Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2668–2678
https://doi.org/10.54941/ahfe1008491

Teaching Human–AI Collaboration: A Task-Based Framework for Inclusive Design Education

Dan Li, Lei Xia, Mudasir Ahmed, and Ling Fan

College of Design and Innovation, Tongji University, Shanghai, China

ABSTRACT

Design education is absorbing generative AI without a shared view of what to teach and when. Two responses dominate: tool-by-tool training, which ages with each model release, and general "AI literacy" modules, which connect poorly to studio practice. Both leave questions of access, bias and skill formation implicit. This paper proposes a task-based framework that separates teaching objectives from specific tools. The framework crosses two established distinctions: Task Nature (Routine/ Abstract), from task-based labor economics, and AI Role (Automation/ Augmentation), from research on human–AI collaboration. The four quadrants, Efficiency Engines, Analytical Partners, Creative Sparks and Co-Creation Studios, describe different task patterns, and each is mapped to a characteristic cognitive mode: convergent, divergent, or an oscillation between the two. We show the framework at work by decomposing one published studio brief, a chair design project run with immersive hardware but without AI, into tasks across the quadrants, and by placing published teaching cases within it. We then read the framework through an inclusive-design lens and identify three tensions that concentrate unevenly: hardware access in Co-Creation Studios, aesthetic bias in Creative Sparks, and skill-base erosion in Efficiency Engines. For each we state the pedagogical implication rather than a single curricular fix. This paper draws on prior empirical studies; curriculum studies are the necessary next step. The contribution is a way for educators to decide what to teach that stays stable as tools change, and that shows where inclusion problems are likely to arise.

Keywords: Human–AI collaboration, Design education, Task-based framework, Inclusive design, Generative AI in pedagogy

INTRODUCTION

Design education is under pressure to absorb generative AI faster than pedagogical scholarship can respond. Two responses have emerged. The first is tool-by-tool training, which ages with each release. The second is general "AI literacy" instruction, which offers principles that students find hard to connect to studio work. Both leave implicit the questions of access, bias and skill formation that decide which students benefit from AI in the studio. A framework that helps educators decide what to teach should also show whose learning is being served.

This paper proposes a task-based framework for teaching human–AI collaboration in design. The unit of analysis is the task: a discrete work

Received September 25, 2026; Revised October 29, 2026; Accepted November 9, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

activity, finer than a course project and coarser than a tool action. We cross two distinctions. Task Nature (Routine/Abstract) adapts the routine/non-routine distinction of task-based labor economics (Autor, Levy and Murnane, 2003). AI Role (Automation/Augmentation) comes from research on how people and AI systems divide work (Brynjolfsson, 2022; Raisch and Krakowski, 2021). We map each of the four resulting quadrants to a cognitive mode from creativity research, so that an educator can say what students practise in it.

To keep the framework concrete we use a running example: a chair design brief with a ten-criterion rubric, taken from our earlier studio study (Xia et al., 2024). That study, run in the year the Apple Vision Pro launched, compared head-mounted with tablet-based sculpting and involved no AI, since testing a new device and AI in one session would have confounded the comparison. For our purpose that is an advantage: the tasks and rubric were fixed before any AI tool entered, so we can ask task by task what an AI tool would change. The study also records, in its own limitations, the access problem behind our first inclusive-design tension: a costly new device shared among a class. The other two tensions are anchored in other published cases (Table 2).

The paper contributes: (1) the 2×2 framework with operational definitions; (2) a mapping of cognitive modes to quadrants; (3) a task audit of one studio brief, with placements of published teaching cases; and (4) a diagnosis of three inclusive-design tensions that concentrate in different quadrants.

TASK-BASED VIEWS OF TECHNOLOGY AND WORK

The task-based approach in labor economics holds that technology changes work by taking over tasks, not whole jobs. Autor, Levy and Murnane (2003) separated routine tasks, which follow explicit rules, from non-routine ones, and showed that computers substitute for the first and complement non-routine problem-solving and communication tasks. A second distinction concerns what the technology is for. Automation hands the task to the machine; augmentation keeps the person in the task and extends what they can do. Brynjolfsson (2022) argues that a focus on human-like automation concentrates gains, and Raisch and Krakowski (2021) argue that in management the two cannot be cleanly separated. Both distinctions were developed for labor markets and firms. The closest quantitative work in design education, a meta-analysis of 27 controlled studies of generative AI in design teaching, sorts interventions by who directs the use, whether the student, the instructor, or both together, and finds the largest effects when instructor and student evaluate AI output jointly (Chiang, 2026). Our axes instead describe the task and the AI's role in it. They are meant to be crossed with such usage modes, not to replace them.

