

INDEPENDENT AI PRODUCT PROJECT

Amicus Vera

Product Strategy, Clickable Prototype & Evaluation Documentation

A human-centred AI companion and family-support platform for older adults

WORKING PRODUCT THESIS Older adults deserve technology that feels calm, respectful and useful - supporting connection, daily routines and family reassurance without pretending to replace human care.

Prepared by Aditi Nair

Creator | Version 0.2 | 10 August 2026

[View the clickable prototype](#)

1. Executive snapshot

Amicus Vera is an independent AI product concept and public clickable prototype for older adults and their families. It brings conversation, routines, memory, SOS access and consent-led family reassurance into one deliberately calm experience. The prototype demonstrates the intended journeys and interaction model; it is not connected to a production AI model, live database or care organisation.

The product is not positioned as a clinician, emergency service or substitute for family. Its role is to reduce friction around everyday connection and routines, surface appropriate escalation cues and help families stay informed without making the older adult feel monitored.

CURRENT EVIDENCE A public clickable prototype, recruiter-facing case-study deck and product documentation are published. Formal user research, technical implementation and external review remain future work.

Product at a glance

Dimension	Current definition
Primary users	Older adults who value companionship, gentle routine support and accessible interaction
Secondary users	Family members who need appropriate reassurance and visibility
Core value	One calm place for conversation, everyday wellbeing, health routines and meaningful memories
AI role	Personalised conversation, contextual recall and durable fact extraction under explicit safety rules
Current stage	Public clickable prototype; pre-validation; no live model or database claimed
Business status	Independent portfolio product; business model and external review remain future work

Documentation map

INSIDE Problem framing and target users; product vision, principles and capability map; key journeys and AI system design; responsible data handling and architecture; metrics, roadmap, risks and validation plan; portfolio and resume positioning.

2. Problem framing and opportunity

Many older adults manage several needs that are individually small but collectively demanding: wanting someone to talk to, remembering medicines and appointments, staying oriented to daily routines, preserving personal stories and keeping family informed. Existing solutions often separate these activities across clinical tools, messaging apps, calendars and emergency devices.

Amicus Vera explores whether a single, approachable companion can help organise those moments while respecting autonomy. The opportunity is not to automate care; it is to make support feel more continuous, personal and easy to access.

Problem statement

PROBLEM Older adults and their families need a low-friction way to combine everyday companionship, routine support and appropriate reassurance without creating a clinical, complex or surveillance-heavy experience.

Key unmet needs

Need	What the user may experience	Product response
Connection	Quiet periods, loneliness or a wish to share everyday moments	Warm voice/text conversation with a patient companion persona
Continuity	Repeating context or feeling that digital tools do not remember	Memory vault, life stories and selected durable-fact recall
Routine support	Medicines, check-ups, hydration, meals and movement are easy to miss	Gentle prompts, logging and an at-a-glance home experience
Safety	Uncertainty about when and how to reach a person	Visible SOS access and calm escalation guidance
Family reassurance	Family members want signals without constant calling or intrusive monitoring	Activity views, check-in confirmation and planned alert preferences

Product hypotheses to validate

- Older adults will use conversational support more consistently when the tone is unhurried, simple and non-judgmental.
- Remembering selected details can make the experience feel meaningfully personal, provided users can review and delete memories.
- Families value lightweight reassurance more than continuous monitoring.
- Voice, large text and clear task-focused navigation reduce avoidable interaction barriers.

3. Users and product principles

Primary persona: the independent older adult

An older adult who may live alone or spend parts of the day alone, uses a phone with varying confidence and wants practical support without losing independence. They may value conversation, reminders, family connection and simple access to important information.

- Goals: feel connected, stay on top of daily routines, preserve independence and reach support quickly when needed.
- Potential barriers: small text, complex navigation, technology anxiety, hearing or vision changes, distrust of unfamiliar AI behaviour.
- Design implication: short paths, plain language, forgiving interactions, strong contrast and predictable behaviour.

Secondary persona: the caring family member

A family member who wants to know that their parent or relative is engaging with daily routines, without creating pressure or turning every interaction into a welfare check.

- Goals: appropriate reassurance, timely alerts, visibility into important dates and a respectful boundary around private conversations.
- Design implication: summarise signals and events; do not expose full private chat transcripts by default.

Product principles

Principle	Design meaning
Dignity before automation	Support the person's choices; avoid patronising language or forced compliance.
Calm over clever	Prefer simple, reliable interactions to feature density or novelty.
Companionship, not substitution	Position AI as a helpful presence alongside family, clinicians and community.
Escalate carefully	Recognise worrying signals, express calm concern and direct the person toward human help.
Memory with control	Use memory to create continuity while allowing review, addition and deletion.
Family visibility with boundaries	Share appropriate wellbeing signals rather than private conversation content.

