BAST

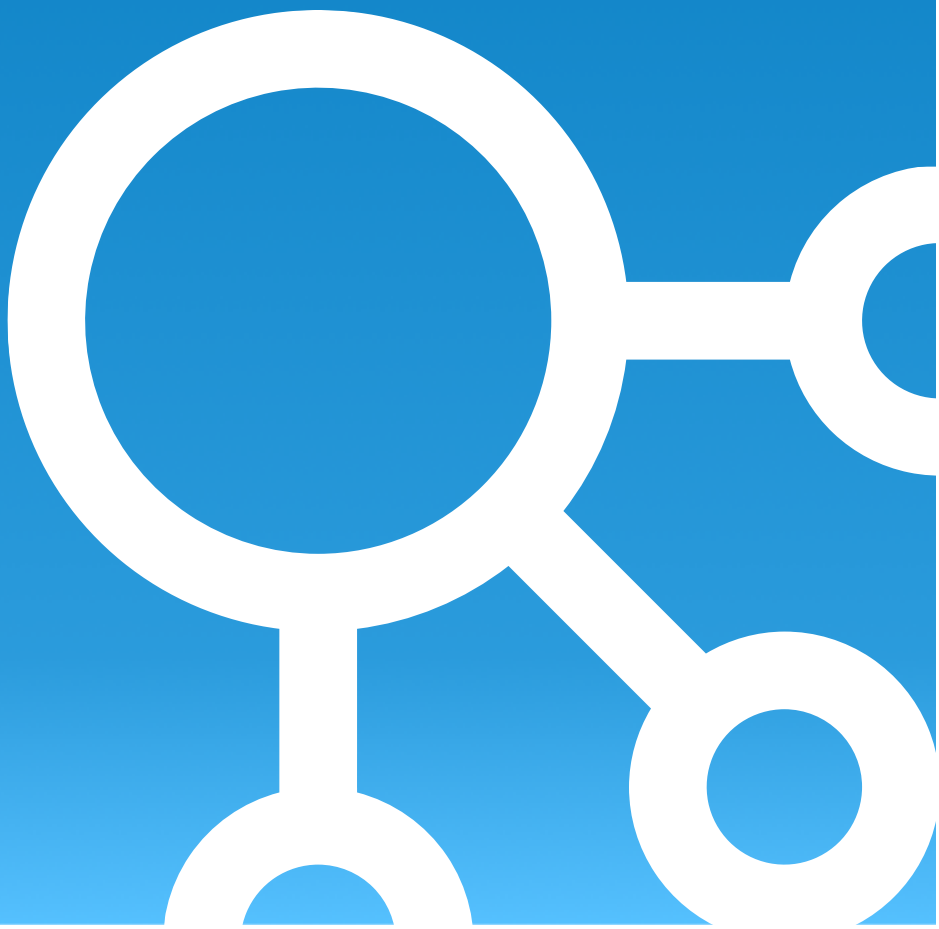

bast.ai

# When Trusted Knowledge Learns to Speak

A look at how institutions can turn expertise into continuous, collective intelligence

## The gap we pretend isn't there.

If you spend any time in healthcare, you start to notice a strange disconnect.

On paper, the system has never been richer in information. Every condition, procedure, and protocol is described somewhere, usually in painstaking detail. Hospitals hold discharge instructions, care pathways, caregiver guides, rehabilitation protocols, pharmacist notes, clinical advisories, and community health resources. Each was shaped by someone who knows what it feels like to carry responsibility for a patient's safety.

And yet, if you really talk to patients and caregivers, you hear something entirely different.

They describe leaving appointments with a plan they do not fully understand. Instructions they are not sure they can carry out. A handful of printed pages that do not quite match the situation waiting for them at home. A sense that they should know more than they do, remember more than they can, and ask fewer questions than they need to.

They are trying to follow guidance while working, caring for children, navigating fear, or healing themselves. They are doing their best in the quiet moments between visits, where things feel fuzzier and lonelier.

