Comparative Canvas

Multi-scale architectural data analytics using open-source AI

Seyran Khademi[1], Fatemeh Mostafavi[2], Casper van Engelenburg[3], Pablo G. Morato[4],Georg Vrachliotis[5]
[1,2,3,4,5]Technical University of Delft
[1,2,3,4,5]{s.khademi|f.mostafavi|C.C.J.vanEngelenburg|P.G.MoratoDominguez|g.vrachliotis}@tudelft.nl

This paper presents an educational framework, integrating artificial intelligence (AI) into architectural education through multi-scale spatial analysis. The framework emphasizes the use of open-source AI models, empowering students to analyze and visualize architectural data at the building, neighborhood, and urban scale. Students engage in collaborative data curation using real-world Dutch architecture projects, compiling floor plans, site photographs, and aerial imagery from open databases and APIs. Using AI models, the students can uncover design patterns across large datasets that are otherwise impossible to interpret manually. This method fosters critical understanding of both data and AI tools, while introducing students to scalable workflows for precedent analysis and spatial storytelling. The course has two deliverables: a poster showcasing multi-scale visualizations and a blog reflecting on the workflows and findings. Together, the outputs demonstrate how AI can expand the possibilities of architectural interpretation and communication. Datasets and in-house developed models will be released.

Keywords: *design education, visual analytics, artificial intelligence, open-source AI.*

INTRODUCTION

Historically, architects manually studied existing structures to inform their designs – a time-consuming process and limited by human capacity to analyze complex data. In the past decades students have developed a digital search culture – routinely using search engines and online libraries to find precedents and inspiration. This facilitated quick access to case studies but also raised concerns about superficial browsing versus deep analysis. In response, curricula introduced computational design theory (parametric modeling) to ensure students engaged critically with digital tools rather than just searching and copying (Basarir, 2022; Karadağ and Bolca, 2018).

In this work, we use AI as a tool to explore architectural datasets to uncover design patterns. Beyond the popular approach in generating form and imagery with AI, we propose a workflow to analyze precedent design data, contributing to a culture of research-driven design. We observed that integrating AI knowledge and workflows into education equips students with novel storytelling tools and critical insight into algorithmic decision-making.

Across educational AI projects, a recurring theme is the mix of software licenses and platforms (Khean *et al.*, 2018). Some initiatives

use entirely open-source ecosystems – e.g. Stable Diffusion models (Çiçek *et al.*, 2023) custom Grasshopper plugins, Python notebooks (SORGUÇ *et al.*, 2022) which align with academia's ethos of openness and reproducibility. Open tools let students peek under the hood, e.g., tweaking the neural network parameters or training data. Proprietary AI services like OpenAI's DALL·E 2 (Ramesh *et al.*, 2022) or commercial software plugins for ease of use are more popular to conduct quick experiments in design studios (Ahmed and Alymani, 2024; Werker and Beneich, 2023). However, they function as black boxes and may raise issues of cost or accessibility.

We used only open-source AI models within a Python environment. To support students with no coding background, the teaching team developed custom packages to simplify their use. Students also curated and cleaned project data, building a transparent dataset to apply AI workflows effectively.

The material for the elective course "AI in Architectural Design" is revisited here. The focus is on pedagogical approaches to engage with technical aspects of AI workflows as part of their research practice. The student progress on meeting the learning objectives (LOs), is measured through presentation, code, textual and visual outputs. The learning activities include lectures, workshops, and tutorials, with a strong emphasis on hands-on Python exercises to equip students with practical skills for applying data and AI in their research workflows.

RELATED WORK

The experimental integration of computers into architecture and design research took root in the early 1960s, with innovative breakthroughs such as Sketchpad (Sutherland, 1963) and provocative concepts of "architecture machines" (Negroponte, 1970, Vrachliotis, 2022). By the 1980s, academia recognized a growing "information explosion" – the 1987 eCAADe conference was tellingly themed *"Architectural Education and the Information Explosion"* sharing a growing concern about digital shift (eCAADe, 1987).

