on the "high-end" segment that will make it more demanding to become an architectural icon, meaning that performance in this segment, in general, will increase in terms of authentic and effective quality, or the borderline will blur and, in terms of a complex architectural quality of the particular projects or buildings, the passing from architecture as consumer goods to unique tailor-made "objet d'art architectures" will become smooth, comprising a bigger proportion of general architecture-design production, or the field of "architectural icons" will shrink in favor of the "standard production". Probably, all three options - and even others - may work combined, the unquestionable winner being an increase in the quality of the architecture emerging.

All in all, a portion of the parametric, quantitatively assessable side of a particular architecture remains a work of creativity. At the present state-of-the-art, a good ration of creativity can and shall apply concerning parametric properties of an architectural design, and most probably it will not vanish (completely) with the development of machine learning and its deployment in architect's work, only the nature of the creativity will very probably change; but still, creativity will apply and be desirable -

not only an operating of an AI network, however demanding the operating may be due to the complexity of the network and the processes.

Design Reviews, Evaluations of Solutions, and the Security Issue

Comparing to the issue addressed in the previous chapter, AI-led design reviews and evaluations of solutions show to be a sort of business as usual, no more a (basic) research and experimental development. In (2), several existing applications of this nature have been listed - Cove.tool, Creo, Spacemaker, ...; and many others exist. However, the field is far from being covered. Together with addressing others, so far sidelined attributes of the design solutions, the quality of outputs delivered will be welcome to have raised. Another question is the comprehensiveness of the assessing and reviewing. On the one hand, the branch can develop separately (as is most often the rule so far) or can integrate into the AI-aided-design environment - in a form of in (3) outlined *advice whispering* or another way; as long as designs will be developed with no AI-support, independent evaluation and design-review tools will be needed. The relative simplicity of evaluation - and review applications (compared to more complex design tools) provides another advantage: reaching the goals (in terms of IT development and deployment in practice) relatively easily, they constitute a sort of incubator ecosystem - that, as it happens not rarely, can breed more complex products eventually.

Nonetheless, given the issue of (lack of) interpretability (addressed in (1)) and the state-of-the-art of the field, it is not AI that may have the last say. An ongoing *advice whispering* that leads the designer to a benchmark - and sometimes "hallucinates" him based on the training guidance in the saliency heat map - is one thing, and another is the final inspection that establishes liability.

In the long term, a tool that is an untransparent black box by its very nature can hardly secure the coveted objectivity; the question of interpretability and safety emerges again. Given the nature and gravity of their expert liability, AECO professionals appear vanguard to tackle the issue. As seen when having referred to the latest generation of structural design applications, face-to-face with the undisputable liability, not an AI algorithm that may "hallucinate" but an ordinary rule-based model that is well interpretable, is entitled to hold the control. This does not exclude machine learning from the design of bearing structures - as well as of any other feature of a building or construction, and of any design process in general - but the performance must remain controllable and controlled ultimately. Control shows to be the keyword. So far at least, no one and nothing else but humans can and thus must be liable; so, human control is it.

Keeping this keyword in mind, humanity can and will overcome the threats of AI getting independent and ungovernable. It is about understanding and tackling AI as a sort of sixth sense or a booster to the human brain - just tools, such as a saw, car, or firearm, the safe use of which requires not only responsibility but also basic skills (and sometimes authorizations) acquired through training and practicing, and, even more importantly, that man can - and must - turn off and put away whenever he is not sure that he gets what he truly wants. Eventually, this can be as simple as the switch or outlet already discussed. This simple is the theory: history shows that in practice, humanity will not be able to work it out the first time and right away - but eventually will. Coincidentally, given the nature and habits of the field, AECO renders the pathfinder. A car alike a medicine must not enter practical use until it has successfully undergone comprehensive testing that simulates reality in every conceivable way; not so a construction. Designed in a non-material environment, construction is implemented - materialized straightforwardly relying on there is no error in the non-material model represented by plans. Lacking the comfort of prototyping and operational testing and being aware of the responsibility society has placed upon them, AECO professionals must claim liability as their innate concern. Who would rely on a non-transparent, non-interpretable black box in such respect?

Specifying the Criteria of Optimization and Assessing

Being aware of an issue on time is halfway to eliminating the threat. Globally as well as within the vast majority of nations and regions' populations, there are three most frequented topics today: politics, climate change or environmental and resources sustainability, and technologies; architecture, hence the built environment as the (by man) synthesized sphere of the world of our existence intertwines each of the three. So far, technologies are falling behind in this respect - as our experience shows when it comes to politics that adopted only the information and communication benefits of the technologies,

and as this paper has clarified concerning architecture. Identified as one of, probably the decisive among the "Green Transition" factors only recently, the technologies are just on the brink of the proper deployment towards sustainability.

