PHICIL-ITATE CHANGE

A PRIMER

SEEDS OF INNOVATION™

A shared decision standard for AI in systems of care, and a way of asking whether **the seed truly fits the soil.**

By Sheila Phicil, MPH, MS, PMP, FACHE
Founder & CEO, Phicil-itate Change

CONTENTS

BEFORE WE BEGIN

A NOTE FROM THE FOUNDER

Why this framework exists, and where it comes from.

I am a **Social Change Futurist™**: I read the patterns in where our systems are heading, and I work to change course before harm becomes permanent, leveraging technology and innovation to usher that change in. I am also a spiritual sovereignty practitioner. I practice, first in my own life, the truth that I am whole and good, created with everything I need to lead and design my life as a sovereign being, and that this same sovereignty belongs to every person. I practice it for myself first, and it carries forward into the systems I help shape. The way I live as a human is inseparable from the way I do this work. The two are one thing.

The SEEDS of Innovation™ framework came straight from lived experience, on both sides of the system. It was forged over a life spent watching what happens when we build things for people without truly listening to them. It was also forged by being one of those people: by knowing, firsthand, what it is to be failed by the very systems meant to care for me. That knowledge shaped a hard-won understanding of how to do it differently.

The first time I understood the problem was in 2007. I was a college student in Haiti, serving as an interpreter for a group bringing wheelchairs, crutches, and prosthetics to families. A few days in, the families started bringing the equipment back. My group decided the people were ungrateful. I asked them myself. What they told me was simple: I have to cross a river, you said to keep it dry, and it is easier for me to crawl than to use a crutch. No one had asked about their lives before deciding what they needed. That was the moment I stopped believing you could fix health one patient at a time, and started understanding that the problem is systemic.

That lesson followed me through every role I ever held. Working on HIV care, I learned that the people who knew how to stay on treatment were patients themselves, so we brought peers into the system and paid them to teach. At the Veterans Affairs hospital, my first attempt to measure physician productivity failed, until I stopped, listened to the doctors, and rebuilt the tool the way they actually understood their work. Within weeks, they had the best data the system had ever seen. The pattern was always the same: when you treat people as experts in their own lives, you can actually help them.

Then I learned the other side of that truth. In 2017, at a major cancer center, I was working as a senior consultant for strategic projects and I was good at my job. So good that I was the one who helped implement a system that decided which patients we could accept for a costly new therapy, based on their insurance. We built a green list, a yellow list, and a red list. Medicaid was on the red list. I had drawn a line that decided who would get care and who would be turned away, and the people on the wrong side of it looked like the people I came from. The tool was efficient. It was compliant. It protected the institution's finances. And it caused profound harm. I have carried that ever since, because it taught me the most important truth in all of this work: **a tool can be excellent,**

profitable, and within the rules, and still be wrong. Good intentions offer no protection. Only better decisions do.

By August of 2023, I was the Director of Innovation for the Health Equity Accelerator at Boston Medical Center, and I implemented the first AI tool at that safety-net hospital, in a geriatric clinic, before most people had even heard of ChatGPT. It worked. It helped our providers and it helped our patients, and the company behind it is thriving today. What made that decision succeed was the lens I brought to it, the same lens I had spent twenty years learning the hard way: understand the system first, design for the people with the most to lose, make sure it fits the whole, protect the people whose data makes it possible, and stay responsible long after launch.

SEEDS of Innovation™ is that lens, written down. It is the discipline I was already using, turned into a standard that anyone can hold, so that the next person bringing AI into a system of care can begin with these lessons already in hand, instead of learning them one painful decision at a time.

I built this from lived experience, and I am sharing it the same way I have learned to do everything else: in the open, with honesty, and in service of the people the system was built to overlook.

— Sheila Phicil, MPH, MS, PMP, FACHE
Social Change Futurist™
Founder & CEO, Phicil-itate Change

BEFORE WE BEGIN

EXECUTIVE SUMMARY

The whole idea, in about a minute.

Artificial intelligence (AI) is entering health and care faster than anyone can fully govern it, and the results so far are uneven. Tools that look impressive have denied people care, failed the moment they left the lab, worked for some groups while failing others, and turned people's most personal data into value they never shared in. These are rarely technology problems. They are decision problems: the wrong choices made before anyone asked the right questions.

SEEDS of Innovation™ is a shared decision standard for AI in systems of care. It gives everyone the same set of questions to answer before they build or buy AI, so they can tell early whether it will work in the real world or fall apart once real people start using it.

SEEDS measures one thing above all: **AI-System Fit™** — whether the real system a tool enters can actually carry it, or will push it back out. A tool can be wanted, funded, and technically excellent and still fail, because the system it joined had no room to hold it. SEEDS is how an organization earns that fit on purpose, before it commits resources.

It works through **five questions, answered in order**:

- **Systemic Signal Assessment** — understand the system you are trying to change, and listen to the people closest to it, before acting.
- **Everyone-Centered Design** — design for the people facing the highest barriers, so the solution works for everyone.
- **Ecosystem Alignment** — make sure the solution holds across the whole living system, from the person at the center out to the environment itself.
- **Data Sovereignty** — keep ownership of data, and a share of the value it creates, close to the people whose lives and work created it.
- **Stewardship** — stay responsible after launch, hold the solution to the reason it was built, and keep the power to stop it.

