From Generative Output to Traceable Design Decisions: A Process-Level Framework for Human–AI Co-Creation in Conceptual Design

Miguel-Ángel Agustín-Fonfría[1][0000−0002−9803−502X], Adolfo-Fernando Ventura-Salom[1][0009−0004−5345−5831], and Ana-Belén Barranco-Sánchez[2][0009−0004−8145−5038]

[1] Universitat Politècnica de València, Departamento de Ingeniería Gráfica, Valencia, Spain
{magustin,advensa}@upv.es

[2] E.S. de Diseño de Aragón, Departamento de Proyectos y Técnicas de Diseño de Producto, Zaragoza, Spain
abbarranco@esda.es

Abstract. Generative visual AI is increasingly used in early-stage design to produce sketches, concept renderings and alternative visual directions. However, visually convincing outputs do not necessarily constitute traceable or justified design decisions. In human–AI co-creation, generated images may remain weakly connected to design intent, evaluation criteria, evidence and design rationale, making the creative process difficult to review, reuse, and, where applicable, assess consistently. This paper proposes a process-level methodological framework for human–AI co-creation with generative visual AI systems. The framework is conceived as a lightweight traceability schema linking design problem, design intent, evaluation criteria, generative input, AI output, selected and rejected features, specifications, evidence, design decisions and iteration records. Its purpose is not to explain the internal behaviour of the generative model, but to document the human interpretation and decision-making process around generated visual artefacts. The proposal is illustrated through a postgraduate conceptual product design case in which generative visual AI was introduced after a systemic intent-structuring phase. The case provides evidence of feasibility for reconstructing the relationship between prior intent, generated outputs, and design decisions, but it is not presented as a controlled validation. The paper contributes to creative AI and human–AI interaction by reframing generated outputs as intermediate artefacts that require explicit links to intent, evidence and accountable design decisions.

Keywords: Generative AI · human–AI co-creation · traceability · conceptual design · design rationale · responsible AI

1 Introduction

Generative artificial intelligence (AI) systems are rapidly entering creative and design-oriented workflows. In conceptual product design, generative visual AI

systems can produce sketches, concept renderings, stylistic variations and alternative visual directions in a short time. This expands the exploratory space available to designers and reduces the representational effort needed to communicate emerging ideas. However, the same capacity introduces a methodological risk: a visually plausible output may be mistaken for a well-founded design concept. A generated image can appear coherent and professionally rendered while remaining weakly connected to the design problem, user needs, functional requirements, technical constraints or semantic intention of the product.

This risk is particularly relevant in human–AI co-creation. In such processes, the relevant unit of analysis is not the isolated image produced by the system, but the interaction between the human designer, the generative model, the generated artefact and the decision that follows. Recent research shows that generative AI can support ideation and creative exploration, but may also affect diversity, fixation, control, authorship and evaluation of creative outcomes [7, 11, 14]. In design research, generative AI has been examined as a support for conceptual design, design cognition and human–AI co-creation processes [5, 9, 20]. These studies suggest that the value of generative AI depends not only on the quality of generated outputs, but also on how designers formulate intentions, interact with the system, interpret results and justify decisions.

The problem addressed in this paper is therefore not whether generative AI can produce useful images for conceptual product design. The central question is how generated visual outputs can be made traceable, interpretable and open to evaluation within a human-led design process. Without process traceability, the transition from image generation to design decision can remain opaque: designers may accept a visual result because it is attractive, surprising or apparently feasible, while the reasons for selecting particular features, rejecting others or reformulating the input remain undocumented. This weakens comparison, evaluation, detection of AI-induced deviations and reuse of knowledge generated during co-creation.

This paper addresses the following research question: *how can generated visual outputs be externally traced to human design intent, evaluation criteria, evidence and decisions in a human-led co-creative process?* The proposed answer distinguishes between internal model explainability and external process traceability. Explainable AI (XAI) addresses the transparency and intelligibility of AI systems and their outputs [3, 10]. In many practical uses of generative visual AI, however, the internal generation process remains inaccessible to the designer. What can be made explicit is the human side of the process: the design intent that precedes generation, the criteria used to assess outputs, the features selected or discarded, the evidence used to justify decisions and the iterations that connect one result to the next.

