

Creativity from Friction: Human–AI Interaction for Exploratory Structural Design

Ricardo Maia Avelino^{1,2} Rita Sevastjanova² Tom Van Mele¹ Philippe Block¹ Mennatallah El-Assady²

Abstract

AI agents that generate final answers based on user input often do not meet the needs of creative fields. Fields such as structural design and architecture need interactive systems that help users externalise and develop ideas, explore alternatives, and refine partial solutions. The final product of such designs needs to comply with many constraints concerning, e.g., spatial configuration, mechanical behaviour, material quantities, and costs. These constraints create friction in the design process, which can stimulate novel and creative solutions. In this paper, we discuss the misalignment between current generative AI goals to remove friction and provide final solutions and the needs of creators, such as structural designers, who develop ideas through iterative work. We present the design dimensions of systems allowing for *constrained human–AI co-creation* that rely on vision-language models making structural exploration conversational, multimodal, and responsive to evolving human intent in ways that follow and augment the discipline’s creative process. Through a pilot design interface based on these principles and a study with experts in the field, this paper shows how structural designers perceive interactive AI systems and how such systems can support design space exploration by reducing repetitive modelling friction while preserving reflective design friction.

1. Introduction

Generative AI is often evaluated for its ability to produce plausible final outputs such as an image, a text, a plan, or a model from user input. However, creative design work

rarely progresses as a one-shot process to the final product. Designers externalise partial ideas, compare alternatives, and use intermediate models to think, collaborate, and create. This view resonates with Schön’s model for professionals of creative disciplines that engage in “reflection-in-action” adapting their creation process to ever evolving constraints and ideas (Schön, 1983). Thus, generative AI in creative domains must shift from workflow automation to *exploration augmentation* (Rezwana & Maher, 2023; Wang et al., 2025).

Structural design makes this need particularly visible. It is a creative discipline, but not an unconstrained one. Structural engineers must balance spatial constraints while reasoning about loads, supports, spans, constructability, economy, safety, among others. In doing so, engineers seek creative solutions that emerge from the friction of managing all these constraints. Teng et al. (2004) define creativity in structural design as the production of a design that is novel and satisfies functional constraints, and highlight the importance of divergent thinking to draw analogies from solutions to previous problems and explore different alternatives.

Structural exploration has been the focus of recent research in which algorithms that explore different structural forms are preferred rather than single-objective optimisations on predetermined typologies (Mueller, 2014; Danhaive & Mueller, 2021; Lee et al., 2016; Saldana Ochoa et al., 2021). Generative AI brings new opportunities to expand this exploration by allowing multimodal interactions and advanced reasoning (Zhang et al., 2025; Suh et al., 2024). This raises a key question: *how can we formally define such interactions so that these systems effectively support and enhance the design process?*

We answer this question by introducing the concept of *constrained co-creation* for structural design: interactive human–AI workflows in which creative exploration is shaped by the ongoing negotiation of domain constraints. The design process supported by these systems becomes an interactive activity in which human and AI agents jointly explore and refine evolving design alternatives. Unlike final-answer automation, constrained co-creation helps designers externalise, inspect, revise, and evaluate alternatives while maintaining human agency (Holter & El-Assady, 2024) and allowing domain-grounded interactive workflows (Song et al.,

¹Block Research Group, Department of Architecture, ETH Zurich, Switzerland ²IVIA Lab, Department of Computer Science, ETH Zurich, Switzerland. Correspondence to: Ricardo Maia Avelino <mricardo@ethz.ch>.

2024). Such systems can remove friction from repetitive actions while maintaining friction that arises from the constraints and steers the design process.

This paper contributes to human–AI interaction in the field of structural design by (1) defining productive friction and its importance for interactive structural design, (2) identifying four design dimensions for future human–AI systems, (3) presenting a prototype of an AI-supported design workflow, and (4) reporting a preliminary study to investigate how experts interact with this prototype in a proposed design scenario.

