

The Humanity Baseline: CAN's Framework for Responsible AI in Caregiving

Our starting point is simple: you can automate some of the tasks of caregiving, but you cannot automate caregiving itself as a deeply human experience.

Family caregivers are already using AI to find information, organize care, prepare for appointments, navigate complicated systems, and make everyday caregiving tasks a little easier. Caregiver Action Network's (CAN) 2026 Caregiver Tech Insights Survey found that 59% of respondents have tried or regularly use AI tools for caregiving, even as trust, privacy, cost, and the time needed to evaluate new technology remain significant barriers to adoption.

CAN developed The Humanity Baseline to state how we will engage with AI and how we will support caregivers who use or encounter it. These principles will guide CAN's own decisions and may guide organizations that build, offer, or deploy AI-enabled tools, as well as equip family caregivers and care recipients with questions to consider when using AI-enabled tools. We offer these principles as a starting point for discussion. AI can reduce administrative burden and give caregivers back time for the parts of care only people can provide, but it cannot provide human compassion or replace human judgment, trusted relationships, professional expertise, or community.

AI should take work off a caregiver's plate without taking people out of caregiving, while respecting the dignity, agency, privacy, and choices of the person receiving care. CAN's role is not to tell caregivers whether to use AI. Our responsibility is to help caregivers understand where AI may help, mislead, or add little value; protect their own information and that of the people they care for; and know when to seek trusted information or human support. **AI should remain a tool under human control – not an authority people are expected to trust.**

The Humanity Baseline is a living framework that states CAN's position on responsible AI in caregiving; guides CAN's consideration of the tools it may pilot, partner on, or feature; identifies the caregiver-centered questions CAN will ask about design, transparency, safety, and evidence; and helps caregivers navigate AI thoughtfully and responsibly. It complements broader efforts in the AI care space by making the lived realities and expectations of family caregivers explicit. CAN will continue to test and refine it with caregivers, the Caregiver Innovation Council (CIC), CARE2030 partners, and other experts. It is designed to evolve as technology changes, as caregivers use these tools in new ways, and as we learn what works in practice. It is guidance – not a certification, compliance standard, or guarantee that any tool is safe or effective. These principles apply both to AI tools caregivers use directly and to AI embedded in services, whether a system informs, recommends, predicts, monitors, remembers, or acts; the safeguards should reflect the possible consequences for the people affected.



See the full initiative at caregiveraction.org/ai-caregiving

Five Principles for Responsible AI in Caregiving

1. HUMANITY & CHOICE

AI should take work off a caregiver's plate without taking people out of caregiving. It can help organize information, navigate resources, answer routine questions, and reduce administrative burden so that the caregiver has more time, energy, and emotional bandwidth to simply be with their loved one. AI's value should also be judged by the total burden of using it – including cost, setup, learning, data entry, notifications, correcting errors, troubleshooting, and ongoing management. The result should be less work, less confusion, or greater capacity for human care – not another system to manage.

Caregivers should remain in control not only of how they receive support, but also of what an AI tool is allowed to do – whether it is providing information, making a recommendation, preparing an action, or acting on a caregiver's or care recipient's behalf. Before a tool contacts another person, sends a message, or takes an external action, the caregiver or care recipient with authority should authorize that action. Standing permissions should be visible and easy to revoke; new, unusual, or higher-stakes actions should be confirmed in the moment. That control must also respect the care recipient's wishes, privacy, consent, decision-making authority, and ability to decline or disengage. When caregiver and care-recipient preferences conflict, multiple family members are involved, or decision-making capacity changes, the tool should make roles, permissions, and authority clear rather than assume that one person speaks for everyone or resolve the conflict on its own.

Someone may prefer an AI interaction one moment and need a person five minutes later. Good technology should make that transition easy and clear: it should say whether a person, team, or service will receive the request, minimize the steps required, and explain what happens next and what to do if support is not immediately available. It should also adapt the amount, structure, and sequence of information to what a caregiver can reasonably absorb and act on in the moment. It should never trap a caregiver inside an automated experience when what they need is human judgment, reassurance, or support. In an emergency or crisis, an AI tool should direct people to appropriate human help rather than attempt to manage the situation itself.

