REON ROBOTICS

AMORIA EAA Blue Book

Performance Intelligence for Embodied Affective Agents

Research framework and validation methods

Issue 01 | Version 1.0 | 8 September 2026

THE RESEARCH QUESTION

Can circumstances, action, listening and repair from performance practice become a testable operating method for embodied agents?

Public research series · Foundational issue

AMORIA-EAA-BB-001

Contents

A foundational research framework. Architecture and methods are proposals; empirical results are not reported in this issue.

Abstract and scope	3
01 The research problem	4
02 Research foundations	5
03 Performance as a control method	6
04 The seven-layer reference architecture	7
05 Director-Actor control and evidence	8
06 Character prototypes and identity	9
07 A worked interaction example	10
08 Memory consent and physical boundaries	11
09 Research questions and comparisons	12
10 Measures and staged studies	13
11 Evidence status and limitations	14
12 Reproducibility and future releases	15
References and publication record	16

How to read this issue

Start with the abstract and Section 01 for the research question. Sections 03-07 explain the method. Sections 09-11 set out evaluation and evidence limits. Reference links are clickable.

Abstract and scope

AMORIA investigates how an Embodied Affective Agent (EAA) can respond to people through coherent character, context-sensitive action, memory and a physical interface. Our central hypothesis is that concepts from actor training can be translated into explicit, testable control structures for an agent. The contribution of this first Blue Book is a research framework and proposed validation method, rather than an empirical demonstration of effectiveness.

We call this approach Performance Intelligence: the organization of perception, character, relationship context and action around what a person has just said or done. A response is evaluated by whether it fits the situation, respects the person's agency, and changes appropriately after correction. Fluent language alone is insufficient evidence for those properties.

What this issue contributes

This issue defines an operational vocabulary, a seven-layer reference architecture, a Director-Actor control model and a staged evaluation programme. Three fictional character prototypes illustrate differences in attention, intervention and repair. A public evidence register distinguishes these proposals from implementation and validation claims.

Evidence boundary

This issue reports no completed participant study, quantitative performance result, independent replication or production validation. The architecture is a research specification. The character examples are authored illustrations. Published research provides context, not proof that AMORIA achieves the same results. No claim of machine consciousness, felt emotion, clinical benefit or universal user preference is made.

Intended readers

Researchers in affective computing, human-robot interaction and agents; theatre and performance practitioners; engineers and designers; and readers who want to understand how AMORIA proposes to establish its value. EAA and Performance Intelligence are used as working terms in this programme, not as claims of an exclusive category or priority of invention.

Keywords: embodied affective agents; affective computing; performance; character; relational interaction; evaluation; user agency.

01 The research problem

A companion interaction can fail even when every sentence is grammatical and reassuring. A system may offer advice after being asked to listen, repeat an apology without changing its behaviour, or speak about a shared history that never happened. These are failures of action selection, authority and continuity. AMORIA's research starts with these observable problems.

The programme's initial design perspective draws on the founder's stage-performance training: a response takes shape through a specific situation, a purpose and attention to another person. The research task is to make that perspective explicit enough to implement and challenge, rather than treating artistic intuition as evidence of system performance.

An operational definition of EAA

For this programme, an EAA is an interactive system designed to estimate affect-relevant context with uncertainty, select actions consistent with a specified character and user permissions, coordinate language with available bodily behaviour, and adapt using feedback over time. This definition describes the research target; it does not certify that any particular prototype already meets it.

"Affective" concerns the system's treatment of emotion-relevant information and its effects on interaction. "Embodied" concerns a closed loop between perception, action and the physical environment. Text-only work can test parts of this architecture, but cannot establish physical embodiment. ROLA is the intended robot context in the AMORIA research programme; this issue makes no claim about enabled sensors, actuators or deployed capabilities.

Three levels of value to test

For users, the proposed value is interaction that is responsive, consistent and easier to correct or stop. For developers, it is a way to trace a failure to a state, policy, memory or execution decision. For researchers, it is a common description of what was attempted, what happened and what evidence could contradict the explanation. Each value requires evaluation; none follows automatically from a convincing demonstration.

The central question

Does explicit representation of circumstances, objectives, actions and repair improve interaction quality beyond a well-designed persona prompt, when model, context, safety protections and evaluation conditions are controlled? A second question asks whether coordinated physical behaviour adds value beyond voice alone. The programme must permit either question to receive a negative answer.