How deeply each question is answered matters as much as the questions themselves. SEEDS measures this depth as **reasoning maturity**, and frames it as a matter of how deeply you listen: to the system, to the evidence, and to the people closest to the problem. Maturity grows over time, through practice and through an honest reckoning with what has failed before.

What makes SEEDS different from the frameworks already in use is where it works. Most frameworks review, comply, or diagnose after a direction is set. SEEDS works one step upstream, governing the decision itself, and it treats the whole system of care, rather than the tool, as the thing to get right. It works alongside the tools an organization already uses, holding them to a shared standard and deciding when each comes into play.

That shared standard is the point. When everyone across a system of care — the people who deliver care, write the rules, pay, fund, build, teach, and receive care — answers the same questions, their separate decisions line up, and the whole system grows stronger together.

SEEDS is already moving into practice, beginning with the Massachusetts League of Community Health Centers (MLCHC), and it is built to be taught and carried by certified practitioners across many organizations over time.

The pages that follow explain SEEDS in plain language, and through a single picture that holds the whole idea together: **a garden, where what you harvest depends less on the seed than on whether the seed truly fits the soil.**

The figure below shows the whole framework at a glance: the five pillars across the top, and the five levels of reasoning maturity rising within each one.

SEEDS of Innovation™

How each pillar deepens through the five levels of reasoning maturity

	Systemic Signal Assessment	Everyone-Centered Design	Ecosystem Alignment	Data Sovereignty	Stewardship
Level 5 · Regenerative	Reads the system so each cycle leaves it stronger	Widens access for everyone, every cycle	Strengthens the whole system, including the environment	Communities govern and share in their data	Keeps learning; harms and gains both feed the next decision
Level 4 · Integrative	Reads how the whole system behaves together	Designs directly with the highest-barrier people	Checks fit across every ring of the system	Shared say and shared value over data	Watches real use; can adjust or reverse
Level 3 · Normative	Uses the accepted process and guidelines	Meets standard accessibility requirements	Aligns to known workflows and rules	Follows consent and privacy policy	Reviews on a set schedule against standards
Level 2 · Protective	Acts to avoid the biggest known risk	Avoids the most obvious harms to some groups	Checks the tool will not break nearby	Guards data against the clearest risks	Monitors for major failures only
Level 1 · Reflexive	Reacts to what is loudest today	Designs for the average user	Looks only at the tool in isolation	Collects data with no say for people	Launches, then leaves it alone

The Foundation

1 WHAT SEEDS IS

SEEDS of Innovation™ is a shared way of making decisions about AI in health and care. It gives everyone the same set of questions to answer before they build or buy AI, so they can tell early whether it will work in the real world or fall apart once real people start using it.

The most important word is **shared**. SEEDS belongs to everyone in the system: the people who deliver care, the people who write the rules, the people who pay, the people who fund it, the people who make the products, the people who teach the next generation, and the people who receive care. When all of them use the same questions, their separate choices line up. The whole system grows stronger together, even when each part is working on its own.

SEEDS gives that shared work a clear shape: five questions to answer, in order, before you act. Answer them well, and you build AI that takes root. Skip them, or answer them out of order, and you build something the real world pushes back out.

The Foundation

2 WHERE SEEDS FITS: THE GARDEN

Growing something in a system of care works the way growing something in soil works. The harvest is rarely about the seed alone. It is whether the seed, the soil, and the conditions all work together. **AI-System Fit™ is that match: whether the real system an AI tool enters can actually carry it, or will push it back out.**

SEEDS fits into the work you are already doing. Your strategy is what you want to grow, and why. Your method is how you plant and tend. SEEDS guides the choices in between: the ones that decide whether anything you plant can truly take root in this soil. A gardener who reads the soil first grows something that lasts. A gardener who plants what worked in someone else's yard starts the season over.

This is where the garden teaches something true about AI, with no technical words needed: **an AI tool succeeds or fails mostly because of where it is planted, the people, the daily work, the trust, the data, far more than because of how clever the tool is.**

The Framework

3 THE FIVE QUESTIONS THAT MEASURE FIT

SEEDS measures **AI-System Fit™** through five questions. You answer them one at a time, in order, because each one prepares the soil for the next. Skip ahead, and you build on soil you have not read yet. Take them in turn, and what you plant has a real chance to take root.

SEEDS PILLAR 1

SYSTEMIC SIGNAL ASSESSMENT

Understand the system you're trying to change, and listen to the people closest to it, before you make a decision.

Before you choose what to build or buy, you study the system you are trying to change, and what it is telling you. That system might be small and specific: the way urgent visits get triaged, how patients move through discharge, how a clinic handles follow-up calls. Whatever it is, it is a living web of people, rules, habits, and timing, all shaping each other. And it sends signals constantly: what keeps working, what keeps straining, who carries the hidden load, where the same trouble returns again and again.