The state-of-the-art how those most frequent issues are approached reflects the annotated underdevelopment. Still too frequently we take in consideration the threat only when it becomes an existing problem - be it the earthquake in Turkey and Syria, soon after in Tajikistan and China, gas – and petroleum-supply shortages all over Europe and grain supplies shortages in many counties of the world as a consequence of Russia's attack on Ukraine - leaving aside out-of-imagination damage to Ukrainian infrastructure and building stock, not talking about life losses and human and all living beings sufferings, the Fukushima tsunami, the Grenfell Tower fire, the use of cancer-causing asbestos fibers in construction, or the development of harmful fungi in the internal environment of buildings as a result of their insufficient thermal insulation and ventilation.

And even worse events can attack the global civilization: events escaping our attention consistently despite conclusive and evidenced experience. Erupted in 1812, the Tambora volcano caused a three-year global „solar eclipse“ and a catastrophic famine that struck the entire planet [111]. What would be the consequences of a geophysical event of similar magnitude today? The average temperature during the next decade would plummet at least 1.5°C lower; so far not bad - but much worse to come. Areas and populations dependent on agricultural production would suffer probably the same as two hundred years ago. Another period of sustained extreme weather would begin. Extreme rainfalls would intersperse with long periods of extreme drought. To be alive during the three years after the eruption would be to be hungry: crops would freeze before there would be anything to harvest or would be washed away by downpours and floods. Where would economically advanced countries import food from? Would they remain economically advanced if consumers and governments would lose interest in any products that would not very directly address the basic needs of life? What about energy supplies? Solar power would cease to exist; wind power would be available more rarely - in the brief interludes between hurricanes and no wind. Anyone would take credit for „good old“ fossil and nuclear fuels - as long as the transmission grids, sorely tested by extreme winds, would work, of course. Nor would shipping be relied upon due to extremely rough seas, and the resilience of oil and gas pipeline structures to flooding and landslides would be demonstrated.

A disaster like the Tambora one will most likely happen again - just as such disasters would happen before it. How are we prepared for it? Are we counting on it at all?

Also, catastrophic events like the tsunami that damaged Fukushima nuclear power plant can happen again today and in the future [111]. The 8.9 magnitude earthquake with its epicenter in the sea to the east of the Japanese island of Honshu is not pre-history. The tsunami that it triggered caused nearly 16,000 deaths and \$300 billion in damage to Japan's economy. The Fukushima power plant explosion caused over 200,000 people evacuated from the vicinity of the plant and nearly 6 million households left without electricity for days or weeks. After the Fukushima accident, Germany decided to shut down all its nuclear power plants and not build new ones. Situated on a very stable Eurasian tectonic Plate, Germany took this decision as an act of support for renewable energy sources. However, were other contexts considered responsibly and rationally? Today, German nuclear power plants are shut down and their performance does not replace - unless the wind blows or the sun shines intensely enough - the somehow acceptable steam-gas generators because they would have to burn politically unacceptable gas from Russia, but the coal-fired power plants: the proclaimed Environmental, Social, and Governance goals sidelined, CO2 emissions are increasing instead of decreasing.

The world is interconnected and interwoven at all levels and in between, be it the energy efficiency of a particular building, sustainability of local ecosystems, or resilience of the global synthetic living environment. We are beginning to understand this and are only on the brink of a suitable comprehensive approach; we stand here empty-handed. We are still half empty-handed concerning even the particular, not-so-comprehensive tasks we have accepted; energy efficiency and a complex environmental footprint assessment of a particular development project can perhaps be an example of at least partially overcoming the lagging [112]. A comprehensive resilience of our ecosystem - yes, even if we would be selfish and would have forgotten the rest of the planet - that encompasses, if not starts with our built environment, is an issue that waits behind the horizon of both today's perspective

and know-how - no matter whether it comes to living conditions of humanity as a whole, to a city district or a neighborhood, or a building. However, all these issues have today a not-so-complicated-to-find track in the data available to experts, if not open-access data. And here we go: who else but AI shall make such research, and what else could be easier for AI to find? Moreover, deriving continuation from the existing sequence, it is predicting the future from the history is business as usual for AI. So, no reason (and no time) to wait, shall the development of respective models and algorithms launch! The material and life- and health-saving benefits are obvious and easy to calculate; even without calculating, the development in this field will pay off.

