Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Strategic Patient Involvement in Medical AI Co-Design: A Pilot Study in Heart Failure

Laura Arbelaez Ossa

laura.arbelaez@ub.edu

Artificial Intelligence in Medicine Lab (BCN-AIM), Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain. https://orcid.org/0000-0002-8303-8789

Machteld Boonstra

Department of Cardiology, Amsterdam University Medical Centre, Amsterdam, Netherlands. https://orcid.org/0000-0001-7550-0489

Fleur Meijers

Artificial Intelligence in Medicine Lab (BCN-AIM), Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain.

Emmah Ng'ang'a

Facultat de Medicina i Ciències de la Salut, Universitat de Barcelona (UB), Barcelona, Spain. https://orcid.org/0009-0007-3835-521X

Richard Stephens

European Society of Cardiology Patient Forum, Sophia Antipolis, France https://orcid.org/0000-0002-0760-2771

Axel Verstrael

European Society of Cardiology Patient Forum, Sophia Antipolis, France. https://orcid.org/0000-0002-6559-0393

Xènia Puig Bosch

Artificial Intelligence in Medicine Lab (BCN-AIM), Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain. https://orcid.org/0009-0001-1148-7921

Folkert W. Asselbergs

Department of Cardiology, Amsterdam University Medical Centre, Amsterdam, Netherlands. https://orcid.org/0000-0002-1692-8669

Maria Maixenchs

ISGlobal, Barcelona, Spain. https://orcid.org/0000-0002-5782-8197

Karim Lekadir

Artificial Intelligence in Medicine Lab (BCN-AIM), Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain. https://orcid.org/0000-0002-9456-1612

Research Article

Keywords: co-design, artificial intelligence, cardiovascular AI, medical AI development, clinical AI, digital health, trustworthy AI, PPIE, patient involvement, participatory design

Posted Date: August 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10808123/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

Background

Artificial intelligence (AI) is increasingly promising to support decisions in acute cardiovascular care, including whether patients presenting to the emergency department can be safely discharged. Ethical and trustworthy-AI frameworks call for involving patients in designing such tools, yet patient involvement in AI has largely been limited to gathering views, acceptability, and interface preferences. This consultative role leaves largely untested what patients contribute when involved strategically, shaping what a tool is for and how it works. What patients contribute when they help design a high-stakes cardiovascular risk prediction tool has rarely been examined.

Methods

We conducted a co-design workshop as a pilot in which patients with lived experience of heart failure (n = 11) participated, within a European project developing an AI-enabled cardiovascular risk prediction tool. The study used a two-stage process: preparatory capacity-building through e-learning and an online webinar, followed by a two-day in-person workshop of 5 facilitated sessions. The workshop comprised five facilitated sessions using participatory design activities structured around trustworthy-AI (FUTURE-AI) design-phase principles. Recordings, co-design artefacts, and surveys were analysed using reflexive thematic analysis.

Results

Patients shifted from a familiar consultative role into an active design role. Working in two independent groups, participants converged on similar designs. Both reframed the tool from clinician-facing to patient-facing and specified: whole-person inputs beyond acute clinical variables; qualitative, and actionable risk rather than a numerical output; the withholding of predictions the model cannot make reliably; and layered, patient-controlled explanation. Because the prediction concerned their own risk of death, participants built emotional safeguards into the design.

Conclusion

Patients involved as strategic partners produce coherent, concrete design specifications that a purely technical process would miss. Realising this contribution requires preparing patients for co-design methods, treating the tool's intended beneficiary as an open question, and building routes by which patient input reaches development.

Trial Registration

Not applicable.

Plain English Summary

Artificial intelligence (AI) tools are being built to help doctors decide how to care for people with heart problems, for example, whether someone who comes to the emergency department can safely go home. These tools are usually designed by technical and medical experts, and patients are rarely asked to help shape them. We wanted to know what patients would design if they were given the chance.

We invited 11 people who live with heart failure to a two-day workshop. Before the workshop, they took an online course and attended a webinar on AI. During the workshop, they worked through hands-on activities based on principles for trustworthy AI, ending with a task where they designed their own ideal tool.

At first, patients were used to being asked only to react to something already made. Once they felt free to design, they went much further. Two groups, working separately, came up with almost the same idea: a tool made for patients, not just doctors; one that sees the whole person, not just test results; that shows risk in plain terms with clear next steps rather than a frightening number; that stays silent when it cannot give a reliable answer; and that lets people choose how much they want to know about their results.

Patients also pointed out that one workshop is not enough, and that funding and rules often stop patient ideas from reaching the final tool. Their designs will now feed back into our project.

Background

Patient and Public Involvement and Engagement (PPIE) in health research is not new, and it has a substantial history in service design and program evaluations across health and social care with established principles, standards, and growing evidence on how involvement should be organized and what its contributions are (1–4). Applying PPIE to Artificial Intelligence, however, is a comparatively young field with its own challenges and requirements (5–7). Unlike a service or a resource that patients can see and respond to, the object being shaped, a trained model, has no perceptible surface: what it learns and how it generalizes cannot easily be inspected or "walked through." Many of the decisions that determine what a model does are made continuously by developers, in a codebase rather than in a room (8). And because a model is retrained, updated, and revalidated over time, the ideal timing and level of involvement is hard to fix. In practice, 85% of participatory AI projects engage stakeholders only once during development (9), suggesting that involvement and continuous engagement remain sporadic. PPIE matters for medical AI because many high-stakes challenges are not purely technical but human-centric: what patients need for their care, how they relate to AI when it is involved in their care, where implementation happens, who can benefit from these AI results.

PPIE in AI research has so far focused on interface preferences, communication, and acceptability, with much literature on patients' attitudes toward AI and only a small but growing set of studies involving patients in the design of health AI (5, 10). While many challenges remain for PPIE in AI, ethical and regulatory frameworks widely endorse PPIE for AI research, positioning it as an ethical aspiration and as a practical requirement for building AI tools that are safe, acceptable, and fit for human-centric applications (9, 11–13). For example, the World Health Organization (WHO) guidance on AI for health emphasizes inclusive, human-centred governance and the need to involve affected stakeholders to address risks and uphold ethical principles (14, 15). Similarly, the FUTURE-AI guideline and trustworthy AI principles position stakeholder involvement as a lifecycle task spanning design, validation, deployment, and monitoring (12).