COGNITIVE MODES IN DESIGN CREATIVITY

Creativity research distinguishes divergent thinking, which generates many possibilities, from convergent thinking, which narrows to the answer the constraints require (Guilford, 1967). In design the two are interleaved.

Goldschmidt (2016) shows designers shifting between generating and evaluating within a single episode of work, and accounts of design as the co-evolution of problem and solution describe the same movement across the design process (Dorst and Cross, 2001). Closer to the present paper, Li, Xia and Fan (2025) mapped occupational tasks in art and design to vision-language-model capabilities and then to divergent or convergent thinking, with curatorial work as the case. We extend that task-to-mode mapping from professional practice to teaching.

INCLUSIVE DESIGN EDUCATION

In this paper, inclusive design is a property of teaching rather than of products: which students a studio pedagogy includes and on what terms. We narrow it to three conditions that a studio using AI can change and that the literature below treats as equity issues: access to tools, the cultural defaults built into them, and the skill students bring. Universal Design for Learning asks teachers to plan for learner variability from the start (CAST, 2024), but we found few applications to university design studios. Design justice moves the question from designing for everyone to power: who takes part in design and how its benefits and burdens are distributed, including the benefits of design education itself (Costanza-Chock, 2020). Work on decolonizing design education adds that curricula modelled on Western schools carry Eurocentric knowledge as if it were universal (Ansari et al., 2018).

Material conditions matter as well. When the Covid-19 pandemic closed UK architecture schools' studios, 29% of 798 surveyed students named lack of access to physical or digital resources as a major concern (Grover and Wright, 2023). In East Africa the studio had worked as an equalizer across socio-economic backgrounds; once teaching moved online, data costs, computers and software licences became barriers (Olweny, Ndibwami and Ahimbisibwe, 2023). None of these literatures says where, in a studio that uses AI, each condition is most likely to bite.

FRAMEWORK: TASK NATURE × AI ROLE

Two Axes

Task Nature runs from Routine to Abstract. We adapt Autor's distinction for teaching. A task is routine when its goal and success criteria can be stated before it starts: remove the background from forty product photographs; produce thirty form variations that satisfy a given brief. A task is abstract when the goal and the criteria are worked out during the task: decide what kind of sitting this chair is for; weigh expressive form against structural support. The distinction is independent of the tool. Generating thirty variations is routine in this sense even when a generative model does it, because the brief is fixed and any single output is cheap to discard.

AI Role runs from Automation to Augmentation. In automation the AI carries out the task and returns an output that the student accepts, rejects or lightly edits. In augmentation the AI takes part in a process that the student directs over several rounds, and each round shapes the next. This is a claim

about the pattern of interaction, not about how advanced the model is: the same model can be used either way. Both axes are continua; we divide them into a 2 × 2 because four named patterns are easier to plan with than two scales.

Four Quadrants

Efficiency Engines (Routine × Automation). Well-specified production steps are handed to the AI: background removal, upscaling, batch rendering of a finished model. The student specifies the step, checks that the output is adequate, and fits it into what comes next.

Analytical Partners (Abstract × Automation). The AI carries out a structured analysis inside a task whose criteria are still open: a first-pass ergonomic check, a synthesis of interview notes, a draft evaluation plan. The AI returns a conclusion; the student tests it against the project's own criteria and decides what it means.

Creative Sparks (Routine × Augmentation). The brief is fixed and the AI widens the set of options within it over rounds of prompting and selection: concept imagery, mood boards, variations in material and color. The student keeps generating, filters, and resists settling early on a direction the tool happens to favor.

Co-Creation Studios (Abstract × Augmentation). The framing itself is in play. Student and AI alternate between proposing and criticizing over a whole project, and problem and solution develop together. This is the closest of the four to traditional studio practice and the most demanding to teach.

