Probabilistic Self-Learning: A Roundup Method for Learning to Navigate the Unknown

Probabilistic Learning: *There is no single right path*.



Roundup uses probabilistic self-learning to repeatedly move a participant from unfamiliarity to recognized contribution across different domains, building a transferable capability to tackle the unknown. The underlying curriculum is the method used to move from unfamiliarity to useful contribution. Software, culinary work, home systems, floral design, construction, healthcare and automotive work are not the curriculum. They are environments in which that method gets exercised, tested and internalized.

The most important educational opportunity created by artificial intelligence may not be that it can answer questions faster, write software faster, or make existing forms of instruction more efficient. Its more profound potential is that it fundamentally changes what it means for a person to encounter something they do not know how to do.

Until very recently, lack of knowledge was often a hard barrier to action. If someone entered an unfamiliar domain without the prerequisite knowledge, tools, vocabulary, or experience, the normal response was to stop and seek instruction. Learning tended to begin with someone else deciding what knowledge was necessary, organizing that knowledge into a sequence, delivering it, and eventually testing whether the learner had absorbed enough of it to proceed.

AI changes that relationship between ignorance and progress. A person can now enter an unfamiliar environment, describe what they are trying to accomplish, begin exploring the system around them, discover the tools and knowledge they need, ask questions as gaps emerge, challenge the answers they receive, refine their understanding, and continue moving toward an

outcome even while significant portions of the environment remain unfamiliar. That does not eliminate the need for knowledge. It changes how knowledge can be acquired and used.

We are calling this approach probabilistic self-learning. The word probabilistic matters for reasons that go deeper than the observation that different learners receive different answers. A learner using this method is operating without certainty that any given next action is correct. They make a reasonable move based on available understanding, observe what happens, update their mental model, and continue. The uncertainty is not merely in what AI produces — it lives in the learning process itself. Two learners beginning with the same destination, the same AI system, and the same physical environment may have substantially different journeys. One may discover an important concept early while another reaches it only after encountering a problem. One may receive too much information and need to learn how to narrow it, while another may accept an answer too quickly and need to learn how to pressure-test it. The educational objective is therefore not to force every learner through an identical path. It is to teach learners how to successfully navigate when the path is not fully known.

The AI Working Environment Comes First

Probabilistic self-learning begins with the creation of a healthy working environment for AI collaboration. That is not a supporting element within the Roundup curriculum. It is the infrastructure through which probabilistic self-learning becomes possible at all. Without it, a participant is simply a young person with access to an AI tool. With it, they have the foundation required to navigate unfamiliar territory.

Giving a young person access to an AI system is not the same as giving them an effective AI learning environment. Access to abundant knowledge becomes useful only when the learner develops the habits and structures required to unlock and apply that knowledge effectively. Before Roundup can teach a participant how to navigate an unfamiliar domain, it therefore needs to help that participant establish an AI working environment that functions as both a knowledge resource and a capability tool.

That means learning how to communicate an objective and provide enough context for an AI system to understand the environment in which the learner is operating. It means learning that an exploratory conversation is different from an execution request, that useful context should be preserved rather than repeatedly reconstructed, and that an answer should often be challenged, refined, compared or examined for gaps before being acted upon. It also means recognizing that AI can confidently produce information that is incomplete, excessive, irrelevant or simply inappropriate for the actual situation.

This should not become another version of teaching children to write perfect prompts. Effective AI collaboration can often look much more like thinking out loud. A learner should become comfortable saying, in their own language: I have never done this before. Here is what I am trying to accomplish. I don't understand this environment yet. Help me understand what I'm walking into, what the important pieces are, what tools people use, and what I need to understand first. The ability to externalize uncertainty that way is itself a skill worth developing deliberately.

The AI working environment thereby becomes an intellectual workbench — the place where the learner can explore an unfamiliar system, organize what they discover, preserve useful knowledge, challenge assumptions and progressively develop enough understanding to act. The deeper capability is not prompting. It is learning how to collaborate effectively with an abundant but probabilistic source of knowledge while retaining responsibility for judgment and direction.

Entering an Unknown Domain

Once that working environment exists, a Roundup begins by giving the participant a destination concrete enough to understand even though the path toward it is not yet visible.

In the first software Roundup, that destination might be simple: by the end of the session, every participant will have created and deployed a real website at a real URL that they can send to someone else. The destination provides direction before the participant possesses understanding. The student does not initially need to know what a code editor is, what a repository does, what deployment means, where a website actually resides, or how a browser retrieves it. Those are not prerequisites for beginning the experience. They are part of the territory the learner is about to discover.