Systemic Signal Assessment is the discipline of reading those signals first. Part of that discipline is drawing the right boundary around the system, deciding what is inside the thing you are trying to change and what sits outside it, so the work is focused enough to actually understand. Then it asks three things:

- **What is really going on here, underneath what we see today?** A single complaint is one signal. The pattern behind it is the truth. You look for the pattern.
- **Whose view counts as evidence?** The people closest to the work, and the people who live with the results, hold the clearest signal. So their experience counts as real proof, and you listen to them first.
- **Is this even a place for AI at all?** Sometimes the honest answer is a better workflow, a change in policy, or more people. Choosing one of those is a full success of this pillar, because the goal is the right call, not the AI call.

This is the gate everything else passes through. When you read the system well, every choice that follows stands on solid soil. When you rush past it, you risk adding AI that helps the system do the wrong thing faster.

IN THE GARDEN

Before you plant a single seed, you observe the soil across seasons. Where does the water gather, and where does it run off? Which corner holds frost into spring? What grows here on its own, and what struggles year after year? You ask the people who have worked this land the longest, because they have seen what one good season would hide.

The soil is always sending signals. The patient gardener reads them first and lets them decide what to plant, where to plant it, and whether to plant at all. That reading is what turns a hopeful guess into a harvest.

A REAL EXAMPLE

In 2013, MD Anderson Cancer Center partnered with IBM to build an AI system that would recommend cancer treatments. The project spent $62 million over four years. Internal testing found the system made unsafe recommendations, and it was never deployed for patient care. The project was abandoned in 2017. The system had been built without first understanding how oncologists actually made treatment decisions, or what the clinical workflow could absorb.

Sources: STAT News, July 2018; Wall Street Journal, March 2017.

SEEDS PILLAR 2

EVERYONE-CENTERED DESIGN

Design for the people who face the highest barriers, so the solution works for everyone.

Once you understand the system and what it is telling you, you start to design. And the question now is *who you design for.*

Everyone-Centered Design means you build for equity and accessibility from the start, by designing for the people who face the highest barriers first. The person who speaks a different language. The person with a disability. The person with the longest commute, the least time, the most reasons to distrust the system. When you design so the solution works for them, it works for everyone, because you have built something strong enough to hold the hardest conditions.

The word "everyone" is the whole point, and it is meant literally. A solution built for an imagined average person tends to serve the middle and strain at the edges, and that strain spreads inward over time. A solution built for the edges holds all the way through. Designing for the highest barriers takes care of the easier cases on its own, with less patching, less special handling, and less breaking down later.

This pillar builds directly on the first. Systemic Signal Assessment is how you learned who faces the highest barriers in this system of care. Everyone-Centered Design is how you put those people at the center of what you build.

IN THE GARDEN

Every garden has a hardest spot. A corner where the frost settles, a strip that goes dry in July, a patch in deep shade. The wise gardener chooses plants that can live *there*, in the toughest soil.

Here is why that works. A plant hardy enough for the frost pocket thrives easily everywhere else in the garden. So when you design for the hardest spot, the whole garden holds, and it holds with less fuss. You water one way. You tend one way. The garden becomes steady and self-sustaining, because you grew it for the conditions at the edges instead of only the easy middle. Plant for the hardest soil, and the whole garden grows strong.

A REAL EXAMPLE

In 2019, researchers discovered that an algorithm used by Optum to identify patients for extra care was systematically biased against Black patients. The algorithm used healthcare spending as a proxy for need, but because Black patients historically had less spent on their care, the tool referred them at far lower rates. At a given risk score, Black patients were significantly sicker than white patients with the same score. The tool had been designed for average cost patterns, leaving those facing the highest barriers invisible.

Source: Obermeyer et al., Science, October 2019.

SEEDS PILLAR 3

ECOSYSTEM ALIGNMENT

Make sure the solution fits the whole living system it joins, working outward from the person at the center all the way to the Earth.

Now you have a solution designed for the people who need it most. The next question is whether it can actually live in the real world it is about to enter.

A solution in a system of care never works alone. It joins a whole living system, and every part of that system depends on the others. The way to see it is to work outward, ring by ring, starting from the same place Everyone-Centered Design started: the person at the center.

At the center is the person closest to the problem, the one who faces the highest barriers, along with the caregivers and family who hold them. Move out one ring, and you reach the people who deliver care and the staff who support them. Out again, and you reach the people who pay and the people who write the rules. And out beyond all of them is the widest ring of all: the Earth itself, the energy and water and resources every solution draws on, including the real environmental cost of running AI. Each ring depends on the ones inside it, and supports the ones around it.

Ecosystem Alignment checks that the solution works across every one of those rings, in real conditions. Because a solution that serves the center while breaking at any ring further out is a solution the whole system will eventually reject. If it helps the patient but burdens the nurse, it breaks. If the nurse can use it but no one will pay for it, it breaks. If everyone wants it but the rules

forbid it, it breaks. And if it works for all of them while draining the planet that holds the whole system, it breaks the largest ring of all. A solution holds only when it holds all the way out.