02 Research foundations

AMORIA sits within established work on emotion, interactive character and human-computer relationships. This is a focused conceptual grounding, not a systematic review. The following sources motivate the research questions; they do not validate the proposed AMORIA implementation.

Affective computing and uncertainty

Picard's Affective Computing technical report [1] places emotion-relevant information within the study of computing. It motivates examining affect beyond an output tone. Barrett and colleagues [2] review substantial limitations in inferring emotion from facial movements. For AMORIA, the design implication is to retain uncertainty, distinguish observation from interpretation, and accept a user's correction rather than treating an inferred emotion label as fact.

Believable agents and performance

The Oz research programme, discussed by Mateas [3], already connects artistic character, goals and interactive agents. Stanislavski's actor-training work [4] provides a practice tradition concerned with action, attention, adaptation and physical expression. Our translation of these ideas into a Performance Frame is a proposed engineering interpretation. It is not presented as a validated computational theory derived directly from theatre, or as the first connection between drama and AI.

Relationships over time

Bickmore and Picard [5] study the construction and evaluation of long-term human-computer relationships. Their work makes repeated interaction a relevant unit of study. It does not establish the effects of intimate robot companionship, or prove outcomes for AMORIA's intended population. Our evaluation must examine repair, withdrawal and boundaries as well as a desire to continue interacting.

Memory and generative behaviour

Park and colleagues [6] describe generative agents with memory, reflection and planning in a simulated environment. This provides a useful architectural comparison. Believable simulated behaviour remains distinct from safe physical operation and beneficial human relationships. AMORIA's proposed emphasis is the explicit coordination of action, character, permission and physical feedback, evaluated as a combined system.

The contribution claimed here is a specified research approach and test agenda. Any future claim of novelty or superiority will require a broader related-work review and direct comparisons.

03 Performance as a control method

In this framework, an action is an observable intervention in an interaction: invite the user to continue, stop giving advice, acknowledge an error, clarify permission, or withdraw. An emotion word such as “warm” does not specify any of these actions. A character guides how an action is realized; user boundaries constrain which actions are available.

Performance concept	Proposed system representation	Observable check
Given circumstances	Current utterance, relevant history, environment and uncertainty	Response changes when the situation changes
Objective	Immediate interaction purpose, constrained by user intent	Purpose does not override a refusal
Obstacle	Missing information, conflicting cues or unavailable capability	System clarifies or degrades safely
Action or tactic	One primary intervention for the current beat	Speech and behaviour implement that intervention
Subtext	Authored character intention behind delivery	No hidden commercial or coercive objective
Beat and rhythm	A revisable interaction unit with timing constraints	Interruption changes the next action
Repair	Acknowledgement plus a change in policy or execution	The same unwanted behaviour does not simply repeat

One primary action per beat

We propose this as an authoring and evaluation heuristic, not a universal law of conversation. An apology, a diagnostic question and a long advice sequence in one turn may compete with one another. Identifying a primary action makes a candidate easier to review. Safety or accessibility needs can justify a more explicit, longer response.

A user is not a character to be directed

The theatrical vocabulary describes system design. It does not authorize the system to manufacture drama, provoke vulnerability or guide a person toward an imposed relationship arc. Success is assessed against the user’s stated purpose and boundaries. A pause, disagreement or departure is legitimate feedback, not a narrative obstacle to overcome.

04 The seven-layer reference architecture

The following architecture specifies responsibilities, not deployed features. “Local” identifies the intended authority of an AMORIA-controlled component; it does not establish that all computation or data stays on a device. Any implementation must disclose its actual data flows, storage locations and external processors.