Rather than beginning with an instructor explaining the entire development environment, the participant uses the AI working environment to begin mapping the territory. If the destination is a real website built with the kinds of tools actual developers use, what environment does that imply? What are the major pieces? What tools will be required? Where will the code live? How does something created on a laptop eventually become accessible through a public URL? Which concepts need to be understood now and which can safely wait?

This is where system-level thinking begins. The mentor does not disappear, but the mentor's function changes. Rather than supplying the map, the mentor helps the learner evaluate the map being developed with AI. If an unnecessary tool appears, the mentor can help the learner question why. If a proposed path is unnecessarily complicated, the learner can ask for alternatives and compare them. The student begins experiencing a fundamental principle of probabilistic self-learning: uncertainty does not require paralysis. It requires orientation.

Knowledge Does Not Have to Precede Action

Roundup does not assume that a participant must fully understand an environment before acting within it. In many cases, action is precisely what creates the context necessary for understanding.

A participant may install a development environment before completely understanding why one is necessary. They may create their first repository before understanding every property of source control. They may deploy a webpage before understanding the complete architecture behind web hosting. But once those things have actually happened, the terminology has somewhere to land. The learner has seen the files, watched the repository change, experienced a deployment and typed a URL into another device to see something they created appear. The abstraction is no longer purely abstract.

This allows learning to move continuously between experience and explanation. Roundup does not need to choose between teaching fundamentals first and simply letting AI handle the execution. Participants can act early enough to experience real progress while repeatedly stepping back to understand the systems through which that progress occurred. The first experience will be imperfect. That is expected. As the learner repeats the process, the environment becomes more recognizable, the questions improve and the learner becomes progressively less dependent on explicit direction.

Systems Understanding Is What Makes AI Steerable

Probabilistic self-learning does not mean outsourcing understanding to AI. It requires enough understanding for the learner to steer AI intelligently.

In a software Roundup, that means developing a basic mental model of computing itself. A young person should understand, at an appropriate level, that computers are physical systems built from components, that processors execute instructions, that programs consist of instructions and data, that memory and storage serve different functions, that information moves among components, that input and output connect computation to the outside world, and that connectivity allows separate computers and systems to communicate. This is not an attempt to produce electrical engineers. It is an attempt to give participants a mental map.

Once that map exists, later concepts have somewhere to attach. Cloud storage is no longer an inexplicable abstraction — it is storage provided by computing systems elsewhere and reached through a network. A smartphone becomes recognizable as a computer containing processing, memory, storage, input and output, sensors and connectivity. A temperature sensor becomes an input into a larger system that can receive information, process it, communicate it, store it and eventually act upon it. That is why a software Roundup can naturally expand from a simple website into local software, sensors, connectivity, gateways, cloud services and increasingly sophisticated applications. The participant is progressively experiencing a larger system while using AI to acquire the knowledge necessary to navigate each new layer. Systems thinking is not being taught as a detached concept. It is being developed through contact with an actual system.

The Artifact Is the Vehicle, Not the Curriculum

The website is not the curriculum. Neither is a home sensor network, a smoked brisket, an air-conditioning diagnosis or a floral arrangement. Those are destinations that give the learning process purpose and direction.

A culinary Roundup could begin with participants being told they are going to successfully smoke a brisket. Another Roundup might ask them to understand why an air conditioner is not performing correctly. A floral Roundup might begin with the responsibility of creating arrangements for a wedding. The domains, tools, terminology and expertise are completely different, but the underlying learning method remains substantially the same. In each case, the learner begins with a meaningful destination in territory they do not fully understand, establishes a working AI environment, maps the domain, acquires knowledge as gaps emerge, acts before

possessing complete understanding, encounters uncertainty through action, returns with better questions and continues.

The repeated movement between uncertainty, discovery, action, evaluation and adjustment is the curriculum. The artifact gives that curriculum somewhere real to go.

Recognized Contribution Gives the Learning Meaning

Roundup begins and ends with recognized contribution. A participant should not simply finish a lesson or satisfy an instructor. They should create something another human being can see, use, appreciate or value.

This is why the real website matters in the first software experience. A participant who begins the day without knowing what a development environment is and leaves with a public URL has experienced something substantially different from completing a tutorial. They can send the link to a parent or grandparent and truthfully say, I built this. The work has escaped the classroom and acquired a witness.

As the experience progresses, the contribution can become increasingly useful. A sensor could help a family understand conditions in its home. A moisture sensor could monitor a garden. A piece of software could solve a recurring household problem. In another domain, a floral arrangement could contribute to a real celebration or a meal could feed actual people. What matters is that the work crosses the boundary between educational exercise and genuine contribution. Recognition then becomes more than praise. It is evidence to the participant that their effort created value.

Repeated experiences of moving from unfamiliarity to useful contribution can progressively change a young person's relationship with the unknown itself. The durable outcome is not simply "I know how to build a website." It is closer to "I have entered environments I didn't understand before, and I know how to begin finding my way."

