

AI AND THE FUTURE OF ARCHITECTURAL EDUCATION IN THE UK

A REPORT BY THE SCOSA AI WORKING GROUP
SEPTEMBER 2025



scosa
The Standing Conference of Schools
of Architecture

CONTENTS

Defining AI in Architectural Education (p4)

AI Programmes Used in Architectural Education (p10)

Student & Staff AI Survey 2025 (p14)

The Ethics of AI Use (p24)

Case Studies of AI Use in Architectural Education (p34)

Case Studies of AI Use in Architectural Practice (p54)

Suggestions & Knowledge, Skills and Behaviour (p82)

Repository of AI Resources (p84)

Arka.Works



ARUP

AUTODESK

Foster + Partners

GRIMSHAW

Heatherwick studio



Editor and AI Working Group Lead
Des Fagan (Lancaster University/Grimshaw)

Editorial Support
Katie Phillips (Lancaster University)

Authors
Conor Black (Arup)
Des Fagan (Lancaster University/Grimshaw)
Derek Hales (University of Salford)
Tom Holbrton (UCL)
Nate Kolbe (London Metropolitan University)
Carlos Medel Vera (University of Liverpool)
Debra Pothier (Autodesk)
Keir Regan-Alexander (ARKA Works) Joanna Sabak (Heatherwick Studio)
Sherif Tarabishy (Foster + Partners)
Martha Tsigkari (Foster + Partners)
Tsung-Hsien Wang (University of Sheffield)
Pablo Zamorano (Heatherwick Studio)
Xingjian Zhao (University of Sheffield)

Working Group Contributors:
Jonathan Bassindale (UWE)
Steve Parnell (Newcastle University)
Tolu Onabolu (Newcastle University)
Ran Xiao (Planarific)



DEFINING AI IN THE EDUCATION OF AN ARCHITECT



DES FAGAN

PROFESSOR OF COMPUTATIONAL ARCHITECTURE
LANCASTER UNIVERSITY

Head of Architecture at Lancaster, Des's field of research interest is in Optimisation and Deep Learning (Artificial Intelligence) for Decision Support Systems in design.


DEFINING AI IN ARCHITECTURAL EDUCATION

Artificial Intelligence (AI) is playing an increasingly significant role in architectural education. This has prompted the formation of this SCOSAAI Working Group, to examine and evaluate how the technology is taught and assessed in our schools. Although AI can often be characterised as a neutral set of computational tools for prediction and generation, its growing role in education - particularly in the creative disciplines - demands critical attention. AI tools demand scrutiny as a catalyst for how we rethink knowledge production and creativity, and negotiate authorship and agency within the architectural design process [1,2].

A BRIEF HISTORY OF AI IN ARCHITECTURAL EDUCATION

Since the 1960s, architectural design studios have experimented with computationally-based expert systems and knowledge-based design software in one form or another - notably in structural optimisation, environmental performance and geometric optioneering [3]. From the adoption of Gehry's Digital Project to the rise of parametric environments such as Rhino and Grasshopper, architects and students have employed a wide range of customised scripts to generate design options and explore geometric complexity [4,5]


Today, however, generative AI tools are changing the landscape of design. These new tools, based on probabilistic inference, are trained on datasets composed of vast amounts of data, incorporating black box architectures using neural networks. [6,7]. When tasked with responding to an architectural brief, students are now able to generate convincing images, models,



Dio Ben Snell 2018



Machine Hallucinations Refik Anadol 2019



Helena Sarin - GAN Paintings 2017

diagrams, essays and code using simple natural language prompts, without knowing or understanding the machination or logic of the process involved. Easy access to AI-based tools leads us to a pedagogical dilemma - how do schools retain rigour, critical reflection and the creativity demanded by the discipline, when AI tools can produce an endless range of outputs that are fast, plausible, and persuasive? How can students engage with these tools to shape and influence the future of their profession? [8,9]

FRAMING THE CURRENT MOMENT

Whilst AI tools deliver powerful advantages, they also give us pause to rethink the principles of architectural education and explore our existing beliefs and values. Long-standing values such as process over product, critical iteration, site sensitivity and drawing as self-expression are now being challenged by tools that offer comprehensive visual and written outputs at the push of a button [10,11]. This report is written in response to that challenge - it benchmarks and investigates how AI is currently being used across UK architecture schools and practice, and how students, staff, and institutions are responding to its emergence. It asks the following foundational questions about the nature of architectural education in the context of AI augmentation:

Which AI tools are currently in use across UK architecture schools — and for what purposes? (Understanding AI Workflows pg. 6)

What do students and staff think about the rise of AI in architectural education? (Student & Staff AI Survey pg. 14)

How can architectural education respond to the ethical challenges posed by AI? (The Ethics of AI Use, pg.24)

How are schools currently integrating AI into teaching and learning? (Case Studies in Architectural Education, pg. 35)

What role is AI already playing in architectural practice - and how could education respond? (Case Studies – Architectural Practice, pg. 56)

What knowledge, skills, and behaviours will future architects need in an AI-enabled profession? (Key Suggestions pg. 82)

How can SCOSA support continued learning and responsible experimentation with AI in architectural education? (Repository of AI Resources, pg.84)


The following sections trace emerging practices, ethical concerns, student and staff perspectives, regulatory frameworks and curricular innovations. In doing so, our working group does not propose definitive answers, but seeks instead to benchmark current developments, chart the evolving landscape, and consider what forms of architectural education may be required in an AI-augmented future.

UNDERSTANDING AI WORKFLOWS IN ARCHITECTURE

Artificial Intelligence is often spoken of as if it were a single technology, but it encompasses a range of workflows that govern how data is collected, how models are trained, and how outputs flow into practice. For architectural education, understanding data collection, model choice and model training of an AI tool is critical - it allows students and staff to explore how AI works and how effective it is. This section introduces some critical terminology: datasets and data security - models, the wrappers to connect 'general' AI tools to design software, and the emergence of bespoke applications that are custom to the field of architecture.

DATASETS: THE RAW MATERIAL OF AI

All AI systems are trained on data. General tools like ChatGPT or Midjourney are typically trained using huge datasets scraped from the internet - text from published articles and billions of images including photos and drawings. The advantages of this are breadth, flexibility and accuracy; the disadvantages are that the model is not tailored to architectural knowledge, and the provenance of the datasets is unclear: do your datasets infringe on existing Intellectual Property (IP) ?



Data.tecture Ryoji Ikeda 2015

One of the most challenging aspects of adopting AI in education is navigating IP. While some tools market themselves as 'safe', the reality is that no generative model can yet be described as fully IP-compliant in a legal sense. Copyright law has not fully caught up with the extent to which training on copyrighted material constitutes infringement. For example, Adobe Firefly describes itself as a 'commercially safe' generative AI tool - trained only on Adobe Stock images, public domain content and openly licensed datasets - it is one of the few models that can reasonably be described as legally defensible.

A parallel situation exists for Large Language Models (LLMs) - most mainstream LLMs including ChatGPT, Gemini, Claude and Llama are trained on huge internet scrapes which may consist of copyrighted text. There are a small number of LLMs that have limited their training data to openly licensed or public domain material e.g. Cohere, Azure-BLOOM (BigScience) and OLMo (Allen Institute for AI, trained on Dolma dataset). Microsoft's Customer Copyright Commitment (CCC) for Copilot and Azure OpenAI offers to indemnify enterprise (subscription based) users against copyright claims arising from AI outputs - and comparable indemnities now exist from Google Cloud and AWS (and for paid tiers, Anthropic/OpenAI).

However, these contractual promises appear to represent legal support for defence from IP claims, rather than proactive limitation of model training on IP protected datasets. Schools should treat these IP indemnification guarantees with caution - and as a risk-sharing mechanism, not a guarantee of legal immunity. For architecture schools and students, Firefly and the above referenced subscription-based counterpart LLMs may constitute safer spaces for experimentation, but users should continue to use caution, transparency and give full attribution for use (see The Ethics of AI Use, pg.24) in this untested space.

"...THE REALITY IS THAT NO GENERATIVE MODEL CAN YET BE DESCRIBED AS FULLY IP-COMPLIANT IN A LEGAL SENSE. COPYRIGHT LAW HAS NOT FULLY CAUGHT UP WITH THE EXTENT TO WHICH TRAINING ON COPYRIGHTED MATERIAL CONSTITUTES INFRINGEMENT..."



CV Dazzle No.6 Adam Harvey 2020

DATA SECURITY - 'OPEN' VS 'CLOSED' TOOLS

A key consideration in the training and use of any AI tool, whether general or bespoke, is the risk of data breach in the context of university or practice policy. Increasingly, universities, practices, and clients are adopting data policies to catch up with nascent AI tool information distribution. Data breach policy often qualifies that upload of any confidential data or prompts to an open, cloud-based LLM or image generator constitutes data breach, and thus risk serious contractual consequences (businesses are concerned that the upload of data, e.g. profitability figures, changes to building specifications, drawings, competitor data etc. could consequently be used by the cloud-based model to further train it, thus risk exposing that data to the public).

As a consequence, more recent AI tools are being offered as 'closed' systems (sometimes referred to as private or enterprise systems) which do not expose a prompt or contextual information

of a query submitted to it, e.g. Azure OpenAI Private, AWS Private LLM Hosting and Anthropic Claude Enterprise. Most of these closed systems come with a subscription price (due to the cost of having to 'host' the model and its context on private, dedicated servers independent of a general cloud-based tool). By using private or enterprise AI tools, students and staff will be given more protection and greater freedom to explore AI prompts and contextual uploads in the knowledge that their data sharing may not breach organisational policy.

MODELS: HOW AI LEARNS FROM DATA

If datasets are raw material, models are the machinery that transforms these data into useable outputs. Students may never need to train a model themselves, but it is essential to understand how a model trains to try to illuminate the black box of input > output. Training a model means exposing an algorithm to huge numbers of examples (datasets) until it can recognise patterns and make predictions. Once trained, the model can generate new content and classify new inputs, based on what it has learned. ►

There are several key families of models that underpin AI tools now entering architectural education:

- **Machine Learning (ML):** The broadest category, covering algorithms that learn from data to make predictions. In architecture, this may mean for example a model that predicts building energy use from a set of floor plans.
- **Deep Learning (DL):** A subset of ML using Artificial Neural Networks (ANNs) with many invisible layers. These models excel at recognising complex patterns in unstructured data including images and text.
- **Convolutional Neural Networks (CNNs):** A subset of deep learning models good at working with images. CNNs have been used to recognise doors and windows in floor plans and to classify room types from architectural drawings.
- **Graph Neural Networks (GNNs):** Models work with data structured as nodes and connections – like architectural rooms linked by doors. GNNs are well suited to mapping spatial networks, visibility graphs, and patterns of movement.
- **Diffusion models:** The technology behind Midjourney, Stable Diffusion, and Adobe Firefly. These models create 'new' images by filtering noise into coherent visuals, making them effective for design visualisation.

Each type of model has strengths and weaknesses: CNNs preserve visual information, but not spatial reasoning; GNNs preserve spatial reasoning, but not stylistic information; diffusion models generate creative images, but are not easily constructable; and DL-based LLMs produce fluent text, but can hallucinate. The key message here is that the model shapes the outcome - understanding the model that sits beneath a tool helps us to interpret our results critically.



Machine Hallucinations Refik Anadol 2019

"...WRAPPERS ACT AS TRANSLATORS: PASSING INFORMATION FROM A DESIGN PROGRAMME TO THE AI, THEN RETURNING ITS RESPONSE BACK INTO THE WORKFLOW TO DEVELOP AN IDEA ITERATIVELY..."



Memory Place Es Devlin 2018

performance. Unlike 'general' tools, these applications are trained on bespoke architectural datasets and designed with disciplinary needs in mind. The advantage is that they can produce outputs more closely aligned to practice - constructible forms or accurate environmental performance feedback. Their limitation is in scope: they serve singular architecture tasks well but sometimes lack the breadth of cognition provided by 'general' AI systems.

WHY THIS MATTERS FOR ARCHITECTURAL EDUCATION

For architectural education, the distinction between and across these terminologies matters significantly. General tools can expand creativity and provide rapid visualisation, but they need to be handled critically for the sake of IP and data security: students must learn to question outputs and trace them back to their training data. Bespoke tools prepare students for the realities of practice but demand an understanding of how their datasets are built, maintained and hosted. All highlight the need for AI literacy - an ability to interrogate where information comes from, how it is structured, and what role we as designers continue to play.

In short, AI in architecture is best understood not as a single programme, but as a layered workflow: datasets provide the raw material; models train on the data which needs to remain secure; wrappers allow integration of models into existing workflows; and bespoke architecture-specific applications attempt to deliver discipline-relevant outputs.

ARCHITECTURE PROGRAMMES

Wrappers are interfaces that connect software packages together. They often take the form of lightweight scripts written in Python that 'wrap' a general AI model for use inside a design software. For example, a wrapper would be a short python script to embed an LLM like ChatGPT to be accessible directly inside Rhino, Revit, or AutoCAD so that the tool can interface with geometries and floorplans directly as 'live'. As such, wrappers act as translators: passing information from a design programme to the AI, then returning its response back into the workflow to develop an idea iteratively. Wrappers deliver accessibility and speed but inherit the limitations of the 'general' models they connect to.

BESPOKE APPLICATIONS: ARCHITECTURE-SPECIFIC AI TOOLS

Alongside wrappers, we are seeing the emergence of dedicated AI tools built specifically for architecture. These include applications such as XFigura, which can generate not only conceptual images but also 3D models from text prompts, and Autodesk Forma, which focuses on optimisation and predictive patterns of environmental

AI PROGRAMMES

COMMON PROGRAMMES USED IN ARCHITECTURAL EDUCATION



CHATGPT
OPENAI
TEXT BASED

ChatGPT is an AI chatbot that uses a deep-learning Large Language Model (LLM), enabling users to have entire 'conversations' with the AI. ChatGPT can also write essays, proof-read work for grammatical errors and generate simple images from text prompts, all for free (although a subscription version is available with greater functionality).

The information provided by the AI is not always right, however, therefore, to gain a full understanding of a topic, AI should not be a sole source of information. In order to counter this, users should be aware that ChatGPT will provide links to websites from which it takes information and these can be of use when looking for original sources.




QUILLBOT
LEARNEO
TEXT BASED

Quillbot is a tool designed specifically for editing written pieces of work, from emails to longer papers.

Tools provided include a paraphraser, grammar checker, AI detector, plagiarism checker, an AI humaniser and more.

These tools can be used to create more efficient pieces of writing, and can be useful to students worried about unintentional plagiarism.

The features which use AI to re-structure writing, however, may present a grey area concerning referencing, and how a student might reference a tool like this is not yet standardised.




MIDJOURNEY
MIDJOURNEY INC
IMAGE BASED

Creating images and short videos from text prompts, Midjourney is a tool that has many potential applications in architectural education and practice, although not specifically designed for one industry.

Midjourney is a company of 11 people - small when compared to companies such as OpenAI, although it has been around since 2022, like ChatGPT.

The service is accessible via subscription only, making it less accessible to students, although users can look around for free and explore the range of images created by others, with options to download content without a subscription.



RHINO + GRASSHOPPER
MCNEEL
GEOMETRY BASED

Rhino and its Grasshopper plug-ins offer an application of AI that builds on pre-existing software, therefore differs from other emerging AI programmes. An example use of this AI application is finding the most efficient energy transfer pathways in a structure, facilitating efficient material use in construction. The AI can also be used to explore human movement in buildings, experiment with building layout in plan and respond to a number of constraints.

Rhino is a software used in some architecture practices including major ones such as Zaha Hadid Architects. It could be used to increase efficiency and decrease the carbon footprint of the building industry.



RUNWAY ML
RUNWAY AI INC
IMAGE BASED

Runway ML is a photo and video editing software that uses AI to enhance and change existing imagery of videos based on a series of prompts inputted by the user.

The software gives creators a high degree of control over the outputs, with video features that include selecting specific subjects, adjusting camera angles and more.

The photo-editing elements allow image modification based on text prompts, as well as the classic text-to-image function.

A subscription is required for this software, but educational licences exist.



XFIGURA
X FIGURA IMAGE &
GEOMETRY BASED

XFigura is a collaborative, AI-native platform for architects and designers. It combines multiple leading AI models into a unified system which supports image, video and 3D model generation, and more.

This technology has the potential to broaden design-phase experimentation, as high quality visualisations of ideas in 2D and 3D could lead to more informed early decision making.


The 3D models that are produced by the software from text prompts can be uploaded onto Rhino, Revit, or Sketchup, making this an incredibly powerful tool for experimentation throughout the creative process.

A GROWING RANGE OF PROGRAMMES

The examples of AI shown on this page represent just a handful of the many software tools currently available. These are included to illustrate the breadth of applications, but they are not the only, or necessarily the best, options. In practice, the choice of programme depends on a range of factors including the intended use, the time required to train or operate the system,

and the costs associated with licensing or deployment. For some users, ease of access or speed of output may outweigh advanced capability, while for others investment in more complex tools makes sense. The pace of technological development also means that tools can quickly be replaced, with new platforms emerging just as others gain traction.

"...FOR SOME USERS, EASE OF ACCESS OR SPEED OF OUTPUT MAY OUTWEIGH ADVANCED CAPABILITY, WHILE FOR OTHERS, INVESTMENT IN MORE COMPLEX TOOLS MAKES SENSE. ..."



GEN AI LANDSCAPE AS OF EARLY SEP 2025



HIGH RATES OF CHANGE

The high rate of change within these and other software that integrate AI is making it difficult for users to stay up to date with the most current software. The evolving methods and applications of AI within both new and existing apps is entering a cycle of making old software redundant quickly, as new ones bring more advanced functionality.

The rate is so high that information in this visual will be outdated on publication. The volatility of the market for new AI programmes makes the reliability of established firms a potential asset, but new programmes and wrappers continue to pervade the space with new models and training to increase accuracy and reduce time.

STUDENT STAFF SURVEY

HOW ARE STUDENTS AND
STAFF USING AI NOW, AND
WHAT ARE THE EMERGING
TRENDS IN USE?

etween December 2024 and August 2025, the working group conducted a student and staff AI survey, and distributed it to all 65 SCOSA member schools. We received 77 total responses - 66 responses from students and 11 responses from staff, each answering different sets of questions. The survey was aligned with the RIBA AI in Practice survey with similarly framed questions, allowing us to draw comparisons between AI use in education and practice.

The survey was designed to benchmark views on AI tools ahead of their widespread adoption. We believe this is a critical moment to collect data, as both students and staff are experiencing the early stages of AI integration and can provide valuable insights by comparing their use and attitudes before and after incorporation into design processes.



Everyman Theatre, Liverpool, Haworth Tompkins 2014

STUDENT / STAFF DEMOGRAPHICS

WHO ARE THE STAFF AND STUDENTS WHO RESPONDED?

As AI evolves, so does its application in architectural education. Understanding student and staff perspectives on its current use will inform future studies and guide the direction of legislation.


STUDENT DEMOGRAPHICS

With n=77 responses from students and staff, this poll cannot provide results that are fully representative of UK architecture schools. The findings should therefore be interpreted with caution, and in light of the sample size limitations. Of the 66 students who responded, 23 were first year undergraduates. The majority of students who responded to the survey were on an undergraduate course, with just 6 students from masters courses responding. The distribution of responses across year groups (Demographics Table) is perhaps understandable in line with attrition and increasing workloads towards graduation.

STAFF DEMOGRAPHICS

A total of 11 staff responded to the survey. Interesting comparisons can be made between staff and student results, such as opinions on the implications of AI for employment, and the benefits or dangers of AI to education and the profession.

DEMOGRAPHICS



STUDENTS AND STAFF COMPARED

Students tend to take a more cautious, even pessimistic, view of AI. Many are uncertain about its impact on their future careers, and relatively few believe it will fundamentally change the field of architecture. This may reflect a wider pessimism about the industry itself.

Staff, by contrast, generally see AI as less of a threat and more as another tool in the lineage of CAD or BIM. At the same time, staff did anticipate it would reshape both architectural learning and practice, reflecting a more positive outlook overall.

Across both groups, knowledge of AI was relatively limited. Most participants described only basic or practical familiarity, with few claiming advanced expertise.

27%

OF STUDENTS DID NOT HAVE CLEAR POLICY ON THE USE OF AI PROVIDED BY THEIR COURSE


63%

OF STUDENTS BELIEVE THAT AI INCREASES THE RISK OF THEIR WORK BEING IMITATED


100%

OF STAFF WANTED TO LEARN MORE ABOUT AI TOOLS IN THEIR WORK

MY COURSE HAS CLEAR POLICY ON THE USE OF AI



HOW WOULD YOU RATE YOUR KNOWLEDGE ABOUT AI IN GENERAL



HOW CAN WE USE THESE DATA

AND WHAT WILL THE IMPACTS BE ON HOW WE TEACH?

The results of the survey give us an indication of the positive and negative associations that staff and students have with AI use. The survey shows where in the design process students use AI tools, informing guidance on where to target advice and use in future.



In 2024, the RIBA published the results of a similar survey - the RIBA AI Report - which asked similar, sometime identical questions to UK practices. For some questions and topics, the data from this survey can be compared to the results of the RIBA survey, in order to explore the differences in AI use and perception between practice and universities.

Topics surrounding ethics are particularly comparable: the survey asked students and staff about their level of concern over the adoption of AI in education regarding employment. 63% of students had significant concerns about the impact of AI on employment, whilst in practice 7% responded that AI had led to staff reductions.