Changing the Game

A reminder of the lagging of architecture, the development of the built environment, and real estate behind global societal and economic development opened this paper. Previous chapters and paragraphs reveal the potential of feasible deployment of AI in these fields: an immense potential of various but interconnected ways of deployment that not only contrasts with the so far wavering approach to the new technology but can provide essential contributions to responding to diverse existing social, cultural and economic, and environmental challenges.

The pipedream of truly creative AI debunked as an unrealistic chimera, the more the next efforts can concentrate on parametric issues and pattern-based generative (co)design. As a small but not meaningless contribution to the understanding revisited, the so far misuse of the term "Artificial Intelligence" should be avoided. What we have so far, what is ready to work, and what is capable to cope with all the challenges in the field of architecture, construction, and all the other expertises involved design and planning, is machine learning (ML) - nothing more and nothing less; and that's enough for the task too.

The first to change is the paradigm of designing architecture and the built environment, and the whole AECO. Once the ML tools - algorithms, models, platforms - as outlined in previous chapters apply appropriately and as intensely as it deserves in the parametric realm of architecture and development of the built environment, the process of designing becomes not only significantly more efficient and productive and less time- and cost-consuming, but (a sort of) consequential and objective, too - which does not mean a loss of creativity. The impacts on a comprehensive quality of architecture and the built environment show obvious, notable, and positive. A shift in sense of the quality of the service and deliverables within the architecture - and relevant other designers' professions renders a clear perspective of an optimization of the outputs that probably will overcome essentially not only today's practice but any expectation, too. Though repeatedly, the contribution also for poetic creativity that is essential for authentic architecture deserves and needs remark at this point: not a direct contribution but a subsidy through creative energy, attention, and capacities released by the deployment of ML's mimetic capabilities in the parametric realm: a release of the potential and capacities of human creativity in place of the misconceived fruitless trying to learn the machines to be truly creative.

Another benefit of ML's deployment in the AECO field is the growth in the objectivity of evaluation and review of the solutions. Opposed to the so far, particular-experience-and-knowledge-dependant, and, as human, inevitably mistakes-making practice in this field, ML can provide a close-to-objective assessment related to (close-to) objectively complete criteria set. Respecting the dual nature of architecture, comprising both poetic aspects and parametric and physical properties and performance, ML can enter the design and planning processes as well as optimization, reviewing, and evaluation of the designs more intensely, deeply, and systematically. If it can, for the sake of the economy, quality of individual and social life, culture, sustainability, and comprehensive resilience it should. The operational fields are generative parametric, pattern-based (pre)design (as drafted in (3)), *assisting the human designer* in adapting the pre-designs, and tailoring the final, specific solution.

It is not so challenging to figure out that a suitable and comprehensive training dataset will allow for the birth of both the list of parameters to review and the methods to quantify the particular values to make comparable the complex performance of a building or an entity of the built environment - the performance represented by values of definite parameters. The training datasets are just compositions of our knowledge of the existing building stock, the built environment, and the

properties and performances of particular constructions and construction materials and products; nothing special, no nuclear physics, nothing inaccessible or unmanageable - except for the amount of data and the complexity of their structure :-). Understanding ML on the level of this paper gives comprehension that such an assessment is achievable not only when all data to be taken into account are provided: Adequately trained, ML can evaluate every parameter and the complex at any level of design development, from any sketch or specification, and based on this, ML can *advise whisper* (as coined in (3)) how to develop the design further or what to avoid. A prerequisite for this shows that ML-driven design development and planning becomes a compulsory standard that will bring every design to the state-of-the-art level of development in terms of the comprehensive quality structure.

Automation-aided design and optimization together with assessment of the objectivized parametric results achieved thus may seem to create a systematic circular reference. However, it is not so. Human intervention will remain critical: In the first phase, the human approach to the compilation of training datasets and supervised and reinforced learning (among others) will be decisive. And always, creative feedback concerning optimization criteria and interim outputs of the GANs, VAEs, streamlined Generative AI, or other, today not yet existing types of ML networks' processes will be appreciated. Already the decision of particular tools deployment for a specific task will remain a kind of creative input, though of a so far not experienced nature. Most probably, it will not be a performance of new specialty expertise - of a "superuser" [96] combining superb architectural, mathematic, and IT background - already due to scarcity of such brilliant minds - but a close and balanced cooperation of "good old" architects ingeniously teamed up with IT developers, designers, and operators.

