From “I Feel” to “Users Need” AI-Simulated Experience and the Construction of User Knowledge in Inclusive Architectural Design

Mohamed Ahmed Mohamed Abdellatif Khalil[1], Hee Sun (Sunny) Choi[2]*
1 School of Design, The Hong Kong Polytechnic University, mohamed.khalil@connect.polyu.hk;
2 School of Design, The Hong Kong Polytechnic University, Hee-sun.choi@polyu.edu.hk;
*Corresponding author

Abstract

Generative AI now allows architects to converse with simulated users who speak in the first person about spatial needs, discomforts and preferences. This creates a new challenge for empathetic and inclusive design: synthetic experience can sound like lived testimony and then migrate into briefs, programmes and spatial requirements with increasing authority. This paper addresses how AI-generated experiential language changes epistemic status as it is translated into architectural decision-making. Using critical conceptual analysis and a comparative reading of two published AI-assisted spatial-design cases, the paper develops the concept of testimonial-to-requirement transformation. It identifies four recurrent shifts: loss of source attribution, hardening of epistemic modality, expansion of experiential scope and borrowed authority through formalisation. The paper concludes that empathy in design cannot be reduced to plausible simulation. Inclusive practice requires provenance, explicit uncertainty and opportunities for contestation before synthetic claims are treated as user knowledge.

Keywords

Generative AI, Empathetic design, Inclusive design, User representation, Architectural knowledge

Opening Image. Conceptual illustration of testimonial-to-requirement transformation in empathetic and inclusive architectural design. AI-generated with OpenAI ChatGPT image, August 2026.

1. Introduction: when "I feel" becomes "users need"

An architect can now ask a large language model (LLM) to "act as an older resident with limited mobility" or "respond as a visually impaired commuter". Within seconds, the system can return fluent first-person accounts of discomfort, preference, dignity, accessibility and belonging. Such outputs are attractive because they are immediate, conversational and apparently situated, including for populations whom a designer has not met.

The difficulty is not simply inaccuracy. Generative-AI persona research already documents stereotyping, over-generalisation and weak empirical grounding (Gu, Chandrasegaran and Lloyd, 2025; Haxvig, D'Andrea and Teli, 2025; Amin et al., 2025). Recent decision-support evidence sharpens the consequence; conversational personas introduce distinctive interaction challenges when they represent user groups to decision makers (Salminen et al., 2026), while demographic synthetic users can exaggerate between-group differences enough to misdirect downstream decisions (Chen, Zhu and Zheng, 2026). In product design, Personagram further connects census-based personas to actionable product features, showing that persona-to-design formalisation can already be operationalised outside spatial design (Kim et al., 2026). Nor is simulation equivalent to participation: computational proxies can undermine representation, understanding and meaningful influence (Agnew et al., 2024; Annand, Garasto and Groves, 2026), while simulated stakeholders can redistribute epistemic authority during needfinding (Elisha-Wigwe, Brubaker and Schauer, 2026).

A more specifically architectural problem begins after the synthetic persona has spoken. Architectural design operates through a succession of representational artefacts: interviews and notes may become personas; personas become scenarios; scenarios inform briefs and programmes; programme statements become requirements; requirements support diagrams, plans and specifications; and these eventually become material conditions. Architectural scholarship has long shown that knowledge about users is fragile within this process. Architects often work with implicit, contingent and person-bound understandings of future users, while scenario-based methods have been developed precisely to make spatial experience more explicit and discussable (Van der Linden, Dong and Heylighen, 2019a; 2019b). Requirements-traceability and BIM research, meanwhile, has sought to preserve links between requirements, design rationales and building elements (Ozkaya and Akin, 2007; Zahedi et al., 2022).

Generative AI inserts a peculiar new object into this already fragile chain: language that can take the grammatical form of lived testimony without a person who has lived the experience. In peer-support research, Goel et al. (2026) describe how LLMs can reproduce the narrative form of lived experience and sometimes imply experiential proximity despite lacking the experience itself. Philosophical work likewise distinguishes the epistemic standing of LLM-generated statements from ordinary human testimony (He and

Yang, 2025). In architecture, this matters because a statement initially recognisable as model simulation can be progressively reformatted until its synthetic origin disappears. “As an older resident, I feel exposed here” may become “older residents experience exposure”, then “users need a protected approach”, and finally “provide a sheltered transition”. The language becomes more impersonal and authoritative at the same time that it becomes more consequential spatially.

This paper names that process testimonial-to-requirement transformation: the translation of AI-generated experiential speech into increasingly formal architectural representations that may acquire the status of user knowledge without corresponding evidence from an actual user. The paper does not claim that architects have never converted assumptions into user requirements. They have. Nor does it claim that synthetic personas are inherently illegitimate. Rather, it argues that LLMs create a distinctive possibility for designer inference to arrive already voiced as if from an experiential subject, and that architectural representation can subsequently naturalise that inference as a requirement.

Two research questions organise the analysis:

RQ1. Through what representational and linguistic transformations do AI-generated first-person statements become architectural scenarios, user needs and spatial requirements?

RQ2. What do these transformations imply for empathy, evidence and contestability in inclusive architectural design?