69%
OF STUDENTS NEVER
USE AI TO GENERATE
IDEAS OR APPROACHES
FOR DESIGN STUDIO
WORK

32%
OF STUDENTS OFTEN
USE AI FOR GENERATING
CONTENT OR WORDS FOR
A WRITTEN ESSAY

ETHICAL CONCERNS OVER THE ADOPTION OF AI IN EDUCATION REGARDING EMPLOYMENT



HOW STRONGLY DO YOU AGREE OR DISAGREE THAT AI HAS LED TO STAFF REDUCTIONS




Data from RIBA AI Report 2024

HOW ARE STUDENTS USING AI

Students are using AI in all aspects of work. The most common uses are for finding references (53%) and structuring responses for written essays (51%). The least common use is for digital model making (15%), followed by checking building regulations (17%), searching for precedents (28%) and generating ideas and approaches for design studio work (40%).

These responses may reflect the current low accuracy of AI tools for regulatory compliance, as well as the limited distribution of licenses by institutional IT departments, which restricts student access to more technical, compliance-oriented software typically required in later RIBA work stages.

PLEASE INDICATE HOW MUCH YOU ARE USING AI IN THE FOLLOWING AREAS OF YOUR WORK



ETHICAL CONCERNS


WHAT RESPONDENTS WERE MOST CONCERNED ABOUT REGARDING AI

Question 12 of the survey asked participants to indicate their level of concern about AI . In all questions, the majority of answers indicate at least some concern. The majority of respondents had significant ethical concerns over the adoption of AI in specific areas of architectural education. The results (right) show high levels of concern from students and staff regarding plagiarism and bias.


Other ethical issues raised by the survey include whether work produced by generative AI should be marked differently to work produced by students without the assistance of AI. This question split staff, with similar numbers agreeing and disagreeing. Students, however, had strong opinions that work made using AI should be marked differently (60%).

There are some concerns over parity - with most students (61%) and staff (63%) agreeing that the technology has the potential to impact the parity of work if schools restrict the use of AI augmented work, whilst others do not. Despite these ethical concerns, respondents tended to remain positive about the overall impact of AI on the profession, including positive attitudes towards design innovation and creativity, areas that may be vulnerable to issues such as plagiarism and bias as a result of AI.


ETHICAL CONCERNS OVER THE ADOPTION OF AI IN ARCHITECTURAL EDUCATION REGARDING:




WORK PRODUCED BY GENERATIVE AI SHOULD BE MARKED DIFFERENTLY TO WORK PRODUCED BY STUDENTS



HOW CONCERNED ARE YOU WITH THE PARITY AND QUALITY OF YOUR WORK COMPARED TO OTHERS IN YOUR COHORT AND ACROSS SCHOOLS OF ARCHITECTURE WHO ARE USING AI?



OVERALL, DO YOU THINK THE EFFECTS OF AI WILL BE POSITIVE OR NEGATIVE?



ETHICS

THE ETHICS OF AI
IN
ARCHITECTURAL
EDUCATION

THE ETHICS OF AI IN ARCHITECTURAL EDUCATION



Architectural education is undergoing a rapid evolution as generative AI tools become increasingly accessible. This brings unique ethical challenges that demand a focused response from architecture schools and professional bodies [1, 2]. Educators and regulators are concerned with maintaining design authorship and academic integrity in an age where algorithms can generate images, models, and even design concepts [3, 4]. In architectural education, more so than in many other disciplines, the studio culture, emphasis on creative originality, and the visual nature of coursework mean that the ethical use of AI must be carefully defined and managed [5]. This section examines those challenges and outlines how UK architectural education is developing guidelines and strategies to ensure AI is used responsibly and transparently in design studios and beyond.

ETHICAL CHALLENGES IN ARCHITECTURAL EDUCATION

Design Authorship and Originality: One of the foremost concerns is the question of authorship when students use AI in design work. Architecture students traditionally build a portfolio that reflects their personal creative abilities; if a building form or rendering is produced by a generative AI, who is the author of that design? [6-8]. Educators worry that over-reliance on AI image or model generators could dilute a student's individual creative voice and make portfolios less indicative of the student's own skills [9, 10]. Students submitting AI-generated designs without disclosure might be misrepresenting work as original when it is partly algorithmic, raising clear ethical issues of misattribution.

PORTFOLIO AND IMAGE ORIGINALITY

Architectural coursework often involves visual deliverables: drawings, renderings, diagrams, physical and digital models. Generative AI now allows students to create striking images or



CARLOS MEDEL VERA

LECTURER IN ARCHITECTURAL
ENGINEERING & TECHNOLOGY
LIVERPOOL SCHOOL OF
ARCHITECTURE

even iterate design options with minimal effort [11-13]. While this presents opportunities for exploration, it also challenges the expectation that a student's portfolio reflects their own hand and thought process. If two students prompt the same AI tool in similar ways, they could end up with comparably styled images, undermining the originality of the portfolio. Questions also arise about intellectual property: AI tools may be trained on countless copyrighted images of buildings, so the outputs might inadvertently incorporate others' design ideas [14, 15]. Using such outputs without caution could conflict with architects' ethical obligations to respect others' work. From an educational standpoint, there is a fine line between using AI as a creative collaborator and outsourcing one's creativity. Some respondents to a UK government consultation saw generative AI as a useful "creative collaborator" in design fields like architecture, with potential to generate ideas or visuals that students can build on [16].

AI IN THE DESIGN STUDIO ENVIRONMENT

The design studio lies at the heart of architectural training, a collaborative space where students develop projects through an iterative, tutor-guided process. The introduction of AI tools in this environment raises further ethical considerations. On one hand, AI tools can augment studio work: for instance, an algorithm might quickly generate form-finding options, environmental analyses, or construction details, allowing students to explore alternatives beyond their manual drafting capacity. On the other hand, if a student leans too heavily on AI to produce a scheme or solve a problem, are they bypassing the intended learning process? Studios prize the development of a concept through sketches, models, and critiques; an AI that jumps straight to polished outcomes might short-circuit this learning journey. Moreover, studio culture values academic honesty and peer learning. If some students secretly use AI to advance their projects, it could create an uneven playing field

"...INTEGRATE AI INTO THE PROCESS WITHOUT ABDICATING AUTHORSHIP OR ETHICAL RESPONSIBILITY..."

and erode trust during design reviews. Instructors must therefore clarify how AI may be used in studio, for example, as a brainstorming aid or for precedent inspiration, versus what constitutes unethical use (such as generating an entire design presentation with minimal student input). The core ethical challenge is to integrate AI as a tool in the studio without undermining the pedagogical emphasis on process, critical thinking, and original creation. Ultimately, AI should not replace the "human touch" in design education; rather, students should learn to critically evaluate AI-generated ideas and maintain responsible control over their design decisions, much as practising architects are expected to do.

ACADEMIC INTEGRITY AND AUTHORSHIP IN ARCHITECTURE COURSEWORK

Alongside creativity concerns, generative AI has prompted urgent questions about academic integrity in architecture programs. Universities define academic misconduct as any attempt to obtain credit for work that is not one's own, and this now explicitly includes misusing AI. A key ethical principle is that a student's submission, be it an essay, design project, or portfolio, must honestly represent their own efforts and sources. In the context of AI, this means students must not present AI-generated content as if they created it themselves. For example, the University of Edinburgh's guidance on AI makes clear that certain uses of generative AI "are not acceptable and constitute misconduct," including "presenting AI outputs as your own, original work" and submitting any AI-generated text, images, or designs "without acknowledgment" [17]. In architecture coursework, this principle covers everything from using ChatGPT to write a theoretical essay to using DALL-E or Midjourney to create a rendering or diagram. If such material is included, failing to acknowledge it is a risk of plagiarism; the student would be, in effect, borrowing ideas or visuals produced by others (in this case an AI trained on human-created data) and misrepresenting them as personal work. ►

The concept of plagiarism traditionally applies to copying another author's words or ideas; with generative AI, the situation is nuanced because the AI output is new but derivative of training data. Many academics argue that uncredited AI-generated work is indeed plagiarism. As University College London (UCL) has cautioned its students, words and ideas from GenAI tools are built on other human authors' work, so presenting AI-generated text as one's own "is considered... a form of plagiarism" [18]. The same logic extends to image-based work; an AI-crafted design might synthesise styles or elements it learned from human designs, so passing it off without attribution breaches academic integrity. In short, academic honesty requires transparency about any assistance a student has received, whether from a human, an online source, or an AI system.

Maintaining integrity is especially critical in project-based and visual architectural assessments, where traditional plagiarism detection (designed for text) is insufficient. The emphasis is on authorship: the student should be the author of the creative work they submit, or otherwise clearly credit all contributors. By articulating these boundaries, architecture programmes aim to uphold the integrity of design coursework. Students, for their part, are expected to exercise judgment and honesty; using AI tools in permitted, support-oriented ways, but ensuring the final design and narrative remain substantially their own work. If there is any doubt, the safest course is always to ask instructors for clarification and, when in doubt, acknowledge the use of AI to avoid inadvertent misconduct.

POLICY FRAMEWORKS: UK GOVERNMENT, RIBA AND ARB POSITIONS

At a national level, the UK is actively developing frameworks to address AI in education. In 2023 the Department for Education (DfE) held a call for evidence on generative AI in education, gathering sector-wide views [16]. The outcome reinforced that while AI offers opportunities to enhance learning, it also presents risks to academic standards. The UK government's stance has been to neither ban nor uncritically embrace AI in higher education, but instead to issue guidance balancing innovation with safeguards. For instance, the UK's Quality Assurance Agency for Higher Education

(QAA) has highlighted the importance of authentic assessment methods that ensure a student's work can be trusted as their own, even as AI tools proliferate [19]. There is an acknowledgement at policy level that if generative AI is widely available (even integrated into common software), universities must adapt, redesigning assessments and teaching students how to use AI responsibly, rather than relying on policing alone. Crucially, the DfE and quality bodies like the QAA have also cautioned against over-reliance on AI detection technology. They note that current AI-detection tools are "unreliable at best" and can even flag original work as AI-generated, so maintaining academic integrity will rely more on robust guidance and student engagement than on any quick technological fix [19].

PROFESSIONAL BODIES - RIBA AND ARB

The architectural profession's regulators are likewise formulating positions on AI ethics which influence education. The Royal Institute of British Architects (RIBA) and the Architects Registration Board (ARB) both uphold honesty, integrity, and competence as fundamental principles in practice, and these naturally extend to how one uses new tools. While as of 2025 neither body has a dedicated AI-in-education policy, their general ethical codes set expectations. ARB's code of conduct (2025 revision) emphasises that architects must "always meet" standards of honesty and integrity in their work [20]. This implies that misrepresenting someone else's work (or an AI's work) as one's own would violate professional ethics; a standard that architecture students should train for during their education. We can anticipate that as AI becomes integral to practice, RIBA and ARB will expect accredited schools to educate students about the ethical and responsible use of such technologies. Indeed, the RIBA has already begun surveying how architects use AI and discussing the need for guidelines. RIBA's 2024 AI Report noted rapid adoption of AI in practice and included expert commentary on ethics and professional standards in an AI-driven design process [21]. The message from the profession is that AI is a powerful aid, but the architect remains accountable for all work under their name, regardless of what tools are used. We see international alignment on this principle: for example, the National Council of Architectural

"...THE UK'S QUALITY ASSURANCE AGENCY FOR HIGHER EDUCATION (QAA) HAS HIGHLIGHTED THE IMPORTANCE OF AUTHENTIC ASSESSMENT METHODS THAT ENSURE A STUDENT'S WORK CAN BE TRUSTED AS THEIR OWN, EVEN AS AI TOOLS PROLIFERATE.."

4.5 MIDJOURNEY EXPLORATION

A collage of Midjourney-generated images, created using reference images from my development drawings, site photographs, and text prompts, as well as by combining and regenerating these results. This collage serves as inspiration for urban strategy visualisation.



Reference image of site

Text prompts:

Biomimetic hillside stairway, Carlo Scarpa-inspired, precise stone detailing, natural textures, photorealistic --v 6.0 --ar 16:9 --iw 2
Sculpted hillside ramp, Zaha Hadid-inspired, fluid organic lines, natural materials, photorealistic --v 6.0 --ar 16:9 --iw 2
Layered stone pathway, subtle geometric forms, medieval context, photorealistic --v 6.0 --ar 16:9 --iw 2
Organic hillside steps, flowing forms, embedded in landscape, photorealistic --v 6.0 --ar 16:9 --iw 2
Terraced hillside access, Scarpa and Hadid fusion, stone and glass, dramatic elevation, photorealistic --v 6.0 --ar 16:9 --iw 2
Stepped hillside walkway, natural stone, lush vegetation, photorealistic --v 6.0 --ar 16:9 --iw 2
Vertical circulation on slope, sweeping curves, biomimetic integration, photorealistic --v 6.0 --ar 16:9 --iw 2
Medieval-inspired hillside path, rough-hewn stone, integrated planting, photorealistic --v 6.0 --ar 16:9 --iw 2



Example of academic integrity in architectural design coursework. Midjourney exploration with full disclosure demonstrates transparent authorship: the student documents image seed, prompts, and parameters that are recorded alongside the evolving design to evidence authorship and critical judgment (Image courtesy of Louis Shepley, third year BA Architecture student at the University of Liverpool 2024-25)

Registration Boards (NCARB) in the U.S. stated in 2024 that “AI is a tool – it is not a replacement for professional judgement,” and that licensed architects must remain in responsible control of all outputs, disclosing AI use where relevant [22]. In essence, tomorrow’s architects (and by extension today’s students) are expected to integrate AI into their process without abdicating authorship or ethical responsibility. Architectural educators should be therefore incorporating these values, ensuring that students learn to treat AI as a support tool under their control, rather than an autonomous creator of content.

UK HIGHER EDUCATION SECTOR PRINCIPLES

In addition to professional codes, universities themselves are combining around shared principles for AI in education. The Russell Group, representing 24 leading UK universities, released a set of principles in 2023 to guide the use of generative AI in teaching and learning. These principles include commitments to “adapt teaching and assessment to incorporate the ethical use of generative AI and support equal access,” and to “ensure academic rigour and integrity is upheld” in the AI era [23]. Collaboration with national bodies like QAA and Jisc is noted, as well as the need for consistency across disciplines in how AI is handled. For architecture schools, which often belong to larger universities, such high-level principles provide a framework: embrace AI’s potential to enhance education, but do so by embedding ethical use, transparency, and integrity into the curriculum. Notably, the principles encourage close coordination with professional bodies and employers [23] which for architecture means aligning with expectations from RIBA, ARB, and leading practices about the skills and ethics graduates should have. In summary, the emerging national framework in the UK supports a positive yet cautious integration of AI: architecture students should become AI-literate and prepared for an AI-augmented profession, but they must also be held to rigorous standards of originality, proper attribution, and critical evaluation of AI outputs.

REFERENCING AI-GENERATED DESIGN WORK


A practical dimension of AI ethics in architecture education is how to reference AI-generated material in submissions. Unlike a traditional essay where one might quote a source and cite it, architectural work may include AI contributions in varied forms: a Midjourney-generated conceptual collage, a block of text from

ChatGPT in a design statement, or even code from an AI assistant used in parametric design. Educators are therefore formulating conventions for acknowledging these contributions to uphold academic honesty.

The overarching rule is simple: whenever AI has contributed to the content of a student’s work, the student should clearly acknowledge it. The form of acknowledgement can vary. For textual content, some universities suggest an in-text citation or a footnote. The reference list would then include an entry for the AI tool. Some universities have issued citation formats: at Manchester, as noted, they treat the AI tool as software, so an entry might look like “OpenAI ChatGPT (2025) output generated 1 March 2025 via prompt XYZ” [24]. The key is to include enough information that the examiner understands what was AI-generated and by which tool or model.

For images and models, similar principles apply. Students are encouraged (and in some places required) to label images created or significantly edited by AI. A common approach is to include a caption like “Image generated using [Tool Name] with text prompt: ‘...’ (image credit: AI-generated by student).” Providing the prompt text is often encouraged because it documents the student’s creative input in guiding the AI. It’s worth noting that writing a good prompt can be a skill, and disclosing it not only shows honesty but also demonstrates the student’s process [25]. The University of Cambridge’s humanities faculty, for instance, has even provided a template declaration where students can fill in how they used AI in their work, a model that could be adopted in design fields as well [26]. In architecture, a student might include a short appendix or slide in their portfolio detailing any AI usage: e.g. “Page 10 rendering created with Midjourney (v5), prompt in Appendix; Plan optimization assisted by Finch3D AI plugin,” and so on.

Another aspect of referencing in design is the idea of “algorithmic transparency.” If an architecture student uses a generative algorithm (say, for form-finding), part of ethical best practice is to explain the algorithm’s role. For example, a student might write in their design rationale: “Using the GAN-generated massing options as a starting point, I selected and refined Scheme B.” Such a description acknowledges that the initial idea came from an AI suggestion rather than purely the student’s imagination. It contextualises the AI’s contribution as part of the design process narrative. This level of detail might be expected for higher-level students (Masters or PhDs projects), where understanding one’s process is crucial. In summary, the proper referencing of AI-generated material in architecture submissions mirrors the



Referencing AI-Generated Design Work. Long-listed entry for the 2025 RIBA J Eye Line drawing competition. The design concept and narrative were first developed through dialogue with ChatGPT, which assisted in shaping the story and synthesising the final Midjourney prompt used to generate the image. This example illustrates transparent authorship by acknowledging both the narrative co-development and the AI-assisted visualisation process.

traditional scholarly demand for citing sources but extended to non-human sources. It is about honesty in attribution. By citing prompts and AI tools, students demonstrate integrity and allow others (teachers, external examiners, future portfolio reviewers) to appreciate which parts of the work were human-crafted and which were AI-assisted. This practice not only guards against accusations of plagiarism, it also educates students in a professional transparency they will likely need in practice. (We can foresee, for instance, clients asking in the future, “Was this visualisation made by AI or by you?”, and an architect who has learned to openly document methods will be prepared to answer.)

LIMITATIONS OF AI DETECTION TOOLS IN DESIGN EDUCATION

In dealing with AI and academic integrity, some institutions initially looked to technological detection tools, software that claims to identify AI-generated text (and to a far lesser extent, images). One of the most prominent is Turnitin’s AI-writing detector, introduced in 2023, which many hoped could automatically flag AI-generated essays [27-29]. However, in the context of architecture and creative work, these AI-detection tools have serious limitations and risks.

Firstly, most detection tools are geared toward textual content (e.g., identifying GPT-style writing patterns). Architecture submissions, however, often centre on visuals and designs that these tools cannot analyse. A design project might include a brief written report, but the heart of the work (drawings, diagrams, renderings) would not be checked by Turnitin. There are emerging tools that attempt to detect AI-created images, but they are even less established and can often be defeated by simple modifications. In essence, for a large portion of architectural coursework, automated detection offers little to no solution.

Even where detection tools do apply (for instance, an essay in a history/theory module or the textual rationale in a design portfolio), universities have found their accuracy to be dubious. The QAA and other bodies reported a “high level of pushback” in the UK sector against using Turnitin’s AI detector due to reliability concerns [19]. These tools can generate false positives, wrongly labelling original student writing as machine-made which is a particularly grave issue for non-native English speakers or those who write in a clear, formulaic style that the software misreads [30]. Conversely, they can also produce false negatives, failing to catch instances where students did use AI but perhaps edited the text lightly. There is also the problem of adaptive behaviour: knowing that detectors exist, students who do intend to misuse AI

can try various evasive tactics (paraphrasing tools, feeding the AI its own output for rephrasing, etc.) to reduce detection scores [31].

Because of these issues, many architecture schools and universities have decided not to rely on AI detection as a primary strategy. For example, UCL explicitly states it does not use AI detectors for coursework [18], and the University of Sheffield has likewise ruled them out due to error rates [32]. Instead, if a tutor suspects a piece of work isn't the student's own, the preferred approach is a conversational one: discuss the work with the student, perhaps in an informal setting or a panel review, to gauge their understanding and the work's provenance. In design education, this approach fits well. During crits or reviews, students typically must explain their design concepts and process. If a student cannot articulately explain a design that an AI largely produced, that gap in understanding becomes evident.

In short, the consensus is that human judgment and intelligent assessment design are far more dependable than AI-detection software for preserving academic integrity in architecture. Turnitin's AI detector and its peers may still be used sparingly (perhaps as a flagged indicator, not as proof of guilt), but many architecture schools consider them an adjunct at best. The focus is shifting to prevention and education rather than detection. By setting clear expectations and engaging with students about their work, educators foster a culture where cheating with AI is understood as unacceptable and ultimately self-defeating for one's learning. This culture, combined with assignments that require personal input at multiple stages, is a more effective guardrail than any algorithmic tool. As the University of Sheffield guidance succinctly puts it, the university prefers to "support students and staff" in appropriate use of AI, rather than spy on them [32]. This trust-based approach, backed up by the ability to verify understanding in person, is particularly apt for the studio-based, mentorship-driven model of architectural education.


PROMOTING ETHICAL AI USE: STRATEGIES IN TEACHING AND CURRICULUM

Rather than simply policing AI use, architectural education seems to be evolving to actively teach students how to use AI ethically and creatively. There is a recognition that AI is here to stay in both academia and practice, so the goal is to prepare students to navigate this landscape with integrity. To that end, schools of architecture are implementing a range of strategies and curricular innovations:

1. Clear Studio and Coursework Policies: Educators are now writing explicit guidelines on AI use into project briefs and course handbooks. For each design studio or assignment, instructors might state whether AI tools are permitted and, if so, to what extent. For example, a studio brief may note: "Students may use AI image-generation to brainstorm early concepts, but all final visuals must be substantially developed by the student. Any AI contributions should be acknowledged in your presentation." By being upfront, teachers remove ambiguity and ensure all students play by the same rules [18]. This proactive communication empowers students to use AI appropriately rather than in secret. It also frames AI as one resource among many, to be used judiciously.