Clinicians feel this disconnect too. They hear the repeated questions, misunderstandings, and missed steps. They know the gap is not about intelligence. It is about context. They are being asked to communicate nuance through short interactions, fragmented systems, and layers of administrative noise to people who may be overwhelmed, anxious, ashamed to admit confusion, or simply exhausted.

And because no one wants to say it aloud, everyone quietly adapts to a reality in which care plans fall apart for reasons that have nothing to do with medicine.

This is the crisis. Not the dramatic, headline-making kind. The quiet kind that erodes outcomes gently and invisibly. The kind that causes patients to disengage, caregivers to feel lost, and clinicians to carry a burden they cannot quite name.

The problem is not that people are short on information. What is missing is continuity and the context that helps a person use that information today.

## A different kind of conversation.

Picture a caregiver standing in her kitchen, staring at a list of instructions she does not quite understand.

Her son has undergone a complex neurological procedure. He is home now. Physically present, but cognitively different. She does not know what is normal, how much help to offer, or which signs mean call someone instead of stay the course.

Once, she would have waited for the next appointment, called the nurse line, or borrowed reassurance from a friend who once cared for a parent. Today, she might search online. She might scroll through anonymous answers written by strangers. She might panic. She might make a choice that feels right but is not aligned with the clinical plan.

Now picture something else.

She pulls out her phone and asks:

"How do I tell him he is making progress, even though some things are still really hard?"

What comes back is not generic advice. It is not internet wisdom. It is not an LLM guessing at something comforting.

It is guidance retrieved from the institution's vetted materials and adapted to this caregiver, this phase of recovery, and this moment. It is phrased gently, grounded in evidence, and paired with the reasons behind it.

It might say:

"At this stage, it helps to be positive and realistic. Focus on small, specific moments instead of big declarations. You could try something like, 'I saw you buttoned your shirt today. That was really good.'"

The source is visible in the conversation. The caregiver gets more than an answer. She gets steadiness - the feeling that she is supported by people who understand what her son is facing.

This is where Bast begins: at the intersection of information and support, logic and empathy, authority and presence.

## Knowledge does not live in documents

A document is not knowledge. A document holds information - a trace of what an expert knew in one setting at one point in time. Knowledge lives in humans. It emerges when the right information meets a specific person, in a specific situation, at the moment they can understand and act on it. Because knowledge emerges through context and relationship, it cannot be conscripted by an institution, a platform, or a model.

We often discover what we know only when the situation asks it of us. Bast is built for that moment. Before the LLM is asked to speak, Bast determines whether the question belongs within the expert-governed scope. If it does, Bast retrieves only the approved passages most relevant to this person's role, language, phase of care, and immediate question. If it does not, the knowledge base is never queried.

The domain expert governs the sources and the rules. The patient, caregiver, or responsible party governs the personal context they bring, including what may be remembered or shared. The LLM never chooses what knowledge to search or what is true. It receives only the narrow, relevant context it is permitted to use and helps that trusted information arrive in a form the person can use now.

Bast does not turn documents into knowledge. It creates the conditions in which the knowledge held across patients, caregivers, clinicians, and institutions can meet in the moment it is needed.

## Memory as part of care.

Healthcare does not happen in a single moment. It unfolds across weeks, months, and sometimes years. Patients revisit fears. Caregivers reevaluate decisions. Clinicians refine plans. Everyone moves at a different pace, with different emotional and cognitive needs.

A static document cannot honor that journey. A single appointment cannot carry that load. A conversation with continuity can.

With the person's permission, the relevant context can travel forward. A patient can return days later and say:

"I am still confused about the thing we discussed last week. Can we go over it again?"

The system can respond without judgment, without requiring the person to repeat the entire story, and without losing the institution's clinical voice.

Memory is not medicine by itself. But feeling recognized can help a person stay engaged with care. Continuity makes room for honesty, curiosity, and hope.

## The Craig Hospital story (in progress)

Craig Hospital, a nationally recognized leader in brain and spinal injury rehabilitation, shared something powerful with Bast: an outpatient procedure manual.

No massive dataset. No EHR integration. No private healthcare information.

Just the distilled information created by clinicians who have spent decades guiding patients and families through neurological recovery.