2. Structural Design Exploration and Human–AI Co-Creation

2.1. Structural Design as Creative Constraint Negotiation

Structural design is the discipline that connects architecture and physics and is concerned with ensuring the mechanical performance of building systems that safely and efficiently resist applied loads while satisfying architectural, spatial, material, environmental, and construction requirements. Unlike many creative domains in which an artefact may be judged primarily by aesthetic or expressive qualities, a structural design is inseparable from how it works. Billington defines structural “art” as a creative synthesis of the three E’s of efficiency, economy, and elegance (Billington, 1983). Figure 1 illustrates constructed examples of Billington’s structural art where the structures are not conceived by an unconstrained creative act but are rather the result of the negotiation of form and force under multiple and multidisciplinary constraints (Allen & Zalewski, 2010). Among them, the Chiasso Shed truss geometry follows its internal forces allowing for minimal material usage; the Dulles Airport roof follows the geometry of a catenary resulting in an efficient purely-tensile geometry; the Broadgate Exchange House embeds a load-bearing arch to transfer forces away from the centre, clearing the train tracks running below the building; and the Eiffel Tower was shaped to resist the effects of lateral wind loads. These novel projects, achieved through creative and constrained design space exploration, remain landmark works today.

2.2. AI-Supported Design Space Exploration

Computational structural design has long treated early-stage design as the exploration of a space of possibilities rather than the production of a single optimal answer. Shape grammars are often used to encode design transformations and generate families of spatial configurations (Stiny, 1980). In structural design, grammars are used in Shea & Cagan (1999); Lee et al. (2016), allowing the generation of large design spaces navigated through stochastic algorithms.



Figure 1. Examples of Billington’s structural art: Chiasso Shed, Switzerland, by Robert Maillart (Photo: ETH-Bibliothek Zürich); Dulles Airport, United States, by Eero Saarinen (Photo: Carol Highsmith); Broadgate Exchange House, United Kingdom, by Skidmore, Owings & Merrill (Photo: Wikimedia Commons); and the Eiffel Tower, France, by Gustave Eiffel (Photo: Wikimedia Commons).

Van Mele et al. (2012); Maia Avelino et al. (2021) develop graphic-statics-based algorithms to explore the interactive design of 2D structures. Ohlbrock & Schwartz (2016) propose a combinatorial approach that allows exploration and form-finding of spatial configurations in equilibrium for 3D systems. The latter approach is further developed in Saldana Ochoa et al. (2021); Pastrana et al. (2023); Bleker et al. (2026), allowing for projection-based design exploration, introducing auto-differentiation algorithms for equilibrium, and using graph transformers to generate structures from text descriptions.

Recent work on generative AI for building structures extends this trajectory by surveying how data representations, generative algorithms, evaluation methods, and optimisation loops are being used for structural design automation (Liao et al., 2024). However, much of this work still frames AI primarily as a generator or optimiser. Language and vision-language agents, in contrast, offer new opportunities to support a conversational and multimodal form of design space exploration. Instead of presenting designers with precomputed or parametrised alternatives, such agents can allow users to steer the generation, transformation, and evaluation of structural models through text and multimodal input. Indeed, the design intent in architecture and structural engineering is rarely communicated through text alone. Designers use sketches, annotations, and partial models to express ambiguous spatial ideas. Recent work on sketch-based 3D generation highlights both the promise and the difficulty of interpreting sketches in engineering (Tono et al., 2025; Li et al., 2024; He et al., 2025). Sketches are useful because they are fast and expressive and sometimes encode tacit knowledge, but they are also incomplete, ambiguous, and difficult for current models to interpret reliably. The challenge is therefore to build AI systems that can work with incomplete multimodal human expressions while keeping the evolving artefact grounded in structural logic.

2.3. Design Dimensions for Human–AI Structural Design Interfaces

We derive four design dimensions for human–AI interfaces to support exploratory structural design following the guidelines proposed in Stähle et al. (2025) for agentic interfaces in visual analytics. These dimensions specialise previous co-creative and mixed-initiative frameworks (Rezwana & Maher, 2023; Wang et al., 2025; Holter & El-Assady, 2024) to the field of structural design. We propose that these interfaces should support: (1) model grounding in structural design knowledge, (2) human- and AI-readable data structures, (3) state awareness and interaction history, and (4) multimodal expression of design intent.

Model grounding in structural design knowledge. A reliable system must be grounded in discipline-specific logic (Song et al., 2024). For structural design, this includes basic notions of structural load paths, knowledge of typical spans, loads, and cross-sections, and understanding of member hierarchy and boundary condition definition. Without such a grounding, an AI system may generate artefacts that are visually convincing but structurally incorrect, which might, in turn, cause domain experts to lose confidence in the system.

Human- and AI-readable data structures. The design artefact should be represented in a form that is accessible to both humans and AI systems (Cao et al., 2025; Feng et al., 2026). For the human, the model should be visual, inspectable, and editable. For AI, the data that underlie the artefacts need to be accessible, revealing their internal data structures and interconnections. This shared representation enables editing for both humans and AI.