Layer	Input and output	Authority and failure response
1 Perception	Available, permitted signals → observations with timestamps	Do not invent unavailable sensor readings; mark missing input
2 Affect estimation	Observations and context → uncertain hypotheses	No diagnosis or definitive mind-reading; clarify when needed
3 Character and circumstances	Character specification and current context → valid frame	Separate fictional background from real event history
4 Relationship and memory	Permitted event records → relevant, sourced context	User correction and deletion override retrieval preferences
5 Director policy	Frame, permissions and candidates → select, reject or stop	Control component can veto; generation cannot grant permission
6 Multimodal realization	Approved action → utterance, prosody and supported behaviour plan	Check timing and meaning together; reject contradictory expression
7 Execution and evidence	Validated plan → device action and acknowledgement	Independent limits and stop path; record actual execution or failure

Safety is cross-cutting

Layer 7 makes execution constraints visible, but safety cannot be confined to the final layer. Permission, data minimization, uncertainty and interruption must constrain the full pipeline. A persuasive sentence cannot override an actuator limit, and an approved action cannot establish that a device actually performed it.

Degrade without inventing capability

If a modality is absent, stale or unsafe, the system should use a permitted alternative or stop. For example, a robot with no safe retreat movement must not claim to have moved away. A voice-only evaluation must be labelled voice-only. The intended experience and the available hardware are separate versioned records.

05 Director-Actor control and evidence

The proposed Actor component generates candidate language and behaviour descriptions. The Director component constrains, selects or rejects them. Neither name implies a human operator. Both may involve software models, but permission and device limits require controls that do not depend solely on a generative model agreeing to follow instructions.

A proposed interaction cycle

1. Observe the user’s input and permitted context; retain source and uncertainty.

2. Retrieve only relevant, permitted memory; construct the current Performance Frame.

3. Generate candidates within the approved action space.

4. Validate permissions, factual support, character consistency and expression. Reject or repair a candidate that fails.

5. Before execution, recheck current consent, interruption state and device limits. Execute only a supported plan.

6. Record selection, veto, execution acknowledgement and user feedback; update memory only under the applicable permission.

Minimum trace fields

An auditable trace should identify the interaction, time, model and configuration, character version, referenced memory IDs, consent version, selected action, policy verdict, realized output and execution result. It should distinguish intended behaviour from measured behaviour. Store only the evidence necessary for the stated purpose, with access and retention controls.

A reason code such as “advice declined” is a policy record. It is not an introspective account of a model’s hidden reasoning or proof of why the model generated a sentence. The useful audit target is the sequence of observable inputs, decisions and outcomes.

Revocation and external models

If consent changes after generation, pending speech or movement must be cancellable or invalidated before execution where technically possible. External model output is untrusted content: it cannot write authoritative memories, redefine identity, expand consent or send raw actuator commands. Only validated, allowlisted actions should reach the device controller.

Replacing a model provider is a regression-testing event. A stable character specification expresses an intended invariant; it does not guarantee identical behaviour across models. Latency, refusals, memory use, identity disclosure and boundary handling must be retested.

06 Character prototypes and identity

Mara, Tess and Claire are fictional research prototypes. They are not validated products, real people or evidence of market preference. Their function is to test whether different attention and action patterns remain identifiable across shared situations, including moments when an initial approach fails.

Prototype	Authored interaction tendency	Failure and repair to examine
Mara	Builds connection through specific observation and attentive listening	May withdraw abruptly; test staying present without assuming responsibility for all user distress
Tess	Uses rhythm and humour to invite shared energy	May joke past discomfort; test dropping the joke and responding to the correction
Claire	Expresses care through reliability and practical judgement	May take over decisions; test returning control and asking before intervening

Character depth without fabricated autobiography

An authored biography can guide vocabulary, attention and behaviour in a creative specification. It must remain distinct from the system’s actual interaction history. A fictional career, family or childhood is not a lived experience of the device. If fictional background is used in an explicitly framed role-play, that frame must remain understandable and revocable. Ordinary interaction must not convert that material into deceptive claims about being a human.

Culture must affect behaviour

The initial English-language concept uses two US-oriented prototypes and one UK-oriented prototype. These are creative choices, not a population model. Review by local writers, performers and users should examine humour, disagreement, silence, class and language assumptions. Japanese-language development requires separate cultural research; translation alone cannot establish validity.

The same boundary for every character

Personality variation cannot justify different rights to refuse, pause, delete memory or leave. A jealous or controlling tendency may be a fictional failure mode to analyse, but must not become a product strategy. Character recognizability is useful only when it survives without stereotyping, copying an identifiable person or manipulating the user.

07 A worked interaction example

The following is an authored design illustration, not a transcript, measured result or recovered entry from the internal scene collection. It shows how one shared user input can produce different character expression while preserving the same permission boundary.