The contribution of the paper is threefold. First, it formulates the traceability gap that appears when generated visual outputs are used as design evidence without documenting their relation to intent, criteria and decisions. Second, it defines a process-level methodological framework—a lightweight traceability schema—that links design problem, design intent, generative input, AI output,

feature selection or rejection, specification mapping, evidence, decision and iteration. Third, it illustrates the framework through a postgraduate conceptual product design case based on a documented pilot experience in which generative visual AI was introduced after a systemic conceptual structuring phase [1]. The case is used to demonstrate feasibility of process reconstruction and analytical usefulness, not as a full experimental validation or a claim of improved design performance.

The paper is organised as follows. Section 2 summarises the relevant background and identifies the gap. Section 3 presents the traceability framework. Section 4 illustrates it through the application case. Section 5 discusses implications and limitations, and Section 6 concludes the paper.

2 Related Work and Gap

2.1 Human–AI Co-Creation and Generative Design

Human–AI interaction research emphasises that AI systems should support appropriate expectations, feedback, controllability, recoverability and human judgement throughout the interaction cycle [2]. These principles are especially relevant in creative settings, where outputs are not evaluated against a single correct answer but against situated combinations of intention, novelty, usefulness, constraints and interpretive value.

In creative and design-oriented contexts, the relation between humans and AI is better described as co-creation than as automation. Generative systems can provide variation, unexpected combinations and rapid externalisations, but the human participant remains responsible for framing the task, interpreting outputs and deciding what should be retained, rejected or transformed. This view is consistent with an augmentation-oriented understanding of AI, in which the aim is to expand human capability rather than replace human reasoning [4].

Recent studies on generative AI in design reinforce this process-oriented view. Fang et al. review generative-AI-enhanced human–AI collaboration in conceptual design and highlight the need to understand interaction across the design process rather than evaluating outputs in isolation [9]. Chen et al. analyse human–AI co-creation through cognition, process, methodology and outcome dimensions [5]. Wang et al. show that creativity in human–AI co-creation depends on design experience, task framing and the way generated alternatives are interpreted and developed [20]. Related work on sketch-based and multimodal generative exploration also shows that design workflows increasingly move across sketches, prompts, images and prototype representations [8, 21].

A common implication is that generative AI should be understood as part of a chain of representational transformations. A design process may move from problem statement to brief, from brief to prompt, from prompt to generated image, from image to selected features, and from those features to specifications or decisions. Each transformation may add, omit or distort information. If these transformations are not documented, the generated image becomes an impressive endpoint rather than a reviewable record of how the design decision was reached.

2.2 Traceability, XAI and Documentation

XAI research responds to the opacity of AI systems and the need to make their behaviour more understandable to human users [3, 10]. Miller further argues that explanation in AI should not be treated only as a technical description of model behaviour, but also as a human-oriented process of making reasons understandable in context [17]. This distinction is important in generative visual co-creation. A designer does not only need to know why a model generated a given image; the designer needs to document why that image, or part of it, was considered relevant to a design problem.

Responsible and trustworthy AI frameworks emphasise transparency, accountability, human oversight and documentation [6]. Documentation practices such as model cards show that AI artefacts should be accompanied by structured information about intended uses, evaluation conditions and limitations [18]. Although model cards document models rather than creative processes, they illustrate a principle that is transferable to human–AI co-creation: AI-mediated artefacts should be accompanied by information that makes their use reviewable and accountable. In this paper, traceability is treated as the documentary condition that allows a process to be reconstructed, whereas accountability is understood as a possible consequence of being able to identify decisions, reasons and responsibilities.

Conceptual product design is a suitable domain for process traceability because it requires continuous transformation between intentions, constraints, visual representations and decisions. Systemic design models structure this transformation through objectives, input variables, action variables, information variables, essential variables and feedback loops [13]. A related systemic model for organisational processes explicitly emphasises the traceability of objectives through actions, information, monitoring and feedback [12]. In this paper, this vocabulary is used only to structure traceability relations; the proposed framework does not require adoption of the full systemic design method.