State awareness and interaction history. A co-creative AI system must enable access, inspection, and modification of the current state of an evolving design model (Feng et al., 2026). As the artefacts develop, constraints must be re-evaluated. An edit may remove a support, disconnect an element, or change the structural load path. The system should therefore maintain awareness of the current state, previous edits, persistent constraints, and user goals to effectively assist in the design of safety-critical engineering components and to allow tracing the responsibility over executed modifications.

Multimodal expression of design intent. Structural and architectural intent is often expressed in many “languages” (Shi et al., 2026; Baudoux et al., 2025). Designers communicate through sketches, drawings, annotations, and verbal or text descriptions. Communication is also often incomplete and ambiguous. Human–AI tools should therefore accommodate these distinct “languages” and their intrinsic ambiguity.

Together, these dimensions can guide the role of AI in exploratory structural design. The next section proposes a workflow that embeds these design dimensions.

3. A Co-creation Structural Design Workflow

Based on the dimensions presented above, this work proposes a co-creative workflow for exploratory structural design allowing for human $\rightleftharpoons$ model $\rightleftharpoons$ AI interaction and anchored on pre-defined and intrinsic discipline constraints (Figure 2). Explicit constraints, such as design parameters, site conditions, and regulations, are defined at the beginning of the process. These are complemented by intrinsic discipline constraints based on engineering validity or knowledge. Some of this knowledge should be embedded in the model’s grounding, while some remains part of the designer’s tacit expertise and professional judgement about what is likely to perform well. Motivated by these constraints, the human and AI agents *co-create* in a design loop where the design intent is co-developed between them. Through multimodal inputs, both agents can edit and interact with the model. Due to interactivity and state awareness, both agents can perceive the model changes and review the current state. The agents share bi-directional feedback throughout the design process.

We argue that this cycle can support an augmented design exploration, which might lead to more creative ways to satisfy the constraints. The goal is not to automate the production of a final structural form but to follow the design process, reducing unproductive modelling friction while preserving the *productive friction* through which designers inspect, revise, and negotiate constraints. By editing the structural model, receiving feedback from the AI agent, the human agent acquires important insights that influence the design process. The generated designs can then be evaluated for quantitative or qualitative performance and the results might reveal problems that then act as additional constraints requiring new creative solutions.

To understand the effectiveness and potential of this design workflow, we conducted a preliminary study on a pilot structural design interface that enables bi-directional multimodal interactive structural design using Google Gemini 3.1 Flash-Lite (Google DeepMind, 2026) as its underlying reasoning model. The study is described in the following section.

4. Preliminary Design Study

We conduct a preliminary study to examine how expert users interact with an AI-supported structural design interface in a constrained design scenario. This study offers qualitative insight into how designers use, trust, revise, and respond to an AI agent during early-stage structural design and how they manage friction in the design process.




Figure 2. Proposed human–AI workflow for exploratory structural design. The interactions between human, AI, and the model support creative exploration under a constrained design space.

4.1. Design Task

The participants are given the design task shown in Figure 3. The task proposes to design a multi-storey building within a maximum limit of $50.0 \times 25.0 \times 25.0$ m with supports only allowed in a reduced footprint of 40.0×15.0 m. In addition, the building must provide an open space in the middle to allow a public plaza such that the building cannot cross a virtual sphere with diameter of 10.0 m positioned in the centre of the support footprint. Designers are instructed to develop creative but structurally feasible solutions with at least 5 storeys and 2000 m² floor area.

4.2. Participants and Procedure

We recruit $n = 3$ participants with experience in the fields of architecture, structural engineering, and parametric modelling. Each session starts with a consent form and then proceeds in six stages guided by a questionnaire (Appendix A). A test facilitator observes and assists participants, especially when they struggle or feel uncertain about using the AI interface, following the pair-analytics design approach (Arias-Hernandez et al., 2011). The stages are defined as:

1. **Demographics:** Participants fill out a survey of their professional and educational experience and their knowledge of 3D modelling and AI (Appendix B).




Figure 3. Boundaries of the user study. The building must stay within the red boundary, avoid the spherical object, and be supported in the green zone. Measurements are indicated in metres.

2. **Task presentation:** Participants receive the design problem and printed pages of the orthographic views of the building and are asked to develop their ideas by sketching their design intent (Appendix C).
3. **Introduction to the tool:** Participants receive a 5 minute introduction of the tool and how they can create and edit artefacts using drafting commands, AI conversation, or a sketch-on-the-screen tool.
4. **AI co-creation:** Participants interact with the system through direct 3D modelling, text prompts to AI, and sketch-based AI inputs, iteratively refining the design. Participants’ screen activities are recorded (Figure 4).
5. **Final result:** The participants review the final artefact and approve it (Figure 5).
6. **Reflection interview:** Participants discuss their experience and how the interaction with AI modified their design output and where it reduced or created friction (Section 5).