PPIE methods, including co-design, are not a single method but an orientation that can be operationalized in varied ways (9, 16, 17). We use PPIE to mean AI research carried out with or by patients rather than to, about, or for them (3, 7). However, PPIE can operate at very different levels of influence. At one end, patients are consulted on predefined aspects of an existing tool, its usability, interface, or language; at the other, they hold strategic influence over what the tool is for, who it serves, and under what conditions it should act. Delgado et al. describe this range as a spectrum from consulting to ownership, defined by how much decision-making power is transferred to patients (9); we define strategic involvement as the high-decision-making-power end, where patients shape the problem rather than react to a solution. Co-design sits at this end of the spectrum: it positions participants as active partners throughout the design process, contributing to problem framing, acceptable use cases, implementation requirements, and governance decisions (18–20). Rather than inviting patients to provide feedback, co-design fosters collaboration between patients and AI developers, engaging them in decisions about the tool itself (9, 16, 21).

Co-design builds on participatory-design and user-centred traditions and shifts the roles of designers and participants, so those affected by a design contribute creatively to what is made rather than only reacting to it (22). However, enabling strategic involvement through co-design is demanding, and the field is still developing shared frameworks for how it should be conducted and evaluated (22). Other factors may also constrain strategic patient involvement in AI research. For example, tensions may exist between participatory aspirations and practical constraints, such as resource or time limitations (23, 24). Some scholars position AI knowledge as a prerequisite for patient involvement, although limited evidence shows that AI training produces strategically different contributions (5, 6).

Acute heart failure is among the most common causes of emergency hospital admission worldwide, carrying substantial in-hospital mortality and high rates of adverse post-discharge events (25, 26). A central clinical challenge in the emergency department (ED) is determining which patients can be safely discharged and which require admission or close monitoring. Studies comparing these decisions against validated risk scores reveal a mismatch between risk and disposition: many lower-risk patients are admitted, while some higher-risk patients are discharged (27, 28). A potential application for AI, therefore, is AI-enabled risk prediction tools to support these decisions and predict major adverse cardiovascular events at 30 days. However, despite promising performance in research settings, these tools have not yet translated into routine clinical practice, partly because design decisions have been driven by what is technically feasible and clinically measurable rather than by what is meaningful and actionable for the people these tools are meant to help care for.

During discharge decisions, clinicians are recommended to consider and address potential barriers that could jeopardize patients' successful self-care, including health literacy, frailty, impaired cognition, financial resources, transportation, symptom monitoring, and social and caregiver support (28). However, these considerations are seldom reflected in predictive cardiovascular risk AI, where models are often built around an idealized patient, a clinical profile in an acute moment, stripped of life context, care history, information preferences, and agency over their own risk. Without patient involvement, tools may be optimized for clinical accuracy but misaligned with patient needs, preferences, and care realities.

All of this makes co-designing cardiovascular risk prediction tools an important case for PPI, given that an AI tool that predicts an individual's risk of death increases the affective and ethical considerations that lower-stakes tools do not, such as how much a patient wants to know about their risk, whether a tool should decline results, who bears the costs of errors, and what happens when these results are wrong.

Aim:

Against this background, we conducted a co-design workshop, as a pilot, within DataTools4Heart (29), a Horizon Europe project developing an AI-enabled cardiovascular risk prediction tool to support ED discharge and admission decisions. Our aim was to examine what patients with lived experience of heart failure specify, build, and recommend for such an AI when involved as strategic partners in its co-design, and what their contributions, and the limits they themselves identify, reveal for involvement in AI design. We structured the co-design activities around trustworthy-AI (FUTURE-AI) principles (12) as a backbone for the ethical design considerations the tool should address.

Methods

Study design

We conducted a pilot, two-day, in-person co-design workshop combining facilitated group discussions with structured participatory activities common in design thinking (journey mapping, discussion of user personas, role-playing, and

paper prototyping). We selected co-design because it supports developing solutions with affected stakeholders and can surface practical constraints, values, and implementation requirements early (30). We used structured facilitation and multiple participation modes (*e.g.*, plenary discussions, breakout groups, and written content) to support inclusive participation and avoid common participatory group pitfalls (22, 31). We report the involvement in accordance with the GRIPP2 long-form checklist (Annex 1), which is designed for studies whose primary focus is patient and public involvement (2).

Involvement took place in two stages. A preparatory phase familiarised participants with the project and with AI in healthcare through a short self-paced e-learning module (https://www.elementsofai.com/) and an online webinar that focused on AI design and AI design requirements. The in-person workshop then ran over two days (Barcelona, October 2025). At the start of the workshop, the tool aim was briefly introduced: a clinician-facing ED risk-prediction mock up interface shown as a prompt for discussion rather than a fixed design to be rated. In the introduction session, an AI developer met participants and answered questions about the model and data. After that, two facilitators ran the workshop, comprising five co-design sessions plus welcome, recap, and closing segments. Each session operationalized one or more FUTURE-AI design-phase principles into a participatory activity (Table 1).

Table 1
Delivered workshop structure, activities, and FUTURE-AI principles

Session	Aims (as delivered)	Participatory process	FUTURE-AI principle(s) grounding	Outputs and data generated
Preparatory	Familiarise participants with the project and with AI in healthcare; establish a shared starting point.	Self-paced e-learning module + online webinar ahead of the in-person days.	Cross-cutting (stakeholder engagement across the lifecycle).	Baseline + post-learning surveys (confidence, comfort, expectations).
Welcome	Set a safe, open tone; introduce the tool and the ED use case; meet the AI developers and team.	--	--	Field notes.
Session 1	Where does the tool fit? What are patients' general needs across the care pathway; who should use it and why.	Usability (define intended use, user requirements); Universality (intended setting and variation).	Vignette of a patient + journey mapping and facilitated discussion.	Blackboard journey map (co-produced); audio/video recording.
Session 2	Who should the tool serve, who might be excluded, and how to avoid exclusion; what fairness means from a patient view.	Fairness (engage stakeholders to define sources of bias); Universality (heterogeneity across groups).	Persona cards (six characters spanning age, ethnicity, digital access, geography) + a blank participant-authored persona; paired discussion then plenary.	Annotated persona cards; participant-authored persona; audio/video recording.
Session 3	How should the tool fit into care: communication, consent, decision-making power, and the patient–clinician–AI relationship.	Explainability (user understanding/communication); Usability (human oversight); autonomy & consent.	Co-constructed role-play (participants directed the scene and characters), enacted twice from different angles, followed by discussion of the enacted interaction; needs surfaced as post its in the recap.	Priorities post its (participant-written, facilitator-clustered). audio/video recording.