The objective of making design decisions explicit and reviewable has a well-established precedent in design rationale research. Issue-based approaches such as IBIS were proposed to structure argumentative problem-solving processes later adopted in design rationale research, while QOC (Questions, Options and Criteria) notation documented the space of alternatives considered and the criteria used to select among them [15, 16]. Later surveys of design rationale systems examined approaches, representation, capture and retrieval in engineering contexts [19]. The framework proposed in the present paper shares this general motivation, but addresses a specific condition of human–AI co-creation that was absent from pre-AI design rationale systems: the need to trace how a machine-generated visual artefact, produced through an internal process that remains inaccessible to the designer, is interpreted, filtered and incorporated into a human-led design decision. In this sense, the proposed traceability schema adapts the design rationale tradition to generative visual AI workflows.

The gap addressed in this paper lies at the intersection of these areas. Existing work shows the relevance of generative AI for design, the importance of human–

AI interaction, and the need for transparency and documentation in AI systems. Less attention has been paid to the specific process-level traceability schema needed to connect generated visual outputs with human design intent, evidence and accountable decisions. The following section defines such a layer.

3 External Traceability Framework

The proposed framework is conceived as a process-level methodological layer between the generative model and the final design decision. Its purpose is not to inspect the internal mechanisms of the model, but to make explicit the human reasoning that surrounds the use of generated visual artefacts: what was intended, how the system was prompted, what was produced, what was selected or rejected, what evidence supported the decision and how the next iteration was defined.

Operationally, the framework should be read at three levels: a set of traceability entities, a set of relations among those entities, and a compact record that instantiates those relations for each co-creation episode. Thus, the proposal is not only the sequence shown in Eq. (1), but the combination of entity definition, relational mapping and documentary record illustrated in Tables 1–3.

The framework is based on three assumptions. First, a generated visual output is an intermediate artefact, not a design solution. Second, the value of that artefact depends on its relation to a previously formulated design intent and explicit evaluation criteria. Third, traceability is produced by documenting relations among entities, not by storing images alone. Table 1 summarises the core entities of the framework.

Traceability emerges from explicit relations among these entities. Intent–input relations document how the design intent is translated into a prompt, sketch or reference. Input–output relations identify which generative input produced each output. Output–feature relations identify what is extracted from a generated image. Feature–specification relations translate selected features into design implications. Evidence–decision relations justify why a feature, output or direction is accepted, rejected or modified.

Together, these relations form a minimal traceability chain:

$$P \rightarrow I \rightarrow C \rightarrow G \rightarrow O \rightarrow F \rightarrow S \rightarrow E \rightarrow D \rightarrow R. \tag{1}$$

The chain is not intended as a rigid sequence. Designers may return from output evaluation to intent reformulation, from feature analysis to prompt redesign, or from decision review to new criteria. The notation only indicates the minimum set of relations required to reconstruct how a generated artefact becomes part of a design decision.

The framework can be connected to systemic design variables. Contextual requirements and prior criteria operate as input variables; prompts and generative strategies operate as action variables; generated outputs and annotations operate as information variables; evaluation criteria, evidence and decisions operate as essential or control variables; and iteration records close the feedback loop.

Table 1. Core entities of the external traceability framework.

Entity	Role in the traceability chain
Design problem (P)	Initial need, opportunity or situation that motivates the design process.
Design intent (I)	Human-formulated conceptual direction expressing what the design should achieve or communicate.
Criteria (C)	Explicit criteria used to assess whether generated outputs are aligned with the intent.
Generative input (G)	Prompt, sketch, reference image or multimodal input provided to the generative system.
AI output (O)	Visual artefact produced by the generative system.
Feature (F)	Visual, formal, functional, semantic or interaction-related element identified in the output.
Specification (S)	Translation of a feature into a requirement, constraint or actionable design property.
Evidence (E)	Observation, argument, comparison, user need or technical criterion supporting a decision.
Decision and iteration (D, R)	Human selection, rejection, modification or acceptance, and the record linking one cycle to the next.