This is where many good ideas meet the real world. A tool can read the system well and serve the highest-barrier people beautifully, and still need this test: does it hold across the whole living system, ring by ring, from the person at the center to the Earth at the edge?

IN THE GARDEN

A plant lives with everything around it. The soil it draws from, the other plants nearby, the bees that visit it, the water it shares, and the people who tend it, harvest it, and clear it away at the end of the season. And around all of that is the land itself, the climate, the water table, the wider world the garden sits inside.

A plant that grows fast by draining the soil, crowding out its neighbors, or poisoning the water looks like a triumph for one season and becomes a problem the next. So the wise gardener looks at the whole living world before planting. Will this plant feed the soil or starve it? Will the bees it draws help the rest of the garden? Can the people here actually tend it with what they have? And does it give back to the land it grows in, or take until the soil is spent? A plant that fits the whole living world, all the way out to the earth that holds it, gives you a harvest that lasts, season after season.

A REAL EXAMPLE

Epic Systems deployed a widely used sepsis prediction model across hundreds of hospitals. A 2021 study found the model missed two out of three sepsis cases and generated an 18% false alert rate, producing roughly 109 alerts for every case it actually caught. The model performed well in controlled testing but failed when it met the real conditions of clinical workflows, alert fatigue, and patient populations it had never been validated against.

Source: Wong et al., JAMA Internal Medicine, July 2021.

SEEDS PILLAR 4

DATA SOVEREIGNTY

The data belongs close to the people whose lives and work created it, along with the value it generates.

This is the boldest part of SEEDS.

AI runs on data. It takes in the data people give and turns it into information, insight, and value. In a system of care, that data comes from human lives and human work: a patient's story, the signals from a body, a nurse's notes, the record of a thousand daily decisions. People are the original intelligence. The machine learns from what they give it.

For a long time, healthcare has treated that data as something the system owns once it is collected. Systems gather it, hold it, and build value on it, while the people who gave it stay outside the benefit. Data Sovereignty corrects that. It holds that ownership belongs close to whoever generated the data, and so does a share of the value it creates.

In practice that means a few things held together. People know their data is being used, and they agree to it freely, each time it is put to a new purpose. They keep a say in how it is used. They share in the value it generates. And they keep the power to stop. The same principle holds at every level: the person closest to the data, whether a patient sharing their story or a clinician whose work generates a record, stays closest to the say over it.

In the age of AI, data is power. SEEDS puts that power back where it begins, with the people whose lives and work create it. That is why this pillar is the most radical of the five: it changes who is in charge.

This pillar builds on the ones before it. By now you understand the system, you have designed for the people who need it most, and you have checked that the solution fits the whole living system. Data Sovereignty asks the question those people would ask if they were in the room: whose is this, and who benefits?

SEEDS draws on established principles here: FAIR (Findable, Accessible, Interoperable, Reusable) for how data is managed, and CARE (Collective Benefit, Authority to Control, Responsibility, Ethics) for who governs it and who benefits (Wilkinson et al., 2016; Global Indigenous Data Alliance, 2019). Together, they anchor Data Sovereignty in both technical rigor and ethical accountability.

IN THE GARDEN

The people who know the land are the reason anything grows at all. Their knowledge of the soil, what it needs, when it turns, what it has survived, is the true source of every harvest.

So when a harvest comes in, the question is who it belongs to. Imagine people who tend the soil with their own hands and hold generations of knowledge about it, and then watch every basket carried away by someone who never knelt in the dirt. Something has gone wrong, no matter how rich the harvest. And there is a deeper cost: push out the people who know the land, and the knowledge that grew the garden walks away with them. Next season, far less grows.

The wise and fair gardener keeps the people who know the land close, and keeps them sharing in what the land brings in. A garden tended by the people whose knowledge grew it keeps growing, season after season. A garden grown by taking gets one harvest, and then the knowledge, and the gate, both close.

A REAL EXAMPLE

> The clearest example is 23andMe. Millions of people paid the company to learn about their ancestry, and many agreed to let their genetic data be used for "research." In 2018, the drug maker GlaxoSmithKline invested 300 million dollars and gained the exclusive right to mine that database to develop new medicines, and the two companies agreed to split the costs and the profits between them. The people whose DNA made the whole thing possible received no share, and most learned of the arrangement only after it was announced. The value moved from the people to the corporations, and stayed there. Genetic data makes the stakes sharpest of all, because a person can change a password, and can never change their DNA.

SEEDS PILLAR 5

STEWARDSHIP

Stay responsible for the solution long after launch, and hold it to the reason you built it.

Stewardship is the discipline of staying responsible the whole way through, a commitment to ethics: to watch what the solution actually does in real use and to keep the power to adjust it, improve it, or stop it when the signals call for that.

You have read the system, designed for the people who need it most, checked that the solution fits the whole living system, and kept the data close to the people who generated it. Now comes the part that lasts the longest.

Launch is the beginning. The work continues from there. A solution keeps working after it goes live, and the world around it keeps changing. New people arrive, conditions shift, and what worked at first can struggle later. Stewardship is the discipline of staying responsible the whole way through: watching what the solution actually does in real use, and keeping the power to adjust it, improve it, or stop it when the signals call for that.