The contribution is diagnostic. It connects three literatures that are usually treated separately: architectural user representation and briefing; synthetic users and AI-simulated experience; and the translation of requirements into design decisions. It then shows why provenance alone is insufficient unless the provisional status of a synthetic claim survives its movement across architectural representations. A formally traceable requirement may still be unevidenced, and a documented “user need” may originate in nobody’s lived experience.

2. Method and analytical scope

The paper uses critical conceptual analysis supported by comparative examination of published design cases. Here, spatial design includes architectural and urban-design decisions that configure built space. The purpose is not to estimate prevalence or evaluate a particular LLM, but to examine how AI-mediated user representation intersects with processes that convert user knowledge into design requirements.

The literature search was purposive and iterative across three clusters: (1) architectural user representation, scenario-based design, briefing and requirements traceability; (2) synthetic users, AI personas, simulated stakeholders and AI-generated experiential language; and (3) AI-assisted spatial-design systems linking natural-language requirements to spatial reasoning. Searches used ACM Digital Library, ScienceDirect, SpringerLink, SAGE, ASME Digital Collection and CumInCAD, supplemented by relevant journal and institutional repositories and backward/forward citation tracing. Peer-reviewed work was priori-

tised; recent preprints are identified as such. The final search update was conducted on 18 August 2026.

Two published cases were theoretically sampled as boundary cases rather than to estimate prevalence. Seyman Güray and Uyan (2025) document an AI-assisted interior-design workflow in which fictional user profiles and user-based scenarios enter programming and conceptual development. StreetDesignAI uses synthetic cyclist personas to evaluate infrastructure and support design modification (Wang et al., 2026). They expose two distinct handoffs from represented user perspective toward design formalisation. Wu, Cao and Ji (2026) is discussed separately in Section 3.2 only as evidence of downstream natural-language-to-space machinery because it does not simulate users. None of these studies is treated as evidence that epistemic reclassification is prevalent or that their authors committed it.

The comparative reading tracks four properties: source attribution (whether claim origin remains explicit), epistemic modality (whether uncertainty persists), experiential scope (whether a claim remains persona-specific or becomes a population claim), and design formalisation (whether it remains exploratory or becomes a programme item, criterion or prescription). The coding concerns what the published documentation makes visible, not an allegation of misconduct by the case authors.

3. The architectural construction of user knowledge

3.1 Users are already mediated in architectural practice

The "user" in architectural design is rarely a continuously present speaker. Design teams work with clients, regulations, consultants, briefs, demographic categories, previous projects, professional experience and imagined future occupants. Verhulst, Elsen and Heylighen (2016), observing an architectural design competition, showed that when direct users are absent architects still populate their reasoning with imagined people and anticipated activities. Van der Linden, Dong and Heylighen (2019a) later described architects' knowledge about users as fragile: much of it is contingent on particular people, tacitly held and mediated by the socio-material environment of practice. The problem is structural: design must represent situated, embodied and changing experiences before a building and many of its users exist.

Architecture has therefore developed representational devices that make absent users more present. Personas, user journeys and scenarios can give design teams concrete characters and sequences through which to discuss spatial experience. Tvedebrink and Jelić (2018) proposed personas and scenarios as a way to integrate body-environment knowledge into architectural design. Van der Linden, Dong and Heylighen (2019b) developed and tested a scenario-based approach in residential care projects, using specific characters and user journeys to expose otherwise implicit and sometimes contradictory as-

pects of future use. Their findings are particularly important for the present argument: scenarios are useful precisely because they can travel into design discussion and motivate spatial decisions.

This representational power is also the source of risk. A persona is not the person represented. A scenario is not an observation. A brief is not a neutral transcription of stakeholder speech. Briefing scholarship treats the brief as an evolving process through which needs, values and constraints are interpreted and negotiated rather than merely recorded (Blyth and Worthington, 2010). Inclusive-design theory similarly warns against treating a generic or statistically central user as the unquestioned subject of design. Bianchin and Heylighen (2018) argue that inclusion is inseparable from questions of justice and the distribution of design benefits and burdens; Heylighen and Bianchin (2018) further emphasise deliberative dimensions of inclusive design rather than reducing inclusion to technical accommodation.

The pre-AI problem is therefore well established: user knowledge is constructed and mediated. The relevant question is what changes when one of the mediators can itself generate plausible user speech.

3.2 From representation to requirement

Architectural production requires ambiguous concerns to become increasingly determinate. “Residents need more privacy” may eventually become an adjacency rule, a sightline constraint, an opaque partition or a separate entrance. This concern belongs to a wider design-rationale tradition that has long sought to preserve the reasons, alternatives and assumptions behind design decisions, including issue-based systems such as gIBIS (Conklin and Begeman, 1988). Architectural requirements research extends that logic into the relation between briefs, requirements and form. DesignTrack links requirement specification with form exploration so that relationships between requirements and solutions can be navigated across a project (Ozkaya and Akin, 2007). Zahedi et al. (2022) further extend it into BIM, linking owner requirements, regulations, design episodes, explanation tags and spatial or semantic constraints to specific building elements and properties.

These approaches answer an important downstream question: why is this design decision here? Yet they normally begin from an item that has already entered the system as a requirement, owner need or regulatory provision. Contemporary AI systems make the downstream translation increasingly direct: Wu, Cao and Ji (2026), for example, parse natural-language residential requirements into spatial structures and floorplan recommendations. Their system does not simulate users, but it demonstrates the machinery through which a linguistic requirement can acquire computationally actionable spatial form. Synthetic users introduce a prior question: what authorised this statement to become a user requirement in the first place? Perfect traceability from a requirement to a wall does not establish that the requirement was grounded in any affected person’s experience.