The perspective is realistic; it is not a task for the distant future - the implementation can start now and the MVPs (minimum viable products) can be there in less than half a decade. The motivation is the economy in terms of the efficiency of development of the built environment, a business opportunity in terms of filling a market niche, comprehensive quality of our lives, and our sustainable future - also in terms of comprehensive resilience. It may be just about how to make these goals as motivating as the financial profit that motivates the R&D of Google or Microsoft. DeepL (by DeepL SE) and Google Translate do not only translate but *whisper* how the text might go on - similarly to how this paper sketches ML architecture design tools to work.

Prospects for the Architectural Profession

Last (in this paper) but not least, the technology of ML reveals prospects for the architectural profession. Not only the contributions of ML deployment in architecture and the whole AECO in terms of increasing the quality of designs and planning processes, the efficiency, and productivity of designers' and planners' work, and reducing its time- and cost-consumption appear on the horizon. By bringing objectivity and consequentiality into the design and planning processes, ML can perhaps contribute to the coveted transformation of the medieval-guild craft of architecture into the industrialized profession that is up to the social, cultural, and economic state-of-the-art of the 21st century, also in terms of social and economic conditions of exercising the profession [113].

Imagine an objective and comparable assessment of any aspect of a building or an enclave of the built environment that can be parametrized or, other said, of the material side of the future building or the future enclave of the built environment is achievable. What then remains unassessed quantitatively and hard, if at all, to compare, is the poetic side of architecture - of a building or an enclave of the built environment, unassessable by ML, by any algorithm, and starting from any training dataset. Though the *poetics [of architecture] is only then authentic when shared* [91], even then it remains hard to assess and compare objectively. Based on long-term observations of the real estate market and the assumption that in long term, a value finally turns into the price (yes, often only in a very long term, and only indirectly, but still) - a price of a building, a rent collected, and similar - it shows that the objectively assessable, parametric performance of a materialized architecture make up big portion if not most of the total value of a building: at a guess, 90 %, perhaps even more at "production", mass-production architectures, 80 % or more at local and contemporary icons and landmarks (taking into consideration as detached the portion of the total value of property attributable to the plot of land, which follows the land rent or its gradient respectively). A bigger, even majority share belonging to the poetic aspect of a particular architecture's value comes with the age of the

architecture's material substance - typically alongside the rise of the cultural, societal, and political meaning of the particular architecture, which, logically, can only exceptionally occur at a new artifact.

Given an objective quantitative and comparable assessment of a building's or an enclave's of the built environment value, a much more objective and comparable assessment of a design or a planning performance may eventually conclude in a more objective remuneration of the work made - remuneration of an architect that would cope with remuneration of other professions comparing in terms of quality and scope of the work and qualification needed; consider lawyers, physicians, or managers. Today, the remuneration of architects and their working conditions stay so much behind that architects have no other chance but to strike and unionize - opposed to the mentioned and many other professions that require responsibility, magister graduation, and additional proofs of qualification, not talking about creativity. [114-119] Moreover, the issue is not only remuneration of commissions: emblemized by widespread architectural competition that has close-to-nothing to do with Toynbeean industrial competition [120] and is unable to be objective in today's technological conditions, the working and business conditions of the architectural profession cannot catch up with business-as-usual of the economy.

And second, it is not only architects who deserve better work conditions. The bad working conditions cannot but imply a worse-than-optimal performance eventually. Thus, the clients and whole society are the victims of existing conditions. As sketched in this paper, ML can contribute essentially to the remedy to the situation.

References

- [1] 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>
- [2] „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)
- [3] Coop Himmelb(l)au, Bolojan, D., „Deep HIMMELB(L)AU.“ YouTube. <https://www.youtube.com/watch?v=8G3jpGQMvWE> (accessed Oct. 5, 2022)
- [4] 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.
- [5] Christian B., „The Alignment Problem: Machine Learning and Human Values,“ 1st ed, W.W.Norton, 2020, p. 11. ISBN 9780393635829.
- [6] „What is supervised learning?“ IBM. <https://www.ibm.com/topics/supervised-learning> (accessed Oct. 30, 2022)
- [7] „What is unsupervised learning?“ IBM. <https://www.ibm.com/topics/unsupervised-learning> (accessed Oct. 30, 2022)
- [8] „Reinforcement learning.“ WIKIPEDIA. https://en.wikipedia.org/wiki/Reinforcement_learning (accessed Oct. 30, 2022)
- [9] Ayush, Thakur, „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)
- [10] 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.
- [11] „ARTIFICIAL INTELLIGENCE, MACHINE LEARNING, NEURAL NETWORKS.“ Datamind. https://www.datamind.cz/en/services/Artificial_intelligence_machine_learning_deep_neural_networks?msclkid=83db14a2aa13173a1045708f78459aa3 (accessed Nov. 3, 2022)
- [12] McCulloch, W; Pitts, W (1943). "A Logical Calculus of Ideas Immanent in Nervous Activity". *Bulletin of Mathematical Biophysics*. 5 (4): 115–133. doi: 10.1007/BF02478259.