Here is the heart of it. You hold the solution to the very thing you set out to address, the signal you named in the first question. The test is whether it still delivers the outcome you intended. Is it doing what you hoped? Better than expected, in a way worth growing? Falling short? Causing harm? You measure it against the goal you set, not against how well the tool performs on its own terms. A tool can be fast, accurate, and efficient and still miss the point. If the intention was to keep more patients out of the emergency room, then that is the measure, and a tool that hits every technical mark while that number stays flat has more work to do.

Stewardship also means holding the power to act on what you see. When the signals show real harm, you need a real way to stop: a named person, a clear threshold, and a working mechanism to pause, change course, or remove the solution entirely. The ability to reverse a decision is part of the decision. A solution you cannot stop is a solution you never fully controlled.

This is the test a solution can pass at launch and still fail later. It can be profitable, follow every rule, stay popular, and still fail stewardship by being left alone too long. The work of SEEDS holds only when someone keeps holding it.

IN THE GARDEN

A garden is never finished on planting day.

You planted for a reason: to feed your family through the winter. So the real work is keeping that reason in view. The plants are growing, but are they growing the food you actually need? The wise gardener watches for the first sign of blight, notices the vine spreading too far, catches the dry week before it becomes a dry month, and sees the surprise, the volunteer plant that came up on its own and turned out to be the best thing in the bed. A plant can look healthy and green and still give you nothing you can eat.

And here is what matters most when something goes wrong: you can pull the plant out. If a plant is choking the others, drawing pests, or bearing fruit that makes people sick, the gardener does not hope it improves. They pull it. The power to remove what causes harm is what keeps the whole garden safe.

A garden left alone after planting goes to seed, no matter how well it started. A garden tended through every season, watched against its purpose, and cleared of what causes harm, is the one that feeds people year after year. The planting is a promise. The tending is how you keep it.

A REAL EXAMPLE

Duke Health launched Sepsis Watch in 2018 with a different approach: the team spent months studying how nurses and rapid response teams actually worked before building the tool. They kept clinicians in the decision loop, tracked outcomes continuously, and expanded only after proving the model held in real conditions. By 2024 it had been running for six years, contributed to a 27% reduction in sepsis mortality, and expanded to three hospitals.

Sources: Sendak et al., JMIR Medical Informatics, July 2020; Duke Today, August 2024.

The Framework

4 HOW DEEPLY YOU LISTEN

The five questions are only half of SEEDS. The other half is *how deeply you answer each one.* Two people can ask the very same question and stand in completely different places, because one is listening to what is loudest today and the other is listening all the way down to how the system really works.

The deeper you listen, the more you draw on: stronger signal, fuller proof, and the voices that are easy to miss. The same question can be answered at five levels, and each level is a real way of working, with each one listening more deeply, and standing on clearer signal, than the last.

- **Reflexive.** You answer from what is right in front of you, moving fast when speed matters. Your signal is a single event, a complaint, something happening today. The question to ask out loud: *What information was unavailable at the time of this decision?*
- **Protective.** You answer to stay safe against the biggest known risk. Your signal is the records and reviews that hold up under scrutiny. The question to ask out loud: *What risk were we optimizing against, and whose risk was it?*
- **Normative.** You answer using the accepted process. Your signal is the official guideline, the standard everyone shares. The question to ask out loud: *Where does lived experience contradict what our documents say?*
- **Integrative.** You answer by listening to how the whole system behaves together. Your signal is lived experience, how the work truly runs, and the patterns that show up over time. The question to ask out loud: *Which system behavior are we reinforcing, intentionally or otherwise?*
- **Regenerative.** You answer so that each choice leaves the system stronger for next time. Your signal is gathered across many cycles and held where the whole system can keep learning from it. The question to ask out loud: *How does this decision make future decisions easier and safer?*

And the levels are fluid. An organization can listen deeply on one question in one season and find itself back at a shallower level when pressure, leadership changes, or urgency shift the conditions. The levels describe how you are reasoning right now, in this decision. They are positions you move through, revisit, and earn again as conditions change.

Two things make this matter. First, maturity lives inside each question, not in a separate step. Your maturity *is* how deeply you listened on each one. You can listen deeply on one question and barely scratch the surface on another. Second, the depth of your listening sets the size of the move you have earned. A fast read backed by a single signal earns a small, reversible step. A deep read backed by real proof earns a larger commitment. That is how SEEDS keeps anyone from betting the whole system on a guess.

No organization starts at Regenerative. The goal is to know honestly where you are, and to move deliberately deeper as your experience builds.

This is why failure matters so much. In a system of care, failure is one of the clearest signals there is, and it is also the one most often buried, especially when people are eager to show that an innovation worked. SEEDS treats failure as something to face honestly and use, a catalyst that informs the very next decision rather than a verdict to hide from. This honesty is most important of all with AI, because so much about how AI behaves in the real world is still unknown. We do not know everything we are walking into. Growing through the levels, with discipline and an honest reckoning with what went wrong, is how we learn what we could not have seen at the start.

