

# Care AI — Bring Human Wisdom Within Reach

Joe Gleason · Advance review draft 3.7 — October 2026

## At a Glance

**The central idea:** AI should help people make better decisions—not quietly make the decisions for them.

**What you'll explore:** How AI can help people organize information, prepare questions, explain approved guidance, develop content, and preserve useful experience. Follow Sam as he prepares for a family meeting without allowing a summary to erase the differences between Pat's and Ellen's perspectives.


**Design and AI:** Give an assistant a defined purpose, appropriate knowledge, clear rules, and a place within a human service. See how a staff member's explanation can become a reviewed article, video, question, or planning resource while retaining its meaning and authorship.

**Put it to work:** Create an **Assistant Role Card** and a **Reviewed-Knowledge Workflow** that identify sources, permissions, reviewers, human response routes, and how corrections become improvements.

**Evidence and evaluation:** Examine source-grounded generation, early caregiver-chatbot research, clinical drafting experience, and responsible-AI guidance. Evaluate accuracy, usefulness, review effort, privacy, access, and actual support—not the fluency or volume of generated answers.

*“One of the most important potential outgrowths of AI in medicine is the gift of time.”*

— Eric Topol, MD, physician-scientist, in *Deep Medicine*, quoted by Scripps Research, March 12, 2019.<sup>69</sup>



**Figure 10 · A draft that keeps the sources separate. Concept illustration of preparation for human review, not a product screenshot.**

Pat's account remains a distinct source in the draft. Ellen's account remains a distinct source rather than being merged with Pat's. Staff observations retain their source and status in the draft. Clinical guidance remains distinguishable from observations and family accounts. Approved knowledge supplies relevant grounding for the draft. A responsible human reviewer checks the draft before it is used. The reviewer corrects unsupported statements, omissions and lost differences. AI prepares rather than makes the care decision. In this story Sam is the social worker; clinical responsibility stays with the appropriate clinicians.

#### **Pat's account**

What Pat reports.

#### **Ellen's account**

What Ellen reports.

#### **Staff observations**

What staff actually observed.

#### **Clinical guidance**

Approved clinical information.

**AI-assisted, source-linked draft → Human review → Conversation and chosen next steps**

Keep differences visible. An assistant does not decide whose account is right.

### **In Context — Getting ready to listen**

Before the next meeting, Sam looked again at the question Ellen had raised.  
Some evenings were becoming harder to organize.

### In Context — Getting ready to listen

Pat had described it differently. He thought they were both tired of talking about what needed doing.

The Day Center had information about a different part of the day: activities Pat had enjoyed, occasions when he had chosen somewhere quieter, and the support the team had provided.

There was also the unresolved request for an additional Tuesday. The latest information available to Sam did not contain a confirmation. That needed checking, not an assumption.

Each item mattered. None explained all the others.

Sam wanted the meeting to begin with the family's concerns, not with everyone reconstructing the record.

He explained the preparation process to Pat and Ellen.

"With your agreement, I can use an assistant to help organize the information we have chosen for this conversation. I will check the draft before we use it. It won't decide what the evenings mean or what support you should accept."

"Will it make us agree?" Pat asked.

"No."

"Then it has a reasonable chance."

Ellen smiled.

"I just want my question to be there when we start."

"That is what we are preparing for," Sam said.

The useful task was not to produce a comprehensive account of the family.

It was to help the people arrive ready to listen.

### A gift when the purpose comes first

The original CarePhysics framework introduced Care AI as assistance intended to make care more personal and reduce repetitive work. This chapter develops that intention into specific tasks, responsibilities, and tests of usefulness—not a promise that adding AI automatically improves care.

**AI should help people make better decisions—not quietly make the decisions for them.**

That boundary leaves a great deal of room for useful work.

An assistant could help a social worker prepare for a conversation, a family member

formulate a question, a nurse draft an explanation, or a Day Center team turn a practical insight into training. It could help an organization compare possible approaches before investing in new content or software.

The opportunity is not simply faster writing. It is making relevant knowledge easier to work with.

When properly applied—with appropriate sources, guidelines, rules, and accountable people—AI can be a gift to a care organization. It can take on some of the repeated searching, assembling, formatting, and drafting that surrounds care, leaving people with a better starting point for the work only they can do.

That is an opportunity to test and protect.