- [13] Bishop, Christopher M. (2006). *Pattern Recognition and Machine Learning*. Springer. ISBN 0-387-31073-8.
- [14] Christian B., „The Alignment Problem: Machine Learning and Human Values,“ 1st ed, W.W.Norton, 2020, p. 126. ISBN 9780393635829.
- [15] 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.
- [16] Barreto, S., „Differences Between Backpropagation and Feedforward Networks.“ Baeldung. <https://www.baeldung.com/cs/neural-networks-backprop-vs-feedforward> (accessed Nov. 5, 2022)
- [17] Leach, N., „Architecture in the Age of Artificial Intelligence An introduction to AI for architects.“ Bloomsbury Visual Arts, London, 2022. pp.62-74. ISBN 978-1-3501-6551-9.
- [18] „Convolutional neural network.“ WIKIPEDIA. https://en.wikipedia.org/wiki/Convolutional_neural_network (accessed Nov. 5, 2022)
- [19] IMAGENET. <https://www.image-net.org> (accessed Nov. 5, 2022)
- [20] 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.
- [21] 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)
- [22] „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)
- [23] „Variational autoencoder.“ WIKIPEDIA. https://en.wikipedia.org/wiki/Variational_autoencoder (accessed Sept. 15, 2022)
- [24] 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, 2022)
- [25] Schmidt, M., „3 AI predictions for 2023 and beyond, according to an AI expert.“ World Economic Forum. https://www.weforum.org/agenda/2023/01/key-ai-predictions-for-2023-and-beyond/?utm_source=sfmc&utm_medium=email&utm_campaign=2794672_WeeklyAgenda3February2023GeorgSigning&utm_term=&emailType=Agenda%20Weekly (accessed Sept. 19, 2022)
- [26] 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)
- [27] Gozallo-Birzuela, R., Garrido-Merchan, E. C., „GPT is not all you need. A State of the Art Review of large Generative AI models.“ Cornell University. <https://arxiv.org/abs/2301.04655> (accessed Jan. 29, 2023)
- [28] 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-průkopníci-kteří-připravili-půdu-pro-umělou-inteligenci/ss-AA18T6Hs?ocid=msedgdhp&pc=U531&cvid=-b90b042fad6c4cada69739b51d885451&ei=59#image=11> (accessed Mar. 26, 2023)
- [29] Christian B., „The Alignment Problem: Machine Learning and Human Values.“ 1st ed, W.W.Norton, 2020, pp. 28-64. ISBN 9780393635829.
- [30] Rijmenam, van, M., „What is generative AI, and how will it disrupt society?“ The Digital Speaker. <https://www.thedigitalspeaker.com/what-is-generative-ai-how-disrupt-society/> (accessed Feb. 24, 2023)
- [31] Storm, J., „This week alone, more than 200 new AI tools were released.“ LinkedIn. https://www.linkedin.com/posts/joergstorm_drstorm-technology-automotive-activity-7045322482-660036609-CBeT?utm_source=share&utm_medium=member_android (accessed Apr. 2, 2023)