Session	Aims (as delivered)	Participatory process	FUTURE-AI principle(s) grounding	Outputs and data generated
Session 4	What do patients want from the tool; identify must-haves; how feedback should be used and reported.	Usability (requirements); Traceability (feedback and reporting).	Facilitated recap of Day 1 + prioritisation and clustering of features and needs.	Black-board requirements and AI features (facilitator-clustered); audio/video recording.
Session 5	"Creating my perfect AI tool": patients as designers specify what the tool should be, do, and show.	Cross-cutting / open design (participants free to integrate or exceed the principles above).	Open group prototyping task in two independent groups, group presentations.	Participant-authored design artefacts (the "eMe" concept; low/med/high risk display; layered basic/intermediate/expert explanation); audio/video recording.
Wrap	Groups present their designs; whole-group synthesis; reflect on process; brief on the follow-up survey.	Traceability (how input will be used and feedback).	Plenary synthesis on a shared board; closing reflection.	Post-its on how you would like to be involved (facilitator clustered), audio/video recording.

Participants

Participants were recruited using purposive sampling to assemble a group with lived experience of heart failure, while aiming for diversity in background and perspectives. Recruitment occurred through open calls on social media and targeted outreach via European Society of Cardiology (ESC) patient networks. Interested individuals were screened by a study co-authors, to confirm suitability for an intensive, facilitated, in-person co-design process. During screening, 29 individuals expressed interest via email; 14 were enrolled and completed the baseline questionnaire. Later, 3 withdrew due to language or scheduling constraints, leaving 11 participants who completed the e-learning module, attended the webinar, and participated in the in-person workshop. Eligibility criteria included: self-reported lived experience of cardiovascular disease; ability to participate in person; English language proficiency; and willingness to engage in the iterative multi-stage format. The exclusion criteria were: (i) severe health conditions that were incompatible with travel, and (ii) limited cognitive capacity to engage and provide informed consent. Survey responses are numbered R1–R14 (the enrolled group) and workshop discussants PT1–PT11 (those attending in person).

Data Collection

Workshop sessions were captured using the university's audiovisual seminar room, which recorded full-room audio and video from two cameras at different angles; two additional audio recorders were placed centrally as a safeguard against technical failure (the room capture proved sufficient and the backups were not needed). Screen and session capture used Open Broadcaster Software (OBS), recording locally on the main author's laptop. We transcribed recordings locally with Whisper, and the main author checked transcripts against the recordings. Alongside the recordings, the workshop produced a set of physical artefacts, which we treated as image data and interpreted against the recordings of participants presenting and discussing them: a blackboard journey map (Session 1); annotated persona cards and one participant-authored persona (Session 2); participant-written post its from the Session 3 recap, co-produced blackboard synthesis prioritisation (Session 4) and participant-authored design mock-

ups (Session 5). Brief surveys were administered at four timepoints (pre-learning, post e-learning, post-webinar, post-workshop).

Data Analysis

Artifacts produced during the workshop varied in authorship: The Session 5 mock-ups and the participant-authored persona were produced by participants directly. The post its were written by participants but clustered by the lead facilitator during the workshop. The blackboard journey map and the closing blackboard synthesis were co-produced: the lead facilitator recorded and reflected participants' contributions in real time as the plenary directed. This in-room, facilitated synthesis functioned as an interpretive act during data generation, was discussed with all team members present, and informed the formal analysis conducted afterward.

Analysis proceeded in three linked stages. First, the workshop was designed deductively: activities were structured around FUTURE-AI design-phase principles, so the data-generation was framed by pre-specified topics. Second, in-room facilitated synthesis (the clustering and blackboard work above) organised material collaboratively during the sessions. Third, we analysed the full corpus (transcripts, recordings of group presentations, artefacts-as-images, field notes, and open-ended survey responses) using reflexive thematic analysis [(32, 33), moving iteratively through familiarisation, coding, and theme development and refinement (using software: ATLAS.ti). Coding was led by the first author; codes and candidate themes were reviewed iteratively with co-authors, and developed themes were reviewed and critically discussed with the patient co-authors. Because activities were principle-structured, the analysis remained deliberately open to specifications that fell within, across, or beyond the principles the activities were designed to elicit; where interpretive disagreements arose, they were resolved through team discussion. Final themes were reviewed by the full author group.

Survey Data

Survey data were analysed qualitatively and descriptively and used to complement workshop qualitative interpretation (34). Survey items (including 1–5 Likert items on confidence and comfort, and open-ended questions) were used descriptively and contextually to complement the qualitative analysis. Given the small, attrition sample, we do not report inferential statistics or treat the Likert items as quantitative outcomes; we summarise them narratively and draw on the open-ended responses thematically.

Team Positionality

Our team spanned several disciplines. The first author is a researcher in AI ethics and social scientist for healthcare, with experience in patient engagement and qualitative methods; EN, FM, and MM are social science researchers with expertise in qualitative and participatory methods; MB is a clinician (technical doctor) on the project's clinical side and co-facilitated the workshop and FA is a senior clinician in cardiology. AV and RS are patient representatives from the ESC Patient Forum; and KL is an AI-for-healthcare researcher and focus on trustworthy AI designs and development. The AI-development team met participants and answered questions but did not facilitate sessions or conduct the analysis. We decided to invite developers to meet participants because, for involvement to reach the build, the people who build the tool must be reachable. Because the team included the tool's developers, we took steps to keep the co-design workshop open: we stated that the tool's direction was not fixed, actively invited critical as well as supportive input, and restated the workshop's purpose at several points. The co-design outputs will inform a revised version of the tool, which we plan to bring back to patients, and a possible follow-on project.

This study was itself a PPIE activity: patients with heart failure were active co-designers rather than research subjects. Patient co-authors (AV, RS) from the ESC Patient Forum contributed to the workshop design and to reviewing all participant-facing materials, attended the workshop in an observer and support capacity (not facilitating, as their prior relationships with some participants could have introduced bias), contributed to the in-room synthesis and to interpretive discussion of the findings, and reviewed and critically discussed the developed themes as co-authors. The first author led formal coding. The study design was informed by prior PPIE within DataTools4Heart.

Results

The final workshop group (n = 11) included 5 women and 6 men representing 7 European countries. All reported a current diagnosis of heart failure with heterogeneous underlying conditions. Most had prior experience with patient engagement, research participation, or health advocacy. Motivations for participation centred on improving heart failure care and quality of life, contributing to science, and ensuring AI-enabled tools reflect patient priorities. All respondents reported comfort communicating in English and using a computer for online sessions. Participant characteristics are summarised in Table 2.

We organise the findings into three themes. Theme 1 describes patients' shift to a design role: what they brought that enabled it, and the challenges the shift entailed. Once in that role, participants redesigned the tool and, working in two independent groups, converged on strikingly similar designs. We organize that redesign into what the tool should be (Theme 2) and how it should behave (Theme 3). Throughout, participants repeatedly raised trustworthy-AI concerns (e.g., fairness, reliability, and explainability) in all sessions.