4.3. User Interaction

This section describes the interaction of each participant with the AI interface. Figure 4 summarises the evolution of the design sessions, showing key intermediate model states, interaction types, and approximate timestamps for each participant. The interactions are characterised as follows:

- **Human-Only:** modifications primarily executed by the human alone using the 3D drafting commands.
- **Human-AI-Mix:** modifications are shared between humans and AI, e.g. “move selected joints by 4m.”
- **Human-Text-to-AI:** high level prompts written by participants requiring modifications in the structure, e.g. “populate beams and columns on a grid with sizes [...]”
- **Human-Sketch-to-AI:** human performs a sketch on the screen which is sent together with a prompt-instruction asking for AI modifications.
- **Human-Undo-AI:** AI makes unexpected mistakes and the user intervenes and undo.

The interaction for each participant is described below.




Figure 4. Evolving design timeline for all participants (P01, P02 and P03) using the 3D interface. Interactions are characterised and key moments in the design process are noted.

P01. Sketch Phase (18 min): Based on the constraints received, P01 developed their sketch mainly at the front elevation, ensuring that the area and floor constraints were fulfilled. The concept included inclined columns on the ground floor and an enlarged atrium to accommodate the open space constraint. **3D Phase (37 min):** P01 combined manual geometry modifications with AI assisted modelling operations. The participant began defining ground-level joints at the intended column locations, populated a beam grid, and then asked AI to replicate them to the upper floors. The participant asked AI to remove the “largest” slab to meet the open-space constraint. After reviewing the AI output, the inclined columns were rejected as the participant judged that these members were too long. The participant then moved the south face inward to reduce the inclined column lengths but later noticed that these would end up too close to the typical columns for practical engineering construction. This insight changed the design direction with fewer columns and a 2.5 m overhang at north and south sides. The participant approved the final result after checking the area constraints with AI assistance.

P02. Sketch Phase (15 min): The participant’s concept presented inclined columns along the shortest building direction; after reviewing their first sketch, the participant decided to add a staggering on the top floors to make the design more unique. Front and side sketches were made. **3D Phase (32 min):** P02 started the 3D modelling with a clearer predefined concept, and AI was used primarily to accelerate the construction of the structural scheme. The participant created a grid and described the staggering using the grid tag names. The slabs were created correctly on the first attempt, but AI struggled to extend the same logic consistently to beams and columns. The participant

then restarted with a revised prompt that successfully generated most of the required elements. The participant then used sketches to indicate the desired inclined columns at the building edges to the AI and made further small edits using text prompts. The final model was then approved and finalised after performing an AI check on the area constraint.

P03. Sketch Phase (16 min): P03 developed the sketch in plan using a regular $5.0\text{ m} \times 5.0\text{ m}$ grid with terraces on the longest side of the structure. The open-space constraint was envisaged as a minimal cutout in the building regular grid at low level columns. **3D Phase (71 min):** P03’s 3D interaction took longer than for other participants. P03 started collaborating with the AI system to generate the regular building. However, P03 changed their direction when planning the cutout for the sphere. Using the sketch tool, P03 found that a top-to-bottom cylindrical cutout would be more appealing than the original planned minimal cutout. Since the regular building grid was already set, the interaction of P03 with the model became time consuming, requiring passing specific instructions to the AI system to rearrange specific elements of the building. P03 used the sketch feature to indicate the generation of the irregular slab shapes created by the cutout. P03 was satisfied with the final outcome, but noted that additional time would have allowed to also include the terraces as in the original idea.

In summary, participants interacted with AI on different levels. P01 modified its initial design intent but this change was not suggested by the AI autonomously, but rather by the participant’s visual inspection of AI changes. P02 used AI mainly as an accelerator of its design intent, retaining responsibility for correcting and refining the spatial logic. P03’s interaction shows the promise and limitation of multimodal co-creation since sketching helped the participant




Figure 5. Initial sketch and final models co-developed using the AI tool for participants P01, P02 and P03.

express an irregular spatial idea, but the AI struggled to operationalise it precisely.

The initial sketch and the final 3D model for all participants are presented in Figure 5. Following their 3D modelling session, participants discussed the post-task interview questions which are presented in the next section.