IN THE GARDEN

Every gardener can look at the same patch of soil and see something different, depending on how closely they listen to it.

Reflexive. "A plant is struggling today. Act now, keep it simple."

Protective. "This spot has lost plants before. Choose the safe, proven option."

Normative. "Follow the planting guide everyone trusts for this zone."

Integrative. "Read how the whole yard works together: the water, the shade, the way the seasons move across it."

Regenerative. "Make every choice leave this soil richer and easier to grow in next season."

No gardener begins as a master. You learn the soil by working it season after season. The tomatoes that failed last year teach you more than the ones that thrived, if you are willing to look honestly at why they failed: the wrong corner, the late frost, the watering you missed. A gardener who studies the failed bed plants a wiser garden the next spring. The one who looks away plants the same disappointment again.

So the depth of listening grows over time, and it grows fastest in the gardeners who face what went wrong. The more seasons you have tended, and the more honestly you have read both the harvest and the loss, the deeper you listen, and the bigger and surer the planting you have earned.

In Practice

5 WHAT GOES WRONG, AND WHAT SEEDS CATCHES

It helps to be clear about what SEEDS is protecting against, because these are not hypothetical worries. Technology has been failing in systems of care for a long time, in patterned, documented ways, and AI is the newest and most powerful version of that same story. The patterns repeat, and each one maps to a pillar that would have caught it.

- **The wrong problem, solved confidently.** A health insurer used an algorithm to decide how long patients needed care after a hospital stay, and it cut people off while they still needed help. When patients appealed, the algorithm was reversed in the large majority of cases, which means it was wrong most of the time it was challenged. The tool worked exactly as built. It was simply built to optimize the wrong thing. This is what **Systemic Signal Assessment** catches: it forces the

question of what problem you are really solving, and for whom, before a single line of code is trusted.

- **Works in the lab, fails in the world.** Again and again, models that performed beautifully in testing have fallen apart once they met real patients in real hospitals, because the conditions on the ground were nothing like the conditions they were trained on. This is what **Ecosystem Alignment** catches: it tests whether a solution survives the whole living system, not just the controlled setting where it was born.
- **Built for the average, harmful at the edges.** Bias has shown up in the large majority of studies that examined AI tools for it, with tools working well for some groups and failing others, often the people already facing the most barriers. This is what **Everyone-Centered Design** catches: by building for the hardest case first, it closes the gap that bias slips through.
- **Value taken, people left out.** Data given by patients and communities has been turned into enormous value that flowed to companies and institutions, while the people who gave it received no say and no share. This is what **Data Sovereignty** catches: it keeps ownership and benefit close to the people whose lives and work created the data.
- **Launched, then left alone.** Tools have been deployed and then forgotten, drifting toward harm with no one watching and no clear way to stop them. This is what **Stewardship** catches: it keeps someone accountable for what the solution does over time, and keeps the power to pull it out.

One failure belongs to AI in a way it never belonged to any tool before it, and it deserves to be called out on its own.

> **THE RISK UNIQUE TO AI**
>
> **Cognitive surrender is what happens when people stop doing their own thinking because the AI already did it for them.** Researchers at the Wharton School identified and named this pattern (Shaw & Nave, 2026). As people lean on AI to help them think, many begin to accept the AI's answer and set aside their own judgment, and they often feel *more* sure of the decision, even when the AI is wrong. In their study of more than 1,300 people across nearly 10,000 decisions, acceptance of wrong AI answers reached almost 80 percent in some conditions. No earlier technology did this. A calculator or a records system saves you effort; it does not quietly take over your thinking. AI does, because it reasons in words, the way a trusted colleague would. In a system of care, where decisions carry real weight and depend on real human understanding, that is a sharp and specific danger. It is also the clearest reason a *shared standard* matters. A standard that lives outside any one person's head keeps human judgment in the loop, so that "the AI suggested it" can never be the place the thinking stops. SEEDS keeps people doing the reasoning, with AI as something they question rather than something they obey.

IN THE GARDEN

Each of these failures shows up in the garden too.

You can plant fast and confident in the wrong soil, and lose the whole bed. You can grow something beautiful in a greenhouse that dies the moment it meets real weather. You can tend the easy center of the garden and let the hard corners fail, until the failure creeps inward. You can let someone haul off a harvest that was never theirs. You can plant and walk away, and come back to a bed gone wild.

And there is the particular trap of trusting the seed catalog over your own soil. The catalog is glossy and certain. It says this plant grows anywhere. The gardener who stops reading their own soil and simply believes the catalog plants the same disappointment, season after season, and feels confident doing it. The wise gardener keeps their own eyes on the soil, and treats the catalog as one voice among many, never the final word.

In Practice

6 HOW SEEDS WORKS WITH WHAT YOU ALREADY USE

If you lead in a system of care, you may already use a framework for AI or change, and a fair question is where SEEDS fits alongside it. The answer is that SEEDS sits in a different place than the others, and it works with them rather than against them.