Time saved in preparing a summary should not automatically become an expectation to handle more families. A useful explanation should not become a reason to eliminate human contact. The organization must decide what it wants the assistance to make possible.

For Sam, the purpose is a better-prepared conversation.

For Ellen, it is having her concern represented accurately.

For Pat, it includes being able to disagree without disappearing from the account.

## Care AI is more than a chatbot

A conversational assistant is one way to provide help. It is not the whole idea.

Some assistance happens before anyone opens a conversation: organizing selected notes, locating approved guidance, preparing an article outline, or identifying information that needs confirmation.

Other assistance happens while someone is thinking: explaining a term, comparing appropriate options, rehearsing wording, or helping turn a broad concern into a question.

Still other work follows the encounter: preparing an agreed summary, organizing feedback, or drafting a revision for the responsible team to review.

Summarization, preparation, navigation, accessible explanations, and routes to people are useful tasks to examine. The evidence for each task differs; none should silently expand into autonomous care.

A drafting assistant also differs from a system authorized to take action.

Preparing a message does not send it. Identifying a possible appointment does not book it. Finding a service does not establish eligibility or availability.

Where a configured system can perform a bounded action, the organization must define the permission, confirmation, and record required. The assistant should never announce that something happened merely because it generated the words describing it.

These distinctions make the service easier to understand. They also help the team decide

where AI is useful and where a simpler tool—or a person—is the better choice.

## Give the assistant knowledge with a clear purpose

The quality of the assignment begins before the prompt.

What information is relevant? Who supplied it? What kind of authority does it have? May it be used for this person and this purpose?

At Genus, we assembled the research and working materials behind CarePhysics to help guide our own development, training, content, and support. The dated Day Center export describes that shared reference foundation for teams, wording decisions, and AI-assisted work. This is an internal working approach, not proof that every resulting feature or answer improves care.

Other organizations can build on the same idea without adopting the same services or technology.

CarePhysics supplies questions and design principles. Your organization supplies its current services, professional guidance, terminology, values, and hard-earned experience. The people you serve help establish what fits.

Those materials should not become one undifferentiated collection.

| Kind of knowledge                    | What it contributes                                                     | What must remain clear                                                                                     |
|--------------------------------------|-------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Research evidence</b>             | Findings and models relevant to the question                            | Source, population, study design, findings, uncertainty, and limits                                        |
| <b>Design and practice guidance</b>  | Ways to explain, demonstrate, invite, support, and review               | Whether it is professional guidance, an organizational rule, or a proposed design approach                 |
| <b>Local organizational facts</b>    | Services, contacts, eligibility, hours, capacity, and responsibilities  | Owner, date checked, applicable setting, and what requires confirmation                                    |
| <b>Illustrative material</b>         | Fictional scenes, sample questions, demonstrations, and draft templates | It is an example—not a real person's history, a testimonial, or an outcome                                 |
| <b>Permissioned personal context</b> | The information appropriate to an individual task                       | Who may access it, why it is needed, what may be shared, and when permission or relevance must be reviewed |

This separation is part of the proposed CarePhysics knowledge-bank approach. A

research result should not become a claim about a local program. A fictional scene should not become a fact about a family. A private concern should not become reusable training material.

An assistant does not need everything the organization possesses.

It needs the appropriate material for the job.

## Retrieval is not training—and a source is not a guarantee

One way to support an assistant is to retrieve relevant passages from an approved collection and make them available when it prepares a response. This is often called **retrieval-augmented generation**: finding material and using it to support generated text. Research has demonstrated the approach on language and question-answering tasks; that does not establish clinical reliability.<sup>70</sup>

The distinction is practical.

Making this chapter available for a response is not the same as retraining the underlying model. Correcting an answer does not, by itself, mean the system will remember that correction next week. Whether conversations are stored or reused for other purposes is another question the organization must resolve.

A maintained source collection, an implemented retrieval system, and a tested care service are different things.

Even retrieved material can be wrong for the task. It may be outdated, apply to a different setting, omit an important qualification, or be interpreted incorrectly.

Ask the assistant to identify the sources supporting its statements. Then make it possible for the reviewer to inspect them. A citation that does not support the sentence is not useful evidence.

When sources conflict, the assistant should expose the conflict and refer it to the designated authority. It should not silently combine two incompatible instructions or assume that the most recent informal note overrides a clinician-approved plan.