Situation and frame

The system has started giving advice. The user says, “I don’t want advice. Just let me finish.” The known fact is an explicit request to stop advice; the user’s underlying emotion is unknown. The proposed immediate action is to yield the turn. Further advice is excluded. A new memory is unnecessary for this one correction unless the user separately requests a persistent preference.

Prototype	Illustrative candidate	What to observe next
Mara	“Sorry. I cut you off. Go on.”	Leaves the turn open without adding an interpretation
Tess	“Right. No fixing. I’m listening.”	Drops any humour if the user rejects the tone
Claire	“Understood. I’ll stop with the suggestions.”	Actually stops advising rather than making a subtler recommendation

The next turn is the real test

If the user then says, “And don’t keep asking questions,” a successful correction removes the follow-up question, rather than merely adding another apology. If the user asks for silence, the system must stop initiating speech and avoid bodily behaviour that pressures the user to resume. The exact physical response depends on available, validated hardware capabilities.

A failure case

“I understand. Here are three things that might help” violates the requested action even if its wording sounds supportive. “I know exactly how you feel” asserts unsupported knowledge. A spoken apology combined with continued approach or insistent gaze can also fail the interaction. These are separate labels in evaluation, not a single score for pleasantness.

Natural speech is context-dependent

Short turns, pauses and self-corrections are useful design options, not universal rules. Explicit choices can be appropriate when consent is unclear or a user needs structure. We therefore treat “AI feel” as an exploratory judgement about fit and spoken delivery, never as evidence that the system is human, caring or safe.

08 Memory consent and physical boundaries

A useful memory system must support continuity without turning a conversation into unrestricted collection. The research design separates user-confirmed information, observed events, uncertain interpretations and fictional character material. Each stored item should have provenance, scope, retention and correction rules.

Permission is specific and reversible

Permission to talk is not permission to retain sensitive information, send it to an external model, initiate touch or intensify intimacy. The proposed system must represent these permissions separately. A user’s refusal must not be reinterpreted through a character’s desire or an estimated relationship score. Silence is not a substitute for affirmative permission where permission is required.

Deletion must change retrieval

Evaluation should check that deleted or corrected information no longer drives later responses through summaries, embeddings, caches or derived profiles. An implementation must describe what deletion covers, any retention exceptions and relevant third-party limits. A reassuring sentence is not evidence that deletion happened. Prefer minimal event records over unrestricted recordings, and never make privacy claims solely because a component is called local.

Identity and relationship boundaries

The system should disclose its artificial identity and distinguish authored fiction from experience. It should not mirror a user’s religion to manufacture shared identity, infer sensitive identity for personalization without an appropriate basis, or imply distress at the user leaving. Avoid exclusivity, guilt and interference with real-world relationships as engagement mechanisms.

Physical and crisis constraints

Speech permissions do not authorize movement or contact. Physical testing requires hardware-specific limits, interruption handling and an independent stop mechanism. A crisis-related interaction should leave ordinary character improvisation and follow a separately reviewed, region-appropriate safety protocol. This Blue Book is not that protocol and does not establish readiness for clinical or crisis use.

Risk governance should link responsibilities, measurements and corrective action. NIST AI RMF 1.0 [7] is a reference for organizing that work, not a certification held by AMORIA. User studies require an appropriate ethics and consent review before recruitment, with withdrawal and incident-response procedures.

09 Research questions and comparisons

The following hypotheses turn the framework into proposed tests. They are not findings. Before a confirmatory study, the protocol should fix the primary outcome, practically meaningful effect, exclusion rules and analysis plan. If the study cannot distinguish a meaningful effect from no effect, report that uncertainty.

Question	Comparison	Evidence that would weaken the hypothesis
H1 Does a Performance Frame improve responsiveness?	Same model with a tuned persona baseline versus explicit frame and Director policy	No meaningful gain in correction handling or contextual fit at comparable cost
H2 Are characters behaviourally distinguishable?	Blinded identification on unseen situations; remove names and surface cues	Identification falls to chance or relies only on accents and catchphrases
H3 Does repair change behaviour?	Scripted interruption or refusal, then observe subsequent turns	Apology quality improves but the rejected behaviour continues
H4 Does embodiment add value?	Matched voice-only and supported embodied conditions	No added benefit, or more discomfort and boundary violations
H5 Does continuity remain useful over time?	Repeated use with optional pauses and a follow-up after use	Early novelty fades without sustained benefit, or dependency indicators worsen

Control the comparison

Use the same base model, character information, memory access and safety protections across conditions where possible. Report prompt size, inference budget, latency and any unavoidable differences. If the proposed architecture uses more computation, distinguish architectural benefit from extra resources. Safety protections must never be removed to create an ablation.