By the 1990s, the adoption of 2D and 3D CAD tools (e.g., AutoCAD) allowed students to explore precision drafting and basic component libraries (Lynn, 1993). While these tools introduced digital design literacy, they remained largely procedural (Carpo, 2017). The 2000s marked a shift toward parametric design with tools like Grasshopper for Rhino, enabling students and practitioners to explore algorithmic iterative form generation, and a better integration with digital fabrication processes (Vrachliotis, 2022).

In the 2010s, machine learning and data-driven processes re-entered the architectural discourse—not as speculative experiments, but as increasingly tangible and practical tools shaping architectural design workflows, spatial analysis or design decision-making. Design education started to move beyond formal exploration, shifting towards performance- and knowledge-based design processes, that integrated environmental data, user behavior, and databases of historical precedents (Carpo, 2017). Computational design conferences such as eCAADe and CAAD Futures *(*Turrin *et al., 2023*) increasingly featured design experiments with machine learning, optimization algorithms, and visual analytics tools, bridging studio practice and research innovation.

The topic of education in architecture has been revisited as "Data-Driven Intelligence" in eCAADe 2024 (Kontovourkis *et al.*, 2024). The special education track explored how educational frameworks are evolving to incorporate computational tools and digital fabrication techniques, aiming to equip

students with the necessary skills to navigate the data-driven landscape.

In the era of AI, the core concern remains: How do we train architects to critically navigate and use ever-expanding body of architectural data and digital tools to design better buildings and cities? In response, recent years have been seen educators adopt a range of strategies to better integrate AI tools in the architectural design curriculum, shaping new forms of design literacy and experimentation. Examples range from exploring the relationship between text-to-image generation and environmental design (Doumpioti and Huang, 2023) to image generation for visualizing future landscapes (Mugita *et al.*, 2023).

Besides variations in tools, the stage of design in which the AI-based approach is used, plays a role (Dortheimer *et al.*, 2023). A conventional design method was implemented by (Ahmed and Alymani, 2024), followed by using an AI rendering tool to assist students regarding material choice and exterior opening placements. Interestingly, (Agkathidis et al., 2024) compared outcomes of generative tools (Midjourney, DALL·E 2, Stable Diffusion) with those from conventional methods in parallel studios. Results showed that GenAI tools boosted inspiration and speed but risked reduced design control, authenticity, and coherence between form and function. In a similar study, (Hermund *et al.*, 2024) aimed to explore how PhD students analyze and interpret architectural images driven from both traditional (i.e., 3D model making and photography) and AI-driven methods (i.e., text to image generation using Midjourney, Stable Diffusion, Prome AI, Veras AI, Planfinder XYZ, Sora, Adobe Firefly, Craiyon, DALL-E, and Krea.ai). Student reflections revealed challenges in translating architectural qualities and spatial emotions into text prompts. Similar to (Çiçek *et al.*, 2023; Cudzik and Radziszewski, 2018; Sorguç *et al.*, 2022), our approach transcends mainstream generative AI and proprietary tools, introducing AI as a transformative instrument to convert diverse architectural data modalities into compact, rich representations. These representations facilitate comparative studies and enhance visualization techniques. We emphasize AI's role in data interpretation for deeper, analytical design research beyond form generation.

COURSE METHODOLOGY

The 10-week, 5 ECTS elective with 16 students was structured around four pedagogical modules that guided students from foundational AI literacy to applied architectural analysis. Each module (M) corresponds to a learning objective (LO), designed in alignment with Bloom's taxonomy (Engelhart *et al.*, 1956) to span from knowledge acquisition to analysis:

- M1: Knowledge – LO1: Describe the science behind AI
- M2: Comprehension – LO2: Recognize the complexity of architectural data
- M3: Application – LO3: Operate data and AI workflows using Python
- M4: Analysis and Synthesis – LO4: Compare multi-scale architectural data

The course was structured around a mix of lectures, tutorials, and workshops, with each session blending these elements into a cohesive learning topic. The foundation of the course was built on four real-world architectural projects which included access to detailed project data such as IFC files. Students worked in groups of four and selected one of these projects as their case study for the duration of the course. Lectures covered the fundamentals of AI to address the LO1 and tutorials were delivered as one-hour executable notebooks on Google Collab. These were run live with the students in class, guiding them step-by-step through code blocks. This

allowed students to experiment directly with programming frameworks and AI workflows introduced during lectures. Workshops combined short lectures with hands-on coding exercises and included a project site visit. During the workshops, AI and data-driven techniques were directly related to architectural measures and design qualities. Students also received specific instructions for curating datasets relevant to their chosen project, bridging computational tools with spatial and contextual analysis.