The goal is not an impressive stack of references.

It is a traceable explanation that the responsible person can check.

## Prepare a summary without settling the disagreement

### In Context — The difference belongs in the draft

For the meeting preparation, Sam selected the information appropriate to the task. Pat and Ellen knew what would be included. Private material outside that agreement was excluded.

The proposed assistant produced a short draft, with links to the selected records and their dates.

| Source                                                      | What the draft should preserve                                                                                                                                        |
|-------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ellen's question</b>                                     | Evenings feel harder to organize. She wants help explaining what remains to be done and considering support.                                                          |
| <b>Pat's perspective</b>                                    | He is tired of evenings being occupied by conversations about care tasks. He wants more ordinary time together.                                                       |
| <b>Reviewed Day Center information</b>                      | Pat has enjoyed reading and discussion and has sometimes chosen a quieter setting. These observations concern the program day, not evenings at home.                  |
| <b>Relevant clinical material selected for this meeting</b> | No new clinical explanation of this particular evening concern is established in the selected material. Questions requiring assessment belong with the clinical team. |
| <b>Additional Tuesday request</b>                           | The latest selected record contains no confirmation. Sam must verify the current position rather than infer what happened.                                            |

### In context

One sentence in the draft read:

“The family agrees that Pat’s evenings are becoming difficult.”

Sam removed it.

The family had not agreed on that description. Ellen had raised a concern about organizing the evenings. Pat had described something he wanted less of.

Sam replaced it with:

“Pat and Ellen describe different concerns about evenings. Discuss both before proposing a change.”

He also checked that the Day Center’s observations were not being used to explain something its staff had not observed.

The draft was now more useful because it held less certainty.

When the meeting began, Sam did not read the table as though it were a verdict.

“Ellen, is this still the question you want to start with?”

“Yes,” she said. “By evening, there are still little arrangements to make.”

**In context**

“And Pat?”

“I would like an evening that isn’t about arrangements.”

“Then perhaps we should look at which work could happen earlier, or be shared differently,” Sam said. “And whether there are separate questions for the clinical team.”

The assistant had not solved the issue.

It had helped keep the important parts available for the people who could discuss it.

**Human review is work, not a button**

A summary can be grammatically correct and still be wrong about a person.

Review must examine the sources, not merely polish the output. Has the draft attributed statements correctly? Preserved uncertainty? Included a concern that changes the meaning? Invented agreement? Turned a suggestion into an accepted responsibility?

The reviewer also checks what should not appear.

The Day Center export describes meeting preparation from staff-selected evidence, with private notes excluded by default and drafts reviewed before use. That is a documented workflow description, not independent evidence that every review catches every error.

Permission needs to be enforced before inappropriate material enters the task. Do not send private notes to an assistant and rely only on a sentence telling it not to reveal them.

A technical control that allows a note to be included does not itself establish that inclusion is authorized or appropriate.

Review also requires time and expertise. Sam can review the support-planning summary within his role. Lena confirms program information. A clinical explanation needs the relevant clinical authority.

WHO identifies **automation bias** as a risk: people may overlook errors or delegate difficult judgments because a system presents a persuasive answer. Its guidance calls for defined tasks, stakeholder involvement, and continuing assessment.<sup>1</sup>

Make disagreement with the draft easy. A rejected sentence is not a failed worker or a disloyal user. It may be the review process doing exactly what it should.

**Turn a useful explanation into shared knowledge**

AI can also help preserve knowledge that would otherwise remain in one conversation.

A staff member finds language that makes a difficult idea clearer. A family points out

why an invitation felt pressuring. A volunteer discovers an effective way to demonstrate a task. The team can examine that contribution and decide whether it belongs in a reviewed resource.

The opportunity is to capture the lesson—not automatically copy the personal conversation.

### **In Context — Lena's words become a starting point**

After several discussions about activities, Lena explained her approach to the team.

“An invitation should leave room for a different answer. Someone can enjoy listening without wanting to lead.”

A colleague asked whether that idea belonged in the family guide.

“It might,” Lena replied. “But I don’t want it to sound as though we can offer anything at any time.”

That qualification mattered too.

With Lena’s agreement, the team prepared a generic brief. It contained her explanation, the center’s actual activity options, the relevant CarePhysics design guidance, and the route for families to ask questions.