With that one source, a new kind of support became possible. Caregivers could ask questions that were never designed to fit neatly into a pamphlet:

- "Is it okay that he is sleeping more this week?"

- "How do I talk about fatigue without making him feel guilty?"

- "What should I encourage today if yesterday was hard?"

Bast could retrieve Craig's guidance and adapt it to the language, role, and moment of the person asking, while preserving Craig's philosophy of care.

This is still a case study in motion. Bast and Craig are experimenting, learning, shaping, and assessing. The outcomes take time, yet the signal is clear: a trusted source can become a living companion when it knows how to meet a human being in context.

Craig is proving a simple truth. Institutions do not need more data to support people. They need better pathways to deliver the information they already trust - and a way to govern the value created when that information becomes useful.

## Start with the library.

We can see the larger opportunity by stepping outside healthcare for a moment.

A librarian holds knowledge about their community that exists in no national repository. They know which questions people are afraid to ask out loud. Which explanation lands with a grandmother raising two grandchildren. Which services the county suddenly needs in February. Which words help someone take the next step toward a job, a benefit, a diagnosis, or a safer life.

Librarians built that knowledge over decades, and every word of it was produced in service of the public.

Today, the internet treats that work as raw material. Community institutions create knowledge. Platforms extract it. Model companies turn it into products. Then the institution is asked to buy access to its own intelligence in a new form.

The problem is not that knowledge moves. Public knowledge should move. The problem is that control and value move in only one direction.

The correction is simple: the steward keeps a seat at the table.

The institution decides how its knowledge is used, which uses are public, which are commercial, what must remain private, and how the value created from that knowledge returns to the people who made it possible.

That is not a case for locking knowledge away. Access and extraction are not the same thing. A library can make knowledge widely available without surrendering its authority over how that knowledge is packaged, represented, or sold. A health system can share expertise without turning patients into inventory.

## The institution hosts the model.

Imagine libraries hosting a language model inside a boundary it governs.

The model is grounded in the library's own service materials, local history, rights-cleared collections, community resource directory, and carefully maintained knowledge of the community. A resident can use it to understand a public benefit, prepare for a job interview, find a trusted health resource, learn a new skill, or simply ask the question they were embarrassed to ask at a desk.

The answers cite their sources. The model knows when to stop and invite a librarian in. The library can see where its materials are out of date, where people are getting stuck, and where a new program could help.

The library does not need to train a frontier model or operate a hyperscale data center. It needs a model whose sources, retrieval rules, permissions, provenance, usage rules, and commercial terms it controls. The underlying computing infrastructure can come from a technology partner. The intelligence remains an institutional asset.

For residents, the model can be a public service: dependable, accessible, and free at the point of need.

For approved commercial users, the same information can become a governed product. A public agency, software company, research group, or AI builder might pay to query a trusted local information service, test a model against an expert-curated benchmark, or license a clearly bounded body of institution-owned or rights-cleared content. Every use is visible. Every source is attributed. Access can expire or be revoked. The underlying corpus does not have to be handed over.

The revenue is not generated by selling residents' questions. It comes from licensing institution-authored information and controlled access to the contextual intelligence built around it.

That is the revGen model: community use is treated as mission; commercial use is priced; and the proceeds sustain the mission. Revenue can fund librarians, digital literacy, new collections, longer hours, stronger governance, and the next generation of community knowledge.

## The same architecture belongs in healthcare.

Every health system already has sources: discharge instructions, care pathways, caregiver guides, rehabilitation protocols, pharmacist notes, nurse education, and community health resources.

A hospital hosts a model grounded in the guidance it already trusts. Patients and caregivers use it between appointments. Clinicians use it to extend their reach without surrendering their voice. The institution decides what the model may remember, what it may learn from, what can leave its environment, and which outside systems may connect to it.

The public-service layer supports the people the health system exists to serve. A separate commercial layer allows approved partners to pay for governed access to the institution's expertise: versioned care guidance, evaluation tools, licensed knowledge services, and lineage-rich APIs that can support other products without exposing raw patient data.

The distinction matters. A hospital's authored clinical guidance can be an institutional asset. A patient's story is not. A steward's authority is always bounded by the rights of the people whose lives produced the data.