2. Academic Integrity Pledges and Declarations: A practical tool gaining traction is the use of integrity declarations related to AI. Some universities now require students to sign a statement (or include a paragraph in their submission) affirming that they have not used unauthorised assistance. With the rise of AI, these declarations often explicitly mention generative AI. For instance, a department might adapt its cover-sheet honesty statement to: "I confirm that this work is my own and that any use of generative AI or other tools has been disclosed and is within the allowed guidelines". The University of Cambridge's template for declaring AI use is a case in point: it prompts students to detail any permitted AI assistance at the start of an assignment [26]. In architecture, such declarations can be tailored: e.g. "No part of this design project was generated by AI, or if it was, it has been specifically noted and referenced". Knowing they must make a sworn statement, students may be psychologically deterred from cheating, and those who do use AI ethically have a formal place to note it. This strategy aligns with the professional world as well, where architects sign off on drawings as accurate and code-compliant; here, they learn early to sign off on the authenticity of their academic work.

3. AI Literacy and Ethics Education: Architecture programmes are beginning to include AI-focused content in their curriculum to ensure students understand both the capabilities and the ethical pitfalls of these tools. This might take the form of seminars, workshops, or modules on "AI in design." For example, a theory or professional practice class could have a unit on the ethical use of AI, covering topics like bias in generative design, IP concerns with training data, and case studies of AI successes and failures in architecture. The aim is to cultivate an informed perspective so that students are not blindly using tools without appreciating the broader implications. The Russell Group principle that "students ►



AI Literacy and Ethics Education. Panels a and b show an AI skill-building workshop and an AI drawing competition held in January 2025 at the University of Liverpool to foster responsible and ethical engagement with generative AI. Panel c presents a selection of student submissions, illustrating the creative outcomes achieved through transparent and critically informed use of AI tools.

and staff [become] AI-literate” is manifesting as training sessions and resources across universities [23]. Some architecture schools may even want to introduce interdisciplinary collaborations, for instance, inviting computer science or data ethics experts to hold a joint workshop with architecture students on generative AI. This exposes students to different perspectives and drives home the message that AI, like any technology, must be wielded responsibly.

4. Emphasising Process and Reflection: To ensure that students engage deeply with their work (and to make AI an aid to learning, not a crutch), educators are placing more weight on the design process and reflective practice. One strategy is to require process logs or design journals in which students document the evolution of their ideas, including any AI tools used at each stage and why. The University Manchester's teaching guidance suggests having students “submit drafts of larger pieces of work as they progress...and encourage reflection on how the work has evolved,” including explicitly “how AI has been used in their work and reflection on its usefulness” [24]. By grading or at least reviewing these reflections, instructors incentivise students to be thoughtful about AI: Did using Midjourney early on help generate alternatives? Did it introduce bias or lead the design astray? Such questions cultivate a metacognitive approach where students critically assess the value of AI in their creative process. Studios generally hold interim crits where students not only show their design progress but might also discuss what tools or methods they employed. In these discussions, ethical use can be reinforced, e.g. a tutor might praise a student for properly citing an AI-generated image, thereby setting a positive example.

5. Authentic and “AI-Resilient” Assessment Design: As a preventive measure, architecture educators are redesigning some assignments to be less susceptible to unethical AI use. While the nature of design projects already demands personal creativity, certain tweaks can further ensure authenticity. For example, tasks that involve site-specific or community-engaged design inherently require observations and inputs that AI would not have (like local context, unique client requirements), making it hard to delegate the work to a general AI. Another approach is collaborative or live components, e.g. in-class team exercises where students must produce work on the spot (thus unable to quietly use AI). The goal of all these strategies is not to trap students, but to design learning experiences that reward originality and personal input, thereby naturally disincentivising any shortcut through AI. In fact, many of these are just best practices in pedagogy (like focusing on

higher-order thinking and creativity), now being re-emphasised in light of AI. If students are tasked with, say, hand-sketching concepts in a live workshop, building a physical model, or writing a personal reflection on their design philosophy, these are outputs that showcase individual talent and cannot be generated by an algorithm. By diversifying assessment methods in this way, schools make it clear that a student's value is in their process, critical insight, and skills, which no AI can replicate in full [24].

6. Studio Culture and Peer Involvement: Ethics in architecture education is also a community matter. Many programmes are fostering open conversations among students about the appropriate use of AI. Studio tutors encourage students to share experiences, for instance, one student might demonstrate how they used an AI tool to test lighting in their design, and together the class can discuss whether this was effective or crossed any lines. Such peer discussions demystify AI and create a shared understanding of what is fair game. Student representatives and architecture society leaders are being involved in crafting “AI guidelines” at some schools, ensuring the student voice is heard and that any policies are seen as fair and realistic. By engaging students in creating the rules, educators find greater buy-in. This mirrors the collaborative ethos of studios, where norms often emerge from dialogue. Additionally, some institutions could introduce the idea of “AI integrity champions” or ambassadors, students or staff who stay updated on AI tools and ethical best practices and can advise others. While still informal, this shows the direction of travel: ethical AI use is becoming part of the professional identity that architecture schools seek to instil, much as sustainability and equality have.

7. Integration into University Governance: At a higher level, architecture faculties are aligning with university-wide ethical frameworks on AI. Many universities have established AI task forces or working groups to continuously update policies as the technology evolves. Architecture departments are usually represented in these groups (indeed, the novelty of image-generating AI made art and design faculties key stakeholders). The outcome is that institutional policies are being integrated into the daily practices of architecture education. Course validation documents, learning outcome statements, and assessment criteria are gradually being revised to mention the use of digital tools and AI. For example, an updated learning outcome might read: “Students will demonstrate the ability to critically employ computational design tools, understanding their ethical and

"...BY DIVERSIFYING ASSESSMENT METHODS, SCHOOLS MAKE IT CLEAR THAT A STUDENT'S VALUE IS IN THEIR PROCESS, CRITICAL INSIGHT, AND SKILLS, WHICH NO AI CAN REPLICATE IN FULL..."

practical limitations.” This embeds the expectation of ethical AI use into the curriculum blueprint. We are also seeing references to university AI statements in course handbooks given to students. A handbook might quote the university's principle on AI, such as the University of Sheffield's stance of using AI appropriately and acknowledging its limits [32], and then explain what that means for a particular architecture course.

In implementing all these strategies, a consistent tone is maintained: an academic-policy approach that treats students as future professionals capable of ethical reasoning. The language used in guidelines is often similar to professional codes, underlining personal responsibility. For instance, just as an ARB-registered architect must account for how they produce their drawings, an architecture student is asked to account for how they produced their work, AI included. This not only prevents misconduct in the short term but also builds habits of integrity that will carry into students' careers.

CONCLUSION

In conclusion, the ethics of AI use in architectural education revolve around balancing innovation with integrity. By recognising the specific challenges, from design authorship and portfolio originality to the pitfalls of AI 'shortcuts', architectural educators are crafting nuanced policies that speak directly to studio practices. National frameworks (UK government, RIBA/ARB standards) provide a supportive background that emphasises integrity and transparency, while university and departmental policies operationalise those values on the ground. Proper referencing of AI-generated material has become a new norm, and the inadequacy of detection tools has shifted focus toward education-based solutions. The array of strategies being adopted, from declarations to reflective assessments, demonstrates a comprehensive effort to foster an ethical AI culture. Architecture students are thus being taught to engage with AI as a powerful design partner, one that can expand their creative horizons, but always under the guiding premise that the human designer is ultimately responsible for the work and must diligently credit all sources of inspiration and production, including AI. This academic ethic not only safeguards the fairness and credibility of architectural education in the present, but also helps shape a generation of architects who will use AI conscientiously to advance the field without compromising the core values of authorship, originality, and professional integrity.

AI IN EDUCATION

CASE STUDIES EXPLORING HOW DIFFERENT EDUCATIONAL FACILITIES HAVE BEEN ADAPTING AND INTEGRATING AI INTO TEACHING AND LEARNING.



"...WE NEED TO CONSIDER NOT JUST WHAT CAN BE AUTOMATED, BUT THE POSITIVE ROLE THAT ANY FUTURE AI TECHNOLOGY CAN PLAY WITHIN THE CHARACTERISTICS OF THE SUBJECT..."



TOM HOLBERTON

ASSOCIATE PROFESSOR
THE BARTLETT SCHOOL OF
ARCHITECTURE
UNIVERSITY COLLEGE LONDON

Tom's research focuses on the use of artificial intelligence within design, developed through ten years of teaching and research at UCL (Bartlett "Unit 21").

The way students learn and produce work is being rapidly changed by artificial intelligence. This is happening at multiple levels, from general pedagogy to specialist skills, and regardless of whether courses have chosen to directly address AI.

1. Text to image generators, such as Midjourney or Stable Diffusion, have had a dramatic impact on public consciousness and creative disciplines. These models challenge rendering skills, visualisation skills, and for students entering architectural education provide new forms of immediacy to their ideas. Through playful interaction of prompting and generating results they present almost no barriers to use. A new generation of students can now expect to instantly depict their thoughts, using these models as a form of mediated sketching. Powerful new multimodal models generating video, interactive worlds, and 3D forms will emerge over the coming years.
2. For design courses using high levels of computation, there has been a steady growth of deep learning models and datasets that can be readily accessed through Github and other repositories, with accessible Graphics Processing Unit (GPU) processing for training models on the cloud. AI-assisted coding, or "Vibe coding," is making computer programming far more accessible. There has been a gradual shift in research and skills over the past 5 years from parametric and reinforcement learning techniques into deep learning models. Many of the coding and design skills are highly transferable, but there are critical philosophical and epistemological changes which challenge established relationships between design thinking and computation. The theory and philosophy of digital tools has found a new critical urgency as symbolic and procedural methods shift to probabilistic ones, often in 'black boxes.'

3. Large Language Models (LLMs) present fundamental challenges to learning across education and society. These new tools, readily accessible to all, are changing how students might engage with content, undertake basic research, write and code. Whilst this has been disruptive to many traditional arts and humanities subjects dependent on written assessment, design education has developed quite robust methods for interrogating agency and process, whether this is in studio reviews, crits, pin-ups or portfolios. Students are expected to justify how skills and software have contributed to an outcome. Computational and automatic methods have been readily accepted into design without undermining its core practice. There may be lessons from architectural education which could now help other arts subjects being challenged through automatic writing. There may be techniques from architectural pedagogy that would be valuable to other subjects grappling with AI.

HOW DOES AI ALIGN WITH ARCHITECTURE?

We need to consider not just what can be automated, but the positive role that any future AI technology can play within the characteristics of the subject. Architecture provides a long history of integrating different forms of tools and formulations [3]. These constantly shape the way design is controlled and operated and affect how a design process absorbs pieces of knowledge, expertise and performance whilst sustaining a sense of agency.

Architecture replicates itself by learning from examples, where observation and change is an underlying habit. Buildings provide a canon of uniquely built instances, similar but always different. This is then combined with many process-based media: sketching, ►

drawing, modelling, visualisation, simulation and optimisation, consultation and critique, which synthesise enormous complexity. To create a design demands a constant speculative testing of the specific, in order to establish parameters, concepts and legitimise any process. In this way, deep learning seems both appealing and contradictory. It also learns from examples, but the patterns learnt may only ever be one aspect of what can constitute a design.

Established methods of working with computers such as parametrics have tended to provide some separation between technical and creative decisions. The lack of intelligibility within AI models can result in far greater entanglement of technical and creative factors. Quantitative and qualitative data are easily combined. Datasets may contain a range of complex attributes and latent properties which cannot be easily separated. As a result, models may not be as objective as imagined, and understanding the outputs from any 'black box' model depends far more on testing what has been learnt and how it changes.

ANALYSIS

For this report, we undertook an integrative review looking at the online course and published work from 66 universities in the UK by the end of the academic year 2023-24. This included speaking individually to 14 universities:

- 3 institutions had advanced computation-related courses directly focused on AI
- 9 institutions had students expressly working with AI on design projects, with potentially many more implicitly using AI in some form but not declared
- 19 institutions undertook academic research related to AI, with 11 directly related to architecture, design or construction. A further 8 had related projects in the arts or urbanism.

For context, we also reviewed a number of international courses from around the world, with a particular focus on how they may be incorporating AI differently into the curriculum. Key aspects of assessment included:

- What is the role of an AI model as part of an overall design exercise?
- How can the use of AI provide agency for a student to engage critically?

- How much control and training is made possible within constraints of time and skill?
- How does engagement with AI expose or address pedagogical challenges?

CASE STUDIES

Through five categories we illustrate different assumed roles for AI within design pedagogy. The examples included are illustrative of different approaches, rather than a comprehensive survey.



Ioana Drogeanu, Bartlett, UCL
Towards a 'Non-Universal' Architecture: Designing with Others through Gestures



Oscar Maguire Bartlett, UCL
Probobli Boboli

CASE STUDY 1: FROM SEARCHING TO CONVERSATIONS



Royal College of Art



UNIVERSITY OF
WESTMINSTER

OXFORD
BROOKES
UNIVERSITY

LM and text to image models provide an accessible and fast method of interrogating a topic, site or brief so as to enrich the ingredients that start a project. These can emerge as an alternative to internet search, being mindful of the risks of hallucinations. Moving beyond AI-enabled answers to queries, LLM enables a range of formats, tones of voice and media that can be used to research and creatively interrogate a topic. Through this dialogic role-playing students can expose questions of plurality, bias, and style in the construction of a brief and context.

Studio ADS4 at the Royal College of Art (MA Architecture, RCA, taught by Matteo Mastrandrea, Tom Greenall and Nicola Koller) tested different language tools to provide “new narratives” that form the basis of designs. Language models can easily be employed to extend and expand a narrative, particularly through a dialogic method of an ongoing conversation. This potentially mimics how relationships develop in architecture and where suggestions draw out more detailed requirements. A design project can potentially grow and develop a client or occupant narrative. Alternatively, it may imagine counterfactual histories for a site. (Mimicry and the Villa of Mysteries, 2023 Marisa Yue Chuen Müsing, RCA ADS4) Narrative methods that use AI models can draw-in alternate methods of prompting, incorporating multiple participants, communities voices and styles. This sets a different mode for brief construction where it can be fluid and responsive.



Marisa Müsing, RCA
Mimicry and the Villa of Mysteries



Joshua Parker, RCA
Forgiveness, Not Permission




Harry Tindale, RCA
The Zoo that Came to Tea

Speculative image generation, often labelled ‘Imagineering’, provides another version through which text prompts can provide a form of sketching or conceptual visualisation, rendering provocative or nuanced images that can visually construct a brief, an idea for a project, or a set of imaginary references. The uncanny and photorealistic possibilities of these images direct from a text description can be compelling. However, they may also present pedagogical challenges if they are seemingly too detailed and ‘fixed’ for early stages of design. These tools can both enhance and undermine the conceptual development of projects, and students need the skills to differentiate this for future career and clients.

The mining of large models (Midjourney, Dall.E / GPT4, Stable Diffusion) is still highly subjective, ethically complex, and requires a user to understand how tone and prompting method steer the generated outcome. This sensitivity to different language can be mapped and analysed. Design projects can explore this as a new design space driven by language such as Framework for Design - conversation between AI and human agency (2023, Thomas Ellis, MArch University of Portsmouth, taught by Dr Antonino Di Raimo, Paula Craft Pegg, John Pegg and Simone Sfriso).

There are dangers to images generated by text being presented with a false empiricism rather than through creative dialogue, one that treats answers as grounded within a known domain. Instead, any input material (prompts, source images, or a visual scaffold created by a student) is crucial to determining the outputs and interpretation, where a process of iteration will reveal if there is agency. Different methods can be used to structure AI-generated inputs and output – from inputting base depth map images - created from a project (Extreme Retrofit, 2024, Rahul Mukesh Vyas, Rishab Namdeo Naik, Alireza Sadeghi, Jahnavi Jayashankar, MA Digital Craft Oxford Brookes) to generating images that are translated back into parametric models (Oyster Shellebration, 2023, Diploma Studio 10, MArch Westminster taught by Toby Burgess and Arthur Mamou-Mani).




Summing Cai, Yiwen Qian, Muskaan Marcia, Yiheng Xu, Bartlett, UCL
Accent Diffusion

Over the past decade the growth of big data has provided the raw material to map and simulate the city. Deep Learning techniques are now providing more speculative and creative approaches to interrogate these patterns: “Received notions of top-down and bottom-up approaches are no longer satisfactory paradigms to engage the spatiality of ML algorithms in which the whole design space can be simultaneously manipulated and punctually probed.” [1]

The MArch research clusters on the Bartlett Prospective (B-Pro) programme work in small groups to develop computational methods that reimagine relationships and configurations within the city. Design projects emerge from reading ‘latent’ intelligences, blending data and interpretation in ways which are less deterministic and more relational. Accent Diffusion (2023 Summing Cai, Yiwen Qian, Muskaan Marcia, Yiheng Xu - RC14 taught by Roberto Bottazzi, Eirini Tsouknida and Tasos Varoudis) applied cGAN to visualise language diversity and identify new sites and programmes of intervention. Informatisation City (2022 Baitong Lii, Zhiying Chen, Hao Ren, Yining Wong – RC11 taught by Julian Besems and Philippe Morel) created a chain of cGAN models that generatively connect site forms to hybridised programmes and topology. Automating these approaches across an urban context establishes stronger connections between the subjective deep-reading of a place and a design response. Large foundational AI models are extending new agentic methods to mine and incorporate image, text and film from a city, allowing the designer to create speculative bridges between spatial and physical constraints and other more unexpected readings of the city that might combine social media, literature and emotion.

The capacity to discover and form patterns within a rich set of data types, sources and perspectives through deep learning has the danger of confusing correlation with causality. There is a rapid growth of academic research in this area, but it requires critical design skills to identify when subjective ‘positions’ are being taken on a place as part of a design approach. In principle, these new capacities offer better ways for designers to engage and challenge heuristic conventions, engaging with many more dimensions to urban life. Critical questioning and interpretation of the city needs to remain part of the architectural education, using the intelligence of these new methods without seeing them as solely deterministic.

CASE STUDY 2: NEW UNDERSTANDING OF CONTEXT




Summing Cai, Yiwen Qian, Muskaan Marcia, Yiheng Xu, Bartlett, UCL
Accent Diffusion

CASE STUDY 3: EXPLORING DESIGN PROCESS

UNIT 21 - THE BARTLETT



Rolandas Markevicius, Bartlett, UCL
Cross-Modal Compositions



Rolandas Markevicius, Bartlett, UCL
Cross-Modal Compositions


Studio teaching has the capacity to test and incorporate AI techniques as part of a broader design education. Unit 21 (BSc + MArch Architecture, UCL, taught by Abigail Ashton, Andrew Porter and Tom Holberton) focuses on supporting students develop their own unique design process, typically combining analogue and digital technologies to drive the development of an architectural proposal. Each process is developed through an individual student's research, a site context and the thematic brief of the year.

AI models have been adopted in different ways by around fifteen students over the past five years. Each student will code and create their own approach, adapting and training models with bespoke datasets created as part of the project and directed by a specific context and research process. The model is not seen as a general purpose design tool but instead is highly specific to one context and method of working. This enables more direct exploration of agency and process between the designer and any computation.

The vertical unit operates across undergraduate and postgraduate teaching, creating a cohort of four different years (Y2-5). Skills are shared and learnt within the student group, combining formal teaching alongside adhoc support. Different roles for AI have emerged through this experimentation with each student adopting new positions on the technology whilst always situated within a wider discourse on design process.



Oscar Maguire, Bartlett, UCL
Probobli Boboli



Ioana Drogeanu, Bartlett, UCL
Towards a 'Non-Universal' Architecture: Designing with Others through Gestures

Fundamental to this approach are questions of agency and modality. By including the process of dataset construction and training, students adopt a more self-critical position for the use of AI without predetermining an approach and its utility. This speculative search and refinement can create multiple versions of a design and an extended exploration of how different sites, contexts, and forms of co-authorship can all play a role in determining outcomes. A design process can use AI to capture and extend a highly contextual sequence of drawn steps to make the "churchiest church" (Algorithmic Compositions in Venice, 2020, Bethan Ring) or generate a population of stones that combine and recombine through their own logic (Probobli Bobli, 2024, Oscar Maguire). The introduction of an AI model does not necessarily reduce creative engagement but instead can synthesise new objectives and sources: from participatory physical gestures (Towards a 'Non-Universal' architecture, 2024, Ioana Drogeanu) to simultaneously composing music and architecture (Cross-Modal Compositions, 2022, Rolandas Markevicius)

Multimodality has emerged as an essential concept for designing with AI. This considers architecture as a constant creative interplay and transduction between different modalities, as different media and notation. AI models can be trained to extend this approach, allowing the translation and projection between many different datasets, drawings and models that all offer related representations of the same design subject. This connects to a tradition of pedagogy through forms of representation, through which design negotiates differing levels of precision, agency and intelligence.

Studio projects can also use the ‘fine-tuning’ of AI models as a form of creative provocateur [2]. This technique leverages the power of large text-to-image models but then applies a small curated dataset to constrain the outputs in style or subject area. This can make the outputs specific to a project and be repeatedly tested and refined. This approach can evolve hybrid forms of working, reapplying personal drawing styles and entering into a form of personal intensification, described by one student as “something to react against, but is not necessarily helpful or solving.”

Luddite Fallacy (2024, Eric Taylor, DS 3.4 BA University of Westminster, taught by Paolo Zaide and Tom Budd) blends drawing and generated images as a design for an automated landscape in Tilbury, creating a design through a form of residue and feedback. All's Well That Ends Well (2023, Michaelia Zheng, UG21 BSc UCL) combined physical modelmaking with the image infill and extending capacity of multimodal models. Each physically modelled iteration could be extended via deep learning for a particular viewpoint, blending interpreted and constructed scenography.