It contained no private family notes.

The assistant’s assignment was not “write something warm about participation.” It was more specific:

Prepare a short article and demonstration outline explaining how participants can discuss activity preferences with our team. Preserve the options and limits in our approved information. Make clear that contribution is optional. Identify questions our team must answer before publication.

The first draft said that participants could always leave any activity and go elsewhere independently.

Lena corrected it. The center could offer choices, but staff still needed to provide appropriate assistance and work within the service’s arrangements.

The revised wording made the support visible instead of presenting choice as the absence of responsibility.

That is a useful form of knowledge capture: a human insight, its context, its limits, and a reviewed way to share it.

The dated export describes team-refined examples and explicit drafting rules for family summaries. It provides a basis for this approach, but it does not establish that every correction automatically changes future model behavior. That requires a defined improvement process.

## Let one idea work across several formats

Once the idea is reviewed, different formats can serve different needs.

A short article might begin:

## There is more than one way to take part

You can tell our team what interests you and what does not. Joining a discussion, listening, trying part of an activity, or asking about a quieter option are all things we can talk through.

You do not have to lead an activity to belong here.

For example: “I would like to listen first. Could someone help me find a comfortable place?”

Staff will explain the options available and any assistance needed. Families can also share questions through the center’s stated contact route.

A video could demonstrate the conversation. Show someone asking to listen, a staff member explaining the available choices, and the person choosing. Do not let the demonstration promise an option the program cannot provide.

A feedback question might ask whether the explanation made those choices clear. It should identify who will review the answer and provide a way to request a response.

A feature-planning brief could ask how the interface distinguishes an offered activity, an expressed preference, and a support arrangement that still needs confirmation.

The same knowledge now informs content, training, and software design without becoming a product catalog.

The recognizable pattern remains: why it matters, the main idea, a demonstration, a chosen next step, questions or connection, and what follows. AI should help carry that design consistently—not make every article or message sound identical.

Clinical instructions and validated instruments require a different discipline. Do not rewrite their essential wording, conditions, or scoring merely to make them fit the template.

And sometimes the useful answer is to create less. A shorter message with a dependable human contact may serve the person better than another complete educational package.

## Personalize through what people tell us

Personalized assistance begins with relevant, permissioned knowledge—not a confident guess.

Ask about language, preferred depth, format, culture, beliefs, the people someone wants involved, and what practical support is available. A person may welcome faith-related material, prefer secular guidance, or want neither introduced into this particular discussion.

Do not infer these preferences from a name, address, age, or family role.

The CarePhysics writing guidance explicitly separates cultural humility from demographic assumptions and calls for learning how this person and family want care organized.

An assistant could help adapt a reviewed explanation for a person who requests less detail, a different format, or another language. The appropriate reviewers still check meaning and suitability. Generated translation should not be treated as qualified interpretation for consequential care conversations.

Keep the facts stable while adapting the expression.

Ellen's preferred brief summary should not omit an unresolved concern. Maya's request for more detail should not grant access to information that was not shared with her. Pat's interest in reading should not become a standing assignment to lead.

Purpose comes from the person. Personalization should help them pursue it, not make the organization's preferred action more difficult to decline.

## Questions do not keep office hours

A family member may remember a question after the service has closed. They may want to revisit an explanation without feeling embarrassed about asking again.

A properly configured assistant could make approved routine information available around the clock. It could help prepare questions, explain the organization's process, or offer a relevant resource within its defined scope.

That can extend support between conversations. It can also help someone arrive at the next conversation knowing what they want to ask.

Consider Ellen using an approved question-preparation tool:

"Help me explain the evening problem without making it sound as though Pat is the problem."

The assistant might offer:

"By evening, I still have arrangements to finish. I would like us to look at what could happen earlier or be shared, while keeping time for us to enjoy together."

Ellen can change it, reject it, or bring it to Sam in her own words. The draft does not automatically become a message to her children.

Or she might ask why Pat has become tired in the evening. The assistant should not diagnose from the conversation. It can explain that the available information does not establish a cause and help her contact the appropriate professional.

Around-the-clock information is not around-the-clock professional care.

State when people respond, provide the approved route for urgent concerns, and make human help available without requiring a person to persuade the chatbot that their question deserves attention.