This changes the economics. The institution no longer has to choose between protecting knowledge and creating value from it. Good governance is what makes the knowledge valuable. Sources remain current. Permissions are explicit. Uses are traceable. Commercial access

helps pay for patient education, navigation, community health work, and the people required to keep the system trustworthy.

The model becomes a revenue-generating asset without making the community the product.

## The data is the encounter, and it flows.

The most qualified data is not another copy of the document. It is the structured record of the encounter between a human and information: how everyday language maps to an expert concept, how the same question changes for a patient or caregiver, what matters in week three instead of year two, what someone asks next, what they correct, and what finally helps them act.

That is human-generated, expert-grounded, context-rich narrative data. For a specialized model, it can be more useful than another indiscriminate sweep of the internet because it records not only what the source said, but what a person needed in order to understand it.

This direction is not fringe. Andrew Ng, Michael I. Jordan, and Margaret Mitchell describe this from different angles: data quality, provenance, context, intended use, and limitations matter as much as model scale.

Bast classifies intent and traverses the ontology before generation, sending only relevant, validated source segments to the LLM. The domain expert controls the corpus. The LLM handles the narrowest part of the work: translating trusted information into useful language.

That already reduces compute. The moonshot is larger than lower token usage. It is to make brute-force inference against a general model unnecessary for most questions. Information should be localized, governed, and retrieved where it is relevant. The end state is simple: stop building data centers to compensate for context we failed to preserve.

Broad statistical patterns still matter. But in high-stakes care, when qualified current evidence shows that the broad prior does not fit the person in front of us, local context must control the decision. This is not personalization for convenience. Healthcare is not maintenance. It is often survival. No one should be controlled by a distant statistical average built from the very data their life helped create.

The incentive model changes with the architecture. Institutions earn by maintaining trusted, relevant knowledge and licensing governed access to it - not by feeding every human interaction back into a global machine.

## Collective Intelligence is the utility.

Collective intelligence is not a bigger brain or a larger database. It is a community's capacity to sense, connect, learn, and respond across different contexts.

Dave Snowden's work begins from the reality that people often know more than they can say, say more than they can write down, and discover what they know in the context of need. Nora Bateson's Warm Data keeps another truth in view: a living system cannot be understood by removing information from the relationships that give it meaning. The family, the clinic, the culture, the institution, and the moment all shape what can be known and what can be done.

Bast is connective tissue for that collective intelligence. It brings relevant, sourced information into a live human situation. It lets questions and corrections return to domain experts as governed signals. It helps an institution notice patterns without requiring every story to be centralized or every contradiction to be flattened into one answer.

The models are not the community's intelligence. The community is. The models help that intelligence find itself when it is needed.

Calling this a public utility does not mean making every dataset public or giving everyone unrestricted access.

It means the intelligence is operated as dependable civic infrastructure. People can reach it when they need it. They can see whose knowledge grounds it. They can challenge an answer, reach a human, and understand the rules under which it operates. The institution maintaining it is local, legible, and accountable.

A library board meets in public. A nonprofit hospital reports on community benefit. Both have duties that a distant annotation vendor or general-purpose model provider does not. Neither is perfect, but each can be called to answer by the people it serves.

That accountability is part of the product and nothing in that loop requires selling a person's conversation.

## Consent is the architecture.

This model only works if it keeps institutional knowledge and personal experience in their proper places.

BastCare is where an individual exercises that boundary. Record a visit with everyone's permission. Receive a plain-language summary. The audio is deleted. Bast retains event IDs, token counts and timing. Then review exactly what would be shared, decide who can see the summary, and revoke access later.

Sharing nothing is always valid.

We humans yearn for products that help us without harvesting our data. BastCare is built around that desire. A person can receive the benefit of institutional intelligence while retaining specific, visible, and reversible authority over how their experience is used.

*Consent is not a blanket checkbox that turns a life into a dataset. It is specific, visible, reversible authority over a defined use.*

## We did not invent the principle, we remixed and made it relavent.

The clearest articulation of this idea comes from Indigenous data sovereignty.