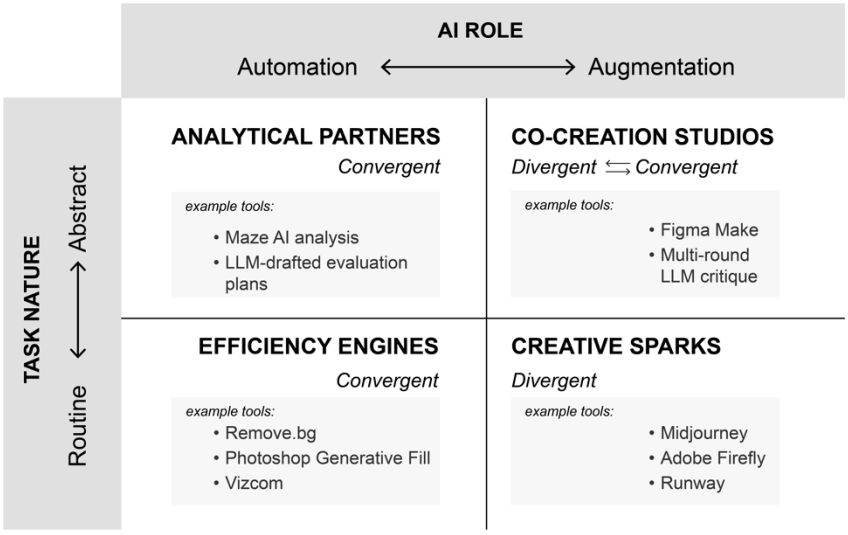


Figure 1: The task-based framework. Task Nature (Routine ↔ Abstract) is crossed with AI Role (Automation ↔ Augmentation); each quadrant is labelled with its predominant cognitive mode. Tools named are dated illustrations (2026), not fixed members of a quadrant.

Cognitive Modes by Quadrant

Each quadrant exercises a predominant mode. In Efficiency Engines and Analytical Partners it is convergent: the AI has produced something, and the student judges whether it is adequate or whether its conclusion holds. In Creative Sparks it is divergent: the point is to have more varied options than the student would reach alone. In Co-Creation Studios it is the oscillation between the two that Goldschmidt (2016) observed in designers, now spread across a person and a system. What the mapping gives an educator is a learning objective that does not name a tool: students practise judging automated output against a rubric, or students practise sustaining divergence over several rounds before selecting.

APPLYING THE FRAMEWORK

In the project reported by Xia et al. (2024), teams of design students aged 18 to 22 modelled two chairs in a two-hour session, and five experienced designers scored each chair on ten criteria, from support and structure to shape expression and fulfilment. Table 1 breaks the project into tasks and asks what an AI tool could do in each and what the student would then be practising. The placements are ours and are meant as a template for an educator's own audit, not as findings.

Table 1: A task audit of a chair design brief.

Quadrant	Tasks Within the Brief	What the Student Practises	Typical Failure
Efficiency Engines	Cleaning reference photographs; mesh clean-up; rendering a finished model in several materials	Specifying a step; checking output for adequacy	Accepting flawed output unchecked
Analytical Partners	First-pass check of seat height and posture against anthropometric data; comparison of precedent chairs	Testing an automated conclusion against the project's criteria	Treating the analysis as the decision
Creative Sparks	Generating many form directions from the fixed brief; color and material variations	Sustaining divergence; filtering; noticing the tool's defaults	Early fixation on a tool-favored look
Co-Creation Studios	Reframing what kind of sitting the chair serves; iterating form with the AI as critic; trading expression against support	Moving between proposing and judging; owning the framing	Letting the AI set the problem

Two things follow. A single two-hour session spans all four quadrants, so the teaching decision is which quadrant a session is for, not which tool to teach. And the rubric is unevenly covered: criteria that can be checked against a standard, such as support, sit in the automation quadrants, while criteria that must be traded off sit in Co-Creation Studios.