If a request is transmitted, report that accurately. Do not say that Sam is handling it until he or the responsible service has accepted it.

The intended outcome is a useful answer or human connection—not a longer conversation with the assistant.

## Give the assistant a role people can understand

Before introducing a new assistant, complete an **Assistant Role Card**.

This is an organizational planning tool, not a certification. It makes the proposed job inspectable by the team and understandable to the people affected.

| Role-card element                       | Example: Family Meeting Preparation Assistant                                                                                                                   |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Human purpose</b>                    | Help people arrive ready to discuss their questions, without reconstructing every record during the meeting.                                                    |
| <b>Authorized users and information</b> | Named staff roles, approved guidance, verified local information, and selected personal material authorized for this purpose.                                   |
| <b>Permitted work</b>                   | Retrieve, organize, draft questions, identify missing information, and prepare source-linked summaries.                                                         |
| <b>Review and decisions</b>             | Sam reviews support-planning material; Lena confirms Day Center facts; clinical questions go to the relevant clinician.                                         |
| <b>Action limits</b>                    | No independent diagnosis, treatment change, capacity judgment, family invitation, disclosure, booking, or declaration that a service occurred.                  |
| <b>Privacy and sharing</b>              | Exclude information outside the task's permissions. Produce separate versions only when appropriate. Explain storage, access, retention, and any secondary use. |
| <b>Uncertainty and help</b>             | Preserve disagreements, identify unsupported statements, and provide the actual human response route and urgent-care instructions.                              |

| Role-card element               | Example: Family Meeting Preparation Assistant                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ownership and evaluation</b> | Name the service owner, backup, review date, configuration version, benefit measures, burden checks, and conditions for stopping or redesigning. |

A separate family-facing assistant needs its own card. Permission to prepare an internal draft is not permission to answer every family question directly.

The proposed role-card approach follows the book’s source specification: define the task, information, permitted output, reviewer, boundaries, and response when something may be wrong.

Explain the role plainly to users. They should know when they are interacting with AI, which organization stands behind the service, and what a person has or has not reviewed.

**Make corrections travel through a deliberate process**

An error corrected today can help improve tomorrow’s work—but only if someone carries the lesson forward.

The team needs to distinguish correcting one record from changing a reusable instruction.

Sam’s correction preserves this family’s different perspectives. A broader design lesson may be that the assistant needs a clearer rule against converting multiple viewpoints into agreement.

The owner can revise that instruction, prepare suitable test cases, and examine the new behavior. The updated version should be checked against both ordinary examples and cases involving disagreement, missing information, or restricted material.

Keep contributors credited. Record what changed, who approved it, and what happened when it was tested.

Do not automatically turn a family’s private exchange into an example for other organizations. A generic or carefully authorized example may communicate the lesson without carrying the private story.

Likewise, “the system learns from feedback” should describe a real mechanism. It might mean revised guidance, changed retrieval, an updated template, a corrected local fact, or a separately governed model-training process. Those are not interchangeable.

The dated Day Center materials propose a register of findings, decisions, changes, and remeasurement—and explicitly say that register is maintained outside the product described there. A promise to learn is not evidence that learning occurred.

## Give people the means to provide oversight

An assistant needs an organization around it.

Someone must maintain the sources, review the work, cover absences, respond to questions, investigate errors, and decide whether a task remains appropriate.

Budget for those responsibilities. Include training, language and accessibility review, technical support, security, evaluation, and the cost of corrections—not only model access.

Give reviewers practical training using fictional or otherwise appropriately authorized examples. Include omissions and plausible mistakes, not just obviously wrong answers. Test whether people can find the source and reject the draft.

Do not introduce fabricated errors into actual family records as a staff test.

Agree with technology suppliers on where information goes, who can access it, how long it remains, and whether it may be used for other purposes. Verify those arrangements for the actual workflow rather than assuming that one general privacy statement covers every use.

A more capable model or a changed source collection can change the experience. Recheck affected tasks before treating the updated configuration as equivalent to the previous one.

NIST's AI Risk Management Framework organizes this ongoing work through **Govern, Map, Measure, and Manage**. It is a risk-management framework, not proof that a particular assistant is safe or effective.<sup>71</sup>

The same approach can serve different settings. A health system might begin with clinician-reviewed message drafting. A home-care organization might prepare training from approved procedures. A Day Center might organize meeting questions. A community network might maintain reviewed resource descriptions.