The CARE Principles for Indigenous Data Governance stand for Collective Benefit, Authority to Control, Responsibility, and Ethics. They were developed for Indigenous Peoples and their right to self-determination; they should not be flattened into generic corporate language. But the challenge they pose to every data system is essential: making data findable and reusable is not enough. We must also ask who has authority, who carries responsibility, and who benefits.

That is the standard this architecture should earn, not merely claim. Regulation is moving in the same direction.

The EU AI Act's Article 50 transparency duties have applied since August 2026, including the requirement that people be informed when they are interacting directly with an AI system. The European Health Data Space entered into force on 26 March 2025, with key provisions and most rules governing the secondary use of health data beginning to apply in March 2029. Together, these frameworks turn transparency, data control, secure processing, and accountability into operating responsibilities.

The governance work is coming either way. Bast gives institutions a way to make it a community asset instead of treating it only as a compliance cost.

## Owned locally, powered by world-class infrastructure

The largest health institutions are already separating ownership of intelligence from ownership of infrastructure.

In November 2025, Mayo Clinic launched Platform_Insights, drawing on a network encompassing 26 petabytes of clinical information. In June 2026, Mayo Clinic and Microsoft announced a frontier healthcare model combining Mayo's expertise and de-identified data with Microsoft's AI and cloud infrastructure. The model will be owned by Mayo Clinic.

## Our moonshot is not one model that knows everything about everyone.

It is a network of institution-owned models: libraries, hospitals, rehabilitation centers, universities, and community organizations, each grounded in the information it is trusted to steward and the human relationships that make that information meaningful.

Each becomes a local node in a larger collective intelligence. Each can serve its community directly, connect to others by choice, and earn revenue from approved commercial uses without surrendering its corpus or its accountability.

The governance is local. The infrastructure can be global. The value returns home.

## The care that happens between appointments.

Most care happens when no one is watching.

It happens in the late-night search. The hesitant text to a friend. The sixth reading of a discharge sheet. The quiet argument in the car. The small doubt no one admits. The moment when fear beats clarity.

This is where Bast sits: in the space between appointments, where human lives actually unfold.

By allowing questions, corrections, clinical expertise, and lived experience to meet without stripping away their context, we create something larger than a chatbot and more useful than a static repository.

We create the conditions for collective intelligence.

## An invitation to co-author the future.

Bast is looking for collaborators who want to shape this model with their clinicians' voices, patient journeys, caregiver challenges, institutional values, and community needs.

We can begin with one or two conditions where the need is greatest. We co-design the conversational pathways, learn where guidance needs warmth, and discover what helps people move forward.

Care deserves a companion. Institutions deserve to be heard. Communities deserve to share in the value of what they know together.

This is how LLMs become public utilities: intelligence remains governable, useful, and accountable to the communities that created it, while private data remains private. ***That is how trusted knowledge learns to speak.***

## Source Notes:


Mayo Clinic, "Mayo Clinic launches Mayo Clinic Platform_Insights to advance digital innovation and quality improvement across healthcare," 3 November 2025.

Microsoft, "Mayo Clinic and Microsoft collaborate to develop a frontier AI model for healthcare," 2 June 2026.

European Commission, "European Health Data Space Regulation (EHDS)," and "Guidelines on Transparency Obligations for Providers and Deployers of AI Systems," 20 July 2026.

Carroll, S.R., Garba, I., Figueroa-Rodriguez, O.L., et al., "The CARE Principles for Indigenous Data Governance," Data Science Journal 19(1):43, 2020.

Andrew Ng / Landing AI, "Data-Centric AI: A Data-Driven Machine Learning Approach."

Michael I. Jordan, "Artificial Intelligence - The Revolution Hasn't Happened Yet," Harvard Data Science Review 1(1), 2019.

Margaret Mitchell et al., "Model Cards for Model Reporting," 2018.

Margaret Mitchell et al., "Measuring Data," 2022.

Dave Snowden, "What We Cannot Say," The Cynefin Co, 2026, and "Creating a Knowledge Strategy," 2014.

Nora Bateson, "Warm Data," 2017, *Combining* (Triarchy Press, 2023).