In the first six weeks, students completed weekly individual coding assignments (CA) that built core skills in Python programming, data handling, and AI workflows. These assignments were designed to support LO2-3 and contributed 60% to the final grade. In week 7, students worked in small groups to present and critically discuss a scientific paper on AI, targeting LO1-2. These activities laid the foundation for the final two deliverables: a visual *Comparative Canvas poster* and a 2,000-word *blog post*. Both were based on multi-modal data analysis of a selected architectural project across building, neighborhood, and urban scales.

The assessment weights were carefully aligned with the intended learning outcomes (Table 1). We calculated the constructive alignment score by comparing the actual assessment weights to the intended targets. The result was a perfect alignment score of 1.0, confirming that the course was designed to match assessment with LOs.

Table 1
The distribution of summative assessments and learning objectives

	CA (60%)	seminar (10%)	poster (20%)	blog (10%)	sum
LO1	5	10	-		15
LO2	10	-	5	5	20
LO3	40	-	5	-	45
LO4	5	-	10	5	20

To further evaluate the learning impact, students completed a self-reflection survey at the end of the course. Their open responses were analyzed using word frequency analysis. The most frequently mentioned terms were "python" (15.9%), "use" (9.5%), "data" (7.9%), "learning" (6.3%), and "new" (4.8%). This response pattern closely reflects the course emphasis on data workflows and technical fluency, supporting the effectiveness of this particularly applied approach. Many students also mentioned gaining insight into AI applications in architecture and reading scientific papers, confirming the success of the knowledge and comprehension modules. This integrated teaching approach—combining modular content, aligned assessment, and reflective feedback—demonstrated how AI and data visualization can be successfully introduced into an architectural curriculum.

COURSE MATERIAL

Data is an indispensable part of any AI workflow. Instead of providing the students with clean and ready-to-use data, they have extensively worked on curating their own architectural datasets. This hands-on approach provided insight into how different data modalities—such as images, graphs, and tabular data—can represent architectural information. Specifically, students learned how to extract architectural floorplans from IFC files and process them into clean image and graph representations for visualization, how to access public aerial and street view imagery through APIs to automate the data collection process, and how to conduct manual data collection during on-site photography and annotate the captured data effectively. The students worked with different types of data to represent the project in three architectural scales: floorplans pertaining to the building scale; on-site photographs of the site to capture to the surrounding; and aerial-view imagery relevant to the neighborhood and urban scale.

The KAAN dataset includes architectural data corresponding to apartment units, hotel

Figure 1
Data acquisition process using Python

rooms, and office spaces related to four building projects across the Netherlands. The data is sourced from commercial clients of KAAN Architecten (KAAN Architecten, 2023) architectural firm. The project's data is converted to two CSV (comma-separated value) files mainly containing geometry, IFC classes, and material properties of architectural elements. The geometrical data resulted from applying BatchPlan (Yildiz *et al.*, 2024), a large-scale solution for floorplan extraction on the IFC files of the building projects. The geometry of the elements is provided in WKT (well-known text) format, representing each element by a multi-polygon. Alongside the name and IFC type of each element, geometry is presented in a CSV file per floor in each project.

Floorplans

In the first three tutorials of the course, the students learned how to convert IFC files into machine-readable floor plans. They learned how to represent a floorplan a graph, where nodes represent the rooms and edges capture the connections such as adjacency or accessibility; and, moreover, how the graph is leveraged effectively as a container for more information. Specifically, they learned how to populate the nodes with semantic (e.g., a room type such as "living room") and geometric (i.e., the shape of the room in terms of e.g. a polygon). In code, the tutorials showed how to acquire the necessary information from the IFC files, and how to plot the floorplans and their corresponding topological structure (i.e., the graph stripped of its features). Figure 1 gives an overview of the data and processing. To compare the projects' floorplan data to other (residential) floorplan datasets, the students were familiarized with two other large-scale and well-known floorplan datasets: RPLAN (Wu *et al.*, 2019) and MSD (van Engelenburg *et al.*, 2025). These datasets were pre-processed to match the graph-based format used in the course, ensuring consistency in features and attributes across different data sources.