Table 2 places eight published studies of AI use in design teaching in the framework. Candidates came from a search of publisher and proceedings pages and two rounds of database search (Consensus; 2022 onward; peer-reviewed only; ten fixed queries; about 220 records screened on title and abstract; criteria and screening record available on request). From about 60 eligible studies we chose eight purposively: at least one per quadrant, a positive and a negative ideation result, and cases whose authors report an access, bias or skill observation. This is not a systematic review; for one see Chiang (2026). Three observations follow. Most published cases are short ideation exercises in Creative Sparks; student-facing cases in Analytical Partners are fewer and mostly concern synthesising user research with a chatbot. A divergent task does not guarantee a divergent result: students ideating with a chatbot produced fewer original ideas than those working without it (Bunn et al., 2024), so the quadrant names the mode the student should be practising, not one the tool delivers. And the quadrant belongs to the use, not the tool: given one brief and the same tools, students in Stanimirovic et al. (2026) chose three different roles for the AI.

Table 2: Illustrative published cases of AI use in design teaching.

Study	Setting and Tools	What Students Did	Quadrant	Issue Reported by the Authors
Batley and Glithro (2024)	UK; 11 product design graduates (8 used the tool); Generative Fill	Created or extended poster backgrounds and added content	Efficiency Engines	Tool learned informally, outside the curriculum
Bartlett, Mills and Hagins (2024)	USA; product design undergraduates; Vizcom	Turned their own hand sketches into renders	Efficiency Engines	Need for hand-sketching skill not reduced; licences supplied by the vendor
Aziz, Kharrufa and Johnson (2026)	UK; third-year interaction design module; students' own chat tools	Compared AI-suggested evaluation plans with their own	Analytical Partners	Oversimplification; trust; less critical thinking
Brook et al. (2024)	UK; 34 first-year product design students; Firefly, Runway, ChatGPT	Built speculative worlds with AI, then designed toys from them	Creative Sparks	Tools limited to free or university-licensed ones; device access; homogenized outcomes
Bunn et al. (2024)	USA; 35 undergraduates; ChatGPT 3.5	Team ideation with and without the chatbot	Creative Sparks	Lower fluency, flexibility and originality with the chatbot
Silva et al. (2024)	Portugal; 25 first-semester design students; Midjourney	Kept a seven-day diary on bias in generated images	Creative Sparks	A third noticed bias in one or none of seven sets; more than a third in five or more

(Continued)

Table 2: Continued.

Study	Setting and Tools	What Students Did	Quadrant	Issue Reported by the Authors
Derksen and Bansal (2024)	USA; final-year print-design students; Midjourney, ChatGPT	Iterated concepts by re-seeding AI output and used AI to critique them	Co-Creation Studios	Prompt writing the main difficulty; some felt a loss of authorship
Stanimirovic et al. (2026)	Serbia; architecture workshop, 34 works; ChatGPT, Gemini	Chose AI as final renderer (26), generator (2) or sparring partner (6)	Three quadrants	Sparring: most consistent quality, lowest sense of authorship; generic "traditional" pastiche

INCLUSIVE IMPLICATIONS

The framework's teaching implications are not neutral with respect to equity. Each quadrant changes the relation between access to tools, prior skill and outcomes. We examine three tensions, each concentrated in a different quadrant.

Hardware Access and Co-Creation Studios

Co-Creation Studios asks the most of students and of equipment. Sustained, multi-round work with AI tends to depend on paid tiers, capable hardware and, increasingly, immersive devices. The chair project, run in the device's launch year, shows the constraint in a comparatively well-resourced studio. Sixteen students shared two head-mounted displays and two tablets, and the authors list the difficulty of sharing a head-mounted display among their limitations (Xia et al., 2024). That display was priced from USD 3,499 at launch (Apple, 2024). The results were mixed across the ten criteria and, with four chairs in total, exploratory; but the cheaper device was not simply worse. Educators already restrict projects to free or university-licensed tools so that all students have equal access (Brook et al., 2024), or choose the device students already own: Welsh et al. (2021) delivered a hands-on product disassembly exercise through smartphone augmented reality to 312 online learners in 44 countries. The implication is to separate the mode from the device: if the objective is oscillation between proposing and judging, ask which is the least expensive setup that still allows it, plan shared access for anything more costly, and avoid assessing outcomes that depend on equipment students do not all have.