Table 2
Sample characteristics (n = 11)

Characteristics	n (%)
Sex	
Women	5 (45%)
Men	6 (55%)
Age group (years)*	
18–29	1 (9%)
30–39	2 (18%)
40–49	3 (27%)
50–59	2 (18%)
≥60	3 (27%)
Country of residence	
The Netherlands	3 (27%)
The United Kingdom	2 (18%)
Ireland	1 (9%)
Germany	1 (9%)
Poland	1 (9%)
Belgium	2 (18%)
Spain (Catalonia)	1 (9%)
Educational level	
University / higher education	10 (91%)
Vocational training	1(9%)
Heart failure–related hospitalisation history	
None	3 (27%)
Once	3 (27%)
Multiple times	4 (36%)
Not reported	1 (9%)
Prior involvement in research / panels / patient advocacy	
Yes	9 (82%)
No / not yet	2 (18%)

1. Shifting to a Design Role

Patients arrived already feeling equipped to contribute, with consistently high confidence across all survey timepoints. The real barrier for patients was not capacity but familiarity. Most participants had only prior experience with reactive, consultative participation and initially struggled to engage with an open design space. Once reassured that open-ended contributions were genuinely wanted, that shift happened.

1a. Confidence to participate came from lived experiences

Across all four survey timepoints, participants reported consistently high “confidence to share ideas" and "comfort speaking up if an AI tool did not make sense" even before engaging with the preparatory materials.

The e-learning module received the widest spread of ratings, with some finding it valuable and others describing it as poorly matched to their needs. Several questioned the relevance of mathematical AI content to the participatory work they were being asked to do. When asked about their potential contribution, participants saw themselves contributing to what counts as safe and effective outcomes (9/11), how the tool should be used in practice (8/11), and how results should be explained (8/11). One participant, after struggling with the mathematical content, articulated clearly where their value lay:

" I think my added value mainly lies in ensuring from the patient's perspective that all factors are taken into account, that the right information is included, and that AI is fed with the right input. Also, making sure that everything is complete and considered in agreements. But I prefer to leave the technical part to the professionals who have expertise in that area " (R7, post-workshop survey)

1b. Co-design required a new involvement mode

Arriving as capable contributors did not mean participants felt comfortable designing the AI. In early sessions, uncertainty emerged about what kind of contribution was appropriate. When unsure how to proceed, participants reverted to familiar consultative patterns, requesting prototypes to evaluate their interfaces and seeking the framing they were accustomed to receiving:

"In the beginning we had to get used to the idea that we could go any direction and nothing was set upfront." (R6, post-workshop survey)

"I was wondering if you share an idea already or a prototype where we can pretend to react to because it's very, um, a big idea so far from different perspectives maybe if there is a map already we can shoot on that." (PT9, workshop)

Once facilitators clarified that all directions were welcome and that there were no wrong contributions, participants shifted, producing substantive and innovative design proposals.

1c. Barriers to AI co-design

As the workshop concluded, participants reflected on what would happen to their contributions after they left, identifying structural conditions such as regulatory, legal, funding, cultural, and sustainability, that work against patient involvement:

"The project is really ambitious. In my opinion, maybe too much. But I hope and I really wish that this ambitious becomes into reality and all of us can enjoy the tool in a non far future. When I say too much ambitious I think in the big troubles of

data protection laws, the relationships between countries and also the complexity to erase false data to avoid mistakes in the predictions." (R8, post-workshop survey)

"My main comment is about the development: it seems like creating such AI tools will require significant funding. I'm curious whether there's a business model behind this, or if the tools will be freely available to patients, yes or no, and how they will be maintained over time." (R3, post-workshop survey)

Several participants explicitly noted that a single workshop is not enough for meaningful co-design and one drew a broader conclusion:

"It also clarified that patient-involvement and co-design needs a lot more academic work in order to come up with rigorous templates and frameworks if we aim to standardize this sort of engagement." (R1, post-workshop survey)

Redesigning the tool: two groups with similar designs

Once in a design role, participants did not evaluate the tool as presented; they redesigned it. In the final task, two groups worked independently, and both arrived at similar tools, a convergence participants noticed themselves:

“We started off with the possibility to use [it] for doctors at the emergency room, and I’m happy to see … both groups came up with a patient app, because I believe the patient should be the center of it all … the patient is the one who has to give the input and also has to live with the condition.” (PT9, workshop)

Both groups independently designed a patient-facing mobile application that drew on whole-person inputs rather than acute clinical variables alone; expressed risk in qualitative, banded terms rather than as a raw number; paired risk with an actionable recommendation; and built in consent and data governance. They also converged on layered, patient-controlled explanation and on a shared patient–clinician data flow. Figure 1 shows the two designs. We present the shared design below along two themes: what the tool should be (Theme 2) and how it should behave (Theme 3).

2. Re-designing the Tool to Become Patient-centric

Once comfortable with the open design space, participants strategically questioned the tool’s purpose, scope, and intended user, producing design proposals that departed significantly from the original concept.

2a. Reframing who the tool is for

The workshop started from a mock-up of an ED cardiovascular mortality risk prediction tool designed to support clinician decisions about discharge or admission. Participants engaged with this concept but quickly identified a fundamental problem: as patients, they would never see it, and it would never reach patients. They challenged the narrowness of the original decision focus, arguing the tool could serve a much broader range of decisions and users or be integrated in a more consistent way into their care:

" …it would be much more valuable if it could be used in a broader spectrum of cases. Because if you present an ED, you've got an acute problem. You know, you've either got a heart attack, cardiac arrest, stroke, you know, if you're in the emergency department, the decision-making is black and white … [about keeping somebody in or discharge them]… It's a very, narrow spectrum of decision. But this tool, I think, would be much more valuable if it was, you know, and there's so many inputs that, you know, it just seems that it could be used in a wider spectrum of decision-making. (PT10, workshop)

Not all participants agreed. One explicitly preferred the original clinician-facing design, arguing that a patient-facing application would generate anxiety rather than empowerment. However, in the final designs, both groups independently produced patient-facing applications.