Most existing frameworks do important work at a specific stage:

- Some help you **manage risk and stay compliant** once you have decided to build or buy. They give you a shared language for documenting what a tool is for and where its limits are.
- Some help you **understand why a tool succeeds or fails** when you try to put it into practice. They are strong at studying the friction between a new tool and the way work really happens.
- Some set the **ethical principles** a health system should honor, the values that protect dignity and rights.
- Some help a hospital **review and screen** the flood of tools vendors are selling, sorting the low-risk from the high-risk.

Each of these is useful, and each does its work *after* the most important decision has already been shaped: the decision of whether to act at all, on what problem, and on whose terms. They review, they comply, they diagnose, and they sort, mostly once a direction is already set.

SEEDS works one step upstream. It is a decision standard, so it governs *the decision itself*, before resources are committed: is this even the right problem, who is closest to it, will this hold across the whole system, who owns the data, who stays responsible over time. SEEDS does not replace your risk

framework, your implementation model, your ethics principles, or your review process. It decides whether and how those tools should be activated in the first place, and it holds them all to the same standard.

There is one more difference worth naming. Most frameworks treat the tool as the thing to get right. SEEDS treats the *system of care* as the thing to get right, and the tool as one part serving it. That is why SEEDS can sit above whatever you already use: it is asking a larger question than any single tool was built to answer.

WHERE SEEDS SITS NEXT TO THE FRAMEWORKS YOU MAY KNOW

These are real frameworks used in health and care today. Each does valuable work at its stage. SEEDS works one step upstream of all of them, governing the decision before these tools are activated.

- **NIST AI Risk Management Framework** — manages AI risk and supports compliance once a tool is chosen. *SEEDS decides whether and how to act before this begins.*
- **CFIR (Consolidated Framework for Implementation Research)** — explains why a tool succeeds or fails when put into practice. *SEEDS asks whether it was the right tool for the real problem in the first place.*
- **WHO Ethics & Governance of AI for Health** — sets the ethical principles a system should honor. *SEEDS turns principles like these into decisions you can actually make.*
- **FAIR-AI** — helps a hospital review and screen the AI tools vendors are selling. *SEEDS governs the larger question of whether to pursue a tool at all, and on whose terms.*

SEEDS does not compete with these. It holds them to a shared standard and decides when each one comes into play.

IN THE GARDEN

Think of everything a serious gardener relies on. A pest guide for spotting trouble. A soil test for what it needs. A planting calendar for timing. A trusted nursery for healthy stock.

Every one of those is valuable, and a wise gardener keeps them all. But none of them answers the first question: *should I plant this, here, at all?* The pest guide assumes you have already planted. The soil test tells you what is there, not what belongs. The calendar tells you when, not whether.

SEEDS is the gardener's judgment that comes before all of it, the thinking that decides what to grow in this soil and why, and then puts every other tool to work in service of that choice. You keep all your tools. SEEDS is what makes them add up to a garden that thrives, instead of a shed full of good equipment and a bed that fails.

In Practice

7 ONE STANDARD, ACROSS THE WHOLE SYSTEM

A system of care has many hands at work in it. The power of SEEDS is that all of those hands use the same five questions, so their separate choices line up instead of pulling apart.

SEEDS sees the whole healthcare system as one living thing, made of the parts where decisions get made and where the results land. Each part owns a different decision, and every decision flows to the others.

- **The people who deliver care** decide whether a tool helps real care or burdens it, and whether to use it.
- **The people who write the rules** decide whether a policy fits how the system really works.
- **The people who pay** decide whether to put money behind a tool.
- **The people who fund it** decide whether to invest in building it in the first place.
- **The people who make the products** decide what to build, how to build it, and what to release.
- **The people who teach the next generation** decide what future leaders learn about building and using AI.
- **The people who receive care** decide whether to trust a tool, share their data, and take part, on their own terms.

When each part answers the same five questions, something powerful happens. A builder designs for the people facing the highest barriers, the same people a provider is fighting to protect, the same people a payer wants to keep out of the emergency room. The questions give them a shared language, so a decision made in one part of the system strengthens the others instead of surprising them later. That is what a *shared standard* does. It turns separate, siloed choices into one system pulling in the same direction.

Here is what the same five questions sound like, in the words of each part of the system.

- **If you deliver care:** SEEDS gives you the standing to say no to a tool that will burden your people, with your reasons documented, and the confidence to say yes when a tool truly serves. You stop inheriting tools that were never going to work in your setting.
- **If you write the rules:** A policy is something you build, and it can miss the real problem the same way a product can. SEEDS gives you the same five questions to test a rule before it becomes law, so you write rules that hold up where people actually live.

- **If you pay for care:** Every tool you fund is money and reputation on the line. SEEDS gives you a clear, documented way to show the decision was made with real rigor, so you back the tools that hold and keep the ones that fail off your books.
- **If you fund innovation:** You have seen good technology die in the gap between a promising pilot and real-world scale. SEEDS is a way to tell, early, whether a tool can survive the system it is entering, so your capital backs what lasts.
- **If you make the products:** People can want your tool and you can still be shut down when the harm shows up or the staff route around it. SEEDS gives you the design standard to build to from the start, so your tool survives the real world and reaches people a tool built for the average user never could.
- **If you teach the next generation:** What you teach becomes how the whole system thinks tomorrow. SEEDS gives you a shared standard to put in front of future builders, providers, and leaders, so responsible decision-making is built in from the beginning rather than learned after the failures.
- **If you receive care:** Your trust, your data, and your willingness to take part are what make everything else possible. SEEDS gives you the standing to ask how a tool uses your information, who benefits from it, and whether anyone is watching what it does over time. Your voice is primary, and the standard protects it.