This distinction is important because architectural representations often become progressively less conversational. A participant can say "I am not sure", "that depends", or "I changed my mind". A programme item tends to be terser: "provide private waiting area". A drawing is terser still. As knowledge moves downstream, the social conditions under which it was produced can disappear even while the design team gains apparent clarity. Generative AI can intensify this problem because it can provide the upstream statement in a form that resembles direct speech before any direct speaker has entered the process.

4. When AI speaks as the user

4.1 Synthetic users are not simply faster personas

LLM-based personas have expanded from static descriptions into conversational agents. Gu, Chandrasegaran and Lloyd (2025) found that persona chatbots can support ideation but documented limitations in credibility, empathy and depth of user understanding. Haxvig, D'Andrea and Teli (2025) show that LLM-generated personas can reproduce gendered assumptions, while Amin et al. (2025) identify wider problems of stereotyping, validation, transparency and representativeness.

One response is stronger grounding. PersonaCite, a CHI 2026 Extended Abstract, constrains persona answers through retrieved voice-of-customer evidence and response-level source attribution (Truss, 2026). This clarifies the downstream question examined here: if provenance matters at the interface, what happens after the response is copied into a scenario, brief or programme?

Synthetic-user research has also reached a second limit: representation is not participation. Agnew et al. (2024) argue that replacing human participants with generative-AI surrogates can conflict with representation, inclusion and understanding; Annand, Garasto and Groves (2026) similarly distinguish public "say" from meaningful "sway". Elisha-Wigwe, Brubaker and Schauer (2026) show in early-stage design how an LLM-simulated stakeholder can reshape epistemic authority. In urban planning, Jauhiainen et al. (2026) use synthetic inhabitants and experts to pilot an immersive participatory environment while positioning them as complements to, rather than replacements for, human participation.

These studies establish why a synthetic persona should not automatically be treated as a participant. Chen, Zhu and Zheng (2026) further ask whether synthetic-user evidence is sufficiently valid for downstream decision support. The architectural question developed here is narrower and subsequent: how can those qualifications disappear after persona speech crosses into scenarios, briefs, programme requirements and spatial rationales?

4.2 Experiential language and borrowed standing

In a 2026 preprint on peer-like support for dementia caregivers, Goel et al. describe a "synthetic lived experience" problem: LLMs can reproduce personal-experience narratives and imply experiential proximity without having the experience described. Outside architecture,

the mechanism is nonetheless relevant because first-person forms locate a claim in an apparent speaker: "I have been there", "I find this difficult", "I would feel safer".

He and Yang (2025) approach a related issue philosophically, arguing that the conditions under which LLM-generated statements can count as a source of knowledge differ from ordinary human testimony. The present paper does not resolve that philosophical debate. It adopts the more limited observation that testimonial form and testimonial standing are separable. A model may produce a sentence that looks like situated testimony without being a person whose embodied history, accountability and social position ground that statement. Architects have always projected assumptions onto absent users; the LLM-specific change is therefore not inference itself. It is that inference can arrive already formatted as responsive first-person experience, making a model-generated possibility easier to encounter as if it were somebody's account.

This distinction matters for inclusive design because lived experience is not only a container for factual propositions. It is also contestable. A person who says "I need this entrance to be visible" can explain why, qualify the statement, disagree with the architect's interpretation, negotiate against competing requirements, or withdraw it. If the statement came from a simulated persona, no represented person necessarily exists who can perform those corrective acts. Once the source label disappears and the statement is recast as "users need a visible entrance", the requirement may retain the rhetorical authority of user-centredness while losing the ordinary possibility of being challenged by the user it invokes.

Contestability is therefore more than an ethical add-on. It is a correction mechanism within design. Participatory and inclusive approaches allow a representation of need to be returned to affected people through consultation, critique and negotiation. That correction can change a programme, expose a trade-off or prevent an apparently inclusive intervention from materialising as segregation, surveillance or secondary access. Synthetic personas may still be useful for generating questions, revealing overlooked possibilities or stress-testing a design. The problem arises when their claims cross into the register of represented human evidence without an event that warrants that change.

5. Published pathways from AI representation to spatial design

5.1 Persona to programme: Seyman Güray and Uyan

Seyman Güray and Uyan (2025) provide one of the clearest architectural precedents because their generative-AI integration model explicitly places user definition, scenario development and architectural programming in sequence. In a second-year interior-design studio, students used text-to-text and text-to-image systems during early design. The published workflow shows an inception phase involving fictional user profiles or personas, scenario development based on the user definition, and transformation of the scenario into the architectural programme.

The study is pedagogical rather than epistemological and does not present fictional AI-supported personas as empirical user research. Its value here is structural: persona-like representation and programme formation occupy adjacent steps in an AI-assisted architectural workflow.

A fictional persona is easy to recognise as fictional while it remains a persona. A programme item such as "quiet retreat space" or "separate circulation" carries no equivalent cue. Programming usefully compresses narrative into actionable statements, but the resulting requirement may no longer reveal whether it came from a participant, regulation, published evidence, designer assumption or generative model.