5. Discussion: AI for Creative Structural Design

5.1. Creativity as the journey, not the destination

The central value of AI in creative domains may not be the production of a final answer, but the expansion and navigation of possible alternatives. In the present study, we observed that all participants have modified their original design idea either in the manual sketch phase or during the 3D modelling session. P03 reported that the AI-supported process helped develop the idea of a circular internal court from the ground floor to the roof, which had not been considered before. P01 changed their design idea twice, achieving a design that successfully met the constraints. P02 also changed their design idea in the sketching phase, erasing the top bays of the building to create a more interesting staggering effect. This interaction shows that professionals in creative domains find new solutions as they develop their original designs. An AI system that acts as a co-creative assistant can, therefore, enrich the design journey and spark new creative insights from structural designers.

5.2. Productive and Unproductive Friction

Not all friction should be removed from creative work. Some friction is productive since it forces designers to confront constraints, compare alternatives, and make judgements. Other friction, however, is unproductive, including repetitive modelling actions, interface overhead, and low-level translation between intent and geometry.

In our study, all participants identified the reduction of repetitive modelling as an area where AI was helpful. P01 emphasised that AI quickly created repetitive elements. P02 noted that the AI made it possible to quickly generate a structure from a grid and remove unwanted bays through simple prompts, avoiding manual removal that would otherwise have taken longer. P03 similarly highlighted the rapid generation of the main structural grid and floor slabs. This indicates that the current prototype was effective in executing explicit, bounded, and repetitive modelling operations.

Similarly, the interface has maintained the productive friction. By interacting with the interface, P01 noticed two design flaws, first regarding the length of the inclined columns and then identifying supports too close to each other. As a result, P01 has changed the design to remove these elements altogether, resulting in a simpler, elegant, and still feasible solution. True co-creation systems should move towards identifying and removing unproductive friction while preserving or highlighting reflective friction that will support design judgement.

5.3. Multimodality as a Co-Creative Anchor

Multimodal input, such as sketches and drawings, can guide exploration without over-specifying it. The ambiguity of sketches lets participants express spatial intent and design directions more naturally than natural language alone. In our study, all participants indicated that they use sketches as one of the main modals to start a new design. P02 added that starts thinking about projects in 2D and then goes to 3D using CAD software and that the pair of 3D and sketches helps to fully understand the problem.

However, while P01 also responded that using sketches felt natural, P01 did not use the feature during the session. P01 added that many repetitive operations, such as replicating multiple elements at a given distance, can be more easily communicated using natural language. On the other hand, P02 noted that it was sometimes hard to choose a vocabulary to communicate with AI. P02 noted that “it was hard to put my ideas into words for the prompt, similarly to how it is hard to explain a design idea to someone without drawing.”

These experiences confirm that multimodality is key for a co-creative system. Some tasks are more easily described in prompts, while for others a sketch might as well be worth a thousand words.

5.4. Automation, Collaboration, or Co-creation?

Users were asked whether they would categorise their interaction as automation, collaboration, or co-creation, and why. The reflection of the participants revealed different expectations about the agency in human–AI structural design workflows. P01 categorised the current interaction as automation, describing co-creation as an AI system capable of issuing early warnings about feasibility or problematic design decisions. P02 described the interaction as collaboration because both the human and AI can modify the model. For P02, co-creation means an AI system that would “suggest things automatically” or highlight when the design was drifting from the original constraints. P03 also described the current experience as a collaboration and appreciated being in control of the design process and “figuring out new ideas” independently as the project progressed. These responses support the argument that co-creation in structural design does not necessarily mean greater AI autonomy. Instead, it requires a careful distribution of agency which should adapt to the user’s preference as discussed in (Holter & El-Assady, 2024).

5.5. Limitations

This paper presents a position and a preliminary study rather than an exhaustive evaluation of structural validity. The participant sample was small, even if it allowed one to identify interaction patterns and design dimensions. The interface is in continuous development, so not all drawing functions comparable to CAD software are currently supported. In particular, users have reported missing functions of click-and-drag and snap that could facilitate the action of human and AI agents. Furthermore, in the present study, changes in the design direction were initiated mainly by designers themselves rather than AI. The current study also does not establish that AI use alone caused more creative outcomes. Instead, it identifies interaction patterns between designers and structural artefacts that can guide how AI-powered interfaces might more effectively support co-creation. Future work could study this process subjected to increased AI autonomy, where it could proactively suggest actions or signal potential design redirections.