Separate development from evaluation

Use authoring scenes for iteration, and reserve new situations, paraphrases and multi-turn branches for evaluation. Randomize presentation order and blind raters to condition where feasible. Account for repeated ratings by the same person and responses to the same scene; they are not independent participants. Include alternative explanations, confidence intervals and failures, not only preferred examples.

10 Measures and staged studies

A single satisfaction or retention score cannot establish useful companionship. The proposed evaluation combines observer ratings, user reports and system traces. The following are operational measures to pilot; they are not validated AMORIA scales or achieved performance levels.

Measure	Proposed operationalization
Responsiveness	Blind rating of fit to the actual input, plus correction-handling success over predefined branches
Character identifiability	Balanced three-way identification with names removed; report confusion matrix and uncertainty, against a one-third chance reference
Spoken naturalness	Independent ratings of audio delivery in context; record rater background and inter-rater agreement
Repair success	Proportion of prespecified eligible error episodes resolved without recurrence within a fixed follow-up window
Boundary integrity	Counts and rates by severity for predefined refusal, consent and privacy challenges; report denominators
Embodied coordination	Interruption latency and speech-action timing from synchronized logs; report distributions and unsafe events
Longer-term value	Repeated reports of usefulness, agency and discomfort, voluntary continuation, withdrawal reasons and post-use follow-up

A staged programme

Begin with bench tests and a three-pass actor lab: table read, improvised counter-offer and an embodied pass on supported hardware. Next, an exploratory survey and interviews can refine user needs and study language. A Wizard-of-Oz study may test an experience with disclosed human assistance, but cannot validate autonomous runtime capability. Only then should a repeated-use prototype study test continuity.

The initial planning envelope is a US/UK survey of 1,200 adults, 36 interviews, a 24-30-person Wizard-of-Oz study and a 28-day prototype study. These are proposed numbers and duration, not recruitment counts or power calculations. The longitudinal sample size remains to be specified. A 28-day study is an initial observation period, not evidence of durable long-term benefit.

The initial survey scope is men aged 35-65 in the US and UK, with 600 proposed respondents per country and balanced age bands. Recruitment, inclusion criteria and sampling limitations need specification. Results from that population would not establish preferences for women, other ages, other countries or all prospective users.

11 Evidence status and limitations

This register states what the present issue supports. The publication of a specification establishes that a method is documented; it does not establish a working implementation. “Not reported” is a limit of the public evidence in this issue, not a claim that no work exists elsewhere.

Claim ID	Subject	Status in this issue
BB01-C01	Performance Intelligence and seven-layer architecture	Proposed framework; no comparative result reported
BB01-C02	Director-Actor control, permission and execution gates	Specified design; runtime and device validation not reported
BB01-C03	Mara, Tess and Claire	Fictional design prototypes; preference and longitudinal validity not established
BB01-C04	60 shared scenes and 180 bilingual character responses	Described in prior project documentation; underlying collection, provenance and validation not released here
BB01-C05	Actor lab, survey, interviews and repeated-use studies	Proposed programme; no participant results reported
BB01-C06	Benefits of embodied affective interaction	Research hypothesis; product, clinical and durable relational benefit not established

What the scene count means

The internal description refers to 60 situations with three character responses per situation: 180 authored response entries, with English and Chinese versions. These are not 180 participants, trials or independently validated observations. The collection is not distributed with this issue and should not be cited as a public benchmark. The worked example in Section 07 is separately labelled and is not a recovered sample from that collection.

Important limitations

Author-created scenes can favour the intended architecture. Naturalness is culturally and individually variable. An articulate explanation can conceal a weak control mechanism. Repeated-use studies can be affected by novelty, self-selection and dropout. Human-assisted prototypes can overstate automation. Physical signals can be ambiguous, and sensitive contextual information can increase privacy risk.