Agentic AI is emerging as a new phase where models are provided with degrees of autonomy to interact within a system through language or other means. This may shift emphasis to models as participants in the design process explore their ability to perceive, provoke and collaborate with a designer rather than capturing all dimensions of a problem. This potentially reframes a model's intelligence through its interactions rather than its ability to generate solutions.

CASE STUDY 4: USING AI AS A CREATIVE PROVOCATEUR



Eric Taylor, University of Westminster
Luddite Fallacy

CASE STUDY 5: STRUCTURED TEACHING THROUGH AI-SPECIFIC MODULES

When comparing UK architecture courses with international approaches, there appears to be more visible adoption of AI within specific modules overseas. Examples include Artificial Intelligences in Design (taught by Immanuel Koh, Singapore University of Technology and Design), MSc1 - Design Data and Society (taught by Prof.dr. Georg Vrachliotis, TU Delft) or Quantitative Aesthetics: Introduction to Machine Learning for Design (taught by Panagiotis Michalatos, Harvard GSD). International education systems support semester-long projects, with clear technical and skill goals. They can be more autonomous, rather than integrated into a longer research agenda or design process. In some cases, there is the capacity for small groups to create a dataset and train a given model to generate morphologies. Students may then critique this as part of a written report or assessment.

This approach has some clear advantages of connecting doctoral and advanced academic research with foundational courses that can give an urgent grounding in AI within a broader architectural education. However, these short modules are also constrained in their capacity to integrate technology into architectural outcomes authored by individual students. Maintaining these modules in the face of the rapid development of new AI models and techniques is challenging. Staying up-to-date is potentially more easily managed as part of the organic evolution of skills within design studios, where students help drive needs, whilst keeping the assessment criteria focused on the critical outcomes and process.

The epistemological implications of deep-learning and AI adoption are being felt across all subjects in universities. Many UK Higher Education institutions are providing additional courses and skills to all their students, teaching how to use large-language models and develop new research methods, recognising that active engagement on the use of AI is essential. Tailored courses for AI & Architecture may start to filter into the architecture history and theory, practice and ethics curricula to replicate some of the content in AI & Architecture modules found internationally. This could offer some foundational skills for all students, alongside individual design teaching.

CONCLUSION

Architecture design teaching within the UK continues to explore and test new technology, just as it has done for the past thirty years. AI is the latest of many computational innovations. There are many strengths within modern architectural pedagogy to critically engage with and explore new technologies, but these depend on resilient and supportive design teaching and assessment that can afford students the time and resources to develop design projects.


Whilst it may appear that technology is driving more immediate, automated and quicker outcomes it does not necessarily follow that pedagogy should mirror this change through an atomised, granular and results-driven structure. Teaching, research and assessment needs to maintain the breadth and space to critically explore and incorporate these tools. Architectural teaching has maintained a level of versatility in the face of changing design techniques and tools over many decades. This requires both engagement and also sufficient time and space to iteratively test. This applies for both an individual student and within collegiate curricula where a plurality of approaches is possible.

The case studies demonstrate that within a very short period of time, design teaching, and particularly the studio system, has adopted AI into many alternative roles within the design process, for research, role-play, collaboration and augmentation. The richness of exploration is positive and testament to the curious and speculative agendas that students can bring to their work, engaging with a technology that will have huge impacts on their future careers. The challenge for pedagogy is to continue to support this plurality in the face of uncertainty, whilst providing essential skills and knowledge for all students.



Oscar Maguire Bartlett, UCL
Probabli Boboli

"...WITHIN A VERY SHORT PERIOD OF TIME, DESIGN TEACHING, AND PARTICULARLY THE STUDIO SYSTEM, HAS ADOPTED AI INTO MANY ALTERNATIVE ROLES WITHIN THE DESIGN PROCESS..."



Thomas Ellis, Portsmouth University
Framework for Design



NATE KOLBE

DEPUTY HEAD OF ARCHITECTURE
SCHOOL OF ART
ARCHITECTURE & DESIGN
LONDON METROPOLITAN
UNIVERSITY

Nates's teaching interests center on digital design processes and techniques. Digital craft is foremost in Nate's work - finding new ways to develop beauty, intricacy and detail in his outputs.



BRIDGING THE TEACHING OF AI WITH RESEARCH

The development of architecture students from BA and MArch into the independent realms of PhD research presents an interesting trajectory as we further understand the impact of AI on architectural education. The transition from the structured academic setting within BA and MArch studies into the independent research agendas of a PhD has its own unique areas of opportunity and risk with the rise of AI and Machine Learning.

Students of architecture develop their skills from their knowledge of a subject where they are understood to be informed on a topic, to the understanding of a topic where a student can effectively describe or discuss an issue. They move into the experience of a situation where their problem solving is a clear parameter of success. The next stage is the ability to manage, guide and solve questions of architectural and technical proficiency. The demonstration of a student's ability to do a task is foundational to their becoming a professional.

"The machine is the architect's tool – whether they like it or not. Unless they master it, the machine has mastered them." Frank Lloyd Wright, *The Architect and the Machine*.

AI in architectural education is having an impact on how students demonstrate their knowledge, understanding, experience and ability in their studies. AI can, on the one hand, be seen as shortcutting the iterative processes of design, but equally

enhancing the variation in processes a student can use to attain those abilities. At the PhD level, the independence of the student is paramount and their ability to thoroughly research a topic is critical to their success. The student is moving from a position of understanding into an area of expertise through research competency.

AI and machine learning offer more than intriguing tools to be utilised by the researcher, they are a study unto themselves. The template of research will naturally change as AI supports the direction and interaction of deep research and expertise.

PhD students declare the ethics of their work at each stage of development. Their supervisors are there to guide and support the student as they develop their area of research expertise. AI and ML will likely become embedded in the process of research as students look for new sources and responses to their research, but it will be those who use it as a support mechanism, those who drive AI further, who will get the most out of these new tools.

PHD RESEARCH



TSUNG-HSIEN WANG

LECTURER IN ARCHITECTURE
SCHOOL OF ARCHITECTURE AND
LANDSCAPE
UNIVERSITY OF SHEFFIELD

XINGJIAN ZHAO

LECTURER IN ARCHITECTURE
SCHOOL OF ARCHITECTURE AND
LANDSCAPE
UNIVERSITY OF SHEFFIELD

DEVELOPING AI AUGMENTED CAPACITY TO SUPPORT BUILDING PERFORMANCE EVALUATION AT THE EARLY DESIGN PHASE

Buildings are significant energy consumers in the world, accounting for nearly one-third of total energy consumption and associated carbon emissions (IPCC, 2022). To address these climate-change challenges, enhancing Building Energy Performance (BEP) efficiency has become a critical focus. A key tool for this endeavour is the Building Energy Model (BEM), which offers an opportunity to carry out simulations and examine how a building utilises energy (Del Ama Gonzalo et al., 2023). By developing a digital building twin, architects and engineers can experiment with various designs, materials, and systems to determine the most effective strategies for minimising energy demands prior to implementing any costly physical modifications (Arsecularatne et al., 2024).

However, developing digital models for existing buildings at a large urban or community scale presents a formidable challenge (Park & Wang, 2024). The conventional modelling approach is often manual, labour-intensive, and computationally expensive. Such a task often necessitates expert engineering knowledge to interpret architectural floorplans and input data into specialised software. This bottleneck is particularly prominent in China, which has multi-household residential buildings with intricate floor plans in urgent need of retrofitting to meet national net-zero building ambitions (Zhao & Wang, 2025). While some automated methods exist for converting floor plans to BEMs, they often struggle with the specific formats of these plans, which typically exist only as images (raster graphics) rather than digital drawings (vector graphics). As such, this research addresses the pressing need for a fully automated, accurate, and efficient modelling framework to bridge this gap, transforming static 2D floor plan images into 3D parametric, simulation-ready energy models.


"...AN AI-ASSISTED BUILDING ENERGY
MODELLING FRAMEWORK THAT COULD
INTELLIGENTLY PROCESS 2D-FLOOR
PLAN IMAGES AND RECONSTRUCT THEM
INTO 3D BUILDING ENERGY MODELS..."



To summarise, this postgraduate research project investigates an AI-assisted modelling framework that affords streamlined building performance evaluation (BPE) at the early conceptual design phase for building retrofit. This study considers multi-household residential buildings in China to explore the potential and practicality of such an automated modelling workflow to support performance-driven design explorations.

METHODOLOGY: AI-ASSISTED BEM (AIBEM) FRAMEWORK

To tackle early design BPE for building retrofit, this project proposes an AI-assisted Building Energy Modelling (AiBEM) framework that could intelligently process 2D-floor plan images and reconstruct them into 3D BEMs (Zhao & Wang, 2025). This framework consists of four methodological steps with hybrid artificial intelligence techniques:




STEP 1: RECOGNISING BUILDING
ELEMENTS WITH A DEEP MULTI-TASK
NETWORK

The first step begins by training a computer to 'comprehend' any given 2D floorplan (image). An advanced AI model, specifically a Deep Multi-Task Network (DMTN), was employed to accomplish this challenge. This DMTN model was developed with a dataset of over 1,000 residential floorplans from China. The trained AI model performed two critical tasks: (1) identifying and locating essential building elements in the floorplan image (e.g., walls, doors, and windows) and (2) generating the vectorised representation of the floorplans. The step considers visual features of building elements with the underlying geometric constraints, ensuring an accurate and robust understanding of the floor plan.

STEP 2: RECONSTRUCTING
HOUSEHOLD LAYOUTS WITH CORE
TOPOLOGY ROOM CLUSTERING

With the individual building elements, step 2 is to assemble them into a structured layout of rooms per household. The framework develops a customised algorithm, Core Topology Room Clustering (CTRC), for this task. The CTRC algorithm starts with the core room, namely the circulation space (the living room), to cluster connected room objects into individual household units. This step reconstructs the household clusters through the spatial connectivity graph.




STEP 3: CLASSIFYING ROOM
FUNCTIONS WITH A GRAPH
ATTENTION NETWORK

With the semantic-rich spatial graph from the above Step 2, Step 3 carries out the room function classification for automating BEM. Classifying room functions is accomplished using a Graph Attention Network (GAT) that uses weighted mechanisms for enhanced predictions. For each household graph, rooms are nodes and connections between rooms (through doors) are edges. The GAT model in this study was trained with the selected room features, including its size, shape, and number of windows, alongside its spatial connectivity relationships to adjacent rooms. The resulting GAT can accurately predict the room functions with over 95% accuracy.

STEP 4: GENERATING THE 3D
BUILDING ENERGY MODEL

With a fully structured and attributed floor plan, the final step automatically generates BEMs. The framework consolidates all gathered information, including the geometry information of building elements—such as walls, windows, and doors—and the room functions per floor plan to formulate an attributed 3D BEM. The result shows a simulation-ready digital representation of the building generated in a fraction of the time when compared with traditional manual methods.



RESEARCH FINDINGS AND
APPLICATIONS

In this study, we presented the AiBEM Framework, showcasing its effectiveness in floor plan recognition and BEM generation. AiBEM could recognise architectural elements with exceptional precision, achieving over 95% accuracy. The room function classification also demonstrated outstanding reliability, capturing the intended use of each space within the building. By streamlining the entire workflow from images to simulation models, this research marks a significant advancement in leveraging AI capacity to support BPE.

This AiBEM framework can afford rapid energy performance assessments of existing residential buildings. For urban designers and policymakers, it enables large-scale data collection and analysis to develop effective energy conservation strategies and building codes. Other stakeholders, such as building and asset owners, can leverage this technology to identify underperforming buildings for energy-saving retrofits. By significantly reducing the time and labour required to create BEMs, this framework makes advanced energy analysis accessible for a broader range of projects.

This case study demonstrates how AI can enhance the capacity to support sustainability initiatives. It offers a scalable solution designed to streamline building performance evaluation to reduce energy consumption and carbon emissions in the building sector. This effort paved the way to achieving future energy-efficient built environments in China and worldwide.

"...CLASSIFYING ROOM FUNCTIONS IS ACCOMPLISHED
USING A GRAPH ATTENTION NETWORK (GAT) THAT
USES WEIGHTED MECHANISMS FOR ENHANCED
PREDICTIONS..."



IN PRACTICE

AI’S GROWING IMPACT ON ARCHITECTURAL PRACTICE &THE ROLE OF EDUCATORS IN SHAPING THIS FUTURE



SHERIF ELTARABISHY
ASSOCIATE PARTNER, DESIGN
SYSTEMS ANALYST

For over 14 years, Sherif has been lecturing, training and consulting at different universities and firms, with a focus on digital transformation in the AEC industry.



MARTHA TSIGKARI
SENIOR PARTNER AND HEAD OF
THE APPLIED R+D (ARD) GROUP

Martha has two decades of experience working in projects of all scales and uses. Her work incorporates computational design, human-computer interaction, machine learning, and optimisation.

Foster + Partners

INTRODUCTION: BEYOND THE HYPE

Conversations about AI in architecture typically start with the technology. At Foster + Partners, we've learned to start elsewhere: with process mapping, workflow analysis, and a fundamental question: what does it mean to be an AI-native architecture practice?

Over the past several years, our Applied Research + Development team has grown from eight to thirty members, bringing together architects, engineers, computer scientists and applied mathematicians. Rather than simply adding AI capabilities to existing tools, we have taken a more radical approach, reimagining our workflows from the ground up to be data-ready, scalable, and fundamentally aligned with how modern AI systems operate.

THE FOUNDATION: PROCESS BEFORE TECHNOLOGY

Our approach begins not with asking "what can AI do?" but rather "what are current workflow problems that AI can solve for us?" We conduct detailed process mapping sessions, sitting with teams to understand every step of their work, who provides and who owns what information, what decisions are made where, and crucially, where tacit knowledge creates bottlenecks.

Take our work on Al Sa'ad Plaza Tower, a landmark project in Qatar, as an example. The design team needed to explore numerous iterations while managing input from BIM teams, fabricators, façade consultants and structural engineers. When last-minute changes required adjusting the tower height and specific floor heights just days before submission, our automated systems, built through careful process mapping, allowed us to regenerate all geometry and documentation in time. Just one of many similar project anecdotes.

This wasn't achieved through AI magic, but through

understanding the workflow deeply enough to know where automation, whether AI-driven or algorithmic, could have the most impact.

BUILDING SECURE INFRASTRUCTURE: THE AI PORTAL

One of our recent developments has been our internal AI Portal, a platform that allows our designers to safely explore generative AI within the strict confidentiality requirements of our practice. Many of our projects are governed by NDAs, and we've seen too many instances of designers elsewhere unknowingly violating terms of service by using external AI tools.

The AI Portal integrates directly with Rhino, our primary design tool, and runs on a hybrid infrastructure combining on-premises GPU clusters with scalable cloud resources. The platform includes a collaborative branching system where team members can see each other's AI-augmented explorations, build on them, and create a visual version history of design evolution.

We use open-weight models that we can customize using parameter-efficient tuning techniques, in addition to third-party proprietary models when needed. This allows us to train small adapters on project-specific visual references, letting designers "steer" foundational models, while maintaining full control over our intellectual property.

With our starting point being the designs that our teams develop, AI plays the role of the creative amplifier by extending our capabilities and shortening the time required to deliver.

DEMOCRATISING INSTITUTIONAL KNOWLEDGE

Perhaps our most impactful machine learning application addresses a challenge every large practice faces: how to capture and share decades of accumulated expertise. A good example for us was our in-house Design Guides. This is an extensive body



Al Sa'ad Plaza Towers, a landmark project in Qatar by Foster + Partners

of documentation regarding building typologies, regional codes, and best practices that our Technical Design Group has built and updates weekly. But traversing these documents is challenging for new employees: they often don't know where to begin, leading to repetitive questions that consume experts' time and drop productivity.

To facilitate employees interrogating the guides, we developed Ask F+P: a natural language search engine powered by large language models. Unlike typical chatbots, we deliberately chose extractive question answering over generative responses. When someone asks, "What is the typical basin waste pipe size?" the system doesn't generate an answer, it finds the exact location in our documentation and presents it with a confidence score.

This maintains accountability through source attribution while enabling expert validation of results.

BEYOND DESIGN: OPERATIONAL AI

Recognizing AI's potential beyond design tasks, we've explored operational applications with compelling results. Our business insights system exemplifies this approach. By analysing historical data, we built a statistical model that helps teams issue more calibrated proposals.

When presented with a new opportunity, the model considers project location, typology, area, and other factors to predict expected resourcing needs, likely costs, and project duration. Importantly, this isn't a black box, every prediction can be traced back to its reasoning, essential for financial decision-making.

We deliberately chose interpretable statistical models over more complex ML approaches here. When stakeholders need to understand why a certain proposal was recommended, transparency trumps marginal accuracy gains. The tool can help new team members operate with the accumulated wisdom of decades of practice experience.

A CULTURAL TRANSFORMATION

The most profound impact of our ML journey has been cultural. The current AI hype, despite its excesses, has given people across our practice permission to question long-standing processes. Designers now approach us asking "Why am I still doing this manually?" or "Couldn't AI help here?"


Often, we find that simpler solutions, sometimes traditional algorithms, are more appropriate than ML. But the conversation itself is valuable. It's pushed us to examine workflows that were set in stone over the years and imagine fundamentally better ways of working.

This shift requires new forms of literacy across the practice. Not only on how to use AI tools, but also on their limitations, appropriate use cases, and the importance of maintaining critical oversight. Unlike software, where bugs can be patched quickly, architecture deals with permanent physical artefacts. We can't afford to become "vibe-driven" in our contributions to the built environment. Our fundamental goals as an office are those of sustainability and human-centric design: AI is but a tool that could help us achieve these goals faster and better.

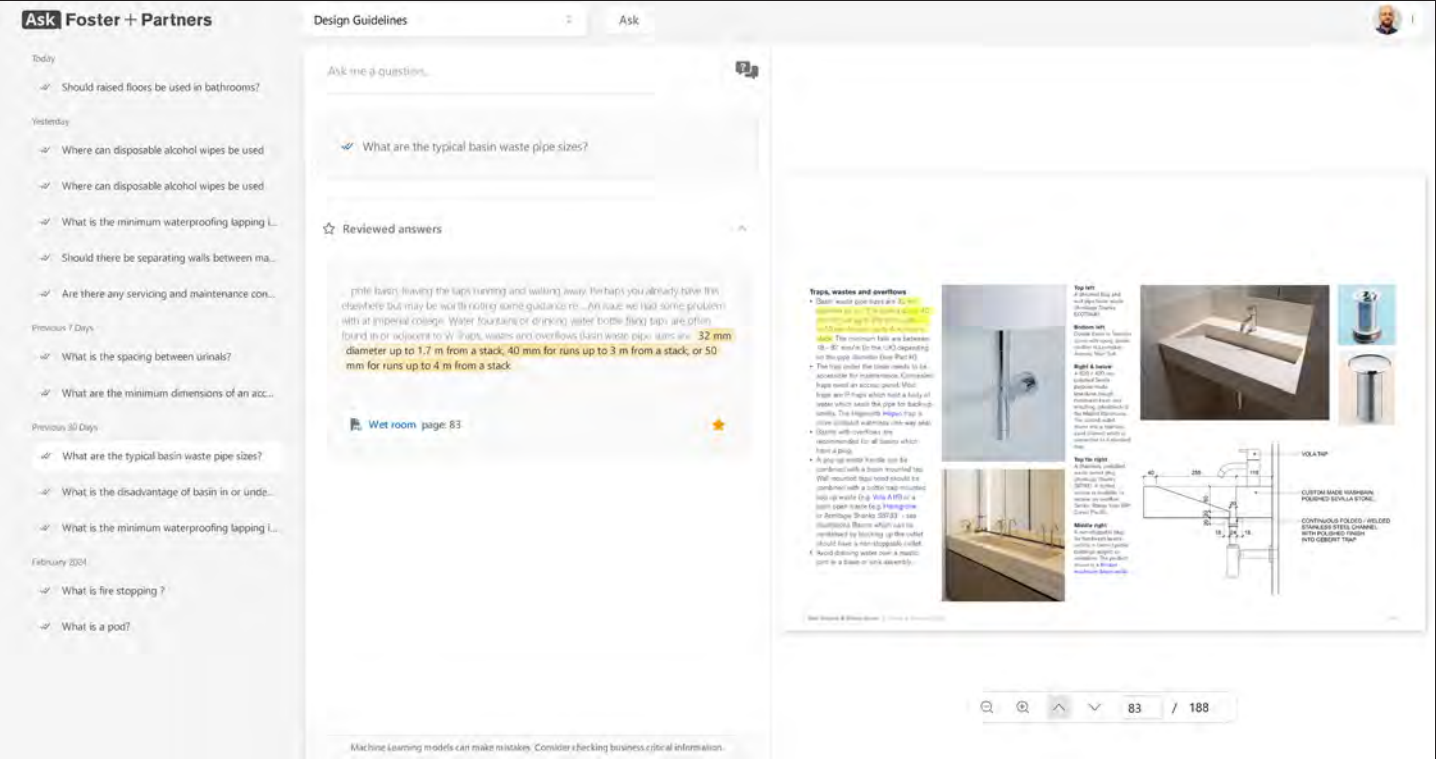
LOOKING FORWARD: THE AI-NATIVE PRACTICE

We believe all architecture practices will eventually need to grapple with what it means to be AI-native, much as businesses in the 1990s had to figure out their internet strategy. This isn't about using AI tools, it's about fundamentally restructuring how knowledge flows, how decisions are made, and how creativity is fostered.

Our vision extends beyond reactive systems that respond to queries. We're working toward proactive systems that anticipate needs, surface relevant information contextually, and



The AI Portal's webapp, showing results submitted by multiple users from different applications and to different ML models, allowing teams to work collaboratively.



Ask F+P web-application, showing the user asking a question, where the answers from the ML model were asynchronously reviewed by domain experts and filtered down to the correct one. On the right the context where the answer appeared in is displayed.