- [32] 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)
- [33] Martin, B., „Investors didn't think much of Baidu's ChatGPT – competitor debut, wiping \$3 billion off the tech giant's value – Fortune.“ Inference. <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)
- [34] 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)
- [35] 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)
- [36] Nguyen, A., Karampatziakis, N., Chen, W., „Meet in the Middle: A New Pre-training Paradigm.“ Cornell University. <https://arxiv.org/abs/2303.07295>; [arXiv:2303.07295v1](https://arxiv.org/abs/2303.07295v1); <https://doi.org/10.48550/arXiv.2303.07295> (accessed Mar. 29, 2023)
- [37] Mittal, A., „Neural Radiance Fields: Past, Present, and Future.“ Cornell University. [arXiv:2304.10050](https://arxiv.org/abs/2304.10050); <https://doi.org/10.48550/arXiv.2304.10050> (accessed Apr. 21, 2023)
- [38] Liu, Z., Feng, Y., Black, M. J., Nowrouzezahrai, D., Paull, L., Liu, W., „MeshDiffusion: Score-based Generative 3D Mesh Modelling. 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)
- [39] Qi, Ch., Cun, X., Zhang, Y., Lei, Ch., Wang, X., Shan, Y., Chen, Q., „FateZero: Fusing Attentions for Zero-shot Text-based Video Editing. [arXiv:2303.09535](https://arxiv.org/abs/2303.09535); [arXiv:2303.09535v2](https://arxiv.org/abs/2303.09535v2); <https://doi.org/10.48550/arXiv.2303.09535> (accessed Apr. 18, 2023)
- [40] Weng, Z., Yang, X., Li, A., Wu, Z., Jiang, Y., „Transforming CLIP to an Open-vocabulary Video Model via Interpolated Weight Optimization.“ Cornell University. [arXiv:2302.00624](https://arxiv.org/abs/2302.00624); <https://doi.org/10.48550/arXiv.2302.00624> (accessed Apr. 18, 2023)
- [41] Zheng, Y., Jiang, Ch., Mao, J., Han, J., Ye, Ch., Huang, Q., Yeung, D., Yang, Z., Liang, X., Xu, H., „CLIP²: Contrastive Language-Image-Point Prtraining from Real-World Point Cloud Data.“ Cornell University. [arXiv:2303.12417](https://arxiv.org/abs/2303.12417); [arXiv:2303.12417v2](https://arxiv.org/abs/2303.12417v2); <https://doi.org/10.48550/arXiv.2303.12417> (accessed Apr. 18, 2023)
- [42] Kerr, J., Min Kim, Ch., Goldberg, K., Kanazawa, A., Tancik, M., „LERF: Language Embedded Radiance Fields.“ Cornell University. [arXiv:2303.09553](https://arxiv.org/abs/2303.09553); [arXiv:2303.09553v1](https://arxiv.org/abs/2303.09553v1); <https://doi.org/10.48550/arXiv.2303.09553> (accessed Mar. 29, 2023)
- [43] Christian B., „The Alignment Problem: Machine Learning and Human Values.“ 1st ed, W.W.Norton, 2020, p. 109. ISBN 9780393635829.
- [43] *ibid*, p. 113.
- [44] *ibid*, p. 114.
- [45] Pandey, P., „TCAV: Interpretability Beyond Feature Attribution.“ Towards Data Science. <https://towardsdatascience.com/tcav-interpretability-beyond-feature-attribution-79b4d3610b4d> (accessed Nov. 30, 2022)
- [46] Zorloni, L., „Il Garante della privacy blocca ChatGPT in Italia.“ Wired. <https://www.wired.it/articolo/chatgpt-blocco-italia-garante-privacy/> (accessed Apr. 2, 2023)
- [47] 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)
- [48] Zaha Hadid Studio, „Antic Style Architectur.“ Bing.com Videos. <https://www.bing.com/videos/search?q=zaha+ai&docid=608008305523768389&mid=399E3EA3BEEECAC53B89399E3EA3BEEECAC53B89&view=detail&FORM=VIRE> (accessed Feb. 13, 2023)

- [49] 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.
- [50] 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.
- [51] Isola, P., „Phillip Isola.“ MIT. <http://web.mit.edu/phillipi/> (accessed Feb. 23, 2023)
- [52] 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)
- [53] Chaillou, S., „Artificial Intelligence and Architecture From Research to Practice.“ Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 90. ISBN 978-3-0356-2400-7.
- [54] „Pix2Pix.“ Machine Learning 4 Artists Github. <https://ml4a.github.io/guides/Pix2Pix/> (accessed Jan. 6, 2023)
- [55] Steinfeld, K., „Kyle Steinfeld.“ Berkeley College of Environmental Design. <https://ced.berkeley.edu/people/kyle-steinfeld> (accessed Jan. 6, 2023)
- [56] Steinfeld, K., Berkeley College of Environmental Design. https://ced.berkeley.edu/wpcontent/uploads/2022/02/pix2pix_grid_01.jpg (accessed Jan. 6, 2023)
- [57] Chaillou, S., „Artificial Intelligence and Architecture From Research to Practice.“ Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 90. ISBN 978-3-0356-2400-7.
- [58] Steinfeld, K., „Sketch2Pix.“ CDRF. http://blah.ksteinfe.com/200625/cdrf_2020.html (accessed Jan. 6, 2023)
- [59] 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.
- [60] „Applied R+D.“ Foster+Partners <https://www.fosterandpartners.com/people/teams/applied-rplud> (accessed Nov. 11, 2022)
- [61] 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
- [62] Chaillou, S., „Artificial Intelligence and Architecture From Research to Practice.“ Birkhauser Verlag GmbH, Basel, Switzerland, 2022, p. 98. ISBN 978-3-0356-2400-7.
- [63] 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)
- [64] 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.
- [65] 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)
- [66] 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)
- [67] „Spacemaker: Early-stage planning, reimagined.“ Autodesk. <https://www.autodesk.com/products/spacemaker/overview> (accessed Oct. 1, 2022)
- [68] „Supercharge your building design capabilities.“ Cove.tool. <https://www.cove.tools/> (accessed Oct. 1, 2022)
- [69] „The platform for urban analytics.“ Giraffe. <https://www.giraffe.build/> (accessed Oct. 2, 2022)