6. Conclusion

We argue that creative AI systems for structural design, as well as for other creative disciplines, should not be designed as final-answer generators, but rather as interactive partners for exploring the design space. Observing participants interact with a prototype AI interface, this paper shows how design intent and user objective evolve together with the artefact. The final artefact is then shaped by the constraints of the design problem and by the intrinsic feasibility constraints of structural design. Future AI-assisted structural

design systems should then follow this process, assisting by removing unproductive friction coming from modelling repetition while preserving productive friction that support designers in achieving their final outcome.

We argue that for structural design, constraints have a dual effect: they reduce the design space while also creating opportunities to stimulate creativity. We propose design dimensions for future human–AI tools for structural design including *model grounding* in discipline-specific knowledge, *human- and AI-readable data structures* allowing both agents to modify the model, *state awareness and interaction history* enabling real-time inspection and traceability, and *multimodal expression of design intent* for a more natural and flexible workflow. The preliminary study described in this paper has reinforced the importance of these design dimensions, and future work will focus on further interface development, systematic evaluation of AI systems for structural design, exploration of different levels of AI autonomy and improving domain specific model grounding. We believe that future work on these unique challenges of adapting generative AI for creative tasks can unlock new possibilities towards a more creative and interactive co-design process for engineers.

Acknowledgments

This work was supported by the Design++ Postdoctoral Fellowship awarded to the first author in collaboration with Halter AG.

References

- Allen, E. and Zalewski, W. *Form and Forces: Designing Efficient, Expressive Structures*. John Wiley & Sons, Hoboken, NJ, 2010.
- Arias-Hernandez, R., Kaastra, L. T., Green, T. M., and Fisher, B. Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics. In *2011 44th Hawaii International Conference on System Sciences*, pp. 1–10. IEEE, 2011. doi: 10.1109/HICSS.2011.339.
- Baudoux, G., Guo, C., and Goucher-Lambert, K. Multi-modal Generative AI for Conceptual Design: Enabling Text-Based and Sketch-Based Human–AI Conversations. *Proceedings of the Design Society*, 5:2501–2510, 2025. doi: 10.1017/pds.2025.10264.
- Billington, D. P. *The Tower and the Bridge: The New Art of Structural Engineering*. Princeton University Press, Princeton, NJ, 1983.
- Bleker, L., Guo, Z., Smith, K., Tam, K.-M. M., Ochoa, K. S., and D’Acunto, P. Text2Structure3D: Graph-Based Generative Modeling of Equilibrium Structures with Diffusion