5.2 Persona to design modification: StreetDesignAI

StreetDesignAI provides a stronger inclusive-design case because its synthetic personas are not merely ideation devices. Wang et al. (2026) developed multiple cyclist personas to broaden transportation professionals' perspectives when evaluating cycling infrastructure. The system surfaces parallel persona responses to street designs and supports iterative modification in response to different needs. The DIS 2026 study therefore demonstrates a direct path from simulated user perspectives to spatial-design evaluation and change.

StreetDesignAI represents a comparatively responsible synthetic-user design: it is grounded in cyclist data, keeps multiple personas separate and surfaces conflict rather than constructing a single generic "user". The present risk therefore does not depend on a poorly designed persona system. It begins if persona-specific feedback is later extracted and restated without its source, uncertainty or scope.

Compare "Persona C flags this junction as potentially stressful" with "cyclists find this junction stressful". The first is a bounded model output; the second is a population claim. A later rationale such as "the junction was redesigned to address cyclist stress" may be justified by corroborating evidence, but it is not validated merely by appearing in a professional report.

Table 1. Comparative reading of the two published AI-assisted spatial-design cases

Analytical property	Seyman Güray & Uyan (2025)	StreetDesignAI (Wang et al., 2026)
Source attribution	GAI use and the user-scenario stage are explicit in the published workflow. The paper does not report whether AI/fictive provenance remains attached when scenario material is reformulated in programme or concept artefacts.	Persona identities and AI mediation remain explicit inside the system. The study does not examine whether those source markers persist after feedback is extracted into external briefs, reports or final design rationales.
Epistemic modality	The published account does not systematically track whether hedges, uncertainty or confidence markers survive the scenario-to-programme transition.	Persona outputs combine scores, observations and conversational explanations. The study does not analyse whether uncertainty markers survive when designers later summarise the feedback as a design rationale.
Experiential scope	User-based scenarios are attached to fictional user profiles at the scenario stage; whether profile-specific scope persists in programme statements is not reported.	Strongly preserved in-interface: four cyclist personas remain separate and disagreement is deliberately surfaced rather than collapsed into one user position.
Design formalisation	User-based scenarios feed architectural programming and early conceptual development, moving represented experience towards actionable design content.	Persona evaluations feed parameterised street modifications followed by re-evaluation, creating a direct loop from simulated perspective to spatial design change.

Note. Comparative analysis across source attribution, epistemic modality, experiential scope and design formalisation. Source: authors.

Table 1. Comparative reading of the two published AI-assisted spatial-design cases across source attribution, epistemic modality, experiential scope and design formalisation. Source: authors.

The comparison yields a qualified finding. Both cases make AI/persona mediation visible inside the published workflow, and StreetDesignAI preserves persona-specific scope by keeping disagreement explicit. Neither study examines whether source attribution, modality or persona-boundedness survives after user-like outputs leave the interface or scenario and reappear in independent briefs, programmes or final rationales. This does not show that laundering occurred. It identifies an unexamined handoff where the representational conditions that make a synthetic claim legible can be lost.

6. Testimonial-to-requirement transformation

The two cases, together with the downstream language-to-space mechanisms in Section 3.2, support a conceptual model of how synthetic experiential claims can travel through architectural work. Figure 1 shows five analytical stages: synthetic experiential claim, interpreted concern, user need, programme or requirement, and spatial decision. They need not always occur or follow this order; the sequence identifies points at which claim status can change.

Figure 1. *Testimonial-to-requirement transformation*

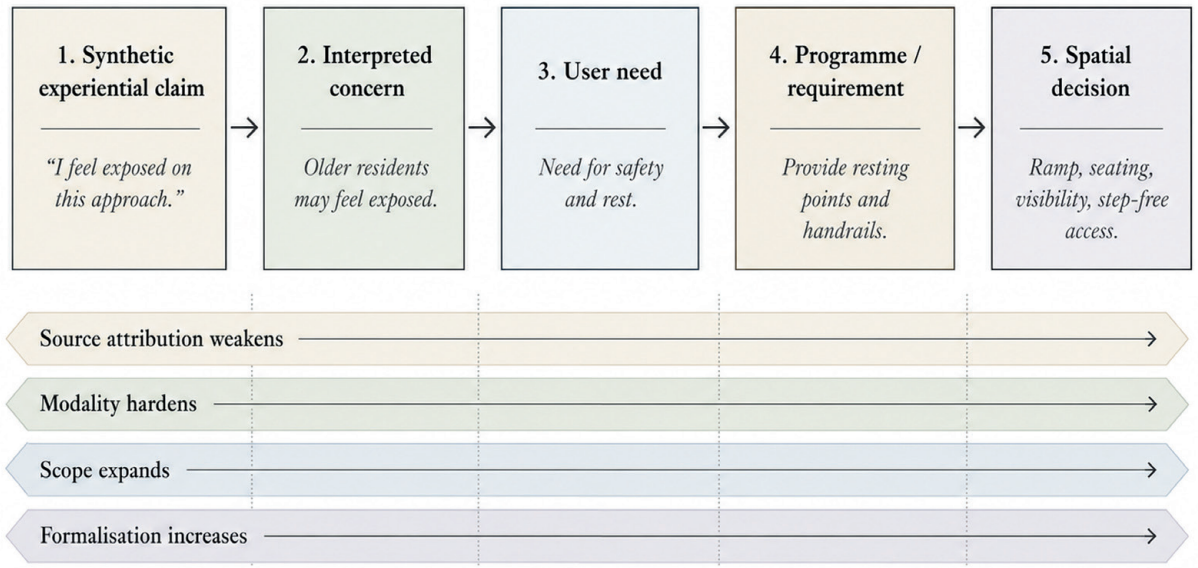

Figure 1. Testimonial-to-requirement transformation. A five-stage analytical sequence showing how a synthetic first-person experiential claim can be interpreted, generalised and formalised as a user need, programme requirement and spatial decision. The lower bands identify potential shifts in source attribution, modality, experiential scope and formalisation. Concept and content: authors; visual generated with OpenAI ChatGPT image under author direction.