This mapping prevents the generated image from occupying the centre of the process. The image is one information-bearing element within a broader system of objectives, actions, evaluation and feedback.

For practical use, each co-creation episode can be documented through a compact traceability record. The record should answer five questions: (1) what was the design intent before generation? (2) what input was used to activate the generative system? (3) which features were selected or rejected? (4) what evidence justified the decision? and (5) what changes in the next iteration? The inclusion of rejected features is deliberate. In generative workflows, discarded outputs may reveal prompt ambiguity, technical infeasibility, ergonomic inconsistency, stylistic bias or deviation from the intended user profile. Recording rejection therefore preserves evidence of human judgement rather than treating failed images as irrelevant noise.

The framework can be evaluated through four dimensions. *Completeness* refers to whether the main entities of the chain are present. *Coherence* refers to the alignment among intent, input, output features and decision. *Traceability* refers to the possibility of reconstructing how a decision was reached. *Accountability* refers to the possibility of identifying who made a decision, on what basis and with what implications. These dimensions do not evaluate the internal quality of the model; they evaluate the quality of the human–AI design process surrounding the generated output.

The framework has two limits. It does not provide a technical explanation of how a generative model produces an image, and it does not guarantee good design outcomes by itself. A traceable process may still lead to a weak decision

if the initial intent is poorly formulated or the evidence is insufficient. Its value is methodological: it makes such weaknesses visible and open to review.

4 Application Case

The framework is illustrated through a postgraduate conceptual product design case derived from a documented pilot experience in the Master's Degree in Design Engineering at the Universitat Politècnica de València [1]. The original intervention involved 25 students organised into project teams and introduced generative visual AI during a three-week block within a project-based design course. Evidence included a student survey, lecturer observation and qualitative review of project deliverables. The case is used here as a proof of concept rather than as a controlled experimental validation.

The relevant feature of the case is that generative AI was not introduced as a free exploratory tool at the beginning of the project. Students first developed a systemic product structuring phase based on a model that organises context, objectives, formal, functional and ergonomic dimensions, variables and feedback relations [13]. This work was operationalised through conceptual panels that condensed product meaning, references, user-related considerations, functional expectations, formal directions and contextual constraints. In framework terms, the panel provided the initial design problem (P), design intent (I) and evaluation criteria (C).

The generative interaction then used sketches, textual instructions, reference images and base images derived from the panel. The tool used in the pilot was Vizcom, but the present paper treats it as one instance of a broader class of generative visual AI systems. The key issue is not tool performance, but the relation between prior intent, generative input and human evaluation of outputs. Generated images were interpreted as intermediate artefacts: they had to be compared with the conceptual panel, broken down into selected or rejected features, and connected to specifications or design decisions.

The survey was answered by 13 of the 25 enrolled students (52%). Although the sample is small, the responses provide useful evidence for assessing feasibility. Perceived professional value was high: 69.2% of respondents gave the maximum score and 23.1% gave the second-highest score. When asked to compare the AI-supported workflow with early compulsory technical 3D modelling, 46.2% rated it as "much better", 46.2% as "better" and 7.7% as equivalent. In addition, 92.3% reported that the tool enabled more variant exploration in less time. These results do not validate the framework experimentally, but they support the claim that the workflow was perceived as professionally relevant and feasible in early-stage conceptual exploration.

This process also made visible a methodological tension. Students valued the acceleration and visual quality of AI-supported exploration, but the usefulness of generated outputs depended strongly on the clarity of the prior conceptual structure. When the panel was explicit, prompts and references were more precise and outputs were easier to evaluate. When the panel was weak or dispersed,

the AI system tended to produce attractive but less coherent alternatives. This suggests that prompt quality is not merely a linguistic skill; it depends on the quality of pre-generative design reasoning.