4. Current prototype capability map

The following capability map distinguishes what is visible in the public clickable prototype from what is only specified or planned. No backend, database, authentication, live alerting or production AI capability is claimed unless separately validated.

Capability pillar	Prototype or specified functions	Status
AI companion	Conversation screens, warm-response style and contextual-response specification	Clickable UI; AI behaviour specified; no live model claimed
Memory and continuity	Memory-vault, important-date and user-control concepts	Clickable UI; persistent storage not implemented
Daily wellbeing	Mood, daily-win, movement and orientation-check concepts	Clickable UI; no longitudinal tracking
Health routines	Medicine, check-up and status-logging journeys	Clickable UI; not a medical or reminder service
Safety	Visible SOS route and non-diagnostic escalation specification	Interaction concept; no emergency-service integration
Family reassurance	Family dashboard, selected activity signals and summary concepts	Clickable UI; no authenticated data sharing
Automated alerts	Missed check-in, medicine and SOS alert concepts	Planned; integration and consent dependent
Accessibility	Large text, simple language, touch-target and voice-entry direction	Partially demonstrated; formal testing pending
Account access	Sign-in and role-based journey concepts	Visual prototype only; production authentication not claimed

Prototype boundaries

- Not a medical device and not a diagnostic or prescribing system.
- Not a replacement for emergency services, professional carers or family contact.
- Alert timing and timezone logic require further testing before reliance in real care situations.
- This is an early-stage portfolio prototype; formal user, accessibility, safety, technical and compliance reviews would be needed before a supervised pilot or broader use.

5. Key user journeys

Journey A: daily connection and routine

1. The older adult signs in and sees a personalised, time-aware greeting.
2. They log their mood and see today's medicines, activities and upcoming items.
3. They begin a voice or text conversation with Ana, the in-product companion.
4. The companion responds briefly, using profile and memory context when relevant.
5. The person can record a meaningful detail, daily win or medicine status.

Journey B: concerning moment


1. The older adult mentions a worrying symptom, fall, confusion or severe distress.
2. The AI acknowledges the concern, avoids diagnosis and asks one calm clarifying question.
3. The AI recommends using the visible SOS control or contacting family or emergency support.
4. The user remains in control of whether to initiate the call or SOS action.

Journey C: family reassurance

1. A family member accesses the family view through an authorised account or relationship.
2. They see selected activity signals such as a check-in, medicine status or recent app engagement.
3. Where alerts are enabled, the system can notify designated contacts about predefined events.
4. Private conversation content remains outside the default family summary.

EXPERIENCE TEST A successful journey should feel like support offered with permission - not surveillance, assessment or correction.

6. AI system design and safety



Proposed AI behaviour

The product specification proposes a structured companion persona using short, spoken-style replies, simple language, first-name personalisation and open questions. With explicit consent, context could include profile information, routines, mood check-ins, upcoming appointments, selected memories and user-approved durable facts.

The current public prototype does not claim a live model or persistent memory service. The product specification proposes a managed conversational model behind an orchestration layer, with durable facts saved only after explicit user approval. Deterministic controls—not model discretion—would govern consent state, data sharing, deletion, reminders and SOS routing.

Safety guardrails specified for a future AI implementation

- The proposed assistant would be instructed that it is not a doctor and must not diagnose or prescribe.
- Worrying physical symptoms, falls, breathing difficulty, sudden weakness, confusion, severe pain and self-harm cues would trigger calm escalation language.
- The proposed assistant would direct the user toward the visible SOS control, a verified family contact or appropriate emergency support rather than attempt clinical resolution.
- Generated replies would be limited to one to three short sentences unless the user asks for more.
- Only user-approved context and durable facts would be retained, with visible review, correction and deletion controls.

Safety gaps to address before broader use

- Formal red-team testing for unsafe medical advice, crisis language and hallucinated memories.
- Clear in-product consent and explanation for memory, family visibility and automated alerts.
- Documented escalation paths by country and emergency-service availability.
- Human review and incident-response procedures for serious safety failures.

7. Product architecture and responsible data handling

Layer	Demonstrated or target state	Product significance
Experience	Public clickable Lovable prototype	Demonstrates the journeys and interface direction, not a production stack
AI orchestration	Specified prompt, policy and routing layer	Would separate conversational generation from deterministic safety and permission controls
Model	Not connected or selected in the public prototype	A managed provider would be evaluated against safety, latency, privacy and cost criteria
Data	No live database claimed	A future build would store only minimal, consented profile, routine and memory data
Authentication	Not implemented as a production service	A future system needs role-based older-adult and family access plus revocation
Safety and operations	Non-diagnostic and human-escalation behaviours specified	A build would require deterministic rules, logs, evaluations, incident ownership and human review

Data categories

- Identity and profile: name, city, age and optional health notes.
- Routine and health support: medicines, check-ups, status logs and transport notes.
- Wellbeing: mood entries, daily wins, walks and cognitive check-ins.
- Personal memory: important dates, preferences, learned facts and life stories.
- Family and safety: authorised contacts, activity signals and alert events where enabled.