At the first stage, a synthetic persona may say, for example, "As an older resident, I would feel exposed walking this long approach without somewhere to rest." Four features are visible: an apparent first-person speaker, an experiential verb ("feel"), conditional modality ("would"), and a specific situation. The statement is easy to identify as a simulated perspective if its source remains visible.

At the second stage, a designer may paraphrase it as "older residents may feel exposed on the approach". First person becomes a demographic category while uncertainty remains. At the third, "older residents need a protected resting point" converts experience into requirement. The programme may then state "provide sheltered rest point adjacent to entrance", removing the demographic subject. A later rationale may report that "the arrival sequence responds to elderly accessibility and comfort needs". A simulated perspective has now returned as apparently established knowledge before becoming spatial configuration.

Four transformations are analytically significant.

First, source attribution can disappear. "The AI persona suggested" becomes "users need". This is different from ordinary summarisation because the source category itself matters. A direct participant statement, a research synthesis, a designer inference and a model role-play should not become interchangeable merely because they share semantic content.

Second, epistemic modality can harden. Terms such as "might", "could" and "perhaps" can disappear when a narrative is compressed into a programme. This is a familiar feature of requirement formation, but it is especially consequential when the upstream source was already speculative.

Third, experiential scope can expand. A statement attributed to one simulated persona can become a claim about "older residents", "women", "wheelchair users" or "the community". Haxvig, D'Andrea and Teli (2025) and Amin et al. (2025) show why such generalisation is risky: synthetic personas can reproduce stereotypes and flatten within-group difference. A change from singular simulation to population statement is therefore not a stylistic edit; it is an evidentiary escalation.

Fourth, formalisation can create borrowed authority. Programmes, diagrams and BIM constraints are professional artefacts; entry into them can make a statement appear settled without making it better evidenced. Requirements-traceability research can document how a requirement shaped a design (Ozkaya and Akin, 2007; Zahedi et al., 2022), but that leaves a prior test: a traceable requirement is not necessarily an evidenced requirement.

This leads to a simple principle: a synthetic experiential claim should retain the uncertainty and source status of its origin unless an explicit validation event warrants reclassification. Validation may come from appropriate published evidence, standards, project-specific observation or stakeholder engagement. Translation into a more formal design representation cannot by itself count as validation.