Table 2 summarises the mapping between the case and the framework, combining the main traceability entities with the preliminary affordances observed in the pilot.

Table 2. Mapping of the application case to the traceability framework.

Framework element	**Case instantiation**	**Traceability value**
Problem, intent and criteria	Product challenge and conceptual panel	Makes intent–output alignment reviewable and, in the educational setting, gradeable.
Generative input and output	Sketches, prompts, references and generated alternatives	Connects visual generation to prior design knowledge.
Selected features	Relevant forms, materials, layouts or interaction cues	Converts visual suggestions into design implications.
Rejected features	Misaligned styles, infeasible forms or unintended cues	Preserves evidence of judgement and detects AI-induced deviations.
Evidence and decision	Panel content, team reasoning, observation and deliverables	Justifies selection, rejection, reformulation or iteration; see Table 3 for a concrete instantiation.
Iteration record	Sequence of inputs, outputs and revisions	Enables review, comparison and reuse of process knowledge.

The case also illustrates why rejected outputs should not be discarded without comment. A rejected image may reveal that a prompt was too vague, that a reference image introduced an unintended stylistic bias, that the model overemphasised a trend, or that the design team had not clarified a relevant criterion. Such rejections are not failures of the process; they are evidence of human judgement. Recording them supports the transition from visual exploration to design rationale.

To illustrate the operational use of the framework, Table 3 presents a retrospective traceability reconstruction for one co-creation episode from the BBQ Experience project, previously reported as part of the pilot experience [1]. The record was reconstructed from available artefacts: the conceptual panel, generated renders and project deliverables. It does not reproduce a real-time log or literal prompts captured during the activity. Rather, it shows that the proposed traceability chain can be reconstructed a posteriori when the prior conceptual structure is sufficiently explicit.

Table 3. Retrospective traceability reconstruction: BBQ Experience project, one co-creation episode.

Entity	Reconstructed content
Design problem (P)	Outdoor cooking experience for social gatherings, combining portability, ease of use and integration of grilling and serving functions.
Design intent (I)	Compact portable BBQ unit with a premium contemporary aesthetic and functionally transparent use. Formal direction: clean volumes, vertical orientation, dark materials and warm accents.
Criteria (C)	Alignment with the conceptual panel: vertical proportions, enclosed body, visible grilling area, avoidance of industrial-kitchen appearance, standing-height interaction and accessible controls.
Generative input basis (G)	Reconstructed input basis derived from the panel: vertical BBQ sketch, descriptors such as compact, premium, dark finish and outdoor social context, plus references to contemporary outdoor furniture.
AI output (O)	Available Vizcom renders showing vertical BBQ units with variations in proportion, surface treatment, chimney configuration and base design.
Selected features (F)	Enclosed dark body, vertical proportions, visible grilling window, minimalist integrated legs and warm wood accent in handle or shelf area.
Rejected features	Horizontal or traditional BBQ proportions, industrial ventilation grilles inconsistent with the premium intent, and oversized chimney configurations dominating the formal composition.
Evidence (E)	The selected direction aligns with panel criteria: vertical proportion, dark material palette, visible but integrated grilling area and contemporary outdoor character. Rejected alternatives moved toward conventional BBQ typologies or industrial cues not specified in the panel.
Decision (D)	Retain vertical enclosed configuration with wood accent as the primary direction. Discard horizontal variants. Refine chimney proportions in the next cycle.
Iteration record (R)	Next cycle: adjusted chimney-to-body ratio, preservation of the same conceptual descriptors and use of the selected render as a reference for further generation.

This reconstruction illustrates three properties of the framework. First, rejected features are as informative as selected ones because they reveal where the generative system drifted from the stated intent. Second, the iteration record shows how a selected output can become a new generative input, closing the feedback loop described in the framework. Third, the reconstruction was feasible because the prior conceptual panel provided explicit criteria against which outputs could be compared. In projects where the panel was less structured, a comparable reconstruction would have been more difficult, which reinforces the framework's dependency on pre-generative conceptual quality.