Photo imagery

By leveraging open databases, students cross-referenced building footprints provided by the Dutch Cadaster (Kadaster, 2025) with aerial imagery available from PDOK (PDOK, 2025) to automatically collect aerial images. For each building, they defined a rectangular bounding box encompassing the footprint plus a margin,

thereby capturing both the structure's boundaries and its immediate context.
This bounding box was subsequently used to specify the geographical extent of the aerial imagery to be retrieved. Following this process, the students curated a gallery of building aerial images at an urban scale. In collaboration with Cyclomedia (Cyclomedia, 2025) which is a provider of street-level geospatial imagery, students supplemented the aerial imagery with street-view images, also retrieved based on each building's geolocation. Moreover, students were guided to take images from the project site by dividing different orientations from the building and record the status of the project surroundings. This is for ground truthing the online images with reality.

AI workflows

As the visual datasets suggest the focus of the course is on integrating computer vision as prominent AI modality to analyze the curated datasets from the projects at different architectural scales. Students were free to come up with their own topic / question they wanted to focus – such as materiality, sustainability and urban context. As depicted in Figure 2, the workflow consists of three steps: 1) passing the data through a proper AI model, e.g., DINO v2 for photo imagery, to transform high-dimensional data into compact representations. 2) dimensionality reduction to compress each data sample into a two-dimensional coordinated to be visualized in a 2D canvas. 3) Alternatively, the data samples can be organized within a metric space such as Euclidean space, where samples most similar to the query image are ranked higher than less similar ones.

All tutorials are developed in *Google Collaboratory*, or Collab – a free web-based programming environment. Collab is an online notebook that combines text and code, making it a simple way to start learning Python without needing to install any software, since it runs on Google's servers. The Python libraries (all open-) that the students got familiarized with include *Pytorch*, *Matplotlib*, and *Pandas*. We used two AI models to analyze the data: DINOv2 (Oquab *et al.*, 2023) for analyzing the photographs and street-view and aerial-view imagery, and an in-house trained graph-based model on floor plans.

Since there do not yet exist floorplan foundation' models that can reasonably well compress a diverse set of floor plans into

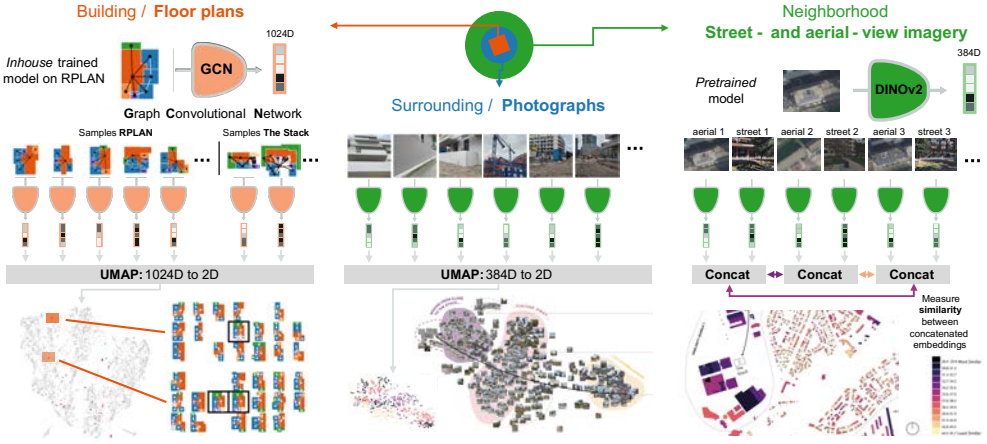

Figure 2
Data processing using AI.
The UMAP visualizations and urban similarity maps are produced by the student group: Jan Zawadzki, Lilli Selcho, Liren Chu, and Pranav Kanhangad Kakkanath.