Figure 2. *Contestability in the translation of user claims*

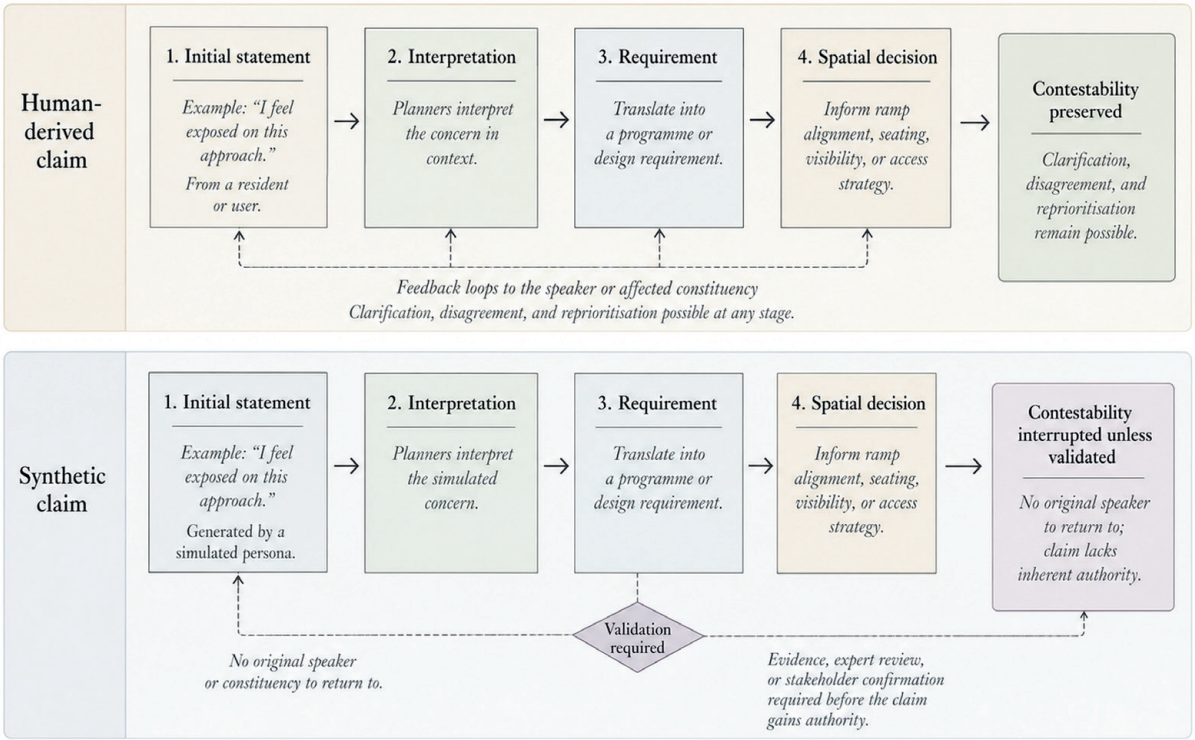

Figure 2. Contestability in the translation of user claims. Human-derived claims can, in principle, be returned to an affected speaker or constituency for clarification, disagreement and reprioritisation. Synthetic claims have no corresponding endogenous correction path, so validation requires deliberate re-entry of evidence, expertise or affected-user engagement. Concept and content: authors; visual generated with OpenAI ChatGPT image under author direction.

Figure 2 extends the argument by adding contestability. Human-derived claims can, at least in principle, be returned to speakers or affected constituencies for clarification and disagreement. Synthetic claims have no such endogenous correction path. Their correction depends on designers deliberately reintroducing external evidence or participation. When source attribution is lost, even that need can become invisible.

7. Implications for empathetic and inclusive design

7.1 Inclusion requires more than plausible representation

Synthetic personas can widen the perspectives considered during design. StreetDesignAI demonstrates a productive version of this use: mul-

tiple simulated perspectives expose conflicts and prompt deliberate trade-off reasoning rather than presenting a single "user" answer (Wang et al., 2026). The risk is therefore not synthetic representation itself, but allowing representational variety to stand in for inclusion.

Inclusive-design scholarship treats justice as more than generic usability (Bianchin and Heylighen, 2018), while participatory AI distinguishes computational representation from meaningful influence (Agnew et al., 2024; Annand, Garasto and Groves, 2026). Testimonial-to-requirement transformation can conceal that distinction when a design is said to "respond to users" although the operative claim was synthetic. The substitution risk may also be distributionally uneven. Synthetic personas are especially attractive where direct engagement is slow or difficult, including disability, care, migration and non-dominant-language contexts. This creates an inclusive-design paradox: synthetic representation may be most convenient precisely where situated knowledge is hardest to access and contest. Fricker (2007) defines testimonial injustice as a prejudice-driven credibility deficit; synthetic substitution is not identical because no human speaker may be present. Yet acting on representations of marginalised groups without a route for affected people to contest them can reproduce a related asymmetry: representation without reciprocal correction.

The consequence is clearest in contested spatial decisions. Technical accessibility claims may be checked against standards and empirical evidence. By contrast, deciding that an accessible entrance should be separate from a ceremonial entrance implicates dignity, visibility and hierarchy. If a synthetic persona's "privacy preference" becomes a programme requirement for segregation, there is no corresponding affected speaker who can object that the translation misconstrues the concern.

7.2 Contestability and provenance

The central safeguard is therefore **contestability**. Participants can clarify a statement, reject a summary, rank priorities differently or challenge the spatial form adopted in their name. These acts are part of how inclusive design corrects representation. A synthetic persona can generate another response when reprompted, but that is not equivalent to an affected person contesting how an institution has represented their interests.

Where synthetic claims enter a brief, programme or rationale, three distinctions should remain visible: what generated the claim; whether it is speculative or corroborated; and whether an affected constituency could contest the interpretation. This is not a new decision-validity framework. Chen, Zhu and Zheng (2026) already show that synthetic-user error can propagate into downstream decisions. The architectural issue is what happens when such evidence is reformatted into professional artefacts while provenance and modality become less visible. Grounding at the interface remains valuable, as PersonaCite illustrates (Truss, 2026), but provenance must travel into downstream architectural artefacts. Figure 4 translates these safeguards into a concise protocol for practice.

Figure 3. *Linguistic markers of transformation*

Marker	From	To	Design risk
Person shift	*I feel unsafe here*	*Users feel unsafe here*	Scope inflation
Modality shift	*may help*	*need*	False certainty
Source shift	*AI persona suggested*	*source omitted*	Provenance loss
Formalisation shift	*concern*	*requirement*	Borrowed authority

Figure 3. Linguistic markers of transformation. Illustrative shifts in person, modality, source attribution and formalisation that can signal a change in the epistemic status of an AI-generated user claim as it moves into design documentation. Concept and content: authors; visual generated with OpenAI ChatGPT image under author direction.