The case demonstrates the feasibility of applying the framework to reconstruct the relation between prior intent, generated outputs and design decisions in a real AI-supported design activity. Its educational setting and small sample prevent causal claims about improved design quality, but the triangulation of survey responses, lecturer observation and deliverable review provides a grounded proof of concept for the proposed process-level traceability schema.

5 Discussion

The proposed framework reframes generative visual AI in conceptual design as a problem of traceable human–AI co-creation rather than as a problem of image

production alone. This shift is relevant because the most visible contribution of generative systems—the rapid production of convincing visual alternatives—is also a source of methodological vulnerability. Generated images may accelerate exploration, but they do not automatically preserve the reasoning that connects a design problem with a justified decision.

The framework contributes to human–AI co-creation by supporting a structured interaction in which the human actor maintains control over intent, evaluation and decision. It is aligned with human–AI interaction principles that emphasise intelligibility, feedback, controllability and appropriate oversight [2]. The traceability record does not modify the generative model, but it helps designers make their interaction with the model explicit and reviewable.

The distinction between external process traceability and internal model explainability is essential. XAI research provides concepts for understanding and managing opaque AI systems [3]. However, the explanation required in a design setting is partly social, contextual and decision-oriented: why a designer considered an output relevant, why a feature was retained, why another was rejected, and how these decisions relate to intent and evidence. The proposed framework is therefore a process-level complement to XAI, not a replacement for model-level explanation.

The framework also extends documentation logic from AI artefacts to AI-mediated creative processes. Model cards show how structured reporting can make model characteristics, intended uses and limitations more visible [18]. The framework proposed here applies a similar principle to co-creation episodes. Its object is not the model or the dataset, but the sequence through which generated outputs are interpreted and transformed into decisions. This may support educational feedback, design review, professional documentation and later analysis of AI influence in creative workflows.

A distinctive aspect of the proposal is its link with systemic design modelling. By translating objectives, action variables, information variables, essential variables and feedback loops into a traceability grammar, the framework locates the generated image within a system of intent, action, evidence and evaluation. This helps preserve continuity between pre-AI conceptual reasoning and post-AI visual exploration, while avoiding the assumption that a polished output is equivalent to a justified concept.

Several limitations remain. The framework depends on the quality of human documentation; superficial records may create bureaucratic effort without real explanatory value. It also adds cognitive and procedural load, which may require different levels of granularity for educational, research or professional use. Finally, it does not solve issues such as training data provenance, intellectual property, bias or environmental cost. Its contribution is narrower: making the human decision process around generated artefacts more explicit, traceable and open to accountability.

6 Conclusions and Future Work

This paper has addressed a methodological problem that emerges when generative visual AI systems are incorporated into early-stage creative design processes. The problem is not only that generative models may be opaque, but that the human process surrounding their use can also become undocumented. When a generated image is treated as a design result without recording the intent that preceded it, the criteria used to evaluate it, the features selected or rejected and the evidence supporting the decision, the process loses traceability.

To address this gap, the paper has proposed a process-level external traceability framework for human–AI co-creation with generative visual AI systems. The framework links design problem, design intent, criteria, generative input, AI output, features, specifications, evidence, decisions and iterations. Its contribution is not to explain the internal behaviour of the model, but to make the human-led chain of interpretation and decision-making explicit.

The application case shows the feasibility of using the framework to reconstruct the relation between prior conceptual structuring, generated outputs and design decisions in a real postgraduate product design setting. The survey results provide preliminary support for the practical relevance of the workflow: respondents perceived high professional value, a clear advantage over premature technical 3D modelling in the conceptual phase, and faster exploration of variants. These findings remain exploratory, but they strengthen the proof-of-concept status of the proposal.

Future work should operationalise the framework through reusable templates, compare traced and non-traced generative design processes, and build annotated collections of human–AI design iterations. The framework should also be tested beyond conceptual product design. By shifting attention from the generated image to the traceable process that gives it meaning, the proposal aims to support a more accountable and human-led use of generative AI in creative practice.