IN THE GARDEN

A community garden is tended by many people. One person knows the soil. One waters. One keeps the pests away. One decides what the plot is for. One brings the seeds. One teaches the new gardeners. One eats from the harvest.

If each of them works from a different idea of what the garden should be, the garden pulls apart: the waterer drowns what the planter chose, the harvest comes in for food no one wanted. But when all of them share the same understanding of the soil and the same questions about what belongs here, every separate task adds up. The garden becomes one thing, tended by many hands, growing what the whole community actually needs.

That is SEEDS across a system of care. Many hands, many decisions, one shared standard, and a system that grows stronger because everyone is reading the same soil.

Looking Ahead

8 WHERE SEEDS IS HEADED

SEEDS is built to grow, in three ways.

It grows through use, and through proof. A decision standard earns its authority by showing, over time, that it leads to better decisions and prevents real harm. SEEDS is building that record in the open: a clear account of how each pillar leads to better outcomes, and a growing library of real cases where the standard changed a decision for the better. The aim is straightforward and honest. Show the work. Let the evidence accumulate where everyone can see it, so that trust in SEEDS rests on what it has actually done, not on what it claims.

It grows through real engagements, beginning now. SEEDS is already moving from idea to practice. The first major example is the Massachusetts League of Community Health Centers, which is embedding SEEDS into its AI playbook, so that the health centers serving communities across the state make their AI decisions through this shared standard. This is SEEDS doing the thing it was built for: helping the people closest to care decide well, before the stakes are locked in. Each engagement like this becomes part of the record, and part of how the next one gets better.

It grows through people, carried beyond any one room. SEEDS is designed to be taught and certified, so that trained practitioners can carry it into their own organizations and their own work. The questions, the discipline, and the way of seeing the whole system are meant to spread from one practitioner to the next. The goal is a system of care where this kind of rigor is widely held, where many people across many roles are asking the same five questions, until making decisions this way becomes simply how good decisions get made.

That is the direction. A shared standard, proven in the open, carried by many hands, growing stronger over time.

IN THE GARDEN

A single good garden feeds a family. But the knowledge of how to grow it, shared and taught and proven across seasons, can feed a whole community for generations. That knowledge spreads the way it always has, from one grower to the next, each one learning to read the soil a little more deeply than the last.

That is where SEEDS is headed. The work now is to tend the first gardens well, to keep honest records of what grows and what fails, and to teach others to read the soil the way a master grower does. A seed becomes a garden. A garden, well tended and shared, becomes a way of growing that outlasts any single season, and any single gardener.

APPENDIX

REFERENCES

All sources verified as of June 2026.

1. Shaw, S. D., & Nave, G. (January 2026). Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. The Wharton School, University of Pennsylvania. SSRN | DOI

1,372 participants; 9,593 trials; acceptance of incorrect AI answers reached 79.8% in some conditions.

2. GlaxoSmithKline & 23andMe (July 25, 2018). GSK and 23andMe sign agreement to leverage genetic insights for the development of novel medicines. GSK press release. GSK

$300M equity investment; four-year exclusive collaboration; costs and proceeds split 50/50.

3. National Institute of Standards and Technology (January 26, 2023). AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1. NIST

4. Damschroder, L. J., et al. (2009). Fostering implementation of health services research findings into practice: a consolidated framework for advancing implementation science (CFIR). Implementation Science, 4:50. DOI

5. World Health Organization (June 28, 2021). Ethics and Governance of Artificial Intelligence for Health. WHO

6. Wells, B. J., Nguyen, H. M., McWilliams, A., et al. (August 11, 2025). A practical framework for appropriate implementation and review of artificial intelligence (FAIR-AI) in healthcare. npj Digital Medicine, 8, 514.

7. Ross, C., & Herman, B. (November 14, 2023). UnitedHealth sued over use of algorithm to deny rehabilitation care to seriously ill Medicare Advantage patients. STAT News. Lawsuit alleges nH Predict, a NaviHealth algorithm predicting patient length of stay, carried a 90% denial-reversal rate on appeal.

8. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (October 25, 2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447-453.

9. Wong, A., Otles, E., Donnelly, J. P., et al. (June 2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Internal Medicine, 181(8), 1065-1070.

10. Sendak, M. P., Ratliff, W., Gao, M., et al. (July 15, 2020). Presenting machine learning model information to clinical end users with model facts labels. JMIR Medical Informatics, 8(7), e22564.

11. Ross, C., & Swetlitz, I. (July 25, 2018). IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show. STAT News.

12. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., et al. (March 15, 2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3, 160018.

13. Global Indigenous Data Alliance (September 2019). CARE Principles for Indigenous Data Governance. GIDA.