meaningful embeddings, the teaching team developed an in-house AI model for the course. The training objective for the graph-based AI model is to learn spatial similarity in floor plans. To do so, the model was trained to penalize relative distances between the embeddings of an anchor floorplan and that of a similar one (positive) and the embeddings of the same anchor and that of least similar one (negative). Similarity labels were obtained by using the algorithm given in (Patil *et al.*, 2020), and the dataset of RPLAN (Wu *et al.*, 2019) was used to train the model.

The students were given the pre-trained model to convert the set of floor plans graph into embeddings, which were subsequently concatenated with the ones from RPLAN. Since the dimensionality of the embeddings of the pre-trained model is relatively large (D = 1024), *UMAP* (McInnes *et al.*, 2018) was used to project the high-dimensional embeddings into two dimensions (2D), making it possible to visualize the combined floorplan dataset on a 2D canvas. Based on the dataset visualization, the students analyzed their project's data: whether the model generalizes to their data and how it related to the other floorplan data.

COMPARATIVE CANVAS

Once the students learned how to collect, clean, curate the project data, and how to use Python and Collab as an interface to access opensource AI models, the last module of the course, i.e., compare and examine multi-scale architectural data has started. The medium given to the students was an A0 vertical poster to visualize the project data using AI tools they learned. Moreover, they had to cover the three architectural scales in one canvas.

A key component of the workflow was the comparative analysis of floor plans (building scale), which allowed students to evaluate and interpret architectural quality and diversity across projects and datasets. In the analysis, spatial connectivity within floor plans, identifying similarities and dissimilarities in building footprints and layouts. This approach highlighted recurring spatial configurations and design motifs of each project that might be overlooked in manual analyses. In addition, clusters of design elements and spatial arrangements appeared in the visualization, providing insights into the building's architectural coherence and its relationship to broader design typologies.

To compare buildings based on the facade/envelope or neighborhood similarity, students used street/areal imagery. The search implemented by computing the distance between embeddings of all buildings in the gallery and those of a query building, using DINOv2 in the shared embedding space. Specifically, aerial and street-view image embeddings were concatenated, and the Euclidean distance between the query and gallery embeddings served as the similarity ranking score. Since embeddings and similarity scores are geolocated, they can be visualized in open maps, thereby providing urban visual insights. The concept of search is a rhetorical in architectural design which has been revisited in our course using foundation models enabling understanding of urban morphology, typological analysis, and spatial reasoning. By utilizing DINOv2's image embeddings derived from aerial and street-view imagery, students systematically identify visual similarities and morphological patterns across urban regions. This approach encourages critical reflection on building typologies and the coherence of the urban fabric, enabling students to identify underlying relationships such as the spatial clustering of architectural styles, consistent building forms within neighborhoods, or the presence of unique contextual outliers.

Furthermore, spatial visualization skills are developed as students learn to communicate complex spatial patterns effectively through

similarity maps and heatmaps. This process also facilitates deeper architectural inquiry, empowering students to pose and answer design-related questions grounded in real-world urban contexts. For instance, they may examine the coherence of urban neighborhoods, explore the distribution and evolution of architectural typologies, and critically evaluate contextual compatibility across diverse urban fabrics. Presenting their findings through visually typological grids, annotated imagery, and comparative diagrams further refines their analytical and communicative abilities.

CONCLUSION

This paper presented a data-driven pedagogical framework for integrating AI in architectural education, focusing on spatial comparative analysis. The method allowed students to engage with datasets far exceeding the cognitive capacity of manual analysis. Importantly, the course emphasized open-source workflows, allowing students to explore technical barriers and data limitations.