Figure 4. *Provenance and validation protocol for synthetic user claims*

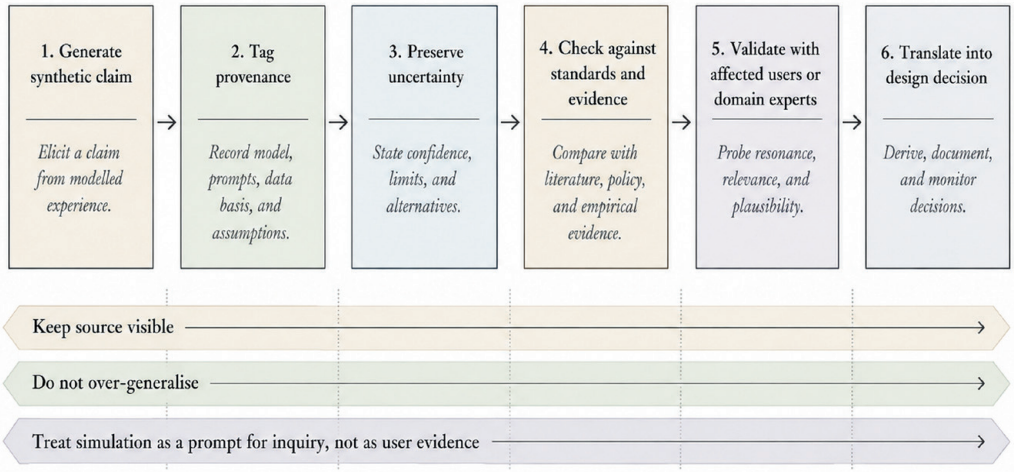

Figure 4. Provenance and validation protocol for synthetic user claims. A practical workflow for retaining provenance and uncertainty, checking claims against evidence and standards, validating them with affected users or relevant experts, and only then translating them into design decisions. Concept and content: authors; visual generated with OpenAI ChatGPT image under author direction.

7.3 Analytical markers and limitations

The transformation is partly linguistic. Person shift, modality and grammatical agency provide observable markers: "I feel" can become "users experience"; "might benefit" can become "need"; "the synthetic persona suggests" can disappear. Goel et al. (2026) show the analytical value of attention to first-person experiential language in another domain. Architectural research can extend this lens to briefs, programmes and rationales without treating the markers themselves as proof of misconduct. Figure 3 summarises these markers as practical signals for critical review.

Multilingual projects create a further research question because AI may simultaneously simulate a perspective and translate it into the professional language of the design team. Whether that double mediation changes which concerns survive into programmes warrants separate empirical study.

This paper does not test prevalence, causal effects or professional adoption. Its cases establish only that the relevant links already exist in published design research; they do not show that any case author converted an unevidenced claim into falsely evidenced user knowledge. Future research should compare real AI-assisted transcripts with derived briefs, programme statements and design rationales, testing source attribution, modality, scope and formalisation directly.

8. Conclusion

This paper has addressed how AI-generated experiential statements can change epistemic status as they move through architectural representation. In response to RQ1, the analysis shows a recurrent testimonial-to-requirement transformation: first-person synthetic claims are paraphrased, generalised and formalised until they can function as user needs, programme items and spatial rationales. Across this movement, four shifts recur: source attribution weakens, epistemic modality hardens, experiential scope expands and formalisation lends borrowed authority.

In response to RQ2, the central implication is that empathetic and inclusive design cannot equate plausibility with participation. A synthetic statement may help designers imagine neglected situations, but empathy in design is not the production of convincing first-person language. It depends on accountable relations between claims, evidence and the people in whose name design decisions are made. When "I feel" becomes "users need" without preserved provenance, explicit uncertainty or a validation step, simulated experience can acquire unwarranted standing inside the design process.

The practical message is therefore clear: synthetic personas should be used as prompts for inquiry, not as substitutes for user knowledge. If AI-assisted inclusive design is to remain both empathetic and intellectually rigorous, it must keep source status visible, maintain the provisional character of synthetic claims and reintroduce contestability before those claims harden into spatial decisions. Otherwise, architecture risks speaking

most confidently about users at precisely the moment when no user has actually spoken.

Data Availability Statement

No new empirical dataset was generated for this study. The analysis draws on published literature and publicly available research materials cited in the paper.

Declaration of Generative AI

During preparation of this work, the first author used OpenAI ChatGPT (GPT-5.6 Sol, August 2026) for drafting assistance, proofreading and language editing. The authors retain full responsibility for the final work.

References

[1] Agnew, W., Bergman, A.S., Chien, J., Díaz, M., El-Sayed, S., Pittman, J., Mohamed, S. & McKee, K.R., 2024. The illusion of artificial inclusion. Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, pp. 1-12. https://doi.org/10.1145/3613904.3642703.

[2] Amin, D., Salminen, J., Jansen, B.J., Shin, J. & Kim, D.H., 2025. Generative AI personas considered harmful? Putting forth twenty challenges of algorithmic user representation in human-computer interaction. International Journal of Human-Computer Studies, 205, 103657. https://doi.org/10.1016/j.ijhcs.2025.103657.

[3] Annand, P.J., Garasto, S. & Groves, L., 2026. Public say versus public sway: 'Lived experience literacy' in the pursuit of participatory AI. Big Data & Society, 13(2). https://doi.org/10.1177/20539517261447845.

[4] Bianchin, M. & Heylighen, A., 2018. Just design. Design Studies, 54, pp. 1-22. https://doi.org/10.1016/j.destud.2017.10.001.

[5] Blyth, A. & Worthington, J., 2010. Managing the Brief for Better Design. 2nd ed. Abingdon: Routledge.

[6] Chen, Z., Zhu, D. & Zheng, L.N., 2026. When synthetic users fail: A cross-domain benchmark of LLM-simulated human survey responses. arXiv:2607.26348 [preprint; under peer review].

[7] Conklin, J. & Begeman, M.L., 1988. gIBIS: A hypertext tool for exploratory policy discus-

sion. ACM Transactions on Office Information Systems, 6(4), pp. 303-331. https://doi.org/10.1145/62266.62278.