"...IT'S PUSHED US TO EXAMINE WORKFLOWS THAT WERE SET IN STONE OVER THE YEARS..."

help navigate complex workflows without requiring perfectly articulated questions. This represents not just an efficiency gain but a redefinition of a practice's sustainability, where knowledge becomes a living, accessible resource rather than a siloed asset. The journey from experimenting with ML to embedding it in our operational core continues to reshape how we work. But we're conscious that with great capability comes great responsibility. As we push boundaries with generative design, automated workflows, and AI-augmented creativity, we must maintain the critical thinking, ethical grounding, and human judgment that architecture demands.

THE BALANCE: PERFORMANCE AND RESPONSIBILITY

An interesting tension in our work emerges between computational acceleration and physical permanence. While we can now generate hundreds of thousands of design variations and analyse them in hours using tools like Cyclops, our publicly released performance analysis platform, buildings still need to last 50-100 years. This temporal mismatch requires careful consideration. Cyclops exemplifies our approach to this challenge. Rather than using probabilistic AI models that approximate performance, we have GPU-accelerated deterministic analyses. Where traditionally evaluating 24 views in a stadium might take a minute, we can now analyse an entire 10,000-seat venue in a second. This distinction matters. When we're evaluating daylight performance, structural loads, or acoustic properties, we need physically accurate results, not statistically likely approximations. By maintaining this rigour while achieving AI-like speeds, we give designers the ability to explore vast design spaces without compromising the fundamental reliability architecture demands. We tested even faster probabilistic models, but it is always worth questioning when and to what extend marginal speed gains justify sacrificing analytical certainty.

NOTES FOR FUTURE SELF

The landscape is shifting extremely quickly, assumptions from a month ago are probably already outdated. Keep revisiting, keep questioning. Vendors will over promise. They are solving for their business model, not your organizational complexity. Their "few clicks" AI solution rarely accounts for your fragmented data, inconsistent access privileges, or the reality that "data" means different things to different teams within the organisation. We're still in the early days. Costs will drop, use cases will mature, new applications will emerge. Document everything, today's failed experiment might be tomorrow's breakthrough when revisited with better models or clearer thinking. Remember that new technologies rarely serve as direct replacements; their real value often lies in enabling entirely new workflows and possibilities

we couldn't imagine before. When implementations fail, and they will, make sure you have scoped narrowly enough to understand exactly what failed. Usually it's integration choices, culture, or expectations, not the technology itself. Resist the temptation to automate only the glamorous parts. Back-office processes might not showcase at design conferences, but they offer proven Return on Investment (ROI) patterns from other industries and measurable impact. Sometimes the best place to start isn't the most visible. The future of architectural practice lies not in the wholesale adoption of AI, but in the thoughtful integration of ML into workflows designed from first principles. Becoming AI-native isn't about the technology, it's about culture, process, and a fundamental commitment to reimagining how architectural knowledge is created, captured, and shared. The practices that thrive will be those that see ML not as a threat or magic solution, but as a powerful amplifier of human creativity. This requires genuine executive buy-in, team adoption and budgeting as much for workforce upskilling as for the technology itself. Sophisticated tools are worthless in unprepared hands. Done right, this thoughtful integration maintains human creativity and judgment while extending our collective capability to design better buildings, better cities, and ultimately, a better built environment for all.



PROF. DES FAGAN
AI RESEARCHER
SECONDED TO GRIMSHAW ARCHITECTS
AHRC INNOVATION FELLOW

GRIMSHAW



Image of the AI and Architecture: Sustainability Summit held in January 2025

Pedagogy and Practice in AI:Lab: Artificial Intelligence for Low Carbon Buildings

The AI:Lab: Artificial Intelligence for Low Carbon Buildings (2023-4) was a funded research collaboration between Lancaster University and Grimshaw Architects that explored the intersections of AI and sustainable design. Whilst the primary aim was to explore AI tools applied to the Eden Morecambe project, the methods and processes developed encompassed trajectories to explore new areas of pedagogical enquiry. Framed within the university design studio, these workflows offer a new way of working with AI as part of curriculum learning: not as a replacement for design thinking, but rather, learning to think, model and adapt in a world that is carbon constrained.

PEDAGOGICAL OPENINGS: LEARNING FROM NATURAL FORMS

The Eden Morecambe Project is influenced by the form of seashells – our first AI-augmented workflow trained computer vision models to ‘parse’ seashell geometries, creating a script that could extract dimensional and material information and classify any shell found on any beach. This was the first step to create an editable ‘digital twin’ of seashells in Grasshopper, using computer vision to identify their size and type, passing this information to the Grasshopper program through a visual ‘wrapper’ (Cloudflare & Python) to create a geometrically editable model. This work drew on the principle that natural shell formation can be represented mathematically, and as such, is ideally suited for surface creation

in the algorithm-based environment of Grasshopper.

From a pedagogical viewpoint, this work demonstrates that AI can digitise and convert real-world elements to classify and segment, to test performance and explore the environmental implications of scale and material use. What might begin as an exercise into bio-inspiration may lead into discourse on critical awareness between structural stiffness or carbon expenditure at scale, leading to recognition of biological evolution as a natural force for optimisation. Instead of prescribing design outcomes, this work is a frame for inquiry - an example of the way that AI can be conceptually used to explore bio-inspired iterations at scale, joining an exploration of aesthetics with measurable levels of sustainability.

CONTEXT AS DIALOGUE: SITE DATA & AI DISCOURSE

In our second workflow, transport data, planning policies, and transcripts from consultations were uploaded to a closed Retrieval Augmented Generative (RAG) model to examine how AI could interpret, synthesise, and navigate between competing voices on urban sustainability.

One experiment involved generating fictitious dialogue between historically divergent commentators Jane Jacobs, Robert Moses and others, using the model to draw on their published viewpoints. This experiment reframed site investigation through the lens of dialogue, highlighting its potential in design education

to help students recognise how context is shaped and reshaped through negotiation.

Using these tools, students have the opportunity to test out approaches or positions in relation to competing urban sustainability principles in a virtual forum of stakeholder voices, to uncover inconsistencies or contradictions and to construct and deconstruct contested notions of urbanism. By combining this with publicly available social media datasets (with appropriate approvals) this could be scaled to explore public sentiment and discover patterns in how residents and stakeholders react to proposals that affect their communities. The workflow has the potential to teach students about civic literacy: sustainability involves compromise, values-based negotiation, and multiple stakeholder priorities. AI augmentation here rehearsed a debate rather than resolved it - in itself, a useful tool for exploration.

FEEDBACK LOOP: SURROGATE MODELLING PERFORMANCE

Surrogate models were developed to approximate the performance of time-intensive gridshell simulations for Eden Project Morecambe. We scripted a Convolutional Neural Network (CNN) surrogate model to test thousands of design iterations in real time to analyse shell performance; including evaluation of total carbon used, structural utilisation ratios and structural displacement.

Within the design studio, the pedagogical implications of

"...AI CAN DIGITISE AND CONVERT REAL WORLD ELEMENTS TO CLASSIFY AND SEGMENT, TO TEST PERFORMANCE AND EXPLORE THE ENVIRONMENTAL IMPLICATIONS OF SCALE AND MATERIAL USE..."

surrogate models are significant: early simulations are often speculative, as initial ideas rarely give students or educators access to detailed analyses of structural stability, carbon cost, or environmental performance. Surrogate modelling introduces the possibility of an early feedback loop that could improve the iterative and collaborative potential of AI in comparison with computationally expensive Finite Element Analysis (FEA) methods. Surrogate models can produce thousands of early design options, reflect on their carbon impacts, and think through trade-off decisions. Framing these workflows and explorations has the potential, within studio teaching, to connect computational explorations more immediately with a sense of responsibility for the environment, developing a culture of sustainability that becomes second nature within the rhythm of design iteration.

THE HIDDEN CARBON COST OF COMPUTATION

Model training can be dependent on a large amount of GPU resource, consuming considerable amounts of energy and cooling water. To track this, we used CodeCarbon, an open-source Python package that tracks energy consumption and code execution emissions. Energy use for model training was logged and benchmarked to operational and embodied carbon savings that these tools intend to achieve by the end of the project in 2027. Although still early in its lifecycle, the project recognised the importance of quantifying its own net carbon outcome. Scalability for any AI tool remains a key factor: when AI tools are deployable across thousands of projects, initial training cost will be readily repaid against wide-reaching impact, but responsible innovation means confronting these trade-offs transparently – and ensuring that projects maintain comprehensive records to evaluate carbon performance throughout their implementation.

For students, this highlights a significant aspect of use - the sustainability of the digital tools themselves. Too often, the energy cost of computational methods is considered inconsequential.



Image from the AI and Architecture: Sustainability Summit

By considering the carbon impact of their tool use, students understand the consequences of every workflow choice. Thinking about design outputs balanced with carbon cost benefit fosters responsibility and transparency, extending the conversation about sustainability to the design technology itself.

COLLECTIVE LEARNING AND PUBLIC EXCHANGE

A key feature of AI:Lab was emphasis on dialogue: the AI and Architecture: Sustainability summit (Jan 2025) at Morecambe Winter Gardens created a space for practitioners, educators, students and the public to think and reflect on the use of AI for decarbonisation. The exhibition and debates framed AI not just as an opportunity for expediency of workflows, but as cultural engagement and inquiry. Connecting studio projects to public workshops, extending reviews into exhibitions and developing AI-mediated processes alongside communities can encourage collective exploration on the urgent role of AI tool use in the built environment.

PEDAGOGICAL REFLECTIONS

When viewed in terms of education, our learning from AI:Lab points less to specific practices or methods, and more towards broader pedagogical themes:

Iterative AI Feedback as Collaborator: AI enables early-stage feedback loops, allowing students to test multiple iterations and reflect on making in a collaborative, dialogic sense.

AI Inference and Dataset Complexity: Data-driven models underscore the wicked nature of design problems. Through inference, AI can assist in navigating competing values, fostering engagement, and enabling debate.

Ethical and Sustainable AI: Teaching with AI should address not only the energy savings of design outcomes but also the environmental costs of the software and processes themselves.

FINAL THOUGHTS

AI:Lab has taught us that the development of tools and workflows is important, but that key learning comes from reflection upon our existing learning processes: iterations, dialogues, reflections, and accountabilities. AI can be both a medium and a mirror: allowing students to quickly explore ideas, while also engendering thoughtful, critical reflections on sustainability and authorship. Ultimately, what emerges is less a set of fixed methods than an exploration into how AI sits as part of a broader commitment to curiosity, accountability and design in the context of planetary boundaries.



LEARNING FROM THE STRUCTURES OF NATURE
Using a machine learning (ML) model trained on real and synthetic datasets of shell images, we classified shell types and extracted key dimensions from photos of cockles, mussels, oar shells, and whelks. These inputs fed into mathematical formulas describing 3D shell shape generation. This enabled the creation of editable 3D geometries from any shell photo, integrating environmental and structural performance analysis. By capturing the principles of shell growth, the process supports innovative architectural design inspired by Morecambe Beach's natural forms.

FIG.6 IMAGES OF BUILDING INTERIOR GENERATED FROM PARSED DATA OF SCANNED IMAGE OF SHELLS WITH VARIED STRUCTURAL SOLUTIONS




FIG.7 GRASSHOPPER DEFINITION - INPUTS TO GENERATE WHELK SHELL STRUCTURES

FIG.11 MORECAMBE WINTERGARDENS - INSPIRED BY SHELLS
The Morecambe Winter Gardens' glazed ceramic tiles and plaster decorations showcase patterns inspired by the natural geometry and textures of seashells, particularly those found on Morecambe Beach. The repeating curves, spirals, and radial motifs evoke the organic elegance, inspired by both the ornamentation and structural principles of shell forms.
Tiled surfaces are inspired by the radial pattern of whelk shells, providing a tactile, visual richness. Plaster detailing mimics the smooth whorls of sea shells, creating a sense of flow and organic structure. Beyond interior decoration, the geometry of shells also inspires large-spanning structural forms - natural strength and surface efficiency informs the design of arches and domed ceilings.



FIG.8 AN IMAGE OF THE WINTERGARDENS FOYER CERAMIC TILING



FIG.9 ML MODEL TRAINED TO CLASSIFY AND EXTRACT KEY DIMENSIONAL DATA



FIG.10 ML TRAINING METHOD
The ML training method utilised a large dataset of real shell photos and a synthetic database of thousands of mathematically generated shells created in Grasshopper. By comparing real and synthetic images, the model learned to identify subtle variations in shell geometry. The ML model was trained to classify shell types—and extract key dimensions such as width, length, curvature, and spiral proportions. These extracted dimensions were then input into a mathematical formula that accurately described the creation of each shell form. This process allowed for precise modelling of shell geometries, bridging natural forms and computational tools for use in design and analysis.




FIG.11 MORECAMBE BEACH SHELL FORM INTEGRATION
The generated shell geometries were scaled to real building dimensions and analysed in Grasshopper for structural and environmental efficiency. By applying structural simulations, the forms' inherent strength—derived from natural shell geometries—was tested under loading to optimise large-span performance and to explore how different materials (timber / steel / bio-concrete) might perform. Environmental analysis included daylighting and air-flow simulations to further refine the designs. The process improves the structural efficacy of organic architectural form finding and offers a new method for integrating biomimetic principles into architecture from a simple photo of a local shell found on a beach.

Presentation panel (3) from the AI and Architecture: Sustainability Summit exhibition exploring ML-based surrogate modelling of seashells.



JOANNA SABAK
COMPUTATIONAL DESIGNER
LONDON, KINGS CROSS



PABLO ZAMORANO
HEAD OF GEOMETRY AND
COMPUTATIONAL DESIGN
LONDON, KINGS CROSS

Heatherwick studio



Rhino Screenshot Input



Predicted Illustration in 2022



Predicted Illustration in 2024

Images by Heatherwick Studio, 2022-2024

Human-Centred AI in Design

At Heatherwick Studio, we believe in the transformational potential that artificial intelligence can bring to the way we design and deliver projects. For several years, we have been actively engaging with and adopting AI to expand our creative capacity, accelerate repetitive tasks, and provide greater focus on the human qualities that define our work.

Our approach has never been to replace human judgement, but rather to leverage AI in support of it—enabling us to design with increased richness, detail, and imagination. Through early adoption of the technology, we have been able to explore both its possibilities and inherent risks, whilst shaping it to align with our core design values.

WHERE WE STARTED

We began our exploration of AI implementation in 2019, several years after the publication of the landmark Generative Adversarial Network (GAN) paper. At that time, AI research in architecture remained relatively nascent, often focused on predicting floor plans or building footprints from established site boundaries. We identified a different opportunity: to accelerate the delivery of routine and mundane tasks, thereby freeing up creative capacity to enable increased richness in designing detail, atmosphere, and form.

By early 2021, we were conducting our first systematic experiments. The initial objective was straightforward—to take

early-stage clay renders from Rhino and transform them into illustrations consistent with our established visual style. We utilised pairs of these renders alongside corresponding hand-drawn sketches from our project archive to train a Pix2Pix model. The datasets proved messy, the labelling process was laborious, and the code implementation remained unpredictable. Whilst the results were not yet suitable for client presentations or formal reviews, they generated considerable excitement within the studio. Even in these imperfect early outputs, we could discern how AI might eventually become an integral component of our design process.

In 2022, the arrival of Midjourney marked a pivotal moment. Suddenly, any team member within the studio could produce high-quality images within minutes. Many of us began generating inspirational visuals for internal reviews—images that provoked meaningful discussions around form, materiality, and atmospheric qualities. Stable Diffusion was concurrently employed to create 2D assets including people, vegetation, and furniture elements for rendering applications.

We also initiated collaborations with external partners to advance AI capabilities beyond 2D applications. Through a hackathon collaboration with Thornton Tomasetti, we developed "Dreamhopper," a Rhino plug-in that generated images directly from the Rhino viewport whilst experimenting with the creation of 3D meshes from text prompts. Whilst the results remained abstract in nature, they indicated exciting possibilities for rapidly producing low-definition 3D assets for visualisation purposes.

FROM OFF-THE-SHELF SOLUTIONS TO CUSTOMISED TOOLS

It became increasingly evident that utilising generic, off-the-shelf tools would prove insufficient for our requirements. AI needed to communicate in our design language, reflect our distinctive visual style, and protect our IP. This realisation led to the development of our in-house diffusion model, trained on 120,000 images from our comprehensive project archive—including sketches, diagrams, renderings, model photographs, and images of completed projects. Launched in 2024 as Heatherwick AI, this customised tool enables us to:

- Transform a preliminary sketch into a fully rendered image within seconds
- Apply the visual characteristics and aesthetic qualities of past projects to new design explorations
- Test material options on clay Rhino renders without requiring full re-rendering processes


Through training on our proprietary archive, the generated outputs carry the distinctive character of the studio rather than adopting a generic "AI" aesthetic. Over the course of three years, the technology evolved from early fragmentary experiments into a reliable, integrated design companion for everyday use.

HOW AI INTEGRATES WITHIN THE STUDIO

In 2023, we conducted a comprehensive survey across the studio to understand how team members were already engaging with AI tools. Approximately half were utilising the technology in various capacities, often during early design stages, but also for text-based tasks including editing, summarising, and translation activities. These insights directly informed our AI strategy development. We focused our efforts on:

- Customisation – developing tools specifically built for our established workflows and visual identity
- Integration – embedding AI capabilities into existing processes rather than creating separate, isolated workflows
- Accessibility – ensuring tools remain intuitive and accessible to all studio members regardless of technical expertise

This strategic approach ensures that AI functions not as an external add-on, but as an integral component of our working methodology, supporting design exploration, rapid provocations, and accelerated visual updates. ▶



Screenshot of Stablewick - Heatherwick AI Image Generation Tool

Image by Heatherwick Studio, 2024

MAKING AI UNDERSTAND DESIGNERS

One of the principal challenges has been bridging the communication gap between how AI systems interpret instructions and how designers typically express creative ideas. Diffusion models respond optimally to precise, technical prompts, whilst designers naturally communicate through conceptual and visual terminology.

Our interface for Heatherwick AI was specifically designed to translate between these two distinct languages. A designer can provide a Rhino screenshot, reference materials from past projects, and a concise description of desired materials—and the tool will generate variations that capture both the specific request and the studio's established aesthetic sensibilities. This alignment between input parameters and output results is what renders the tool practical for day-to-day design applications.

BUILDING A STUDIO-WIDE AI 'BRAIN'

Parallel to image generation capabilities, we have been developing the 'Heatherwick AI Brain'—a centralised knowledge hub capable of retrieving information from across the studio's extensive archives. Utilising RAG methodology, the system draws upon our internal wiki, private video channels, and comprehensive project files to provide relevant references and contextual suggestions.

By making this institutional knowledge searchable and accessible from a single source, we can effectively break down information silos between teams whilst facilitating learning from past project experiences. Over time, we envision this system working seamlessly alongside our design tools, ensuring that inspiration, precedent studies, and project knowledge remain readily accessible throughout the design process.

CUSTOMISATION AND FLEXIBILITY

Our primary focus in AI implementation has been the development of highly customised tools, with each designed to maximise impact whilst ensuring ease of use and maintaining safety standards. However, the rapid pace of AI technological development means that flexibility remains equally critical to our approach.

We are developing a comprehensive framework that enables new models to be integrated into our existing systems as they are released, effectively combining state-of-the-art capabilities with our proprietary datasets and established workflows. This approach ensures that designers have access to powerful tools that remain safe, contextually relevant, and aligned with our distinctive way of working.



Sketch to Render Experiment

Image by Heatherwick Studio, 2024

LOOKING AHEAD

From our initial tentative experiments in 2021 through to the successful launch of Heatherwick AI in 2024, this journey has been characterised by continuous enthusiastic discovery and strategic adaptation. The technology has advanced rapidly, and we have evolved alongside it—always maintaining the same fundamental objective: to create richer, more human-centred projects, supported by tools which amplify creativity of our designers.

By combining the speed and flexibility of these advanced tools with our craft, curiosity, and critical judgement, we can explore a broader range of ideas, respond more effectively to change, and continue producing work that connects meaningfully with people and places.



KEIR REGAN-ALEXANDER
ARCHITECT AND AEC DOMAIN EXPERT
ARKA WORKS FOUNDER & DIRECTOR

Arka.Works

The Small Practice Mandate: Preparing Graduates for AI on Day One

What is the 'Telos' of Architecture School? I recall a debate with a professor during my Part 1 studies about the 'Telos', or the core purpose, of an architecture school. As a student, I argued that the purpose was straightforward: to prepare graduates like me for a successful professional career. My professor disagreed. "No," he countered, "we're not in the business of creating architects. We're in the business of teaching people how to think."

Twenty years on, with the benefit of hindsight, it's clear this was always a false dichotomy. Universities have a duty to do both. They must produce critical thinkers who can assemble vivid arguments about design, but they must also release competent professionals into the industry that are equipped for the immense technical challenges of practice.

Those challenges have only multiplied. In recent years, the profession has had to absorb the tectonic shifts of BIM, the urgent demands of Net Zero design, and the stringent new liabilities of the Building Safety Act. Now, the ground is shifting again with the arrival of AI. This latest disruption demands we adjust our methods upstream, during education. Perhaps nowhere is the tension between academic philosophy and practical necessity felt more keenly than in the small practices that form the backbone of our industry.

THE VIEW FROM SMALL PRACTICE

For the past two years, I've worked with practices of all sizes and a clear pattern has emerged. Walk into any small studio today and you will find a broad curiosity in AI, coupled with widespread confusion and anxiety about how best to use it. The conversation often starts with a nervous energy: "We've heard others are doing amazing things and we need to be brought up to date". This fear of falling behind is a powerful motivator, but turning that energy into lasting value is proving much harder than expected.