References

1. Agustín-Fonfría, M.Á., Ventura-Salom, A.F., Adán-Roca, V., Santiago-Praderas, V.: El concepto antes que la imagen: IA generativa y modelo sistémico en la enseñanza del diseño de producto. In: Libro de actas del XII Congreso de Innovación Educativa y Docencia en Red. Editorial UPV, Valencia (Jul 2026).
2. Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E.: Guidelines for Human-AI Interaction. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–13. CHI '19, Association for Computing Machinery, New York, NY, USA (May 2019).
3. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion **58**, 82–115 (Jun 2020).
4. Brynjolfsson, E.: The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. Daedalus **151**(2), 272–287 (May 2022).

5. Chen, G., Yu, Z., Xie, Y., Liu, Z., Yu, C.: The study of human-AI co-creation design under generative artificial intelligence: Cognition, process, method, and outcome. Journal of Engineering Design **37**(4), 1–42 (2026).
6. Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López De Prado, M., Herrera-Viedma, E., Herrera, F.: Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. Information Fusion **99**, 101896 (Nov 2023).
7. Doshi, A.R., Hauser, O.P.: Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances **10**(28), eadn5290 (Jul 2024).
8. Edwards, K.M., Man, B., Ahmed, F.: Sketch2Prototype: Rapid conceptual design exploration and prototyping with generative AI. Proceedings of the Design Society **4**, 1989–1998 (May 2024).
9. Fang, C., Zhu, Y., Fang, L., Long, Y., Lin, H., Cong, Y., Wang, S.J.: Generative AI-enhanced human-AI collaborative conceptual design: A systematic literature review. Design Studies **97**, 101300 (Mar 2025).
10. Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., Yang, G.Z.: XAI-Explainable artificial intelligence. Science Robotics **4**(37), eaay7120 (2019).
11. Heigl, R.: Generative artificial intelligence in creative contexts: A systematic review and future research agenda. Management Review Quarterly **76**(1), 955–992 (Feb 2026).
12. Hernandis Ortuño, B., Agustín-Fonfría, M.Á., Esnal-Angulo, I.: Modelo Sistémico Para La Gestión de Empresas. ID-Think®, rDis® Red Internacional de Investigación en Diseño Sistémico., Valencia, 1 edn. (2017).
13. Hernandis Ortuño, B., Agustín-Fonfría, M.Á., Esnal-Angulo, I.: Ecodiseño Concurrente. Red Internacional de Investigación en Diseño Sistémico, Valencia, 1 edn. (2019).
14. Hou, J., Wang, L., Wang, G., Wang, H.J., Yang, S.: The double-edged roles of generative AI in the creative process: Experiments on design work. Information Systems Research pp. 1–20 (Oct 2025).
15. Kunz, W., Rittel, H.W.J.: Issues as Elements of Information Systems. Working Paper 131, Institute of Urban and Regional Development, University of California, Berkeley, CA (Jul 1970)
16. MacLean, A., Young, R.M., Bellotti, V.M., Moran, T.P.: Questions, Options, and Criteria: Elements of Design Space Analysis. Human-computer interaction **6**(3-4), 201–250 (1991).
17. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence **267**, 1–38 (Feb 2019).
18. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T.: Model Cards for Model Reporting. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. pp. 220–229. FAT* '19, Association for Computing Machinery, New York, NY, USA (Jan 2019).
19. Regli, W.C., Hu, X., Atwood, M., Sun, W.: A Survey of Design Rationale Systems: Approaches, Representation, Capture and Retrieval. Engineering with Computers **16**(3-4), 209–235 (Dec 2000).
20. Wang, N., Kim, H., Peng, J., Wang, J.: Exploring creativity in Human–AI co-creation: A comparative study across design experience. Frontiers in Computer Science **7**, 1672735 (Sep 2025).
21. Zhang, C., Wang, W., Pangaro, P., Martelaro, N., Byrne, D.: Generative image AI using design sketches as input: Opportunities and challenges. In: Creativity and Cognition. pp. 254–261. ACM, Virtual Event USA (Jun 2023).