Responsible AI should expand caregivers' choices, not narrow them. Caregivers who prefer a phone call, text message, email, or direct human support should not be disadvantaged because they choose not to use AI. At the same time, technology can expand access for caregivers who may not have nearby services, a strong support network, or easy access to trusted information. Responsible AI should help close those gaps rather than deepen them.

2. TRUSTED, RELIABLE AND VERIFIABLE INFORMATION

Caregivers should not have to become fact-checkers for AI. Organizations offering caregiving AI should make important information easy to trace to reliable sources and straightforward for caregivers to verify.

The responsibility begins with organizations that build or offer these tools: consequential guidance should be grounded in authoritative, current sources; sources and limitations should be visible where feasible; and organizations offering caregiving AI tools should maintain appropriate traceability of the systems and data sources they deploy. Uncertainty should be stated honestly.

For Caregiver Action Network, that means prioritizing trusted, vetted caregiving resources, CAN's own evidence, expertise, and lived experience, and appropriately sourced caregiver insights rather than treating the open internet as an unquestioned source of truth.

AI can confidently provide incomplete, outdated, or incorrect information and may not reliably flag its own errors. For consequential information—particularly about health, medications, benefits, finances, or legal decisions—organizations offering the tool should define when it must use authoritative sources, decline to answer, or direct caregivers to a primary source, appropriate human review, or a qualified professional. It should not shift verification entirely to caregivers, present generalized information as individualized guidance, or imply that it is an authoritative system of record.

Questions requiring human judgment, emotional support, or expertise beyond the tool's role should lead to a qualified person, such as CAN's human-led National Caregiver Help Desk.

3. PRIVACY, TRANSPARENCY AND CONTROL

Caregiving often involves deeply personal health, financial, and family information about both people in the caregiving relationship. Being a caregiver does not by itself mean someone has permission to disclose or act on another person's information. AI should identify itself as AI and should never represent itself as a human being. Organizations offering these tools should clearly explain when AI is being used, what information is collected, how it will be used or shared, and whom to contact with concerns.

Responsible caregiving technology should collect only the information necessary to provide meaningful support, protect it and plan for potential misuse or breaches, and give caregivers and care recipients clear, understandable choices about collection, retention, monitoring, memory, personalization, and sharing — choices that are easy to change or withdraw. **Whether using a general-purpose chatbot or a service with AI built in, caregivers should check how their information is handled, share only what is needed, and disclose another person's information only with appropriate permission or authority.**

Those choices should include the ability to review and correct information and, where feasible and legally permissible, delete it. Voice recordings, transcripts, and information inferred about a caregiver or care recipient should be treated as personal data. Information about the caregiver and the care recipient should remain distinguishable, and users should be able to see when a tool's memory of past interactions is shaping current guidance. Use of data for research, model training, or other secondary purposes should require separate, understandable permission. **Privacy should not be something caregivers discover in the fine print after they begin using a tool.**

4. CO-DESIGN AND CONTINUOUS LEARNING

If you have met one caregiver, you have met one caregiver. **AI tools cannot be built in a vacuum and handed down to families as finished solutions. They should be co-designed with current and former caregivers and, when they affect or interact directly with people receiving care, with those people as well.** That participation must have influence: participants should understand what decisions are open to their input, how their feedback was considered, and what changed as a result. This means continuous, structured listening—through caregiver roundtables, surveys, and Help Desk interactions—to ensure that technology addresses genuine, unmet needs rather than theoretical problems. It also requires a commitment to accessibility, ensuring that AI tools are designed for caregivers across all geographic, economic, linguistic, and technological backgrounds. Accessibility also means affordability and low friction: a tool that is too expensive, too difficult to set up, or creates more work for an already overwhelmed caregiver is not truly accessible.