"Inputting data myself and looking at an app, this is Instagram on steroids. This is anxiety waiting to happen." (PT8, workshop)

2b. Widening the inputs to the whole person

Participants argued that the model should draw on social, contextual, and lifestyle factors such as living situation, culture and diet, religious practice and fasting, activity, as conditions on how risk is interpreted:

“Religious practices and belief … if you’re in a religion that doesn’t believe in any invasive medicine … or if you … do fasting at certain times … that could … [be] showing up on you.” (PT2, workshop)

They also raise questions about what the predictions mean for them, for their bodies and their experiences, and whether data can represent them accurately:

“A 62 year old woman like me, but who hasn’t got a heart issue, might have a risk factor of 40 anyway … you need to know so that you’re not comparing me with some fit 23 year old.” (PT2, workshop)

The demand beneath these inputs was to be treated as a whole person rather than a data point:

“We’re looking to be seen as a person … we want to be heard holistically.” (PT2, workshop)

"The most important thing is what the person, not patient, person, wants out of this interaction. Everyone wants something different." (PT10, workshop)

2c. Integrating AI into the shared patient-clinician space

These choices came together in their designs’ architecture: a tool centred on the patient, into which the patient enters personal, lifestyle, and symptom information and to which clinicians add clinical data, shared both ways with consent. One group’s “eMe” captured this integration:

“We call this application eMe … the electronic version of myself … you have to tick the button consent of sharing this data … this information then goes to the … GP … and then the GP can add medical data … and … share it … back to the individual as well as to the hospital.” (PT11, workshop)

Thinking across the care pathway, participants argued the tool should not exist in isolation but connect general practice, the ED, and patient-facing interfaces. One participant, drawing on ten years of diagnostic delay, made this case directly:

"It took me ten years to get diagnosed with all the classic symptoms. That's a huge red flag for me, there's a real gap in that initial diagnostic thing. That would be my priority." (PT2, workshop)

3. Adapting To Patient-centric Outputs

Participants described in detail how AI outputs should be made usable for them: when and how results should be given, when to withhold them, and how much to explain.

3a. Qualitative and actionable outputs

Several rejected a numerical risk score as the primary output on both comprehension and emotional grounds; the affective stakes of mortality were explicit:

“It’s shocking, to be honest … if I was just … an average patient … I would be shocked … information like that is better human to human.” (PT10, workshop)

They preferred banded categories, and wanted the output to say what to do and to prepare them for the clinical encounter:

“Could you possibly do it like … the obesity scale where … you’re in this upper percentile, or … the middle risk percentile?” (PT2, workshop)

“If the results can be translated into some questions as a preparation for your consultation.” (PT1, workshop)

Participants stressed that outputs should reflect different values and life priorities, not just individual clinical risk, for example:

"Some people just want to live as long as possible, whereas other people want to be able to do a lot more active things within ten years." (PT2, workshop)

3b. Withholding predictive outputs

Participants suggested the tool decline a prediction when it could not do so reliably. They expressed that the cost of error is borne by the patient, and it is better to opt out of predicting uncertain predictions:

“If … you are in an underrepresented cohort and the tool cannot give you a decent prediction score … it should also not give you one.” (PT8, workshop)

“I would accept an AI tool which would say, sorry, I can’t give you the risk right now because I’m still learning. I have not enough data.” (PT11, workshop)

and, on the lack of usability of a bare figure:

“Your 10-day mortality risk is only 6.4, very good.” (PT11, workshop, ironical tone and gestures)

3b. Explanations as a layered and patient-controlled output

Participants design explanation to be available in graded depth, chosen by the patient and adjustable to their situation and energy:

“Having that option of knowing about the toolkit … it’s patient’s choice, but let’s at least give the patient an option.” (PT7, Session 3)

“In acute heart failure situation, I don’t challenge much. I don’t have the energy for it.” (PT8, Session 3)

Both groups implemented this in their designs: a language-selectable, patient-facing interface; a low/medium/high output in place of a number; and tiered “basic / intermediate / expert” explanations with more detail and more technical language as the patient chose to go deeper (Fig. 2).

Discussion

Summary of Findings

In this pilot, co-designing an AI tool with patients produced strategic input that fundamentally reframed its purpose, scope, and intended users, proposals that would not have emerged from evaluating a predefined design for usability or its interface. Two independent groups converged on the same design: a patient-facing tool built on whole-person inputs, communicating banded and actionable risk, gated by reliability, and explained in patient-controlled depth. Operationalizing trustworthy AI principles as co-design activities helped surface ethical requirements and safeguards that could inform more patient-centred tool development.

The tool’s affective specificity shaped what patients designed. Because the prediction concerns the patient’s own risk of death, participants built emotional safeguards into the tool such as refusing raw numbers, banding risk, insisting the tool withhold rather than mislead, and controlling how much they wished to know. This suggests that the stakes of an AI tool, not only its subject matter, may indicate the need for patient co-design even if they are not the expected end-users.

Participants also identified that the conditions surrounding AI development systematically work against co-design input reaching decisions that matter. Regulatory timelines, funding ceilings, and the absence of feedback loops between workshops and development pipelines mean that involvement remains a challenge – even within our own project.

Relation to Previous Research

Previous research on co-design in healthcare AI suggests that actively involving patients can yield more patient-centred and sustainable tools (35). Thompson et al., studying PPIE within a large AI and big data programme, found that patient contributors consistently prioritised person-centred outcomes, while researchers remained constrained by data-driven metrics (5). We observed this tension as patients pushed the design beyond clinician-facing, data-driven scope. Classic participatory design scholars distinguish between *having a voice* and *having a say*, namely the difference between being consulted and having real influence over decisions and outcomes (36). For AI specifically, Delgado et al. formalize this distinction across a spectrum from consulting to ownership, making explicit that PPIE varies in the degree of power transferred to stakeholders to own the decisions that shape a tool (9). We use strategic involvement to denote this higher-influence, that of patients shaping what the tool is for and how it works, not only reacting to a predefined design. Our findings suggest patients can contribute at this strategic level: they reframed the tool's purpose and users rather than refining its interface. It is, therefore, important that co-design for AI focus on transferring this design power and enabling strategic involvement, although this is neither easy nor yet operationally standarized.

Our participants’ insistence that the tool withhold predictions it could not make reliably, which is what Verstrael and colleagues describe as epistemic honesty from the patient’s perspective: patients did not want false certainty but wanted the tool to disclose, or withhold, in proportion to what it could reliably know, treating the acknowledgment of uncertainty as a condition of trust rather than a failure (37). That patients reached this position from lived experiences, without technical prompting, suggests trustworthy-AI principles can be expressed in patients’ own terms.

Our findings diverge from previous research on AI literacy. Thompson et al., found that a more heterogeneous sample benefited from introductory AI workshops (5), whereas in our study, confidence was high before any preparatory content and remained stable throughout. This matters because AI design and development frameworks identify patients' digital literacy as a key human factor to be addressed and recommend structured training activities as foundation for meaningful participation (12, 6). Frost et al.'s scoping review of empirical methods in PPIE for AI reveals that this assumption is common: most studies provided background information about AI to participants before involvement, yet many still reported participants' lack of prior knowledge as a limitation (6). They argue many researchers treat patient unfamiliarity with AI as a problem to be corrected before genuine participation, which is unclear in our case, as patients felt confident in their experience even before training.