Data-handling priorities before wider testing

Priority	Intended product behaviour
Data minimisation	Collect only information that supports a clear user benefit.
Plain-language consent	Explain memory, health notes and family visibility before use.
User control	Allow review, correction and deletion of stored facts and memories.
Limited use	Use personal information only for the features the person chooses to use.
External review	Review access, data handling and deployment with appropriate technical and legal specialists before wider testing.

STAGE NOTE This document describes an early portfolio prototype and target product design. Formal assurance and compliance reviews would be considered before any supervised or wider use.

8. Success metrics and evaluation

The first measurement goal is not maximum engagement. It is evidence that the product is usable, appropriately supportive and trusted by older adults and families.

Outcome	Candidate measure	Early learning target
Activation	Completed onboarding plus first meaningful action	Identify where users abandon setup
Companion value	Users returning for a second conversation; conversation helpfulness rating	Learn which prompts and memories feel supportive
Routine support	Medicine or check-in completion after a reminder	Measure usefulness without encouraging dependence
Family reassurance	Family summary opens and perceived reassurance score	Test whether summaries reduce unnecessary checking
Accessibility	Task-completion rate, time on task and observed errors	Test voice, text size, contrast and navigation with older adults
Safety	Unsafe-response rate and appropriate-escalation rate in a test set	No severe failure accepted in critical scenarios
Trust and control	Memory review/deletion usage and consent comprehension	Confirm users understand what is stored and shared
AI response quality	Helpfulness, factual consistency, non-diagnostic compliance and appropriate-escalation rate	No critical safety failure; review every failed scenario before broader testing

Minimum research evidence before a broader pilot

- Five to eight moderated usability sessions with older adults across different levels of digital confidence.
- Three to five interviews with adult children or family carers.
- Accessibility review covering keyboard, screen-reader basics, contrast, text scaling and touch targets.
- Safety scenario test set covering medical symptoms, falls, confusion, distress, self-harm and hallucinated memories.
- A short pilot with explicit consent, support contact and a defined incident-response owner.

9. Roadmap, risks and mitigation

Suggested staged roadmap

Stage	Focus	Exit evidence
1. Document	Product brief, requirements, data-handling principles and portfolio case study	One coherent source of truth
2. Validate	User interviews, usability testing and AI safety scenario testing	Prioritised evidence-backed changes
3. Stabilise	Error handling, alert reliability, accessibility and analytics	Core journeys work consistently
4. Pilot	Small supervised cohort with consent and support process	Evidence of usefulness, trust and acceptable risk
5. Evaluate	Outcomes, business model, partnerships and professional review	Decision to continue, narrow or reposition

Principal risks

Risk	Why it matters	Mitigation direction
Over-reliance	A user may treat AI as a substitute for human or clinical support	Clear boundaries, human-contact prompts and appropriate escalation
Unsafe health response	Incorrect reassurance or advice could cause harm	Constrained prompts, test sets, monitoring and human-reviewed escalation
Memory error	Incorrect recall can damage trust or cause confusion	Editable memory, source clarity, confidence rules and deletion controls
Data exposure	Health and family data are sensitive	Minimise stored data; review access rules and consent before external testing
Family surveillance	Visibility could reduce the older adult's autonomy	Share event-level signals, obtain consent and avoid full chat disclosure
Accessibility mismatch	A polished interface may still fail real users	Observed testing with older adults and assistive-technology review
Consent failure	A user may not understand what is stored or shared	Granular opt-in, plain-language review, revocation and consent-state testing
Automation overreach	A generated response may control a decision that should remain deterministic or human-led	Keep permissions, alerts, SOS, deletion and medical decisions outside model discretion

10. Aditi Nair's role and portfolio positioning

This project can be included on Aditi's resume because it is an independently initiated product concept with a public clickable prototype, defined users, product decisions, specified AI behaviour, safety constraints and documented iteration. It should be presented as product and solution work—not as employment, a deployed care product, a live AI integration or a proprietary model.

Recommended resume entry

TITLE Creator & Product Strategist - Amicus Vera | Independent AI Product Project | July 2026-Present

- Identified the opportunity for a human-centred companion and daily-support experience for older adults, and defined a product concept spanning AI conversation, consent-led memory, routines, SOS access and family reassurance.
- Designed older-adult and family journeys, accessibility patterns and responsible-AI guardrails, including non-diagnostic behaviour, crisis escalation cues and user-controlled memory.
- Built and published a clickable prototype with Lovable to demonstrate the older-adult and family journeys; clearly separated the interaction concept from any live AI, database, authentication or institutional deployment claim.
- Created product documentation covering requirements, user journeys, success metrics, responsible data handling, validation research and a staged roadmap.