Aesthetic Bias and Creative Sparks

Creative Sparks tools lower one barrier: a student no longer needs advanced rendering skill to put an idea in front of others. They raise another. In a study with 540 participants in 27 countries, images generated from prompts that named no country most resembled the United States, and naming the country improved the match only partly (Basu, Babu and Pruthi, 2023). A benchmark covering eight countries found large gaps in models' cultural awareness of

cuisine, landmarks and art (Kannen et al., 2024), and participants in South Asia found outputs that homogenised and exoticised their cultures, which the authors read as an outsider's gaze (Qadri et al., 2023). The bias also enters upstream: an audit of a widely used aesthetic-quality predictor, applied to filter training images, found that it rates realistic Western and Japanese imagery highest (Taylor et al., 2026). All students may drift toward the look the tool favors. Students who work from a non-dominant tradition meet the problem directly: digital art students in Kuwait found that a text-to-image model stereotyped or erased Gulf heritage, and learned to counter-prompt to assert it (Abdulmajid et al., 2026). Students vary in whether they see it: shown images chosen for bias, about a third of first-semester design students noticed it in one or none of seven sets, while more than a third noticed it in five or more (Silva et al., 2024). The implication is to make the tool's defaults an object of study. A Creative Sparks assignment can ask what the tool returns for an unqualified prompt, what it takes to move it, and which references it cannot reach.

Skill-Base Erosion and Efficiency Engines

Efficiency Engines take over production steps that once required craft training: masking, retouching, mesh repair, rendering. Students who arrive without that training can take part in design conversations earlier and present work at the same finish as their peers: in a first-year architecture studio AI refinement improved the buildability of weaker starters' concepts most while several strong starters declined and novelty did not rise (Sektani et al., 2026). The same automation carries the problem Bainbridge (1983) described for process control: the person is left to monitor the automation and to step in when it fails, and both jobs require the skill no longer practised. The student in Efficiency Engines is in that position: the objective is to judge whether automated output is adequate, and that judgment has been built by doing the step by hand. In one product design course, sketch-to-render software did not reduce the need for strong hand sketching (Bartlett, Mills and Hagins, 2024). The two effects fall on the same students: those who gain most from the lower entry cost have the least prior practice to draw on when output must be judged. The implication is to sequence Efficiency Engines use across a curriculum instead of allowing or banning it outright: open access early, so that students can participate, and later teaching of the underlying skills as objects of study, so that judgment has something to rest on.

Diagnostic Scope of the Framework

We have not set out a tension for Analytical Partners. Equity questions exist there, but published teaching cases are few, and we prefer to leave the cell open than to fill it by assertion. Across the three cases the framework's contribution is diagnostic. It does not tell an educator what to do about hardware cost, biased defaults or lost practice. It shows in which kind of task each is most likely to concentrate, so that it can be dealt with when a session is planned.

DISCUSSION AND LIMITATIONS

Would a list of tools show the same tensions? In part. But a list does not show where each problem concentrates and why: cost follows sustained augmentation, biased defaults matter most where the task is to widen options, and lost practice matters most where the task is to judge automated output.

The framework does not replace existing pedagogical frameworks. TPACK, for example, describes what a teacher needs to know to teach with technology (Mishra and Koehler, 2006). It does not describe how work is divided between a student and an AI system, which is what our two axes address. The paper has several limits. Quadrant boundaries are judgments: both axes are continua, and a tool can sit in several quadrants. The task audit is our reading of one brief from one institution, and that project used no AI. The cases in Table 2 are illustrative. Their selection followed a written protocol with fixed queries and inclusion criteria, but the search was in English, on one database whose ranking is semantic rather than exhaustive, and stopped at a purposive eight. A systematic search with pre-registered filters and two independent screeners is the next step for the case table. The framework was developed within one tradition of design education, and the three tensions are not exhaustive: disability, neurodiversity and the language in which students can prompt deserve separate treatment. The framework does make testable claims. Sessions can be assigned to quadrants and compared on what students practise and who benefits, and the predicted concentration of each tension can be checked against course records. Curriculum studies of this kind are the necessary next step.