A recent critical reflection of PPIE in heart failure research similarly documented the tension between scoping a project based on clinical evidence and patients' perspectives, which may prioritise different outcomes such as feasibility within daily life or acceptability in the context of multimorbidity(38). This resonates with our finding that patient involvement needs to start early enough for the scope and purpose of a project to remain genuinely open to revision, rather than asking patients to accept or reject a problem already defined without them. Kirk et al., reflecting on a co-design intervention, similarly found that involvement can become 'the good thing itself'—difficult to critique—and that ceding genuine control to participants was harder in practice than in principle, as the co-design process itself became something the team struggled to hold open (31). Alongside our findings, this suggests the barrier to co-design may not be on participants' capacity but the conditions surrounding the process.

Our findings point to a tension between the aspiration for co-design and the conditions under which AI development is typically resourced. During our workshop, familiarisation with co-design happen relatively quickly when it was structured and guided, but this level of facilitated, multi-day engagement is rarely routinely funded. Additionally, it is possible that what is needed is co-design training for researchers and participants is integrated more fundamentally in AI, though this too depends on funding and timelines. For example, Foster et al.'s analysis of a UK-based PPIE funding scheme found that even when dedicated funding was available to support patient involvement from the earliest stages of grant development, a £600 ceiling and pressure to spend within three months meant that most researchers could consult contributors only once or twice (23). Smith et al., similarly acknowledge that the lack of clarity around how to operationalize and budget for PPIE across the full development cycle remains one of the most consistently cited barriers in literature (39).

Taken together, these findings point to a structural problem: ethical AI guidelines that endorse stakeholder engagement rarely specify the time, budget, or organizational architecture needed to make PPIE meaningful. A recurring assumption in the design of AI clinical tools is that the intended beneficiary and the intended user are the same person. In cardiovascular AI, clinicians are routinely positioned as the end users whose experience matters for design, while patients may be classified as passive stakeholders who are affected by the system but without power to influence its development (40). Where patients have been engaged, the dominant approach has been to build AI literacy and trust rather than involving them in model scope or problem definition (35). The result, as Wang et al.'s systematic review of human-centred design of AI clinical decision support systems found, is that patient-facing perspectives are largely absent from the literature, with the field's attention concentrated on clinician usability and workflow (41). This reflects a broader pattern in healthcare AI where patients are positioned as lacking the capabilities to engage with technical AI decisions, creating conditions in which patients may internalize a passive role and expect to be consulted on something predefined rather than invited to shape what the problem is (42).

For co-design in medical AI, this suggests moving beyond engaging patients only as end users and considering who counts as a stakeholder more broadly, including those affected by a tool but who will never use it directly. Making the

intended beneficiary of an AI tool an explicit, open question may be one of the easiest changes AI co-design practices can make.

Limitations

Several limitations apply as a pilot study. The sample was small, experienced, and advocacy-engaged: these were people well placed to work in an open design space, and the study asked what such participants produce under these conditions, limiting generalisation to less-empowered or less-experienced participants. Second, in-person participation in Barcelona required English proficiency and the ability to travel, constraining who could participate. Third, we have not yet incorporated the co-design outputs fully into subsequent AI development, so we cannot report how these findings functioned in the final tool; implementation, a return of the revised tool to patients, and a follow-on project are the planned next phase. Finally, this was a single cardiovascular mortality-prediction tool; the approach will need to be tried with other participants and other tools, under similar conditions, to see how co-design changes the scope and intentions of such tools.

Conclusions

Structuring patient co-design around trustworthy-AI principles let patients with heart failure engage with those ethical principles and, from them, design a different tool: patient-facing, whole-person, actionable, reliability-gated, and explained on their own terms. Our findings suggest two practical considerations for enabling strategic patient involvement in designing high-stakes AI tools such as cardiovascular mortality predictors.

First, co-design preparation should also familiarise participants with design practices, such as typical design-thinking exercises, prototyping, and alike. Our findings showed the barrier was unfamiliarity with what involvement looks like for AI design. Preparatory components may need to legitimize experiential knowledge, orient participants to the difference between evaluating a predefined tool and shaping a problem, and provide facilitated space for that transition, including time for familiarisation; a dedicated co-design module could help.

Second, the intended beneficiary of the AI tool should be asked as an explicit and open question in co-design when the aim is to design ethical AI. In our workshop, participants independently identified that the tool was clinician-facing and would never be seen by patients, which restructured the entire design conversation. If developers make explicit whose problem the tool solves, for whom its outputs are intended, and whether those are fixed or open questions, patients can interrogate those assumptions from the start.

Our study shows both what such involvement can produce and where its influence meets structural limits. Both are necessary knowledge for anyone seeking to involve patients in building high-stakes medical AI.

Abbreviations

AI
artificial intelligence
ED
emergency department
ESC
European Society of Cardiology
FUTURE-AI
international consensus guideline for trustworthy and deployable AI in healthcare

GRIPP2
Guidance for Reporting Involvement of Patients and the Public (version 2)
PPIE
patient and public involvement (and engagement).

Declarations

Ethics approval and consent to participate

The study received ethical approval from the University of Barcelona (Institutional Review Board IRB00003099, CER072527). All participants provided written informed consent, reaffirmed at the start of the workshop. Involvement was voluntary; participants could withdraw at any time without consequence. Participants received travel reimbursement as compensation for their participation.

Consent for publication

Data were handled confidentially and analysed in de-identified form. Quotes are attributed to respondent numbers (R1–R14 for survey responses; PT1–PT11 for workshop discussions). Participants consented to the use of anonymized co-design quotes and artifacts they produced; the design artifacts reproduced in Fig. 1 contain no identifying information.

Funding

This work received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement No. 101057849 (DataTools4Heart project).

Authors' contributions

LAO conceived and designed the study, facilitated the workshop, led data collection and analysis, and drafted the manuscript. MB co-facilitated the workshop and contributed clinical and design input. FM and EN contributed to workshop organization and qualitative analysis. MM contributed to qualitative methodology expertise and data analysis. AV and RS, as patient co-authors from the ESC Patient Forum, contributed to workshop design, participant-facing materials, in-room synthesis, and interpretation, and reviewed the developed themes. XPB contributed to workshop preparation and manuscript preparation. FWA and KL contributed clinical and AI-development expertise and critically revised the manuscript. All authors read and approved the final manuscript.

Acknowledgements

We thank the patients who generously shared their time and expertise. We acknowledge the support of the ESC Patient Forum and Inga Drossart for recruitment and their contributions to workshop planning.

AI assistance (Claude, Anthropic, Opus 4.8, medium effort) supported the reorganization and editing of manuscript sections. All analysis, interpretation, and substantive content are the authors' own.