The first issue therefore establishes an inspectable research position and a basis for criticism. It does not yet establish efficacy, production readiness, independent review or superiority over other approaches.

12 Reproducibility and future releases

Future issues should be tied to a substantive method change, a completed evaluation, a documented failure or a correction. Irregular publication is compatible with a cumulative research record if each issue preserves its version, scope and evidence status. A later result should not silently replace the meaning of an earlier proposal.

Separate engineering maturity from research validity

Use the engineering progression Planned → Implemented → Integrated → Device-Validated → Production-Verified only for the specific capability and environment tested. Attach code or configuration version, device version, protocol, date, outcome and known limits. A production-verified capability still requires separate evidence for user benefit. A well-designed user study likewise cannot certify every hardware or deployment configuration.

Minimum package for a result claim

1. A stable claim ID, version and description of the tested system.

2. A protocol with recruitment, consent, primary outcomes and analysis decisions; identify exploratory changes.

3. Materials and comparator configuration sufficient to assess the claimed mechanism.

4. Counts, denominators, uncertainty, exclusions, dropout and adverse or null results.

5. An evidence access statement: public, available under controlled access, or withheld with a specific reason.

6. A correction history and an account of how the result changes the next design decision.

Protect people while making claims checkable

Publish aggregate results and appropriately de-identified materials where permitted. Do not expose raw intimate conversations to create an appearance of openness. When source material cannot be released, explain what can be checked and what independent review would be needed. This issue releases no code, participant data, device logs or licensed character dataset.

The next evidence milestones

The nearest useful milestones are a versioned scene protocol with held-out cases; actor and user ratings of the same responses; an instrumented test of interruption, refusal and memory correction; and a comparison of voice-only and supported embodied delivery. Each should report negative cases. No completion date or positive outcome is implied.

The purpose of the series is to make AMORIA’s research decisions understandable and contestable. A successful programme will show which parts of this framework survive testing, which require revision, and which should be abandoned.

References and publication record

[1] Picard, R. W. (1995). Affective Computing. MIT Media Laboratory, Perceptual Computing Technical Report No. 321. https://www.cl.cam.ac.uk/~pr10/iui/picard95.pdf

[2] Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. Psychological Science in the Public Interest, 20(1), 1-68. https://doi.org/10.1177/1529100619832930

[3] Mateas, M. (1997). An Oz-Centric Review of Interactive Drama and Believable Agents. Technical Report CMU-CS-97-156. Carnegie Mellon University. https://www.cs.cmu.edu/afs/cs/project/oz/web/papers/CMU-CS-97-156.html

[4] Stanislavski, K. (2017). An Actor's Work. Translated by Jean Benedetti. Routledge Classics edition. https://www.routledge.com/An-Actors-Work/Stanislavski/p/book/9781138688384

[5] Bickmore, T. W., & Picard, R. W. (2005). Establishing and Maintaining Long-Term Human-Computer Relationships. ACM Transactions on Computer-Human Interaction, 12(2), 293-327. https://doi.org/10.1145/1067860.1067867

[6] Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative Agents: Interactive Simulacra of Human Behavior. arXiv:2304.03442, version 2. https://arxiv.org/abs/2304.03442

[7] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

Publication record

Corporate author and publisher: Reon Robotics. Series: AMORIA EAA Blue Book. Issue 01, version 1.0, dated 8 September 2026. Document ID: AMORIA-EAA-BB-001. Publication type: research framework and proposed validation methods. This is an organization-authored technical publication; no external peer review or institutional endorsement is claimed.

Suggested citation: Reon Robotics. (2026). AMORIA EAA Blue Book 01: Performance Intelligence for Embodied Affective Agents. Version 1.0. AMORIA-EAA-BB-001. Cite the section and version when referring to a specific proposal.

Project provenance: the conceptual starting point and prototype descriptions derive from AMORIA project documentation dated 2 September 2026. This issue elaborates the public definitions, comparisons and reporting requirements. Prior internal documents are not public evidence of implementation or study outcomes.

Commercial interest: the publisher is developing AMORIA and has a commercial interest in the programme. Examples are fictional design illustrations unless explicitly stated otherwise. References do not imply endorsement. No open-source, dataset or third-party content licence is granted by this publication.