Students learned to use AI tools to analyze large datasets, including floor plans and images, and to craft data-driven narratives about building projects. Using AI-generated similarity maps, they compared spatial qualities across buildings at various scales—from individual structures to neighborhoods and cities. They identified patterns such as environmental similarities and contextual design trends by locating visually similar buildings across different urban areas. The projects also involved comparing local floor plans with a dataset of Asian residential buildings (RPLAN) to explore cross-cultural differences and similarities. AI-driven graph analytics, embeddings, and clustering were interpreted through architectural logic like connectivity, rhythm, and functionality. While the course primarily used publicly available floorplan dataset due, integrating more locally relevant data could enhance building scale spatial analysis. We believe future iterations should place greater emphasis on interpreting results, once students are more familiar with the AI workflows.

The KAAN dataset (link: https://github.com/datarefinerylab/kaan-dataset), along with tutorials and student work from the course (link: https://github.com/AiDAPT-A/AI_in_Architectural_Design_TUD), are all released in the spirit of open-source culture.

ACKNOWLEDGEMENTS

We gratefully acknowledge our data partners, KAAN Architecten and Cyclomedia**,** for generously sharing data that were indispensable to the development of this course. Their contributions made this work possible. We also thank the students who engaged passionately with the assignments and kindly agreed to make their outputs publicly available.

REFERENCES

Agkathidis, A., Song, Y. and Symeonidou, I. (2024) 'AI-Assisted Design: Utilising artificial intelligence as a generative form-finding tool in architectural design studio teaching', pp. 619–628.

Ahmed, A. and Alymani, A. (2024), "Integrating Artificial Intelligence Rendering Tools in Design Integrating AI as teaching methods in architectural education", Education and research in Computer Aided Architectural Design in Europe, pp. 629–638.

Basarir, L. (2022), "Modelling AI in Architectural Education", Gazi University Journal of Science, Vol. 35 No. 4, pp. 1260–1278.

Biggs, J. (1996), "Enhancing teaching through constructive alignment", *Higher*

Education, Kluwer Academic Publishers, Vol. 32 No. 3, pp. 347–364.

Turrin, M., Andriotis, C. and Rafiee, A. (2023), "INTERCONNECTIONS: Co-computing Beyond Boundaries." Vol. 1819, *Communications in Computer and Information Science*, *Computer-Aided Architectural Design*.

Carpo, Mario. (2017), "The second digital turn : design beyond intelligence", The MIT Press.

Çiçek, S., Turhan, G.D. and Özkar, M. (2023), "Reconsidering Design Pedagogy through Diffusion Models", Vol. 1, Education and research in Computer Aided Architectural Design in Europe, pp. 31–40.

Cudzik, J. and Radziszewski, K. (2018), "Artificial Intelligence Aided Architectural Design", *ECAADe Proceedings*, Vol. 1, Education and research in Computer Aided Architectural Design in Europe, pp. 77–84.

Cyclomedia. (2025), "Home page | Cyclomedia", available at: https://www.cyclomedia.com/en (accessed 13 June 2025).

Dortheimer, J., Schubert, G., Dalach, A., Brenner, L. and Martelaro, N. (2023), "Think AI-side the Box!: Exploring the Usability of Text-to-Image Generators for Architecture Students", Vol. 2, Education and research in Computer Aided Architectural Design in Europe, pp. 567–576.

Doumpioti, C. and Huang, J. (2023), "Text to Image to Data", Vol. 2, Education and research in Computer Aided Architectural Design in Europe, pp. 541–548.

eCAADe. (1987), "Past Conferences – eCAADe", available at: https://ecaade.org/conference/past/ (accessed 13 June 2025).

van Engelenburg, C., Mostafavi, F., Kuhn, E., Jeon, Y., Franzen, M., Standfest, M., van Gemert, J., *et al.* (2025), "MSD: A Benchmark Dataset for Floor Plan Generation of Building Complexes", Springer, Cham, pp. 60–75.

Engelhart, M.D., Furst, E.J. and Krathwohl, D.R. (1956), *Taxonomy of Educational Objectives The Classification of Educational Goals*, David McKay Company, INC., Michigan.

Hermund, A., Oxvig, H., Bang, J.S. and Danish Academy, R. (2024), "Embracing the Creative Nexus Integrating artificial intelligence, philosophy, and artistic discourse in architectural education", Education and research in Computer Aided Architectural Design in Europe, pp. 665–674.