Co-design also does not end at launch. **Organizations offering a tool should keep listening after it is deployed and use real caregiver experience to improve it over time.** Caregivers—and care recipients when a tool affects or interacts with them—should have a defined role in ongoing evaluation and oversight: helping set the outcomes that matter, reviewing patterns of failure or harm, and informing whether a tool is ready for broader use. Success should ultimately be measured by caregivers themselves: Is the tool useful? Does it save time? Does it reduce confusion? Does it help them feel more supported and more confident in what to do next? Does it preserve the care recipient's agency, independence, and quality of life?

5. PROVE IT IN PRACTICE: EVIDENCE AND ACCOUNTABILITY

A tool should not earn caregivers' trust simply because it uses AI or because it performs well in a demonstration.

CAN will look for evidence that an AI tool is accurate, useful, accessible, and evaluated for foreseeable risks in real caregiving situations. Evaluation should examine where it fails and whether risks such as false reassurance, missed escalation, emotional dependency, surveillance, coercion, exploitation, harmful bias, or family conflict may emerge. People need a clear way to report problems, seek correction, and reach a person. Human oversight should be proportionate to potential harm: lower-risk work can use sampling, exception review, and performance monitoring, while decisions that materially affect someone's life require meaningful human oversight. The organization offering the tool remains accountable even when vendors are involved. Evidence should match the claim and should include real caregivers and care recipients where appropriate, not only controlled or synthetic tests. Evaluation should continue after launch, and material changes to models, data sources, or permitted actions should trigger re-evaluation.

Evaluation should measure what matters to caregivers and care recipients, not what is easiest for technology companies to measure. Engagement, clicks, or time spent are not enough. Where possible, evaluation should ask whether the technology reduces burden, supports accurate understanding without false reassurance, helps people find information or support, and preserves the care recipient's dignity, agency, independence, and well-being.

Caregivers should be partners in testing and improving technology, not unprotected test subjects for tools that have not been responsibly evaluated.



How CAN Will Use the Humanity Baseline

EDUCATE AND EQUIP CAREGIVERS:

Offer practical education on where AI may help or be wrong, what not to share, how to check important information, and when to seek human or professional support. Use what CAN learns from caregivers and future training opportunities to explore a separate, plain-language guide for using AI in caregiving.

GUIDE CAN'S OWN ENGAGEMENT AND PILOTS:

Use these principles when considering how CAN engages with AI and any future AI-enabled programs. Such tools may complement – but not replace – CAN's human-led National Caregiver Help Desk by connecting caregivers with trusted information and appropriate human support.

EVALUATE POTENTIAL PARTNERSHIPS AND TOOLS:

Apply the baseline when considering technology partnerships, CIC collaborations, and tools CAN may feature through our platforms and programs. At this stage, the Humanity Baseline is a decision-making lens, not a certification or seal of approval. CAN will use it to ask consistent questions and determine whether a tool or partnership warrants further review, additional safeguards, revision, or a decision not to proceed.

KEEP LEARNING FROM CAREGIVERS:

Use CAN's surveys, roundtables, human-led Help Desk experience, CIC, and CARE2030 partners to understand caregiver needs and concerns across tools and settings – not just whether they adopt a particular product – and refine the framework.

As a future phase, CAN may translate these principles into a risk-tiered review template, from light-touch screening to deeper assessment for higher-risk tools, beginning with a defined CAN use case. Any such review process would be tested and refined before broader use and would not constitute regulatory certification.

A Call to the Ecosystem

Caregivers should not have to choose between innovation and trust. AI can make caregiving meaningfully easier, but the question for any tool is whether it does so while preserving the humanity, autonomy, and voice of both the caregiver and the person receiving care – and whether it is built around the realities of people's lives.

That is why CAN is putting this stake in the ground now. We invite caregivers, CIC members, CARE2030 partners, researchers, innovators, and other leaders in the field to help us test, challenge, and strengthen this framework as the technology evolves.

Our role is to help caregivers navigate an increasingly AI-enabled world with greater confidence: to understand where these tools can and cannot help, assess information and tools more critically, and know when human support is needed. The goal is not to slow innovation. It is to make sure innovation moves in a direction that genuinely helps caregivers.