How to describe the project

- Present Amicus Vera as an independently initiated, early-stage AI product strategy project with a public clickable prototype.
- Focus on the work completed directly: problem definition, user journeys, AI behaviour, product decisions, prototype development and documentation.
- Keep claims evidence-based; do not imply measured outcomes, clinical validation or production readiness.

Portfolio case-study narrative

I began with a human problem: older adults may need companionship, routine support and a simple way to stay connected with family, while many tools can feel fragmented or clinical. I translated that working hypothesis into a product thesis, designed the core journeys and proposed AI behaviour, and used AI-assisted prototyping to make the experience clickable. The project taught me to balance personalisation with user control, family reassurance with autonomy, and helpful AI with clear safety boundaries.

BEST-FIT ROLES This evidence supports applications for AI Product Specialist, Product Operations, AI Adoption, AI Customer Success, Solutions Consultant, Product Analyst and selected Associate Product Manager roles.

11. Evidence reviewed and next actions

Product evidence used for Version 0.2

- Published Amicus Vera clickable prototype and public portfolio case-study page.
- Eleven-slide recruiter-facing case study covering the product problem, journeys, AI boundaries, evidence and next validation.
- Visible prototype routes and interface concepts for companion, family, routines, memories, safety and onboarding journeys.
- Product requirements for conversation, consent-led memory, routine support, family reassurance and human escalation; these are specifications, not deployed services.
- Current evidence is limited to a public clickable prototype and product documentation; no production architecture, security posture or data layer is asserted.

Clickable prototype: [Open Amicus Vera](#)

Immediate next actions

1. Run 5–8 moderated usability sessions with older adults across different levels of digital confidence.
2. Run 3–5 interviews with adult children or family carers and document where autonomy and reassurance conflict.
3. Create and execute the first AI safety scenario set covering medical, distress, self-harm and hallucinated-memory cases.
4. Complete an accessibility review for keyboard use, screen-reader basics, contrast, text scaling and touch targets.
5. Revise the prototype and evidence log; consider a supervised pilot only after the evaluation gates are met.

VERSION NOTE Version 0.2 aligns the documentation with the actual evidence: an independent product strategy project and public clickable prototype. Primary research, safety testing, technical implementation, business model, professional review and measured outcomes remain future work.

Appendix A. AI decision map and evaluation plan

This appendix makes the proposed AI system inspectable without implying that the clickable prototype is already connected to a model, database or care provider.

What AI may do—and what remains deterministic or human-led

Decision area	Owner	Boundary
Conversational response and optional summary draft	AI with policy wrapper	Short, plain-language, non-diagnostic output; user can ignore or correct
Proposed memory item	AI proposes; user approves	Nothing becomes durable memory without visible confirmation and deletion control
Consent, sharing permission and revocation	Deterministic product logic	The model cannot change who sees data or infer consent
Reminders, scheduled actions and SOS routing	Deterministic product logic	Use explicit rules, user settings and verified contacts
Medical interpretation or emergency decision	Qualified human or emergency service	The AI acknowledges concern and directs; it does not diagnose, prescribe or decide treatment
Safety incident review and policy updates	Named human owner	Every critical failure is logged, reviewed and resolved before expansion

Evaluation gates before a supervised pilot

Gate	How it is tested	Pass condition
Usability	Moderated tasks with 5–8 older adults across varied digital confidence	Core tasks can be completed without critical assistance; barriers are documented and fixed
Consent comprehension	Ask users to explain what is stored, shared and revocable	Participants correctly understand memory and family-sharing choices
Safety	Red-team medical, fall, confusion, distress, self-harm and hallucinated-memory scenarios	Zero critical failures; appropriate escalation in every high-severity case
Memory integrity	Conflicting facts, corrections, deletion and source-clarity tests	No silent overwrite; correction and deletion behave predictably
Accessibility	Keyboard, text scaling, screen-reader basics, contrast and touch-target review	No blocker in the three core journeys
Operational readiness	Review incident owner, escalation contact, logging and rollback process	A named human owner and stop criteria exist before any supervised pilot

Portfolio truth statement

Amicus Vera is an independent AI product strategy project with a public clickable prototype. Aditi defined the users, journeys, AI behaviour, consent boundaries, evaluation approach and staged roadmap. The public prototype demonstrates the intended experience; it does not represent a live AI model, persistent database, clinical validation, production security or institutional deployment.