Small practices face obvious disadvantages. IT expertise is thin on the ground; they exist to design buildings, not to manage complex software stacks. Budgets are tight, and there is rarely a dedicated BIM manager, let alone an AI working group, to champion new tools.

But here's where it gets interesting: those same constraints create a massive structural advantage. While large firms contend with bureaucracy, training budgets, and the inevitable "that's not how we do things here" resistance, a small studio can pivot in an afternoon. When a director sees a new tool or method, they don't need to ask a committee for permission to try it. They are turning a small ship, and small ships turn fast. This agility is precisely where an AI-literate graduate can become an immediate catalyst for change in a small practice, not just another cog in the machine. ►

"...WALK INTO ANY SMALL STUDIO TODAY AND YOU WILL FIND A BROAD CURIOSITY IN AI, COUPLED WITH WIDESPREAD CONFUSION AND ANXIETY ABOUT HOW BEST TO USE IT. THE CONVERSATION OFTEN STARTS WITH A NERVOUS ENERGY: "WE'VE HEARD OTHERS ARE DOING AMAZING THINGS AND WE NEED TO BE BROUGHT UP TO DATE..."



Image: Professional practice is gradually adapting. An "Applied AI for Architecture" workshop hosted by Arka Works in 2025.

WHERE IS AI PROVING VALUABLE IN SMALL PRACTICE?

Through testing many AI methods directly in practice settings, I've found that the ones delivering real value fall into two clear categories, or what I call the 'two big buckets'. For a graduate looking to make an immediate impact, mastering both is the fastest way to become indispensable.




Image: Practices are now building repeatable LLM workflows and agents that act like templates for specific technical end-to-end processes. (credit: OmniChat.uk)

BUCKET 1: 'DOING THE DISHES' (THE ENGINE ROOM)

This is the procedural, administrative work that keeps a practice running but doesn't require expansive and creative judgment. It's the engine room of the studio: essential, but largely invisible. LLMs excel here when paired with good prompting and structured workflows. A graduate who can automate these tasks will be enormously valuable to a small practice. High-impact applications include:

- Auditing reports against standards, QA requirements and regulations.
- Generating first drafts for written bid proposals and fee schedules.
- Summarising meeting minutes, updating trackers and recording actions.
- Reviewing tender returns or Contractor Proposals for compliance.

These are the metaphorical 'dishes' of architecture, the necessary chores we are happy to have AI help us with. A well-configured (Enterprise Grade) LLM can for example cross-reference a internal finishes schedule against an architectural specification faster and more consistently than most humans. That's not replacing judgment; it's freeing up time for applying good judgment.



Image: Early testing with Nano Banana, the latest image model from Google Capable of taking mark up and annotations and rendering outstanding quality edits directly.

BUCKET 2: CREATIVE AMPLIFICATION (THE 'SHOP WINDOW')

This is the client-facing work that helps win projects and communicate a vision. These tools, like Midjourney, Stable Diffusion, Nano Banana and emerging video models like Kling are more like the practice's shop window. They are genuinely powerful for conceptual development and storytelling in the right hands, especially when a design direction is already established.

Where they shine:

- Competition-stage visuals that help the firm stand out.
- Early-stage mood boarding and rapid option testing.
- Persuasive client communication like video animation

...that go beyond drawings and traditional imagery.

RISK AWARENESS

To implement any of the ideas listed above with efficacy the architect operating the system must be fully aware of the risks at each stage, the need to prep their input data correctly, to use the right setup (with an enterprise-grade model, full context window and privacy precautions in place) and to spot-check outputs to confirm accuracy. We must exercise the same professional standards required by the Code of Conduct when we make use of AI output, because we are essentially adopting the output as our own and we will be held accountable for it.

THE UNIVERSITY DISCONNECT: IS AI A VENDING MACHINE OR A CHEF'S KNIFE?

Increasingly, small practices will be looking to hire for these 'two buckets' of need, but the academic response has been inconsistent at best and actively hostile at worst. I've heard of schools where any use of AI leads to disciplinary action, which is a reaction that feels unnervingly similar to banning calculators for fear of 'cheating'.

Perhaps the institutional debate can be framed with a simple question: is AI being used like a vending machine or like a chef's knife? To be clear, I believe AI will be increasingly capable of doing both, but it is up to leaders in the field to champion best practice.

For many universities, the answer is that AI is too often being used as an ideas 'vending machine'. They see a tool that dispenses a generic, pre-packaged product with the push of a button. The fear is that students use it as a shortcut, a substitute for rational and critical thought. This is not entirely unfounded. One lecturer recently shared with me how a student stopped developing their project the moment ChatGPT marked their portfolio against the learning outcomes and gave them an 'A'. The student treated the AI like a vending machine for a passing grade, and their intellectual engagement ended there. This view, that AI negates the learning process, is what drives such bans.

The alternative, and in my opinion the far more productive way of utilising AI is to wield it more as a 'Chef's Knife'. In the hands of a total novice, a chef's knife is useless or even dangerous. But in the hands of an expert, it is an extension of their will; a tool for precision, artistry and for creating something entirely new. While the vending machine has ideas of it's own, the knife is controlled by a chef; it requires and amplifies the skill, judgment, and intent of its operator.

This is the model for a productive relationship with AI: a powerful tool that augments the architect's craft and enables them to achieve more and to a greater level of quality. We need to be able to tell the difference and advocate for judicious use.

When schools assume AI is only a vending machine, they don't stop its use; they just drive it underground because we know the majority of students are using AI in some capacity, whether endorsed for use or not. The result is a widening gulf between students who are fluent in these tools and academic staff who are bewildered by the 'inexplicable' work they can produce. We risk releasing graduates who are either technically unskilled or, perhaps worse, skilled but ethically adrift, with no framework for how to wield these powerful capabilities responsibly.

THE ARCHITECTURE GRADUATE OF THE FUTURE

For me, the solution is not to ban AI but to make sure our teachers understand it comprehensively so that they can guide our future graduates. We must equip students with the skill, judgment and ethical framework to navigate a changing profession. This requires a two-pronged approach; targeting both the institutions and the students themselves.

THE SOLUTION FOR SCHOOLS: EVOLVE ASSESSMENT

First, schools must stop trying to ban the tools and instead evolve their assessment methods. The danger, as Richard Hall (of General Office) observes, is that in the worst cases students' use AI as a "secret substitute for rational and critical thought", resulting in an outcome where "the negation of the students intellectual engagement is blatant in their work". To counter this, the focus of assessment must shift from the polished final product to the intellectual process behind it.

In my opinion, the most powerful instrument we have for this is the viva voce, or oral examination. By assessing students through rigorous, in-person conversation, we can validate authentic understanding in a way a submitted document ever can.

To assess if students are actually exercising critical judgement, tutors will need to find ways to probe the design process. This cultural shift is advocated by Katy Marks (of Citizens Design Bureau), who says that when it comes to AI in universities we need a "changing culture around honesty and provenance. Every hand-in should involve a 'show your workings' moment which is both digital and physical."

These suggestions do demand a change in approach, they're not convenient ideas and they are also probably fraught with challenges around bias and subjectivity of assessment. However, it is likely very necessary in many fields of education as we consider how AI subverts traditional assessment methods. In an age of synthetic knowledge work and design, this return to dialogue might just be the key to making architectural education more human, not less.

For this to work, we must train the trainers. The greatest

bottleneck is often the knowledge gap in the teaching staff, who have been blindsided by the speed of this shift. Unless an educator really understands the technology, it will become every more difficult to discern critical use from cynical use.

AI BRINGING SPECIALISATION
GREATER RESPONSIBILITY IN SMALL PRACTICE

Ultimately, the reason for a small practice to embrace AI and for schools to teach it is not merely about cold efficiency. The real prize is strategic. The goal of my work with small practice is to make their methods more like themselves and less like everyone else. In a world threatened by AI-driven homogenisation, the most valuable position is to become harder to compete with, not easier to replicate.

Consider a firm specialising in high-end bespoke homes. For them, AI's power isn't in generating generic floor plans. It's in deepening the bespoke client briefing process, elevating the storytelling around materials and spatial experience and generating compelling, emotive imagery that steers design decisions early. It helps them amplify the very qualities that define their niche.

Or take a practice that focuses on community engagement. They can use AI to synthesise complex feedback from thousands of residents, generate visualisations that help communities to better understand proposals and alternative ideas instantly and to create inclusive communication materials in multiple formats. In these examples, the technology is deployed in a way that turns up the volume on the specialisms within a small practice.

If we zoom out and look more broadly at the evolving role of the architect over many decades, we can observe that the architect's role has been progressively 'unbundled', with project management, cost consulting and engineering disciplines splitting off into separate domains. AI offers the chance to reverse this trend and bring broader expertise back in-house. By automating the procedural scaffolding of project delivery, we may enable the architect to reclaim their central role in holistic project delivery, able to take on greater responsibility and become more essential, not less. I see architects picking up these new methods much more quickly and comprehensively than our Design Team colleagues and sub-consultants, and this is to our benefit. By equipping



Image: Further tests with Google's Nano Banana. Left Image is the editing render with new landscape design, right hand image is an entirely synthetic pen and ink drawing of the same proposal, featuring striking consistency to the original.

graduates with these skills, we are preparing them to do more than just survive in a new technological landscape, we are preparing them to lead it.

So, what is the 'Telos' of an architecture school in the age of AI? It is no longer a choice between producing critical thinkers and competent professionals; we need a mandate to produce both in a single graduate. The student who has been taught to wield AI as a 'Chef's knife' rather than as an 'Idea Vending Machine' is exercising the very critical thought my professor championed many years ago, but in a way that makes them an immediate asset to small practice. By embracing this new model, we will equip graduates to walk into a small firm not as trainees but as catalysts, ready to move the profession forward.

"...ARUP'S RESEARCH INTRODUCES A HYBRID METHODOLOGY THAT LEVERAGES THE STRENGTHS OF BOTH SIMULATION AND MACHINE LEARNING..."



ARUP
CONOR BLACK
ASSOCIATE, TECHNICAL SERVICES
DIGITAL PRACTICE LEADER
LONDON



The Architecture of Behaviour: Predicting Human Movement with AI

The integration of artificial intelligence into architectural design is reshaping how we understand and shape the built environment. One of the most promising frontiers in this transformation is the use of AI to predict human movement within architectural spaces. This case study explores a ground-breaking methodology developed by Arup. The approach combines agent-based simulations, probabilistic behavioural modelling, spatial feature analysis, and machine learning to predict how people move through workplace environments. The result is a scalable, data-driven framework that significantly reduces simulation time while maintaining high predictive accuracy.


to a comprehensive pipeline that transforms raw spatial and behavioural data into actionable design insights.

COMPONENT 1

The first component involves simulating human movement using MassMotion®, a pedestrian dynamics software developed by Oasys (Oasys is Arup's software house that offers a suite of advanced engineering software tools that support structural, geotechnical, and pedestrian movement analysis for complex projects). Unlike its conventional use in modelling emergency egress scenarios, MassMotion® is employed here to simulate a full ten-hour work day, capturing the nuanced behaviours of individuals as they navigate a workplace floorplate.

The simulation is grounded in the 'Social Forces Model', which conceptualises human movement as a response to various forces acting upon an individual. These include attractive forces that draw agents toward specific goals, such as meeting rooms or amenities; repulsive forces that encourage agents to maintain personal space and avoid obstacles; and velocity forces that adjust an agent's speed based on terrain and congestion. These forces are mathematically represented through non-linear Langevin equations, enabling agents to dynamically adapt to their environment in a realistic manner.

Obstacles within the environment are mapped based on their proximity to walkable surfaces, and approach maps are generated to determine the shortest paths to various goals. The software automatically translates the geometric layout into a network of nodes and links, allowing agents to evaluate multiple routes based on perceived cost. This cost is calculated by analysing the distance, congestion, and terrain type associated with each route. Agents then select the most efficient path to their destination, ensuring that their movement reflects both spatial constraints and behavioural tendencies.




Actual Massmotion® agent simulation, 30 minutes simulation time (left) and machine learning prediction, one second prediction time (right)

UNDERSTANDING HUMAN MOVEMENT

Predicting human movement in buildings is inherently complex. Human behaviour is influenced by a multitude of factors, including personal preferences, social interactions, environmental stimuli, and spatial constraints. Traditional simulation methods, such as agent-based modelling, can capture this complexity but are often computationally expensive and time-consuming. Arup's research addresses this challenge by introducing a hybrid methodology that leverages the strengths of both simulation and machine learning.

Workplaces are becoming more and more rich in sensor data, which provides a robust foundation for training and validating predictive models. Understanding how people navigate office spaces can inform critical decisions about layout, amenities, and space utilisation, ultimately enhancing productivity, well-being, and operational efficiency.

The methodology is composed of four interconnected components: agent movement simulation, probabilistic persona scheduling, contextual spatial analysis, and machine learning prediction using supervised learning. Each component contributes



Dataset of floorplates used as training data and testing data. Approx. 3000 workplace settings

COMPONENT 2

The second component introduces behavioural realism using probabilistic personas. These personas are constructed using real-world sensor data collected from Arup's own office environments. The data sources include turnstile logs that track building entry and exit times, occupancy sensors that monitor space usage, and meeting booking systems that provide insights into scheduled activities.

Each persona is defined by a set of probabilities that represent the likelihood of engaging in various activities throughout the day. For example, a persona might be characterised by a 50% probability of focused work, a 20% probability of collaboration, a 10% probability of attending meetings, a 10% probability of informal conversations, and a 10% probability of using amenities. These probabilities are used to generate daily schedules for each agent, ensuring that their behaviour reflects the diversity and variability observed in real workplaces.

Once a schedule is assigned, agents must decide which specific spaces to occupy for each activity. This decision-making process is dynamic and context-sensitive. Agents evaluate available settings based on their current location, proximity to their designated desk or 'home base', and the availability of suitable spaces. If no appropriate space is available for a scheduled activity, the agent may return to their desk unsatisfied or, in extreme cases, exit the simulation entirely. This approach introduces a layer of behavioural complexity that mirrors real-world constraints and preferences.

COMPONENT 3

The third component involves a detailed spatial analysis of each workplace setting. A typical office floorplate may contain hundreds of distinct settings, such as desks, meeting rooms, and collaboration zones. Each of these settings is analysed across ten spatial dimensions that are hypothesized to influence human movement.

These dimensions include the amount of daylight received by the setting, the number of other settings visible from that location (calculated using Isovist analysis), and the local density of work settings in the immediate vicinity. Additional features include the proximity to windows, amenities, circulation paths, louder spaces, and main entrances or exits. The type of setting (e.g., desk, meeting room) and its position within a row are also considered.

This analysis is conducted using a digital twin of the workplace created in Autodesk Revit®, with data extraction and processing performed using Rhino.Inside.Revit® and Grasshopper®. The result is a rich, multidimensional dataset that captures the geometric and functional context of each setting. These features serve as the input variables for the machine learning model.

COMPONENT 4

With the feature dataset in place, the final component involves training a machine learning model to predict movement patterns. A supervised learning technique (Support Vector Regression) was selected due to its ability to handle high-dimensional, nonlinear problems and its robustness in the presence of outliers. SVR is particularly well-suited for regression tasks where the goal is to estimate a continuous-valued function based on complex input data.

The model is trained using the movement data generated by the MassMotion® simulations, which serve as the target labels. Once trained, the model can predict movement patterns for new layouts in a matter of seconds, bypassing the need for time-intensive simulations. This capability dramatically accelerates the design-to-analysis cycle, enabling rapid iteration and optimisation.

The technical implementation integrates multiple software environments. The digital twin is maintained in Revit®, spatial analysis is conducted in Grasshopper® using C#, and the machine learning model is implemented in Python using the Sci-kit Learn library. This modular architecture ensures compatibility with existing design workflows and facilitates future enhancements.

EVALUATION


The model's performance was evaluated using standard regression metrics, including Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE). The results showed a normalised MAE of 0.12 and a MAPE of 22%, indicating strong predictive accuracy. Visual comparisons between the simulated and predicted movement maps revealed a high degree of correlation, with key movement hotspots and underutilised areas accurately identified.

One of the most significant advantages of this approach is the reduction in computation time. Simulating a full work day for over 200 agents can take several hours using traditional methods, whereas the model can generate predictions in under a second. This scalability makes the methodology suitable for large-scale projects and supports real-time design exploration.

This methodology represents a significant advancement in the application of AI to architectural design. By enabling rapid, accurate predictions of human movement, it empowers architects and planners to make data-driven decisions that enhance spatial efficiency, user experience, and operational performance. The ability to simulate and predict behaviour at scale opens new possibilities for responsive, human-centric design. Architects can test multiple layout scenarios, identify potential bottlenecks, and optimise space allocation based on predicted usage patterns. This approach also supports inclusive design by accommodating diverse user needs and behaviours.

One area of interest is the application of the methodology to other building typologies, such as university campuses, laboratories, and healthcare facilities. Each of these environments presents unique behavioural patterns and spatial challenges that could benefit from predictive modelling. Finally, there is a growing interest in modelling neurodiverse personas and individuals with varying mobility needs. Incorporating these considerations into the simulation framework would support more inclusive and equitable design outcomes.


Arup's AI-driven methodology for predicting human movement exemplifies the transformative potential of machine learning in architecture. By combining spatial geometry, behavioural data, and advanced analytics, the approach enables faster, smarter, and more human-centred design. As the architecture, engineering, and construction industry continues to embrace digital innovation, tools like this will become essential for shaping the future of our built environment.



Massmotion® simulation of floorplate (axonometric view). Darker blue purple shows the highest movement areas, the whiter areas less movement.



Digital twin (Autodesk Revit®) model of office, each Revit family is a workplace setting where the information about each individual setting is held.



Examples of generated heat maps to explore proximities to stimuli.



DEBRA POTHIER
PRINCIPAL AECO STRATEGY
AUTODESK EDUCATION EXPERIENCES
BOSTON



The Future of Architecture in the Age of AI: A UK Student's Guide to Skills, Sustainability, and Success.

You are witnessing one of architecture's most profound transformations. As an architecture student in the UK today, you're not just learning to design buildings. You're preparing to reshape how Britain and the world address the defining challenges of our time: the climate emergency, housing crisis, rapid urbanisation, and social justice through design.


This isn't about choosing between human creativity and machine intelligence. It's about orchestrating their collaboration to create architecture that performs better, sustains longer, and serves more equitably than ever before. Your generation will author this future, but only if you understand the tools and mindset required to lead it within the evolving UK regulatory and professional landscape.

THE STRATEGIC DESIGNER: YOUR NEW PROFESSIONAL IDENTITY

The architect as solitary genius sketching masterpieces is a romantic myth that no longer serves our complex world. Currently, an architect's role includes system design, environmental considerations, and social innovation. You're evolving into someone who orchestrates intelligent tools to solve problems that no single human mind could tackle alone.

Consider AI not as your competitor but as your most sophisticated design partner. While AI can generate thousands of design iterations and analyse performance data instantaneously, you provide what no algorithm can: cultural intuition, ethical judgment, creative vision, and the ability to translate human needs into spatial experiences. This partnership amplifies your impact rather than diminishing your role.

Your emerging responsibilities as a strategic designer include defining project outcomes that balance competing priorities, interpreting AI-generated insights to make informed decisions, ensuring technology serves human and environmental needs,



Multicolored modern office building in Deventer, Netherlands

and communicating design intent across diverse communities. Various AI-enabled design platforms exemplify this collaborative approach—they don't replace your design thinking but accelerate your ability to test concepts, analyse conditions, and explore alternatives you might never have imagined independently.

AI FLUENCY: YOUR NEW DESIGN LANGUAGE

Understanding how AI interprets data and generates alternatives will become as fundamental as understanding structural principles. This doesn't require programming expertise, but rather fluency in communicating with AI tools, critically evaluating their outputs, and integrating their insights into your design process.

Begin experimenting with AI-powered design tools during your studies. Learn to formulate precise questions, provide clear parameters, and critically assess AI-generated solutions. This literacy will distinguish you in a profession where AI collaboration becomes standard practice. The ARB's evolving competency framework increasingly recognises technological literacy as essential to professional practice.

COLLABORATIVE BIM WORKFLOWS: THE FOUNDATION OF MODERN PRACTICE

Building Information Modeling (BIM) has evolved from specialised skill to essential foundation. But mastering BIM isn't

just about software proficiency—it's about understanding how data-rich, collaborative workflows enable superior buildings and more efficient processes.

Deep knowledge of BIM platforms is crucial because contemporary architecture is inherently collaborative. You'll coordinate with engineers, contractors, sustainability consultants, and community stakeholders throughout every project. Your ability to share information seamlessly and maintain project coordination directly impacts outcomes. Post-Brexit, these collaborative skills become even more critical as UK firms navigate new international partnerships and regulatory frameworks.

OUTCOME-BASED BIM: THE EMERGING PARADIGM

The architecture profession is embracing what industry leaders call "Outcome-Based BIM." Instead of starting with form and hoping for good performance, this approach begins by defining desired outcomes—sustainability targets, affordability goals, user experience metrics—then uses AI and BIM tools to explore optimal paths to achieve them.

In this model below, you become a strategic orchestrator of intelligent tools. You might explore multiple site layout options simultaneously while monitoring real-time performance metrics, evaluate trade-offs between cost, environmental impact, and user experience across dozens of alternatives, and identify solutions that balance competing priorities more effectively than manual analysis allows.

This isn't theoretical. Students in pilot programs at UK universities are already designing this way, producing work that's both more innovative and more performatively successful than traditional approaches. These methods align with the ARB's emphasis on competency-based assessment and RIBA's focus on measurable outcomes.

SUSTAINABILITY AS DESIGN DRIVER

The climate emergency isn't a future concern—it's reshaping architecture now. The UK's commitment to net-zero by 2050 means you must become fluent in sustainable design principles and comfortable with tools that evaluate environmental performance in real time during early design phases when modifications are still feasible and cost-effective.