Genus supplies technology. Partners remain responsible for their services, staff, programs, professional decisions, and voice.

AI can help reveal that a resource is missing. It cannot replace unavailable transport, staffing, respite, or clinical care with reassuring prose.

## How would we know it helped?

Begin with one defined task and the process used before AI assistance.

For meeting preparation, measure the whole task: selecting information, producing the draft, checking sources, correcting it, and preparing the approved version. Fast generation is only one part.

The Day Center research materials propose five distinct evaluation layers: draft quality, human-review quality, workflow efficiency, family impact, and safety and fairness. A strong result in one does not establish the others.

For an initial local test, use a defined set of appropriately authorized cases and a stated review period. Record the model, instructions, source collection, and workflow version.

Check whether statements are supported, important information is missing, and disagreements survive summarization. Distinguish errors caught during review from errors that remain afterward.

Then ask whether the output helped. Did the family's questions appear accurately? Did the meeting address them? Was the explanation understandable? Did a requested human response actually occur?

Count the effort required from staff and families. A tool may reduce blank-page effort without saving minutes. It may save staff time while transferring work to relatives. Both differences matter.

Examine access and performance across the languages, formats, and situations the service intends to support. Include people who decline AI assistance and keep appropriate human alternatives available.

A high approval rate is not enough. It could reflect good drafts, rushed review, or overreliance. Sample approved outputs and give a named leader authority to pause or narrow a workflow when serious errors, privacy failures, or unacceptable burden appear.

Separate reach, understanding, first action, adoption, useful participation, outcomes, and attributable impact. Willingness to recommend, actual referrals, and services received by referred people remain different questions.

We are not measuring whether the assistant sounds caring.

We are asking whether its use helps people provide and receive better support.

## Models and Evidence Behind This Chapter

The following research examines components of the proposed workflow: retrieval, clinical drafting, caregiver assistance, and governance. The population and task in each source define what its results can support.

### Source-grounded generation

Give the assistant relevant material to work with

Patrick Lewis and colleagues' 2020 NeurIPS paper combined a language-generation model with retrieval from an external knowledge collection. Across knowledge-intensive language tasks, the system produced improvements over comparison models, including more specific and factual generated language in the tasks studied.<sup>70</sup>

**Where we used it:** Retrieving relevant guidance and selected records before preparing a draft.

**Evidence boundary:** These were computational language tasks, not care encounters. The paper does not show that retrieving documents guarantees faithful summaries, correct

citations, privacy, or clinical safety.

**Local question:** Did the assistant retrieve the right material and use it accurately—not merely include a source label?

### A purpose-built caregiver chatbot

Early acceptability is encouraging, but it is not outcome evidence

Sheung-Tak Cheng and Peter H. F. Ng's 2025 **PDC30 Chatbot** study developed an assistant around a dementia-caregiving guidebook and defined behavioral instructions. The researchers compared responses to 21 common questions and then conducted a two-week acceptability study with **21 family caregivers**.

Participants generally rated it favorably for usability and helpfulness. They had been asked to use it at least daily during the trial.<sup>72</sup>

**Where we used it:** The proposal to combine a defined role with selected caregiving knowledge.

**Evidence boundary:** The small, short acceptability study did not establish reduced caregiver burden, better clinical outcomes, or lasting benefit. A guidebook's prior evidence does not automatically validate a chatbot using it. The research review informing this book classifies direct caregiver-AI evidence as **Emerging**.

**Local question:** Is the support understandable and useful, and what remains unanswered or requires a person?

### AI-drafted clinical messages

Usefulness and time saved are separate outcomes

Patricia Garcia and colleagues' 2024 study evaluated AI-drafted patient-message replies at Stanford Health Care. The **five-week, single-group quality-improvement study included 162 clinicians; 73 completed both surveys**.

Mean draft use across clinicians was about **20 percent**. Survey findings suggested reduced perceived task burden and work exhaustion, but measured reading, writing, and reply times did not improve.<sup>73</sup>

**Where we used it:** Counting review and correction work and examining usefulness separately from speed.

**Evidence boundary:** There was no randomized control group, follow-up was short, and survey participation was incomplete. This was a particular clinical workflow, not a Day Center trial or evidence about Genus.