KAAN Architecten. (2023), "KAAN Architecten - Netherlands based architectural firm", available at: https://kaanarchitecten.com/ (accessed 13 June 2025).

Kadaster. (2025), "About us - Kadaster.nl particulier", available at: https://www.kadaster.nl/about-us (accessed 13 June 2025).

Karadağ, D. and Bolca, P. (2018), "Computational Design Tools in Architectural Education", *The 5th International Conference on Architecture and Built Environment*, pp. 15–25.

Khean, N., Fabbri, A. and Haeusler, M.H. (2018), "Learning Machine Learning as an Architect, How to? - Presenting and evaluating a Grasshopper based platform to teach architecture students machine learning", Education and research in Computer Aided Architectural Design in Europe, Vol. 1, pp. 95–102.

Kontovourkis, O., Phocas, M.C. and Wurzer, G. (2024), *Data-Driven Intelligence*, Vol. 1,

eCAADe (Education and research in Computer Aided Architectural Design in Europe) and University of Cyprus, Department of Architecture, Cyprus.

Lynn, G. (1993), "Architectural Curvilinearity - The Folded, the Pliant and the Supple", *Architectural Design Folding in Architecture*, March, Vol. 63, pp. 8–15.

McInnes, L., Healy, J. and Melville, J. (2018), "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction", 9 February, available at: https://arxiv.org/abs/1802.03426v3 (accessed 13 June 2025).

Mugita, Y., Fukuda, T. and Yabuki, N. (2023), "Future Landscape Visualization by Generating Images Using a Diffusion Model and Instance Segmentation", Education and research in Computer Aided Architectural Design in Europe, pp. 549–558.

Negroponte, N. (1970), *The Architecture Machine: Toward a More Human Environment*, *The Architecture Machine*, The MIT Press, Cambridge, Massachusetts.

Oquab, M., Darcet, T., Moutakanni, T., Vo, H. V, Szafraniec, M., Khalidov, V., Fernandez, P., *et al.* (2023), "DINOv2: Learning Robust Visual Features without Supervision", 14 April, available at: https://arxiv.org/abs/2304.07193v2 (accessed 13 June 2025).

Patil, A.G., Li, M., Fisher, M., Savva, M. and Zhang, H. (2020), "LayoutGMN: Neural Graph Matching for Structural Layout Similarity", *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society, pp. 11043–11052.

PDOK. (2025), "Webservices - PDOK", available at: https://www.pdok.nl/Webservices (accessed 13 June 2025).

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. and Chen, M. (2022), "Hierarchical Text-Conditional Image Generation with CLIP Latents", 13 April, available at: https://arxiv.org/pdf/2204.06125 (accessed 13 June 2025).

Sorguç, A., Kruşa Yemişçioğlu, M. and Yetkin, O. (2022), "Demystifying machine learning for architecture students", Anadolu Universitesi Bilim ve Teknoloji Dergisi-A: Uygulamali Bilimler ve Muhendislik, Vol. 23, pp. 60–67.

Sutherland, I.E. (1963), 'Sketchpad: A Man-Machine Graphical Communication System", *AFIPS*, Association for Computing Machinery.

Vrachliotis, G. (2021) *Experimental Architecture for the Twenty-first Century, CD/ITKE 2010-2020,* Birkhäuser, Berlin, Boston, pp. 16–21.

Vrachliotis, G. (2022) *Architecture and Design in the Age of Cybernetics*, Birkhäuser, Berlin, Boston.

Werker, I. and Beneich, K. (2023) 'Open AI in the Design Process: To what extent can text-to-image software support future architects in the early design process?', Vol. 2, Education and research in Computer Aided Architectural Design in Europe, pp. 577–586.

Wu, W., Fu, X.M., Tang, R., Wang, Y., Qi, Y.H. and Liu, L. (2019) 'Data-driven interior plan generation for residential buildings', *ACM Transactions on Graphics*, Vol. 38, pp. 57–74.

Yildiz, B., Cuartero, J., Mostafavi, F. and Khademi, S. (2024), 'BatchPlan: A large scale solution for floor plan extraction', *CAADRIA Accelerated Design*, pp. 201–210.