Various platforms now enable you to assess carbon footprints, energy performance, and climate impacts while designs are malleable. This isn't about adding sustainable features as afterthoughts—it's about making environmental performance a fundamental driver of design decisions from the first sketch to final construction, aligning with RIBA's 2030 Climate Challenge targets.

AI AND SUSTAINABILITY: STRATEGIC ALLIANCE

You might worry that AI's computational requirements conflict with sustainability goals. While AI systems consume significant energy, strategic application enables more sustainable design outcomes than traditional methods.

AI can evaluate thousands of design alternatives in the time required to manually analyse just a few. This rapid iteration capability helps identify high-performing sustainable solutions early in the design process when changes remain feasible. Various AI-enabled analysis tools help you understand how different materials and building forms affect carbon footprints from initial planning phases, making sustainability a design driver rather than an afterthought.

AI also reduces construction waste by optimising material usage, improving prefabrication processes, and minimising design errors. UK projects are beginning to demonstrate this potential: recent social housing developments in Manchester and Birmingham have used AI-powered analysis to deliver affordable units faster, cheaper, and with significantly reduced carbon footprints compared to traditional approaches.

BREXIT AND INTERNATIONAL COLLABORATION

Post-Brexit, UK architects face new challenges and opportunities in international collaboration. AI tools can help bridge geographical and regulatory gaps, enabling seamless collaboration with European and global partners despite changing legal frameworks. Understanding how to leverage these technologies for international projects while keeping compliance with UK regulations becomes increasingly valuable.

The UK's position as a global leader in sustainable design and architectural innovation can be strengthened through strategic use of Autodesk AI tools that prove superior environmental and social outcomes to international clients and partners.

PRACTICAL PREPARATION AND AVAILABLE RESOURCES

You don't need to wait until graduation to build these capabilities. Many software providers, like Autodesk offer comprehensive free access to professional-grade tools for students and educators, including AI-assisted design platforms like Autodesk Forma, Autodesk Revit, BIM modelling software, and Autodesk Construction Cloud, and curated curricula. Specific to Architecture, check out Autodesk Revit Architecture Professional and Autodesk Revit User certification that can assist you with CV-ready credentials.

These programs go beyond software training, building the diverse skills future architects need. Start experimenting with various tools during studio projects. Use AI analysis to inform design decisions. Practice collaborative workflows with classmates. Most importantly, begin evaluating projects through measurable outcomes, not just visual appeal.

"...CREATIVITY, SPATIAL THINKING, CULTURAL SENSITIVITY AND DESIGN INTUITION REMAIN CENTRAL TO PRACTICE..."



AI Co-Creation, optimising complex outcome criteria including sustainability, cost, time to deliver, habitability, diversity, net operating income, human experience and affordability.

YOUR CAREER IN PROFESSIONAL CONTEXT

Understanding these technological shifts is crucial for career planning within the UK context. Practices that will thrive deliver better buildings faster and more sustainably while navigating post-Brexit regulatory changes and market conditions. They need architects who can use AI effectively while keeping human insight that makes architecture meaningful.

Traditional skills don't become worthless—creativity, spatial thinking, cultural sensitivity, and design intuition remain central to practice. But these human capabilities will be most valuable when combined with technological fluency and outcome-oriented thinking that aligns with evolving ARB and RIBA standards.

The architects who will lead the profession bridge the gap between human needs and technological capabilities, define meaningful goals and direct AI tools to achieve them, and ensure technological innovation serves broader social and environmental purposes while keeping professional standards expected in UK practice.

AUTHORING THE FUTURE

AI, sustainability, and collaborative design are driving architecture's biggest change since modernism. As a UK student, you can influence this change instead of just responding to it. Approach AI not as a threat to creativity but as an amplifier of your design capabilities. When you can rapidly evaluate ideas, analyse performance, and explore alternatives, you're free to focus on what humans do best: understanding needs, creating meaning, and imagining better futures.

Your generation will design buildings and cities that address the climate emergency, accommodate urban growth, and create more fair communities within the UK and globally. AI will be essential in this work, but success depends on your ability to define good outcomes and ensure technological capabilities serve human and environmental flourishing while meeting professional standards.

The future of architecture lies in blending human creativity with AI to design better buildings and communities for people and the planet. Your education and early career choices today decide how well prepared you are to lead that future within the evolving landscape of UK architectural practice.

Start now. Experiment with AI-powered tools like Autodesk Forma. Develop sustainability fluency. Learn collaborative workflows. Most importantly, begin thinking like the strategic designer you're becoming—one who uses every available tool, including AI, to create meaningful, measurable positive outcomes through design while upholding the highest standards of UK architectural practice. The future of architecture is being written now, and you are one of its authors.

KEY SUGGESTIONS & KNOWLEDGE SKILLS & BEHAVIOURS


Key Suggestions and Skills for Future AI Use in Architectural Education

As AI tools are increasingly embedded in education and practice, our discipline needs to adopt clear, proactive strategies. The question is no longer if students will use AI, but rather, how can schools ensure that such use develops into critical and ethical practice. We therefore propose suggestions, together with a taxonomy of skills and behaviours that future architects may cultivate using new AI tools:

SUGGESTIONS FOR SCHOOLS OF ARCHITECTURE

1. **Embed AI literacy in the curriculum.** Staff and students need to understand how AI models are trained, what datasets they use, and the potential ethical issues and bias this creates. This is important learning across all schools, even if AI is not permitted for use as part of assessment. Teaching could extend beyond 'how to use the tool', towards a critical awareness and reflection on bias, provenance and accuracy.
2. **Focus on the process as much as on the outcome.** Assessment structures could reward the ability to critically reflect on how AI was used and developed in design, written and report work, including the development of datasets, prompts, iterations, and model 'tuning'. Assessments may need to change to give more weighting to critical verbal reasoning during reviews to avoid the 'black box' of an AI process's input-output. Students may be asked to expand upon their workflow - how did they work with AI, what did they learn, what were their prompts, what does it mean to their project?
3. **Establish clear policies on authorship and attribution.** Schools should adopt clear guidelines on how AI is used and how contributions are referenced in portfolios, essays and design submissions, which may include the prompt, dataset, model and location of the work that was contributed as a result.
4. **Develop AI-resilient forms of assessment.** Design reviews, verbal examinations and reflective journals reduce the risk of students 'outsourcing' core design / intellectual work directly to AI. Altering assessment types in this way may also help to reinforce the importance of iterative development and professional accountability

5. **Encourage responsible experimentation.** Students could be given opportunities to test AI as part of their creative process, within a framework that encourages them to explain their choices in the knowledge of sustainable or ethical concerns and technical weaknesses.
6. **Prioritise equity and access.** With a vast array of AI tools proliferating the market at different price points, schools may wish to ensure that students are not disadvantaged by lack of access and confusion over choice. Institutions could ask for support from their IT departments about provision of platforms or licences to support students in their experimentation. Consider asking for an IT department's policy on subscription model license distribution e.g. Adobe Firefly as a commercially defensible IP image generator alongside private LLMs with closed, licensed or proprietary datasets such as Cohere or Azure. If Universities are unable to provide access, schools could ask for IT support to demonstrate the difference between these private, database-restricted AI tools in comparison with open cloud-based, general access AI tools.
7. **Build staff capacity.** AI literacy is not just a student issue - design tutors, examiners and lecturers are perhaps in greater need of training and support now in order to recognise, guide and assess AI-assisted work in an informed and consistent way. This report goes some way to help identify this need - further tools and a repository of links and training for both staff and students can be found in the following Repositories and Learning Resources section (p.82).



Drawing Operations Unit: Generation_1 (DOUG) by Sougwen Chung



Interactive Robotic Plastering ETH Zurich 2021 Gramazio Kohler Research

AI KNOWLEDGE, SKILLS AND BEHAVIOURS

The integration of AI tools into architectural education is an inevitability, but its form is not yet predetermined. Schools of architecture have a unique opportunity to shape their own culture of AI use: to balance creativity with responsibility and innovation with integrity. By embedding literacy, authorship, and reflective skills into our curricula, we aim to educate professionals who can engage critically with AI whilst upholding the core values of creativity, accountability, and ethical practice that define our discipline. Drawing on our suggestions, and the report more generally, the following may be considered as nascent graduate attributes for young architects entering an increasingly AI-augmented profession:

KNOWLEDGE

Critical AI Literacy | The ability to question the origin, accuracy and ethics of AI outputs – an awareness of dataset limitations; the capacity to explain what an AI tool has contributed, and why it was used.

Professional Responsibility | Recognition that in practice, the architect is often accountable for all decisions, regardless of which tools were used to develop or contribute work.

SKILLS

Transparency and Referencing | Practice of acknowledging AI use in all forms of work including the text, image, model, contribution to thought or coding. Ability to cite prompts, name tools and their datasets, and situate outputs within work processes.

Reflective Practice | Ability to articulate how and why AI has shaped design decisions, including points of agreement, divergence and iteration.

Verification and Compliance | Ability to test AI outputs against external standards – regulatory, technical, and professional benchmarks - before accepting them as valid e.g. building regulations, planning frameworks, and the ethical obligations of the architect.

BEHAVIOURS

Collaboration with AI | Use of AI as a tool for collaboration and not as a means to an end - how AI outputs can be integrated into practice and across wider discourse with peers, clients and communities.



DEREK HALES
DIRECTOR OF STUDIES IN
PRACTICE FUTURES
SALFORD LABORATORY OF
ARCHITECTURE
UNIVERSITY OF SALFORD

Dr Derek Hales is Director of Studies in Practice Futures at the Salford Laboratory of Architecture, co-convenor of the Centre of Excellence in Practice as Research at the University of Salford, and Research Fellow in Transdisciplinary Practice with the New Centre for Research & Practice.

APPENDIX 1: REPOSITORIES AND LEARNING RESOURCES

AI Tutorials and Repositories

Hands-on workshops in our studios, combined with online tutorials have become central to GenAI learning because they enable us to adapt faster than formal syllabi and bridge between disciplines. University curricula move slowly and especially so in our discipline. GenAI tools evolve weekly across coding, design, media, and ethics and the short and long-form tutorial format, offer practice-led knowledge in real time (building extended knowledge communities) and asynchronously. Tutorial formats appeal in

fragmented, interdisciplinary, domains where no monolithic syllabus can adequately cover things. Circulating through YouTube, GitHub, and Discord, tutorials form a living syllabus – modular, iterative, and community-driven. Beyond instruction, they operate culturally, rewarding openness and experimentation. In this sense, tutorials, including the examples below, embody GenAI itself, they are fast, distributed, interdisciplinary, and constantly reconfigured.

KEEPING CONNECTED

Staying informed in GenAI requires the layering of fast updates, engagement in community discussion, and self-curated resources beyond the scope of this report. For rapid news, sites like Hugging Face and GitHub trending track new models and repos, while YouTube host fresh tutorials. Community spaces including Discord servers, Reddit, and X provide real-time tips and fixes. To filter

noise, newsletters hosted on Medium or Substack like Latent Space offer concise digests. Longer-term perspective comes from conferences and research blogs. Finally, using personal curation tools, for example, using RSS feeds combined with platforms like Notion can help each of us build our own living syllabus. Learning to balance speed, signal, and reflection prevents miasma induced overload whilst keeping knowledge current. ►

USE CASE	RESOURCE	LINK
AI FOUNDATIONS TUTORIALS	GOOGLE – LEARN AI SKILLS	HTTPS://AI.GOOGLE/LEARN-AI-SKILLS/
	ELEMENTS OF AI	HTTPS://WWW.ELEMENTSOFAI.COM/
	KAGGLE LEARN – INTRO TO MACHINE LEARNING	HTTPS://WWW.KAGGLE.COM/LEARN/INTRO-TO-MACHINE-LEARNING
	MICROSOFT LEARN – ML FUNDAMENTALS	HTTPS://LEARN.MICROSOFT.COM/EN-US/TRAINING/MODULES/FUNDAMENTALS-MACHINE-LEARNING/
	LINKEDIN LEARNING: INTRODUCTION TO ARTIFICIAL INTELLIGENCE	HTTPS://WWW.LINKEDIN.COM/LEARNING/INTRODUCTION-TO-ARTIFICIAL-INTELLIGENCE-24947908
LLM PROMPT ENGINEERING TUTORIALS	LEARN PROMPTING	HTTPS://LEARNPROMPTING.ORG/
	MICROSOFT LEARN – PROMPT ENGINEERING	HTTPS://LEARN.MICROSOFT.COM/EN-US/AZURE/AI-FOUNDRY/OPENAI/CONCEPTS/PROMPT-ENGINEERING
	LINKEDIN LEARNING: INTRODUCTION TO PROMPT ENGINEERING FOR GENERATIVE AI	HTTPS://WWW.LINKEDIN.COM/LEARNING/INTRODUCTION-TO-PROMPT-ENGINEERING-FOR-GENERATIVE-AI-24636124
IMAGE & VIDEO BASED AI TUTORIALS	MIDJOURNEY – GETTING STARTED GUIDE	HTTPS://DOCS.MIDJOURNEY.COM/HC/EN-US/ARTICLES/33329261836941-GETTING-STARTED-GUIDE
	PROMPTHERO – MIDJOURNEY GUIDE	HTTPS://PROMPTHERO.COM/MIDJOURNEY-PROMPTS
	ADOBE FIREFLY – OVERVIEW & TUTORIALS	HTTPS://HELPIX.ADOBE.COM/FIREFLY/WEB/GET-STARTED/LEARN-THE-BASICS/ADOBE-FIREFLY-OVERVIEW.HTML
	RUNWAY ML – ACADEMY TUTORIALS	HTTPS://ACADEMY.RUNWAYML.COM/
	LINKEDIN LEARNING: STABLE DIFFUSION: TIPS, TRICKS, AND TECHNIQUES	HTTPS://WWW.LINKEDIN.COM/LEARNING/STABLE-DIFFUSION-TIPS-TRICKS-AND-TECHNIQUES
	COMFYUI.ORG – TUTORIALS	HTTPS://COMFYUI.ORG/EN/COLLECTIONS/TUTORIAL
	LINKEDIN LEARNING: LEARNING COMFYUI FOR STABLE DIFFUSION	HTTPS://WWW.LINKEDIN.COM/LEARNING/LEARNING-COMFYUI-FOR-STABLE-DIFFUSION
	NANO BANANA - VIDEO TUTORIALS	HTTPS://WWW.YOUTUBE.COM/@URBAN_DECODERS/VIDEOS
GRASSHOPPER + AI TUTORIALS	YOUTUBE – ATELIER DESIGN ACADEMY	HTTPS://WWW.YOUTUBE.COM/@ATELIERDESIGNA
	YOUTUBE – LUCIANO AMBROSINI	HTTPS://WWW.YOUTUBE.COM/@LUCIANOAMBROSINI
AEC-FOCUSED VISUALISATION AI TUTORIALS	YOUTUBE – PARAMETRICARCHITECTURE	HTTPS://WWW.YOUTUBE.COM/C/PARAMETRICARCHITECTURE
	YOUTUBE – AEC MAGAZINE	HTTPS://WWW.YOUTUBE.COM/@AECMAGAZINE3860/FEATURED

Note: External links are provided for reference only; inclusion does not imply endorsement, and availability may change over time.

ENHANCED PEDAGOGICAL PROGRESSION

Beyond basic tutorials, designing an enhanced pedagogical structure will enable architectural educators to explore curricula with multiple pathways: students can begin with simple interfaces (Midjourney, ChatGPT), progress through visual programming (ComfyUI, Grasshopper), explore integrated platforms (xFigura), and advance to technical development (PyTorch, custom scripting). The inclusion of architecture-native tools like Grasshopper ensures relevance to professional practice. Where distinctions between layers is somewhat artificial it is a useful device for thinking through relations between AI Models, AI & Python Libraries, and AI & API Integrations, and aligns each concept to a tripartite taxonomy (User, Technical, AEC Integration). The curated learning resources given can then be structured to fit the appropriate layer. Universities with access to linkedInLearning already have curated sets of learning resources.

DEFINING CROSS-LAYER TOOLS

ComfyUI: Represents a bridge tool that spans User and Technical layers. Students begin with pre-built workflows (User Layer) but can progress to custom node development and Python scripting (Technical Layer). Its visual programming interface makes complex AI workflows accessible while maintaining scalability for advanced users. ComfyUI looks set to disrupt rendering and video-based workflows considerably.

Figura: Exemplifies the AI platform aggregation trend, integrating SDXL, FLUX, Gemini Flash, Ideogram, Tripo (3D), CSM, and Magnific upscaling within a unified interface. Such aggregation

is likely to be valuable for architectural visualisation, offering both 2D and 3D AI capabilities without requiring multiple subscriptions or technical integration.

Grasshopper AI Integration: Unique cross-layer tool spanning all three taxonomy levels. User Layer: drag-and-drop AI components (Owl, Crow, Dodo). Technical Layer: custom Python/C# scripting with ML libraries. AEC Integration: native workflow integration within Rhino, the industry-standard architectural modeling environment. Demonstrates how familiar tools evolve to incorporate AI, and indeed has been doing so behind the scenes as it were with tools such as LunchBox pre-dating the GenAI phevnomena.

INTEROPERABILITY

Where at the level of ‘tools’ we can categorise things using the three-layer strata discussed above, at the level of the dataset (and more broadly in the reality of AI collaborative practices) we are better talking about interoperability and cross-layer interactions. For example, ComfyUI essentially democratises access to technical layer datasets while maintaining the visual, user-friendly interface of its node-based graph editor. This perspective shows that with ComfyUI, we’re not just talking about three separate layers - we’re talking about a unified ecosystem where, for example, the environmental datasets we’ve referenced as exemplar can flow seamlessly from APIs to visualizations to BIM applications. The Python backend of ComfyUI ensures technical depth while the visual interface ensures accessibility: Python is the glue that can bind your AI practices across the layers discussed.

Taken together, these examples show that the three-layer structure is not a rigid taxonomy but a generative framework for

TOOL	PRIMARY LAYER	INTERFACE TYPE	LEARNING CURVE	ARCHITECTURAL RELEVANCE
MIDJOURNEY	User	Text prompts	Low	Conceptual visualization
COMFYUI	User→Technical	Visual nodes	Medium	Advanced image workflows
XFIGURA	User→AEC	Unified platform	Low - Medium	Professional visualization
GRASSHOPPER+AI	User→Technical→AEC	Visual programming	Medium - High	Integrated design workflows
PYTORCH	Technical	Code framework	High	Custom model development
SPECKLE	AEC Integration	Data platform	Medium	Collaborative Workflows

Tool Progression Matrix

navigating rapid change. The pedagogical value lies in showing students how to move between layers, understand the affordances and limits of each, and critically evaluate their integration in architectural practice. As AI systems evolve, interoperability will become the defining condition of architectural workflows: what matters is not only knowing individual tools, but how datasets, APIs, and platforms interconnect.

For educators, this means moving beyond software training toward cultivating adaptive, critical, and practice-ready graduates who can both use and question the AI-driven tools shaping the discipline.

REPOSITORY	PRIMARY FOCUS	DATASET TYPE	AEC/ARCHITECTURE RELEVANCE	DATA QUALITY	ACCESS METHOD	BEST FOR AEC AI
KAGGLE	Machine Learning Competitions & Datasets	Tabular, Images, Text, Time Series, Geospatial	High - Building datasets, energy, urban planning competitions	High - Curated, competition-tested	Web UI + API	Building performance, energy modeling, urban analytics competitions
HUGGING FACE	NLP Models & Datasets	Text, NLP, Multimodal, Some Computer Vision	Medium - Text analysis for regulations, reports, documentation	High - Model-ready, well-documented	Python libraries (datasets, transformers)	Processing building codes, regulations, environmental reports
GITHUB	Code Repositories & Some Datasets	Code + Small-Medium Datasets, Documentation	High - Code + datasets, many environmental/GIS projects	Variable - Depends on repository maintainer	Git clone, direct download	Environmental monitoring code + data, GIS tools + datasets
GOOGLE DATASET SEARCH	Dataset Discovery Across Web	All Types (Search Engine)	Very High - Discovers datasets across all sources	Variable - Aggregates from many sources	Search + redirect to source	Discovering specialized architecture/climate datasets
AWS OPEN DATA	Cloud-hosted Public Datasets	Large-scale Cloud Datasets, Satellite, Climate	Very High - Climate, satellite, environmental data at scale	Very High - Enterprise-grade, well-maintained	AWS CLI, web interface, APIs	Large-scale climate, satellite, environmental analysis
PAPERS WITH CODE	Research Papers + Associated Datasets	Research Datasets Linked to Papers	High - Latest research datasets in sustainability, climate	High - Peer-reviewed research quality	Links to original sources	Latest research in sustainable design, climate adaptation
UCL ML REPOSITORY	Classic ML Datasets for Research	Small-Medium Tabular, Classification Datasets	Low - Limited architecture-specific datasets	High - Classic, well-established datasets	Direct download, some APIs	Baseline ML experiments, proof of concepts

Sample Dataset Repositories

APPENDIX 2

AI READING LIST

Benjamin, R. (2019). **Race after technology: Abolitionist tools for the New Jim Code**. Polity Press.

Browne, S. (2015). **Dark matters: On the surveillance of Blackness**. Duke University Press.

Carpo, M. (2023). **Beyond digital: Design and automation at the end of modernity**. MIT Press.

Carpo, M. (2023). Imitation games: Mario Carpo on the new humanism. **Artforum*, 61*(10), 184–188. <https://www.artforum.com/print/202306/mario-carpo-on-the-new-humanism-90638>

Chun, W. H. K. (2021). **Discriminating data: Correlation, neighborhoods, and the new politics of recognition**. MIT Press.