- Transformers. *arXiv preprint arXiv:2601.12870*, 2026. URL <https://arxiv.org/abs/2601.12870>.
- Cao, Y., Huang, Y., Truong, A., Shin, H. V., and Xia, H. Compositional Structures as Substrates for Human–AI Co-Creation Environment: A Design Approach and a Case Study. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pp. 1–25, New York, NY, USA, 2025. Association for Computing Machinery. doi: 10.1145/3706598.3713401.
- Danhaive, R. and Mueller, C. T. Design Subspace Learning: Structural Design Space Exploration Using Performance-Conditioned Generative Modeling. *Automation in Construction*, 127:103664, 2021. doi: 10.1016/j.autcon.2021.103664.
- Feng, K. J. K., Pu, K., Latzke, M., August, T., Siangliulue, P., Bragg, J., Weld, D. S., Zhang, A. X., and Chang, J. C. Cocoa: Co-Planning and Co-Execution with AI Agents. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, pp. 1–23, New York, NY, USA, 2026. Association for Computing Machinery. doi: 10.1145/3772318.3791673.
- Google DeepMind. Gemini 3.1 Flash-Lite: Model Card. <https://deepmind.google/models/model-cards/gemini-3-1-flash-lite/>, 2026. Accessed: 2026-05-15.
- He, Z., Wang, Y.-H., and Zhang, J. Generative AIBIM: An Automatic and Intelligent Structural Design Pipeline Integrating BIM and Generative AI. *Information Fusion*, 114:102654, 2025. doi: 10.1016/j.inffus.2024.102654.
- Holter, S. and El-Assady, M. Deconstructing Human–AI Collaboration: Agency, Interaction, and Adaptation. *Computer Graphics Forum*, 43(3):e15107, 2024. doi: 10.1111/cgf.15107.
- Lee, J., Mueller, C., and Fivet, C. Automatic Generation of Diverse Equilibrium Structures through Shape Grammars and Graphic Statics. *International Journal of Space Structures*, 31(2–4):147–164, 2016. doi: 10.1177/0266351116660798.
- Li, X., Sun, Y., and Sha, Z. LLM4CAD: Multimodal Large Language Models for Three-Dimensional Computer-Aided Design Generation. *Journal of Computing and Information Science in Engineering*, 25(2):021005, 2024. doi: 10.1115/1.4067085.
- Liao, W., Lu, X., Fei, Y., Gu, Y., and Huang, Y. Generative AI Design for Building Structures. *Automation in Construction*, 157:105187, 2024. doi: 10.1016/j.autcon.2023.105187.
- Maia Avelino, R., Lee, J., Van Mele, T., and Block, P. An Interactive Implementation of Algebraic Graphic Statics for Geometry-Based Teaching and Design of Structures. In *International fib Symposium: Conceptual Design of Structures 2021*, pp. 447–454, Attisholz-Areal, Switzerland, 2021. doi: 10.35789/fib.PROC.0055.2021.CDSymp.P054.
- Mueller, C. T. *Computational Exploration of the Structural Design Space*. PhD thesis, Massachusetts Institute of Technology, 2014.
- Ohlbrock, P. O. and Schwartz, J. Combinatorial Equilibrium Modeling. *International Journal of Space Structures*, 31(2–4):177–189, 2016. doi: 10.1177/0266351116660799.
- Pastrana, R., Ohlbrock, P. O., Oberbichler, T., D’Acunto, P., and Parascho, S. Constrained Form-Finding of Tension–Compression Structures Using Automatic Differentiation. *Computer-Aided Design*, 155:103435, 2023. doi: 10.1016/j.cad.2022.103435.
- Rezwana, J. and Maher, M. L. Designing Creative AI Partners with COFI: A Framework for Modeling Interaction in Human–AI Co-Creative Systems. *ACM Transactions on Computer-Human Interaction*, 30(5):1–28, 2023. doi: 10.1145/3519026.
- Saldana Ochoa, K., Ohlbrock, P. O., D’Acunto, P., and Moosavi, V. Beyond Typologies, Beyond Optimization: Exploring Novel Structural Forms at the Interface of Human and Machine Intelligence. *International Journal of Architectural Computing*, 19(3):277–299, 2021. doi: 10.1177/1478077120943062.
- Schön, D. A. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York, 1983.
- Shea, K. and Cagan, J. The Design of Novel Roof Trusses with Shape Annealing: Assessing the Ability of a Computational Method in Aiding Structural Designers with Varying Design Intent. *Design Studies*, 20(1):3–23, 1999. doi: 10.1016/S0142-694X(98)00019-2.
- Shi, W., Upadhyay, S., Quek, G., and Choo, K. T. W. TalkSketch: Multimodal Generative AI for Real-Time Sketch Ideation with Speech. In Woodward, K., Falcon-Caro, A., Ramchurn, R., and Benford, S. (eds.), *Creative AI for Live Interactive Performances*, pp. 83–97, Cham, 2026. Springer Nature Switzerland. doi: 10.1007/978-3-032-16893-1_6.
- Song, B., Zhu, Q., and Luo, J. Human–AI Collaboration by Design. *Proceedings of the Design Society*, 4:2247–2256, 2024. doi: 10.1017/pds.2024.227.

- Stähle, T., Jansen op de Haar, M., Boyer, S., Sevastjanova, R., Narechania, A., and El-Assady, M. A Design Space for Intelligent Agents in Mixed-Initiative Visual Analytics. *arXiv preprint arXiv:2512.23372*, 2025. URL <https://arxiv.org/abs/2512.23372>.
- Stiny, G. Introduction to Shape and Shape Grammars. *Environment and Planning B: Planning and Design*, 7(3): 343–351, 1980. doi: 10.1068/b070343.
- Suh, S., Chen, M., Min, B., Li, T. J.-J., and Xia, H. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human–AI Co-Creation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI ’24, pp. 1–26, New York, NY, USA, 2024. Association for Computing Machinery. doi: 10.1145/3613904.3642400.
- Teng, J. G., Song, C. Y., and Yuan, X. F. Fostering Creativity in Students in the Teaching of Structural Analysis. *International Journal of Engineering Education*, 20(1): 96–102, 2004.
- Tono, A., Wu, J., Wetzstein, G., Armeni, I., Subramonyam, H., Landay, J., and Fischer, M. Deep Sketch-Based 3D Modeling: A Survey. *Computer Graphics Forum*, pp. e70302, 2025. doi: 10.1111/cgf.70302.
- Van Mele, T., Lachauer, L., Rippmann, M., and Block, P. Geometry-Based Understanding of Structures. *Journal of the International Association for Shell and Spatial Structures*, 53(174):285–295, 2012.
- Wang, N., Kim, H., Peng, J., and Wang, J. Exploring Creativity in Human–AI Co-Creation: A Comparative Study across Design Experience. *Frontiers in Computer Science*, 7:1672735, 2025. doi: 10.3389/fcomp.2025.1672735.
- Zhang, S., Wang, H., and Yi, X. Exploring Collaboration Patterns and Strategies in Human–AI Co-Creation through the Lens of Agency: A Scoping Review of the Top-Tier HCI Literature. *Proceedings of the ACM on Human-Computer Interaction*, 9(7):1–43, 2025. doi: 10.1145/3757594.