**Local question:** Does the assistance improve the task, reduce effort, save time, or create some combination of benefits and costs?

### WHO and NIST guidance

Responsibility remains an ongoing organizational task

WHO's guidance on generative health AI identifies potential uses in administration, education, research, and care while warning about inaccurate outputs, bias, privacy risks, and inappropriate delegation. It recommends defined tasks and involvement of the people affected. NIST's AI RMF and Generative AI Profile provide structures for evaluating and managing risks across the system's use.<sup>1</sup>

**Where we used them:** The role card, source permissions, review responsibilities, testing, and authority to stop or redesign.

**Evidence boundary:** These are governance frameworks, not outcome studies or certifications of the proposed assistant.

**Local question:** Who is responsible when the answer is wrong, the information is inappropriate, or the promised human response does not arrive?

## Behavioral and social-learning foundations

Use knowledge to support people's choices

COM-B helps us investigate whether the barrier is understanding, opportunity, motivation, or a combination. Self-Determination Theory keeps choice, capability, and relationships in view. Social Physics directs attention to how useful knowledge travels between people.

The research review informing this book gives these different evidence judgments: strong conceptual support for COM-B, moderate intervention evidence for Self-Determination Theory, and emerging-to-moderate support for the broader Social Physics perspective. These grades do not transfer to an AI application simply because it names the models.

**Where we used them:** Helping Ellen formulate a chosen question, preserving Pat's perspective, and carrying Lena's reviewed explanation into resources others can use.

**Local question:** Does the assistance strengthen human understanding and connection—or merely increase interaction with the system?

## One thing to try

Choose one repetitive preparation task that your team understands well.

Use fictional or appropriately authorized material to prepare a source packet and an Assistant Role Card. Ask for one draft. Have the responsible person check it against the sources, including what was omitted.

Record what helped, what required correction, and the total effort.

Then complete this sentence:

**“The assistant prepares \_\_\_\_\_ so that \_\_\_\_\_ can spend more attention on \_\_\_\_\_.”**

If the final blank describes only producing more output, reconsider the purpose.

## A better question is a beginning

### In Context — “That is what I meant”

During the meeting, Ellen looked at the revised summary.

“That is what I meant,” she said. “Not that every evening is bad. There are just things still waiting for me.”

Pat read his part.

“And I would like fewer meetings after supper.”

Sam kept both statements.

They began discussing which arrangements needed attention, what could be shared, and which questions belonged with the clinical team. They did not assign new responsibilities to Maya or Daniel without asking them.

The additional Tuesday remained a matter to verify through the people already responsible. It did not become confirmed because the summary was neatly organized.

Before they finished, Sam asked what had been useful about preparing this way.

“We got to the question,” Ellen said.

Pat looked at the page.

“And I survived being summarized.”

The assistant had helped with preparation. The conversation still belonged to them.

Whether this approach would continue to help—and whether it was worth the work required—needed more than an encouraging moment. It needed careful listening, appropriate measures, and a willingness to change what did not serve people.

That is the final principle.

## Notes

1. World Health Organization (2024). WHO releases AI ethics and governance guidance for large multi-modal models. January 18. [Source](#)

[Return to text 1](#) · [Return to text 2](#)

69. Scripps Research (2019). Eric Topol pens book on artificial intelligence in medicine. March 12. Institutional account of Deep Medicine. [Source](#)

[Return to text](#)

70. Lewis P, Perez E, Piktus A, et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Advances in Neural Information Processing Systems. 33:9459–9474. [Source](#)

[Return to text 1](#) · [Return to text 2](#)

**71.** National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. DOI: 10.6028/NIST.AI.100-1. [Source](#)

[Return to text](#)

**72.** Cheng ST, Ng PHF (2025). The PDC30 Chatbot—Development of a Psychoeducational Resource on Dementia Caregiving Among Family Caregivers: Mixed Methods Acceptability Study. JMIR Aging. 8:e63715. DOI: 10.2196/63715. [Source](#)

[Return to text](#)

**73.** Garcia P, Ma SP, Shah S, et al. (2024). Artificial Intelligence–Generated Draft Replies to Patient Inbox Messages. JAMA Network Open. 7(3):e243201. DOI: 10.1001/jamanetworkopen.2024.3201. [Source](#)

[Return to text](#)