The artifact proves that they arrived. The recognized contribution gives the arrival meaning.

Why Multiple Domains Matter

The scalability of Roundup comes from separating its universal learning methodology from the specialized knowledge contained within each domain. Young people do not need to become experts in every Roundup they enter. In fact, moving among domains strengthens the underlying learning capability for a specific reason: the deliberate reintroduction of unfamiliarity is itself part of the training.

Once a participant becomes comfortable in software, moving them into culinary work or home systems restores the condition Roundup is designed to address. They are unfamiliar again. The methodology that worked in software must now be applied somewhere entirely different. When it works again, something important happens. The participant accumulates direct personal

evidence that the methodology transfers — that entering unfamiliar territory and becoming capable is not a one-time event but a repeatable outcome. Do that across enough domains and you are no longer merely teaching skills. You are producing evidence, in the learner's own experience, that they can enter unknown territory and find their way through it.

Over time, the participant also begins noticing that unfamiliarity itself has a recognizable structure. New domains contain systems that can be mapped. They contain tools whose purposes can be discovered. They contain terminology that can be learned in context, relationships that can be understood, experts whose judgment can accelerate learning, and constraints that must be respected. AI gives the learner a powerful mechanism for surfacing that previously inaccessible structure, while the domain expert provides the practical judgment necessary to keep the learning connected to reality.

The universal capability Roundup is attempting to develop is therefore not fluency in any particular domain. It is the progressive confidence and practical ability to say, I don't know how this works yet, but I know how to start figuring it out — and to mean it because experience has proven it true.

The Role of the Mentor

The availability of AI does not eliminate the mentor. It makes the mentor's judgment more important while changing the mentor's function.

In a probabilistic learning environment, the mentor is not primarily responsible for distributing information. AI can provide extraordinary quantities of information on demand. The mentor instead helps the participant develop judgment about what that information means in the actual environment. A capable mentor understands the domain well enough to recognize when an important consideration is missing, when an answer is technically plausible but practically wrong, when a learner is following instructions without understanding the underlying system, or when safety, ethics or real-world constraints require intervention.

The mentor also demonstrates something important: expertise does not mean possessing every answer. Experienced people investigate, ask questions, test assumptions, compare alternatives, consult others, make mistakes and revise their thinking. The mentor therefore becomes an experienced navigator accompanying someone who is learning how to navigate — less the person at the front of the room with the answers and more the person beside the learner helping them develop the judgment to find answers themselves. That relationship is especially important because a probabilistic environment cannot be completely scripted in advance. Different participants encounter different problems and pursue different paths. The mentor preserves direction and judgment without removing the productive uncertainty through which the learner develops.

The Intersection With Bobby Hodgkinson's Course Concept

Bobby Hodgkinson's Empowering Future Innovators: An AI-Driven Interdisciplinary Course for Small Business Solutions approaches this territory from a different but highly complementary direction. His proposal places students from different academic backgrounds into interdisciplinary teams working on authentic problems supplied by local businesses. The students are expected not merely to study AI but to build functional prototypes, deploy them into real environments, collect feedback, iterate and produce measurable value for organizations outside the university. His paper also explicitly contemplates extension into workforce retraining and AI literacy for local businesses — treating the university not merely as a classroom but as a mechanism for helping the surrounding community absorb practical AI capability.

That structure is strongly aligned with Roundup's emphasis on recognized contribution. Bobby's students are expected to create for somebody. The problem is real enough that the solution can succeed or fail in contact with an actual user, which produces precisely the kind of uncertainty in which probabilistic self-learning becomes useful. The student cannot turn to the back of a textbook. The customer may not fully understand what solution is needed. The prototype must encounter reality and be revised accordingly.

Probabilistic self-learning adds a different layer to that structure. Bobby's model creates an excellent environment in which uncertain, interdisciplinary, real-world learning can occur. Roundup attempts to make the learner's method for navigating that uncertainty explicit and transferable. Together the two concepts suggest an environment in which real community problems provide the destination, domain experts provide practical judgment, AI provides unprecedented access to knowledge and capability, and the learner develops a repeatable method for moving through what they do not yet understand. Neither concept needs to subsume the other. They address complementary sides of the same opportunity.

The broader Roundup proposition is not to make every young person an expert in software, culinary work, construction, healthcare, floral design or any other particular domain. It is to repeatedly place young people in unfamiliar but navigable territory and give them the working environment, mentorship and methodology necessary to find their way through it.

The artifact proves that they arrived. Recognized contribution gives the arrival meaning. Repetition across domains builds something more durable than any individual artifact: the confidence and practical capability to confront the unknown without waiting for someone else to provide the entire path.