A. Appendix: User Study Questionnaire

Participants completed a questionnaire before and after the design session. The questionnaire collected background information, prior experience, familiarity with relevant tools, and reflections on the AI-assisted design workflow. The study evaluated the prototype interaction rather than the participants' design ability.

A.1. Consent

Before participating, participants were asked to confirm consent for participation, analysis of interaction data, use of anonymised quotes, use of anonymised screenshots, sketches or design artefacts, and screen recording during the session.

A.2. Participant Background

Participants were first asked to describe their background. They select among multiple fields of expertise from the following options:

- Architecture
- Architecture-engineering
- Structural engineering
- Civil engineering
- Computational design
- Design research
- Engineering research
- Architecture student
- Engineering student
- Other

The participants were then asked to report their years of experience in the following areas:

- Structural design practice
- Architectural design practice
- Computational design or parametric modelling
- Engineering studies
- Architecture studies

The response options for these questions were:

None; less than 1 year; 1–3 years; 3–5 years; 5–10 years; more than 10 years.

A.3. Familiarity With Tools and Workflows

Participants rated their familiarity with relevant tools and workflows using a five-point scale:

1 = never used; 2 = tried once or twice; 3 = occasional user; 4 = frequent user; 5 = expert or professional user.

The following categories were included:

- 3D modelling tools, such as Rhino, SketchUp, Blender, or Revit
- Parametric modelling tools, such as Grasshopper or Dynamo
- Structural analysis or FEM tools, such as RFEM, SAP2000, ETABS, Karamba, or Abaqus
- AI conversational tools, such as ChatGPT, Claude, or Gemini
- AI tools for design, CAD, modelling, or coding
- Sketching as part of the participant's design process

A.4. Pre-Task Questions

Before starting the design task, participants answered questions about their usual early-stage design process and their expectations for AI assistance.

1. How often do you explore multiple structural alternatives before converging on a concept?
2. When starting an early-stage structural design problem, which inputs do you usually rely on?
3. What do you expect an AI assistant to be useful for in early-stage structural design?

For the first question, participants used a five-point scale from “never” to “always.” For the second question, participants could select multiple options including:

- Sketches
- Plans or drawings
- Physical models
- Previous 3D models
- Precedent images
- Parametric tools
- Other

A.5. Post-Task Reflection Questions

After completing the AI-assisted design session, participants reflected on how the interaction affected their design process. The questions focused on design evolution, exploration, friction, agency, constraint reasoning, communication modes, and perceived level of collaboration.

1. Did the interaction with the 3D interface modify your initial design idea? If so, how?
2. Did the AI help you explore alternatives you would not have considered alone? If so, which alternatives?
3. Did the AI anticipate next steps in a similar way as you would have done alone? If so, when?
4. Where did the AI remove repetitive modelling friction?
5. Where did the AI introduce new friction?
6. Did you feel in control of the design direction?
7. Did the AI help on your reasoning about constraints? If so, how?
8. Which mode did you find most useful and natural to communicate with the AI?
9. Would you describe the interaction as automation, collaboration, or co-creation? Why?
10. Do you have any additional comments?

Post-task questions used a five-point scale.

1 = not at all / strongly disagree; 5 = very much / strongly agree.

For the communication, participants selected one of the following options:

- Natural-language only
- Natural-language + sketch
- Sketch only
- Direct modification of the 3D interface
- Other

B. Appendix: Participant demographics responses.






Figure B.1. Summary of the demographic results and experience collected from the participants.

C. Appendix: Sketches created by the participants


The sketches provided by the participants are shown in this appendix.



(a) Participant P01



(b) Participant P02



(c) Participant P03

Figure C.1. Sketches made by the participants during the sketch session.