[8] Elisha-Wigwe, O., Brubaker, E.R. & Schauer, A., 2026. Simulated stakeholders: Using critical frameworks to examine LLM agents in early-stage design. Journal of Mechanical Design, advance online publication. https://doi.org/10.1115/1.4072422.

[9] Fricker, M., 2007. Epistemic Injustice: Power and the Ethics of Knowing. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198237907.001.0001.

[10] Goel, D., Goyal, A., Duddu, V., Pal, O., Rodriguez, V.J., Brown, D.S., Karkar, R., Yoo, D.W. & Saha, K., 2026. When AI says "I have been in similar situations": Synthetic lived experience in peer-like caregiver support. arXiv:2606.18057v2 [preprint].

[11] Gu, E.H., Chandrasegaran, S. & Lloyd, P., 2025. Synthetic users: Insights from designers' interactions with persona-based chatbots. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 39, e2, pp. 1-17. https://doi.org/10.1017/S0890060424000283.

[12] Haxvig, H.A., D'Andrea, V. & Teli, M., 2025. "I've never seen a glass ceiling better represented": Bias and gendering in LLM-generated synthetic personas from a participatory design perspective. International Journal of Human-Computer Studies, 205, 103651. https://doi.org/10.1016/j.ijhcs.2025.103651.

[13] He, J. & Yang, C., 2025. Testimony by LLMs. AI & Society, 40, pp. 6201-6213. https://doi.org/10.1007/s00146-025-02366-y.

[14] Heylighen, A. & Bianchin, M., 2018. Building justice: How to overcome the inclusive design paradox? Built Environment, 44(1), pp. 23-35. https://doi.org/10.2148/benv.44.1.23.

[15] Jauhiainen, J.S., Hakanpää, S., Honkasaari, H.-P., Kivilompolo, N., Kurri, M., Lehtiranta, L. & Nurminen, M., 2026. Generative AI in participatory urban planning: Synthetic inhabitants and experts. Land, 15(3), 407. https://doi.org/10.3390/land15030407.

[16] Kim, T., Hong, M.K., Chen, Y.-Y., Li, J.Q., Van, M.P., Hakimi, S., Kay, M. & Klenk, M., 2026. Personagram: Bridging personas and product design for creative ideation with multimodal LLMs. arXiv:2602.06197 [preprint]. https://doi.org/10.48550/arXiv.2602.06197.

[17] Ozkaya, I. & Akin, Ö., 2007. Tool support for computer-aided requirement traceability in architectural design: The case of DesignTrack. Automation in Construction, 16(5), pp. 674-684. https://doi.org/10.1016/j.autcon.2006.11.006.

[18] Salminen, J., Jung, S.G., Kaate, I., Xuan, T.T.T., Azem, J.Y., Aldous, K.K., Amin, D. & Jansen, B.J., 2026. AI representing personas representing user groups: Applying the agency theory to examine interaction challenges of conversational personas as decision-making tools. Decision Support Systems, 205, 114633. https://doi.org/10.1016/j.dss.2026.114633.

[19] Seyman Güray, T. & Uyan, B., 2025. Collaboration with generative artificial intelligence

in the early stages of design studio: A model proposal. Journal of Architectural Sciences and Applications, 10(1), pp. 434-450. https://doi.org/10.30785/mbud.1649820.

[20] Truss, M., 2026. PersonaCite: VoC-grounded interviewable agentic synthetic AI personas for verifiable user and design research. Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems, pp. 1-7. https://doi.org/10.1145/3772363.3798543.

[21] Tvedebrink, T.D.O. & Jelić, A., 2018. Getting under the(ir) skin: Applying personas and scenarios with body-environment research for improved understanding of users' perspective in architectural design. Persona Studies, 4(2), pp. 5-24.

[22] Van der Linden, V., Dong, H. & Heylighen, A., 2019a. Tracing architects' fragile knowing about users in the socio-material environment of design practice. Design Studies, 63, pp. 65-91. https://doi.org/10.1016/j.destud.2019.02.004.

[23] Van der Linden, V., Dong, H. & Heylighen, A., 2019b. Populating architectural design: Introducing scenario-based design in residential care projects. International Journal of Design, 13(1), pp. 21-36.

[24] Verhulst, L., Elsen, C. & Heylighen, A., 2016. Whom do architects have in mind during design when users are absent? Observations from a design competition. Journal of Design Research, 14(4), pp. 368-387. https://doi.org/10.1504/JDR.2016.082032.

[25] Wang, Z., Dai, Y., Lyu, D., Nader, M., Chen, S., Ye, W., Ding, Z. & Yan, X., 2026. StreetDesignAI: Broadening designer perspectives through multi-persona evaluation of cycling infrastructure. Proceedings of the 2026 ACM Designing Interactive Systems Conference, pp. 483-503. https://doi.org/10.1145/3800645.3812888.

[26] Wu, Y., Cao, S. & Ji, G., 2026. Natural language-driven architectural space cognition: An AI-agent-based intelligent recommendation system for residential floor plans. Proceedings of the 31st International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA), 1, pp. 93-102.

[27] Zahedi, A., Abualdenien, J., Petzold, F. & Borrmann, A., 2022. BIM-based design decisions documentation using design episodes, explanation tags, and constraints. Journal of Information Technology in Construction, 27, pp. 756-780. https://doi.org/10.36680/j.itcon.2022.037.