Availability of data and materials

The datasets analyzed in this study are not publicly available due to the risk of participant identification but are available from the corresponding author on reasonable request.

Competing interests

Richard Stephens is a member of the editorial board of Research Involvement and Engagement. The remaining authors declare that they have no competing interests.

References

1. Staniszewska S, Brett J, Mockford C, Barber R (2011) The GRIPP checklist: Strengthening the quality of patient and public involvement reporting in research. Int J Technol Assess Health Care 27(4):391–399. 10.1017/S0266462311000481
2. Staniszewska S, Brett J, Simera I, Seers K, Mockford C, Goodlad S et al (2017) GRIPP2 reporting checklists: tools to improve reporting of patient and public involvement in research [Internet]. 2. 10.1136/bmj.j3453
3. UK Standards for Public Involvement [Internet] [cited 2026 Aug 24]. Available from: https://sites.google.com/nihr.ac.uk/pi-standards/home
4. Brett J, Staniszewska S, Mockford C, Herron-Marx S, Hughes J, Tysall C et al (2014) Mapping the impact of patient and public involvement on health and social care research: a systematic review. Health Expect Int J Public Particip Health Care Health Policy 17(5):637–650. 10.1111/j.1369-7625.2012.00795.x PubMed PMID: 22809132; PubMed Central PMCID: PMC5060910
5. Thompson A, Bartle V, Remfry EA, Reynolds DJ, Barnes MR, Reynolds NJ et al (2025) Public and Patient Involvement in Artificial Intelligence and Big Data Healthcare Research: An Exploration of Issues and Challenges Within the AI-Multiply Project. Health Expect 28(6):e70490. 10.1111/hex.70490
6. Frost EK, Bosward R, Aquino YSJ, Braunack-Mayer A, Carter SM (2024) Facilitating public involvement in research about healthcare AI: A scoping review of empirical methods. Int J Med Inf 186:105417. 10.1016/j.ijmedinf.2024.105417
7. Rubio AP, Dunn J, Borthwick J, Saunders S, Tochel C, Anderson H et al (2026) Beyond the dataset: integrating public voices in data science. Res Involv Engagem 12(1):2. 10.1186/s40900-025-00752-w
8. Reynolds DJ, Mountain S, Bartle V, Remfry E, Barnes MR, Reynolds NJ et al (2025) Targeting everyday decision makers in research: early career researcher and patient and public involvement and engagement collaboration in an AI-in-healthcare project. Res Involv Engagem 11(1):100. 10.1186/s40900-025-00753-9 PubMed PMID: 40830991; PubMed Central PMCID: PMC12366416
9. Delgado F, Yang S, Madaio M, Yang Q The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In: Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization [Internet]. New York, NY, USA: Association for Computing Machinery; 2023 [cited 2026 Feb 6]. pp. 1–23. (EAAMO '23). Available from: https://doi.org/10.1145/3617694.3623261 doi:10.1145/3617694.3623261
10. Hui CY, Shenton AV, Martin C, Weatherill D, Moylan D, Hayes M et al (2024) Patient and public involvement workshop to shape artificial intelligence-supported connected asthma self-management research. PLOS Digit Health 3(5):e0000521. 10.1371/journal.pdig.0000521
11. Siala H, Wang Y (2022) SHIFTing artificial intelligence to be responsible in healthcare: A systematic review. Soc Sci Med 296:114782. 10.1016/j.socscimed.2022.114782

12. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP et al (2025) FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare [Internet]. 5. 10.1136/bmj-2024-081554
13. Ratti E, Morrison M, Jakab I (2025) Ethical and social considerations of applying artificial intelligence in healthcare-a two-pronged scoping review. BMC Med Ethics 26(1):68. 10.1186/s12910-025-01198-1 PubMed PMID: 40420080; PubMed Central PMCID: PMC12107984
14. WHO, ETHICS AND GOVERNANCE OF ARTIFICIAL INTELLIGENCE FOR, HEALTH. WHO GUIDANCE [Internet] (2021). Available from: https://iris.who.int/server/api/core/bitstreams/f780d926-4ae3-42ce-a6d6-e898a5562621/content
15. WHO (2024) Ethics and governance of artificial intelligence for health Guidance on large multi-modal model
16. Messiha K, Chinapaw MJM, Ket HCFF, An Q, Anand-Kumar V, Longworth GR et al (2023) Systematic Review of Contemporary Theories Used for Co-creation, Co-design and Co-production in Public Health. J Public Health 45(3):723–737. 10.1093/pubmed/fdad046
17. Yu J (2025) Participatory Design revisited: framings, key features, and its boundary with co-design. CoDesign 0(0):1–30. 10.1080/15710882.2025.2573020
18. Papoutsi C, Wherton J, Shaw S, Morrison C, Greenhalgh T (2021) Putting the social back into sociotechnical: Case studies of co-design in digital health. J Am Med Inf Assoc 28(2):284–293. 10.1093/jamia/ocaa197
19. Birhane A, Isaac W, Prabhakaran V, Diaz M, Elish MC, Gabriel I et al Power to the People? Opportunities and Challenges for Participatory AI. In: Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization [Internet]. New York, NY, USA: Association for Computing Machinery; 2022 [cited 2025 Dec 3]. pp. 1–8. (EAAMO '22). Available from: https://dl.acm.org/doi/10.1145/3551624.3555290 doi:10.1145/3551624.3555290
20. Bødker S, Kyng M (2018) Participatory Design that Matters—Facing the Big Issues. ACM Trans Comput-Hum Interact 25(1):4:1–431. 10.1145/3152421
21. Vargas C, Whelan J, Brimblecombe J, Allender S (2022) Co-creation, co-design, co-production for public health – a perspective on definitions and distinctions. Public Health Res Pract 32(2):e3222211. 10.17061/phrp3222211
22. Antonini M (2021) An Overview of Co-Design: Advantages, Challenges and Perspectives of Users' Involvement in the Design Process. J Des Think 2(1):45–60. 10.22059/jdt.2020.272513.1018
23. Foster A, Caunt S, Schofield H, Glerum–Brooks K, Begum S, Gleeson P et al (2024) Evaluating a grant development public involvement funding scheme: a qualitative document analysis. Res Involv Engagem 10(1):57. 10.1186/s40900-024-00588-w
24. Syversen A, Umney O, Howell L, Breen J, Briggs E, Hancox Z et al (2025) How can we involve Patients? - Students' perspectives on embedding PPIE into a doctoral training centre for AI in medical diagnosis and care. Res Involv Engagem 11(1):77. 10.1186/s40900-025-00750-y
25. Niazai A, Jamil H, Hameed M, Sheikh S, Nisar MR (2025) Artificial intelligence in cardiovascular diagnostics: a systematic review and descriptive analysis of clinical applications and diagnostic performance. BMC Cardiovasc Disord 25:849. 10.1186/s12872-025-05327-x PubMed PMID: 41315941; PubMed Central PMCID: PMC12661670
26. Razzaq M, Aslam M, Aqeel MB, Malik M, Haroon R, Kalsoom U (2024) The Role of Artificial Intelligence in Predicting Cardiovascular Diseases: Current Trends and Future Directions. Azerbaijan Pharm Pharmacother J 23:1–19
27. Miró Ò, Rossello X, Gil V, Martín-Sánchez FJ, Llorens P, Herrero-Puente P et al (2019) Analysis of How Emergency Physicians' Decisions to Hospitalize or Discharge Patients With Acute Heart Failure Match the Clinical Risk Categories of the MEESSI-AHF Scale. Ann Emerg Med 74(2):204–215. 10.1016/j.annemergmed.2019.03.010