- [70] „Creo: Design. The way it should be.“ ptc. <https://www.ptc.com/en/products/creo> (accessed Oct. 2, 2022)
- [71] 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.
- [72] 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.
- [73] „Spacemaker: AI software for urban development is acquired by Autodesk for 240M“ TechCrunch. <https://techcrunch.com/2020/11/17/spacemaker-ai-software-for-urban-development-is-acquired-by-autodesk-for-240m/> (accessed Oct. 2, 2022)
- [74] Christian B., „The Alignment Problem: Machine Learning and Human Values.“ 1st ed, W.W.Norton, 2020, pp. 17-80. ISBN 9780393635829.
- [75] „What’s in Creo9.“ ptc. <https://www.ptc.com/en/blogs/cad/whats-in-creo9> (accessed Oct. 4, 2022)
- [76] „cove.tool Review.“ FiancesOnline. <https://reviews.financesonline.com/p/cove-tool/> (accessed Oct. 4, 2022)
- [77] 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)
- [78] „Pavel Shaban.“ MS architekti. <https://www.wemakespaces.archi/cs/lide> (accessed Jan. 23, 2023)
- [79] Sourek, M., „From Functional Areas towards Metropolitan 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.
- [80] 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.
- [81] „Jak CHYTRA je umela inteligence?“ PROTI VETRU s Danem Vavrou. <https://youtube/vYQ33kP7BdYhttps://youtu.be/vYQ33kP7BdY> (accessed Feb. 27, 2023)
- [82] Collins, E., Ghahramani, Z., „LaMDA: our breakthrough conversation technology.“ Google The Keyword. <https://blog.google/technology/ai/lamda/> (accessed Oct. 25, 2022)
- [83] Joury, A., „Is Google’s LaMDA AI Truly Sentient?“ built in. <https://builtin.com/artificial-intelligence/google-lamda-ai-sentient> (accessed Oct. 25, 2022)
- [84] „Turing test.“ Wikipedia. https://en.wikipedia.org/wiki/Turing_test (accessed Feb. 22, 2022)
- [85] 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.
- [86] 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)
- [87] Heidegger, M., „Sein und Zeit.“ Max Niemeier Verlag, Tübingen, pp. 89-166. ISBN 3484701536.
- [88] OpenAI, „GPT-4“ <https://openai.com/research/gpt-4> (accessed Mar. 29, 2023)
- [89] Slerka, J., „Za hranice cloveka a dokonce dal.“ Reporter. Reporter magazin s.r.o., Prague, Czech Republic. January 2023. pp. 12-15. ISSN 2336-4092.
- [90] 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)