CONCLUSION

Design educators need a way to decide what to teach about AI that does not depend on which tools exist this year. We have proposed one: cross the nature of the task with the role the AI plays in it, name the four resulting patterns, and state for each the cognitive mode the student practises. Read through an inclusive-design lens, it shows that hardware cost, biased aesthetic defaults and lost practice do not fall evenly across them. The framework does not resolve these problems. It tells an educator where to look for them before a course runs, and it remains to be tested in one.

ACKNOWLEDGMENT

This study was funded by the National Social Science Foundation of China-Art Project (Grant number: 24BG140).

REFERENCES

Abdulmajid, M., Alali, N. and Alsharrah, A. (2026) “Collaborating with machines: Generative AI as a catalyst for creative development in digital art education”, Digital Education Review, 49, pp. 114–132. doi: 10.1344/der.2026.49.114-132.

Ansari, A., Kiem, M., Martins, L.P. de O. and Vieira de Oliveira, P.J.S. (2018) “Three perspectives on decolonising design education”, PARSE Journal, 8. doi: 10.70733/53gdbvmu92sa.

Apple (2024) "Apple Vision Pro available in the U.S. on February 2", Apple Newsroom, 8 January. Available at: https://www.apple.com/newsroom/2024/01/apple-vision-pro-available-in-the-us-on-february-2/.

Autor, D.H., Levy, F. and Murnane, R.J. (2003) "The skill content of recent technological change: an empirical exploration", The Quarterly Journal of Economics, 118(4), pp. 1279–1333. doi: 10.1162/003355303322552801.

Aziz, A., Kharrufa, A. and Johnson, I. (2026) "Integrating GenAI into HCI pedagogy: Supporting students' learning of evaluation methods", EduCHI '26, pp. 1–8. doi: 10.1145/3803869.3803872.

Bainbridge, L. (1983) "Ironies of automation", Automatica, 19(6), pp. 775–779. doi: 10.1016/0005-1098(83)90046-8.

Bartlett, K.A., Mills, J.E. and Hagins, C. (2024) "Comparing understandability of hand sketches versus AI-generated renders for product design", Proceedings of E&PDE 2024, pp. 199–204. doi: 10.35199/EPDE.2024.34.

Basu, A., Babu, R.V. and Pruthi, D. (2023) "Inspecting the geographical representativeness of images from text-to-image models", 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 5113–5124. doi: 10.1109/ICCV51070.2023.00474.

Batley, A. and Glithro, R. (2024) "Exploring the synergy of AI generative fill in Photoshop and the creative design process utilising informal learning", Proceedings of E&PDE 2024, pp. 1–6. doi: 10.35199/EPDE.2024.1.

Brook, E.E., Hanley, C., Campbell Cole, I., Hayes, S. and Mutch, C. (2024) "Dreamworlds: A case study presenting the potential of text-to-image AI in product design education", Proceedings of E&PDE 2024, pp. 175–180. doi: 10.35199/EPDE.2024.30.

Brynjolfsson, E. (2022) "The Turing Trap: The promise & peril of human-like artificial intelligence", Daedalus, 151(2), pp. 272–287. doi: 10.1162/daed_a_01915.

Bunn, B.J., Wright, G., Novoa Munoz, M. and Howell, B.F. (2024) "Exploring ChatGPT's impact on student-team ideation outcomes for new product development: A pilot study", Proceedings of E&PDE 2024, pp. 533–538. doi: 10.35199/EPDE.2024.90.

CAST (2024) Universal Design for Learning Guidelines version 3.0. Available at: https://udlguidelines.cast.org.

Chiang, S.-B. (2026) "Effectiveness of generative AI in higher education design pedagogy: A meta-analysis of usage modes", International Journal of Technology and Design Education. doi: 10.1007/s10798-026-10122-6.

Costanza-Chock, S. (2020) Design justice: Community-led practices to build the worlds we need. Cambridge, MA: The MIT Press. doi: 10.7551/mitpress/12255.001.0001.

Derksen, G. and Bansal, S. (2024) "A design and AI course: In conversation with machines", Proceedings of E&PDE 2024, pp. 527–532. doi: 10.35199/EPDE.2024.89.