28. Fountoulaki K, Ventoulis I, Drokou A, Georgarakou K, Parissis J, Polyzogopoulou E (2022) Emergency department risk assessment and disposition of acute heart failure patients: existing evidence and ongoing challenges. Heart Fail Rev 1–13. 10.1007/s10741-022-10272-4 PubMed PMID: 36123519; PubMed Central PMCID: PMC9485013
29. DataTools4Heart [Internet] (2026) [cited 2026 Mar 9]. DataTools4Heart. Available from: https://www.datatools4heart.eu/
30. Bate P, Robert G (2006) Experience-based design: from redesigning the system around the patient to co-designing services with the patient. Qual Saf Health Care 15(5):307–310. 10.1136/qshc.2005.016527 PubMed PMID: 17074863; PubMed Central PMCID: PMC2565809
31. Kirk J, Bandholm T, Andersen O, Husted RS, Tjørnhøj-Thomsen T, Nilsen P et al (2021) Challenges in co-designing an intervention to increase mobility in older patients: a qualitative study. J Health Organ Manag 35(9):140–162. 10.1108/JHOM-02-2020-0049 PubMed PMID: 33960175; PubMed Central PMCID: PMC9251644
32. Braun V, Clarke V (2014) What can thematic analysis offer health and wellbeing researchers? Int J Qual Stud Health Well-Being 9(1):26152. 10.3402/qhw.v926152 PubMed PMID: 25326092
33. Braun V, Clarke V (2023) Toward good practice in thematic analysis: Avoiding common problems and be(com)ing a knowing researcher. Int J Transgender Health 24(1):1–6. 10.1080/26895269.2022.2129597
34. Braun V, Clarke V, Boulton E, Davey L, McEvoy C (2021) The online survey as a qualitative research tool. Int J Soc Res Methodol 24(6):641–654. 10.1080/13645579.2020.1805550
35. Banerjee S, Alsop P, Jones L, Cardinal RN (2022) Patient and public involvement to build trust in artificial intelligence: A framework, tools, and case studies. Patterns 3(6):100506. 10.1016/j.patter.2022.100506 PubMed PMID: 35755870; PubMed Central PMCID: PMC9214308
36. Kensing F, Greenbaum J (2013) Heritage: having a say. In: Simonsen J, Robertson T (eds) Routledge international handbook of participatory design. Routledge, pp 21–36
37. Verstrael A, Camaradou J, Scheenaerts B (2026) Trustworthiness of artificial intelligence from a patient perspective. Eur Heart J 47(30):4003–4004. 10.1093/eurheartj/ehag303
38. Forsyth F, Hartley P, Mant J, Rowbotham S, Sharpley J, Wood A et al (2026 Apr) A critical reflective analysis of patient and public involvement in a programme of heart failure with preserved ejection fraction research. Res Involv Engagem 17. 10.1186/s40900-026-00867-8
39. Smith L, Keast J, Flowers J, Dambha-Miller H (2025) A novel framework for operationalising patient and public involvement: lessons from designing an AI-informed exercise prescription grant. Res Involv Engagem 11(1):130. 10.1186/s40900-025-00801-4 PubMed PMID: 41163113; PubMed Central PMCID: PMC12574169
40. Miller GJ (2022) Stakeholder roles in artificial intelligence projects. Proj Leadersh Soc 3:100068. 10.1016/j.plas.2022.100068
41. Wang L, Zhang Z, Wang D, Cao W, Zhou X, Zhang P et al (2023) Human-centered design and evaluation of AI-empowered clinical decision support systems: a systematic review. Front Comput Sci 5. 10.3389/fcomp.2023.1187299
42. Ossa LA, Rost M, Bont N, Lorenzini G, Shaw D, Elger BS (2025) Exploring Patient Participation in AI-Supported Health Care: Qualitative Study. JMIR AI 4(1):e50781. 10.2196/50781

Footnotes

1. Robustness was not delivered as a discrete session. Reliability concerns such as that the tool should withhold a prediction when it cannot make one reliably and should communicate uncertainty arose across Sessions 1–5.

Figures

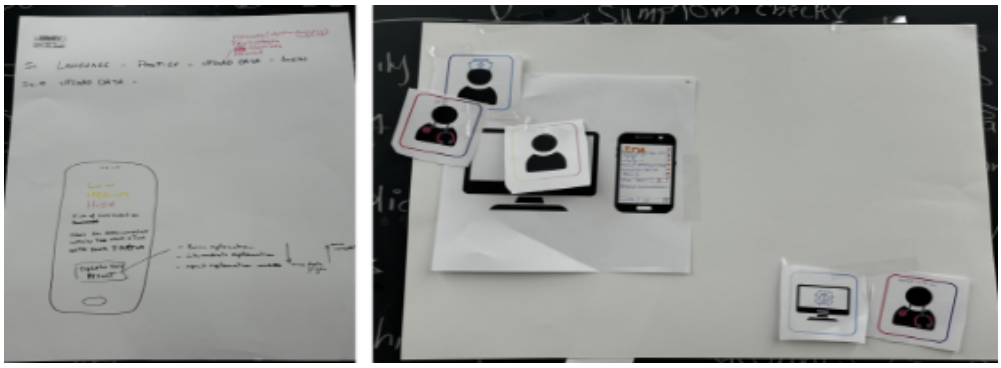


Figure 1

Independently produced designs from the two groups. Both are patient-facing mobile applications with banded (non-numerical) risk, whole-person inputs, actionable recommendations, and built-in consent. One group (left) detailed a language-profile-data flow and a layered basic/intermediate/expert explanation; the other (right) named its design “eMe” and elaborated a shared patient-clinician data model.