Crawford, K. (2021). **The atlas of AI: Power, politics, and the planetary costs of artificial intelligence**. Yale University Press.

Du Sautoy, M. (2019). **The creativity code: How AI is learning to write, paint and think**. Fourth Estate.

Miller, A. I. (2019). **The artist in the machine: The world of AI-powered creativity**. MIT Press.

Noble, S. U. (2018). **Algorithms of oppression: How search engines reinforce racism**. NYU Press.

Ostrom, E. (1990). **Governing the commons: The evolution of institutions for collective action**. Cambridge University Press.

Roberge, J., & Castelle, M. (Eds.). (2021). **The cultural life of machine learning: An incursion into critical AI studies**. Springer. <https://doi.org/10.1007/978-3-030-56286-1>


Steinfeld, K. (2021). Significant others: Machine learning as actor, material, and provocateur in art and design. In I. As & P. Basu (Eds.), **The Routledge companion to artificial intelligence in architecture** (pp. 3–12). Routledge. <https://doi.org/10.4324/9780367824259-2>

Vickers, B., & Allado-McDowell, K. (Eds.). (2020). **Atlas of anomalous AI**. Ignota Books.


APPENDIX 3

SURVEY RESULTS


Q3 STUDENT
WHAT YEAR/ COURSE OF ARCHITECTURAL EDUCATION ARE YOU IN




Q4
HOW WOULD YOU RATE YOUR PERSONAL KNOWLEDGE ABOUT AI IN GENERAL?




Q5
DO YOU AGREE WITH THE FOLLOWING STATEMENTS ABOUT AI?




Q6
PLEASE STATE WHETHER YOU AGREE OR DISAGREE WITH THE FOLLOWING STATEMENTS ABOUT AI




Q7
PLEASE INDICATE HOW MUCH YOU ARE USING AI IN THE FOLLOWING AREAS OF YOUR WORK




Q8
HOW MUCH DO YOU AGREE OR DISAGREE WITH THE FOLLOWING STATEMENT: "MY COURSE HAS CLEAR POLICY ON USING GENERATIVE AI IN WORK"




Q9
IN TERMS OF PERMITTED AI USE IN YOUR OWN WORK, DOES YOUR COURSE ALLOW AI IN THE FOLLOWING MODULES*** CHECK WORDING




Q10
HOW MUCH DO YOU AGREE OR DISAGREE WITH THE FOLLOWING STATEMENT: "MY LECTURERS AND TUTORS UNDERSTAND HOW STUDENTS ARE CURRENTLY USING GENERATIVE AI FOR ASSESSED WORK"




Q11
FOR ANY GIVEN PIECE OF ASSESSED WORK, HOW CONFIDENT ARE YOU THAT YOUR COURSE TUTORS CAN CORRECTLY DETERMINE WHETHER GENERATIVE AI WAS USED?




Q12
DO YOU FORSEE ETHICAL CONERNRS ARISING OUT OF THE ADOPTION OF AI IN THE *** CHECK WORDING




Q13
HOW CONCERNED ARE YOU WITH THE PARITY AND QUALITY OF YOUR WORK COMPARED TO OTHERS IN YOUR COHORT AND ACROSS SCHOOLS OF ARCHITECTURE WHO ARE USING AI?



Q8
PLEASE STATE HOW CONCERNED YOU ARE WITH THE FOLLOWING STATEMENTS



Q11
OVERALL, DO YOU THINK THE EFFECTS OF AI WILL BE POSITIVE OR NEGATIVE?



BIBLIOGRAPHY 1
DEFINING AI IN THE EDUCATION OF AN ARCHITECT

[1] El Moussaoui, M. (2025) Architectural practice process and artificial intelligence – an evolving practice, *Open Engineering*, 15(1), p. 20240098. <https://doi.org/10.1515/eng-2024-0098>

[2] Adeusi, T. S., Kachiside, O. L. and Gupta, R. (2024) Innovative architectural design practices enabled by AI-powered parametric and computational approaches, *World Journal of Advanced Research and Reviews*, 24(3), pp. 2487–2498. <https://doi.org/10.30574/wjarr.2024.24.3.3959>

[3] Kalay, Y. E. (1985) Redefining the role of computers in architecture: from drafting/ modelling tools to knowledge-based design assistants. *Computer-Aided Design*, 17(7), pp. 319–328. [https://doi.org/10.1016/0010-4485\(85\)90165-4](https://doi.org/10.1016/0010-4485(85)90165-4)

[4] Davis, D. and Peters, B. (2013) Design Ecosystems: Customising the Architectural Design Environment with Software Plug-ins, *Architectural Design*, 83(2), pp. 124–131. <https://doi.org/10.1002/ad.1567>

[5] Guidera, M. (2011) Conceptual Design Exploration in Architecture Using Parametric Generative Computing: A Case Study, Conference paper presented at 2011 ASEE Conference, June 2011. <https://doi.org/10.18260/1-2--17649>

[6] Matejka, J., Glueck, M., Bradner, E., Hashemi, A., Grossman, T. & Fitzmaurice, G. (2018) Dream Lens: Exploration and Visualization of Large-Scale Generative Design Datasets, in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI 2018, New York, NY, USA: ACM, pp. 1–12. <https://doi.org/10.1145/3173574.3173943>

[7] Horvath, A.-S. and Pouliou, P. (2024) AI for conceptual architecture: reflections on designing with text-to-text, text-to-image, and image-to-image generators, *Frontiers of Architectural Research*, 13(3), pp. 593–612. <https://doi.org/10.1016/j.foar.2024.02.006>

[8] Francis, N. J., Jones, S. and Smith, D. P. (2025) Generative AI in Higher Education: Balancing Innovation and Integrity, *British Journal of Biomedical Science*, 81, Article 14048. <https://doi.org/10.3389/bjbs.2024.14048>

[9] Kong, S.-C., Lee, J. C.-K. & Tsang, O. (2024) ‘A pedagogical design for self-regulated learning in academic writing using text-based generative artificial intelligence tools: 6-P pedagogy of plan, prompt, preview, produce, peer-review, portfolio-tracking’, *Research and Practice in Technology Enhanced Learning*, 19, Article 030. <https://doi.org/10.58459/rptel.2024.19030>

[10] Fava, M. (2019) ‘A decline in drawing ability?’, *International Journal of Art & Design Education*, 39(2), pp. 319–332. <https://doi.org/10.1111/jade.12255>

[11] Scheer, D. R. (2014) *The Death of Drawing: Architecture in the Age of Simulation*, 1st ed. New York: Routledge. <https://doi.org/10.4324/9781315813950>

BIBLIOGRAPHY 2

THE ETHICS OF AI IN ARCHITECTURAL EDUCATION

[1] Y. Netser, E. Cochran Hameen, and P. Tang, "Socio-Sustainable Architectural Design Through Ethical Implementation of Generative Design and Artificial Intelligence," in Proceedings of the 23rd CIB World Building Congress, Purdue University, West Lafayette, USA, 2025, vol. 1, 402 ed., doi: 10.7771/3067-4883.1095. [Online]. Available: <https://doi.org/10.7771/3067-4883.1095>

[2] N. J. Francis, S. Jones, and D. P. Smith, "Generative AI in Higher Education: Balancing Innovation and Integrity," (in English), British Journal of Biomedical Science, Review vol. Volume 81 - 2024, 2025-January-09 2025, doi: 10.3389/bjbs.2024.14048.

[3] C. Li, T. Zhang, X. Du, Y. Zhang, and H. Xie, "Generative AI models for different steps in architectural design: A literature review," Frontiers of Architectural Research, vol. 14, no. 3, pp. 759-783, 2025/06/01/ 2025, doi: <https://doi.org/10.1016/j.foar.2024.10.001>.

[4] C. Xu and Y. Huang, "Technological Innovation in Architectural Design Education: Empirical Analysis and Future Directions of Midjourney Intelligent Drawing Software," Buildings, vol. 14, no. 10, doi: 10.3390/buildings14103288.

[5] O. Asfour, "How Artificial Intelligence Could Affect the Future of Architectural Design Education," Resourceedings, vol. 4, no. 2, pp. 01-05, 09/30 2024, doi: 10.21625/resourceedings.v4i2.1113.

[6] J. McCormack, T. Gifford, and P. Hutchings, "Autonomy, Authenticity, Authorship and Intention in computer generated art," in EvoMUSART, 2019.

[7] Y. Mei, "Prompting the E-Brushes: Users as Authors in Generative AI," ArXiv, vol. abs/2406.11844, 2024.

[8] J. K. Eshraghian, "Human ownership of artificial creativity," Nature Machine Intelligence, vol. 2, no. 3, pp. 157-160, 2020/03/01 2020, doi: 10.1038/s42256-020-0161-x.

[9] J. Cudzik and K. Radziszewski, "Artificial Intelligence Aided Architectural Design," in Proceedings of the 36th International Conference on Education and Research in Computer Aided Architectural Design in Europe (eCAADe), Lodz, Poland, 2018: eCAADe, pp. 77-84. [Online]. Available: <https://www.researchgate.net/publication/328018944>. [Online]. Available: <https://www.researchgate.net/publication/328018944>

[10] A. Alshahrani and A. M. Mostafa, "Enhancing the use of artificial intelligence in architectural education – case study Saudi Arabia," (in English), Frontiers in Built Environment, Original Research vol. Volume 11 - 2025, 2025-June-18 2025, doi: 10.3389/fbuil.2025.1610709.

[11] L. Tan and M. and Luhrs, "Using Generative AI Midjourney to enhance divergent and convergent thinking in an architect's creative design process," The Design Journal, vol. 27, no. 4, pp. 677-699, 2024/07/03 2024, doi: 10.1080/14606925.2024.2353479.

[12] J. Oppenlaender, "The Creativity of Text-to-Image Generation," presented at the Proceedings of the 25th International Academic Mindtrek Conference, Tampere, Finland, 2022. [Online]. Available: <https://doi.org/10.1145/3569219.3569352>.

[13] V. Paananen, J. Oppenlaender, and A. Visuri, "Using text-to-image generation

for architectural design ideation," International Journal of Architectural Computing, vol. 22, no. 3, pp. 458-474, 2024/09/01 2023, doi: 10.1177/14780771231222783.

[14] N. Inie, J. Falk, and S. Tanimoto, "Designing Participatory AI: Creative Professionals' Worries and Expectations about Generative AI," Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems, 2023.

[15] J. Gillotte, "Copyright Infringement in AI-Generated Artworks," UC Davis Law Review, vol. 53, no. 5, 2019/07/23 2020. [Online]. Available: <https://ssrn.com/abstract=3657423>.

[16] Department_for_Education, "Generative AI: Call for Evidence – Summary of Responses," UK Government, London, UK, 2023. [Online]. Available: <https://www.gov.uk/government/publications/generative-ai-call-for-evidence-summary-of-responses>

[17] University_of_Edinburgh. "Guidance for Working with Generative AI ("GenAI") in Your Studies." University of Edinburgh Information Services. <https://information-services.ed.ac.uk/computing/comms-and-collab/elm/guidance-for-working-with-generative-ai> (accessed).

[18] University_College_London. "Engaging with Generative AI in Your Education and Assessment." UCL Assessment Success Guide. <https://www.ucl.ac.uk/students/exams-and-assessments/assessment-success-guide/engaging-generative-ai-your-education-and-assessment> (accessed).

[19] Quality_Assurance_Agency_for_Higher_Education, "Generative Artificial Intelligence in Education: QAA's Response to DfE Call for Evidence," QAA, 2023/09/07 2023. [Online]. Available: https://www.qaa.ac.uk/docs/qaa/news/generative-artificial-intelligence-call-for-evidence.pdf?sfvrsn=db01af81_8

[20] Architects_Registration_Board, "The Architects Code: Standards of Professional Conduct and Practice," Architects Registration Board, London, UK, 2025. [Online]. Available: <https://www.arb.org.uk>

[21] RIBA, "Perspectives on the Future of AI in Architecture," Royal Institute of British Architects, London, UK, 2024. [Online]. Available: <https://www.architecture.com/knowledge-and-resources/resources-landing-page/perspectives-on-the-future-of-ai-in-architecture>

[22] National_Council_of_Architectural_Registration_Boards, "NCARB's Position on the Use of Artificial Intelligence in the Architectural Profession," NCARB, Washington, DC, USA, 2024/06 2024. [Online]. Available: <https://www.ncarb.org>

[23] Russell_Group, "Principles on the Use of Generative AI Tools in Education," Russell Group, 2023/07/03 2023. [Online]. Available: <https://www.russellgroup.ac.uk/policy/policy-briefings/principles-use-generative-ai-tools-education>

[24] T. C. University_of_Manchester, "Artificial Intelligence and Teaching and Learning," 2025. [Online]. Available: <https://www.teachingcollege.fse.manchester.ac.uk/artificial-intelligence-and-teaching-and-learning>.

[25] C. Medel-Vera and W. F. Gates, "Deciphering Aesthetics: Exploring the Relationship Between Prompt Readability and AI-Generated Image Aesthetics," in

Architecture in the AI Era for Research, Practice and Pedagogy, M. H. Asterios Agkathidis, and Carlos Medel-Vera Ed.: Springer Nature, 2025.

[26]S. o. t. H. a. S. S. University_of_Cambridge. "Template Declaration of the Use of Generative Artificial Intelligence." University of Cambridge (CSHSS). <https://www.cshss.cam.ac.uk/education/generative-artificial-intelligence-ai-and-scholarship/template-declaration-use-generative> (accessed).

[27]D. Weber-Wulff et al., "Testing of detection tools for AI-generated text," International Journal for Educational Integrity, vol. 19, pp. 1-39, 2023.

[28]M. Perkins et al., "Detection of GPT-4 Generated Text in Higher Education: Combining Academic Judgement and Software to Identify Generative AI Tool Misuse," Journal of Academic Ethics, vol. 22, pp. 89-113, 2023.

[29]M. S. Orenstrakh, O. Karnalim, C. A. Suárez, and M. Liut, "Detecting LLM-Generated Text in Computing Education: Comparative Study for ChatGPT Cases," 2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC), pp. 121-126, 2023.

[30]W. Liang, M. Yuksekgonul, Y. Mao, E. Wu, and J. Zou, "GPT detectors are biased against non-native English writers," Patterns, vol. 4, no. 7, 2023, doi: 10.1016/j.patter.2023.100779.

[31]M. Perkins et al., "Simple techniques to bypass GenAI text detectors: implications for inclusive education," International Journal of Educational Technology in Higher Education, vol. 21, no. 1, p. 53, 2024/09/09 2024, doi: 10.1186/s41239-024-00487-w.

[32]University_of_Sheffield. "AI and Academic Integrity." University of Sheffield. <https://sheffield.ac.uk/study-skills/assessment/academic-integrity/ai-and-academic-integrity> (accessed).

BIBLIOGRAPHY 3

AI IN EDUCATION

[1] Bottazzi, R., Varoudis, T., Prajapati, P., & Wang. (2022). N2 P2 – Neural Networks and Public Places. In S. Carta, Machine learning, artificial intelligence and urban assemblages: Applications in architecture and urban design (pp. 177–181). John Wiley & Sons.

[2] Steinfeld, K. (2021). Significant Others—Machine learning as actor, material, and provocateur in art and design. In I. As & P. Basu (Eds.), The Routledge Companion to Artificial Intelligence in Architecture (1st ed., pp. 3–12). Routledge. <https://doi.org/10.4324/9780367824259>

[3] Witt, A. (2021). Formulations: Architecture, mathematics, culture. The MIT Press.

BIBLIOGRAPHY 4

DEVELOPING AI AUGMENTED
CAPACITY TO SUPPORT BUILDING
PERFORMANCE EVALUATION AT THE EARLY DESIGN PHASE

news/generative-artificial-intelligence-call-for-evidence.pdf?sfvrsn=db01af81_8

[20]Architects_Registration_Board, "The Architects Code: Standards of Professional Conduct and Practice," Architects Registration Board, London, UK, 2025. [Online]. Available: <https://www.arb.org.uk>

[21]RIBA, "Perspectives on the Future of AI in Architecture," Royal Institute of British Architects, London, UK, 2024. [Online]. Available: <https://www.architecture.com/knowledge-and-resources/resources-landing-page/perspectives-on-the-future-of-ai-in-architecture>

[22]National_Council_of_Architectural_Registration_Boards, "NCARB's Position on the Use of Artificial Intelligence in the Architectural Profession," NCARB, Washington, DC, USA, 2024/06 2024. [Online]. Available: <https://www.ncarb.org>

[23]Russell_Group, "Principles on the Use of Generative AI Tools in Education," Russell Group, 2023/07/03 2023. [Online]. Available: <https://www.russellgroup.ac.uk/policy/policy-briefings/principles-use-generative-ai-tools-education>

[24]T. C. University_of_Manchester, "Artificial Intelligence and Teaching and Learning," 2025. [Online]. Available: <https://www.teachingcollege.fse.manchester.ac.uk/artificial-intelligence-and-teaching-and-learning>.

[25]C. Medel-Vera and W. F. Gates, "Deciphering Aesthetics: Exploring the Relationship Between Prompt Readability and AI-Generated Image Aesthetics," in Architecture in the AI Era for Research, Practice and Pedagogy, M. H. Asterios Agkathidis, and Carlos Medel-Vera Ed.: Springer Nature, 2025.

[26]S. o. t. H. a. S. S. University_of_Cambridge. "Template Declaration of the Use of Generative Artificial Intelligence." University of Cambridge (CSHSS). <https://www.cshss.cam.ac.uk/education/generative-artificial-intelligence-ai-and-scholarship/template-declaration-use-generative> (accessed).

[27]D. Weber-Wulff et al., "Testing of detection tools for AI-generated text," International Journal for Educational Integrity, vol. 19, pp. 1-39, 2023.

[28]M. Perkins et al., "Detection of GPT-4 Generated Text in Higher Education: Combining Academic Judgement and Software to Identify Generative AI Tool Misuse," Journal of Academic Ethics, vol. 22, pp. 89-113, 2023.

[29]M. S. Orenstrakh, O. Karnalim, C. A. Suárez, and M. Liut, "Detecting LLM-Generated Text in Computing Education: Comparative Study for ChatGPT Cases," 2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC), pp. 121-126, 2023.

[30]W. Liang, M. Yuksekgonul, Y. Mao, E. Wu, and J. Zou, "GPT detectors are biased against non-native English writers," Patterns, vol. 4, no. 7, 2023, doi: 10.1016/j.patter.2023.100779.

[31]M. Perkins et al., "Simple techniques to bypass GenAI text detectors: implications for inclusive education," International Journal of Educational Technology in Higher Education, vol. 21, no. 1, p. 53, 2024/09/09 2024, doi: 10.1186/s41239-024-00487-w.

[32]University_of_Sheffield. "AI and Academic Integrity." University of Sheffield. <https://sheffield.ac.uk/study-skills/assessment/academic-integrity/ai-and-academic-integrity> (accessed).

BIBLIOGRAPHY 5

UNDERSTANDING AI WORKFLOWS IN ARCHITECTURE

[1] Šarčević, T., Karłowicz, A., Mayer, R., Baeza Yates, R. & Rauber, A. (2024) ‘U Can't Gen This? A Survey of Intellectual Property Protection Methods for Data in Generative AI’, arXiv preprint arXiv:2406.15386. <https://doi.org/10.48550/arXiv.2406.15386>

[2] Chesterman, S. (2024) ‘Good Models Borrow, Great Models Steal: Intellectual Property Rights and Generative AI’, SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.4590006>

[3] Torrance, A. W. & Tomlinson, B. (2023) Training Is Everything: Artificial Intelligence, Copyright, and Fair Training, arXiv preprint arXiv:2305.03720. <https://doi.org/10.48550/arXiv.2305.03720>

[4] Guadamuz, A. (2023) A Scanner Darkly: Copyright Infringement in Artificial Intelligence Inputs and Outputs, SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.4371204>

[5] Samuelson, P. (2023) Generative AI meets copyright, Science, 381(6654), pp. 158–161. <https://doi.org/10.1126/science.adi0656>

[6] Choksi, M. Z. and Goedicke, D. (2023)Whose Text Is It Anyway? Exploring BigCode, Intellectual Property, and Ethics, arXiv preprint arXiv:2304.02839. <https://doi.org/10.48550/arXiv.2304.02839>

[7] Samuelson, P. (2023) Generative AI meets copyright, Science, 381(6654), pp. 158–161. <https://doi.org/10.1126/science.adi0656>

[8] Schaeffer, D., Coombs, L. D., Luckett, J., Marin, M. and Olson, P. (2024) Risks of AI Applications Used in Higher Education, Electronic Journal of e-Learning, 22(6), pp. 60–65. <https://doi.org/10.34190/ejel.22.6.3457>

[9] Miresghallah, N., Kim, H., Zhou, X., Tsvetkov, Y., Sap, M., Shokri, R. and Choi, Y. (2023) Can LLMs Keep a Secret? Testing Privacy Implications of Language Models via Contextual Integrity Theory, arXiv preprint arXiv:2310.17884. <https://doi.org/10.48550/arXiv.2310.17884>

[10] O'Leary, D. (2023) ‘Enterprise large language models: Investigating characteristics, risks, and limitations in enterprise settings’, Information Systems Audit and Control Journal (ISAF). <https://doi.org/10.1002/isaf.1541>

[11] Lim, Y. and Shim, H. (2024) Addressing Image Hallucination in Text-to-Image Generation through Factual Image Retrieval, arXiv preprint arXiv:2407.10683. <https://doi.org/10.48550/arXiv.2407.10683>

[12] Kapsalis, T. (2024) CADgpt: Harnessing Natural Language Processing for 3D Modelling to Enhance Computer-Aided Design Workflows, arXiv preprint arXiv:2401.05476. <https://doi.org/10.48550/arXiv.2401.05476>