Dorst, K. and Cross, N. (2001) "Creativity in the design process: Co-evolution of problem–solution", Design Studies, 22(5), pp. 425–437. doi: 10.1016/S0142-694X(01)00009-6.

Goldschmidt, G. (2016) "Linkographic evidence for concurrent divergent and convergent thinking in creative design", Creativity Research Journal, 28(2), pp. 115–122. doi: 10.1080/10400419.2016.1162497.

Grover, R. and Wright, A. (2023) "Shutting the studio: The impact of the Covid-19 pandemic on architectural education in the United Kingdom", International Journal of Technology and Design Education, 33(3), pp. 1173–1197. doi: 10.1007/s10798-022-09765-y.

Guilford, J.P. (1967) The nature of human intelligence. New York: McGraw-Hill.

Kannen, N., Ahmad, A., Andreetto, M., Prabhakaran, V., Prabhu, U., Dieng, A.B., Bhattacharyya, P. and Dave, S. (2024) "Beyond aesthetics: Cultural competence in text-to-image models", Advances in Neural Information Processing Systems 37, pp. 13716–13747. doi: 10.52202/079017-0439.

Li, D., Xia, L. and Fan, L. (2025) "Aug-Creativity: Framework for human-centered creativity with vision language models", in Kurosu, M. and Hashizume, A. (eds.) Human-Computer Interaction. HCII 2025. LNCS 15770. Cham: Springer, pp. 86–101. doi: 10.1007/978-3-031-93864-1_7.

Mishra, P. and Koehler, M.J. (2006) "Technological pedagogical content knowledge: A framework for teacher knowledge", Teachers College Record, 108(6), pp. 1017–1054. doi: 10.1111/j.1467-9620.2006.00684.x.

Olweny, M.R.O., Ndibwami, A. and Ahimbisibwe, A. (2023) "Online architectural education: Reflections on COVID-19 emergency remote learning in East Africa", E-Learning and Digital Media, 20(6), pp. 579–597. doi: 10.1177/20427530221117329.

Qadri, R., Shelby, R., Bennett, C.L. and Denton, R. (2023) "AI's regimes of representation: A community-centered study of text-to-image models in South Asia", FAccT '23, pp. 506–517. doi: 10.1145/3593013.3594016.

Raisch, S. and Krakowski, S. (2021) "Artificial intelligence and management: The automation–augmentation paradox", Academy of Management Review, 46(1), pp. 192–210. doi: 10.5465/amr.2018.0072.

Sektani, H.H.J., Miran, F.D., Abdullah, H.K. and Sherwani, A.B. (2026) "Generative AI and the reordering of design creativity in architectural education: Evidence from undergraduate students", Architecture, 6(3), 133. doi: 10.3390/architecture6030133.

Silva, V., Buzzo, D., Hernández-Ramírez, R. and Ayanoglu, H. (2024) "Understanding how undergraduate students perceive biases in AI-generated images, a research-through-design probe", Proceedings of E&PDE 2024, pp. 258–263. doi: 10.35199/EPDE.2024.44.

Stanimirovic, M., Wang, J., Zupancic, D., Petronijevic, A.M. and Ciric, I. (2026) "Dialogic artificial intelligence in the architectural design studio: An empirical study of the AI sparring model", Buildings, 16(14), 2878. doi: 10.3390/buildings16142878.

Taylor, J., Agnew, W., Sap, M., Fox, S.E. and Zhu, H. (2026) "The algorithmic gaze of image quality assessment: An audit and trace ethnography of the LAION-Aesthetics Predictor", FAccT '26, pp. 6383–6402. doi: 10.1145/3805689.3806462.

Welsh, E., Li, D., Hart, A.J. and Liu, J. (2021) "Scaling hands-on learning principles in manufacturing through augmented reality disassembly and inspection of a consumer product", 2021 ASEE Annual Conference & Exposition Proceedings. doi: 10.18260/1-2—37699.

Xia, L., Li, X., Qin, Y., Li, D. and Fan, L. (2024) "Enhancing design education through spatial computing: A comparative study of traditional and immersive technologies in chair design projects", 2024 IEEE Frontiers in Education Conference (FIE), pp. 1–8. doi: 10.1109/FIE61694.2024.10893086.