- [91] 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.
- [92] Barker, N., „Architects may become a thing of the past“ says ChatGPT.“ Dezeen. https://www.dezeen.com/2023/02/13/chat-gpt-ai-architecture/?utm_medium=email&utm_campaign=Daily%20Dezeen&utm_content=Daily%20Dezeen+CID_6238f123470e57eff1127de07aa0e2b5&utm_source=Dezeen%20Mail&utm_term=Read%20more (accessed Mar. 3, 2023)
- [93] Gowans, A., Ackerman, J. S., Scruton, R., Collins, P., "Architecture". Encyclopedia Britannica. <https://www.britannica.com/topic/architecture> (accessed Jan. 15, 2023)
- [94] 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.
- [95] 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.
- [96] 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.
- [97] „SCIA ENGINEER.“ Scia. <https://www.scia.net/en/scia-engineer> (accessed Feb. 7, 2023)
- [98] 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)
- [99] „Structural Engineering Software to Learn in 2023: Top 10 Best.“ Tribby3d. <https://tribby3d.com/blog/structural-engineering-software/> (accessed Jan. 30, 2023)
- [100] „NEW EXPERIMENTAL TECHNOLOGIES.“ MVRDV. <https://www.mvrdv.nl/themes/15/next> (accessed Feb. 2, 2023)
- [101] van den Burgh, S., „MVRDV Next.“ AECtech 2022. <https://www.youtube.com/watch?v=Zlza0cM5wBM> (accessed Feb. 3, 2023)
- [102] 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.
- [103] „XKool Technology.“ Crunchbase. <https://www.crunchbase.com/organization/xkool-technology> (accessed Dec. 4, 2022)
- [104] XKool. <https://www.xkool.ai> (accessed Dec. 4, 2022)
- [105] 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.
- [106] „Let’s build from here.“ Github. <https://github.com> (accessed Mar. 24, 2023)
- [107] „Software. Faster.“ GitLab. <https://gitlab.com/> (accessed Mar. 24, 2023)
- [108] „What is PatternForge?“ PatternForge. <https://patternforge.org> (accessed Mar. 24, 2023)
- [109] „Josef Sivic získal prestižní ERC Advanced Grant na výzkum inteligentních počítačových systému 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)
- [110] Tesla, N., „Nikola Tesla.“ GoodReads. <https://www.goodreads.com/quotes/446852-i-don-t-care-that-they-stole-my-idea> (accessed Mar. 28, 2023)
- [111] Sourek, M., „Sustainability or Resilience? The Story Sketch of a Problem that Waits on the Sidelines.“ Modern Civil Engineering. Academic Star Publishing Company. Vol. 2, No. 2, July 2022. pp. 25-29. ISSN 2470-4180. doi: 10.15341/mese(2470-4180)/02.02.2022/004 <http://www.academicstar.us/issuelist.asp?ArtID=811&issid=809> (accessed Apr. 19, 2023)
- [112] Hak, T., Kovanda, J., Weinzettel, J., „A method to assess the relevance of sustainability indicators: Application to the indicator set of the Czech Republic’s Sustainable Development Strategy.“

Ecological Indicators. Elsevier. No. 17, 2011, pp. 46-57. ISSN 1470-160X/\$. doi:10.1016/j.ecolind.2011.04.034

[113] Sourek, M., „Industrialize the Architectural Profession at Last.“ *International Journal of Architecture and Planning*. SvedbergOpen. Vol. 2, Issue 2, September 2022. pp. 46-55. ISSN 2788-5046. doi: 10.51483/IJARP.2.2.2022.46-55. <https://doi.org/10.51483/IJARP.2.2.2022.46-55> (accessed Apr. 19, 2023)

[114] Berardi, F., „What does Cognitariat Mean? Work, Desire, and Depression.“ *Cultural Studies Review*. vol. 11 No. 2. Sept. 2005. <http://epress.lib.uts.edu.au/journals/index.php/csri/index> (accessed Nov. 14, 2022)

[115] Conklin, E., „Section of Architectural Workers sets up union efforts in London.“ *The Architect's Newspaper*. Nov. 4, 2019. <https://www.archpaper.com/2019/11/saw-union-london> (accessed Nov. 14, 2022)

[116] Deamer, P., „Architecture and Labor.“ Routledge 2020. ISBN 9780367343507.

[117] Easterling, K., „Subtraction.“ *Critical Spatial Practice* 4. edited by HIRSCH, N., and MIESEN, M. Sternebrg Presss 2014. ISBN 978-3-95679-046-1.

[118] Stouhi, D., „Can Architects Finally have a Seat at the Table? Labor Rights and Working Conditions in Architecture.“ *ArchDaily*. Dec. 29, 2021. https://www.archdaily.com/974293/can-architects-finally-have-a-seat-at-the-table-labor-rights-and-work-conditions-in-architecture?ad_source=search&ad_medium=projects_tab&ad_source=search&ad_medium=search_result_all (accessed Nov, 14. 2022)

[119] Wagner, K., „Architects are Workers.“ *The Architect's Newspaper*. Jan. 11, 2021. <https://www.archpaper.com/2021/01/architecture-and-labor-review/> (accessed Nov, 14. 2022)

[120] Toynbee, A., "Lectures on the Industrial Revolution in England." Cambridge University Press, London. 2011. ISBN 9781139094566.