Continuity Before Language?

A Mayorga Mnemosyne Hypothesis on How Intelligence Inherits a Past and Builds a Future

Francisco J. Mayorga, Jr.

Public Edition 1.2
Revised August 12, 2026 | First publicly published August 13, 2026
Independent Research Essay | Zenodo Archival Copy

Abstract

This essay asks a deceptively simple question: what allows intelligence to build on what came before? Language clearly matters, and memory clearly matters, but neither by itself guarantees that a person, an institution, or an artificial agent will inherit the past accurately enough to use it well. The essay calls that larger problem Continuity and develops it as a ladder, beginning with simple persistence and rising through stateful, adaptive, semantic, epistemic, and governed forms of inheritance.

The argument is intentionally broader than a technical paper and narrower than a claim that Continuity alone created civilization. It places the Mayorga Mnemosyne hypothesis in conversation with Merlin Donald, Bernard Stiegler, John Dewey, Michael Tomasello, Terrence Deacon, Daniel Dor, Jerry Fodor, Lev Vygotsky, research on wisdom and organizational memory, and recent AI memory architectures. Its central proposal is that cumulative intelligence requires more than the survival of information. The past must remain ordered, revisable, interpretable, and trustworthy enough to guide what happens next. Language transformed what humans could carry across minds and generations; contemporary AI makes visible how easily fluency and inheritance can come apart. The essay closes by asking what must survive if increasingly capable intelligence is to acquire not merely a memory, but a trustworthy history.

Keywords: continuity; cumulative intelligence; language; cultural evolution; AI memory; longitudinal agent state; epistemic continuity; governance; wisdom; Mnemosyne

Publication Status

Independent research essay, Public Edition 1.2. This edition underwent structured human and AI adversarial review, primary-source adjudication, source-integrity verification, and a reader-centered revision. The controlling prose was finalized August 12, 2026 and first published publicly on Medium on August 13, 2026. This Zenodo deposit preserves the full scholarly edition with front matter and numbered references. The work forms part of the continuing Mayorga Mnemosyne Continuity Framework™ research program.

Suggested Citation

Mayorga, Francisco J., Jr. “Continuity Before Language? A Mayorga Mnemosyne Hypothesis on How Intelligence Inherits a Past and Builds a Future.” Public Edition 1.2. First publicly published August 13, 2026. Zenodo archival deposit, August 14, 2026.

Copyright and Version

Yuval Noah Harari offers a compelling answer to one of the oldest questions we can ask about ourselves: how did an animal that was, for most of its existence, neither the strongest nor the fastest creature on the landscape become capable of building religions, markets, laws, states, sciences, corporations, bureaucracies, and technologies powerful enough to alter the planet? His candidate for the decisive breakthrough is *language*, especially the human ability to use words to do more than describe what is present: to coordinate around stories, rules, institutions, and imagined realities that no individual could sustain alone. In that sense, language became something like the operating system of collective life.[1]

The argument becomes even more provocative in the age of artificial intelligence, because we are now building engineered systems with extraordinary linguistic capability, systems that can enter the textual machinery of civilization at machine scale and operate inside institutions whose daily life is already made of documents, instructions, contracts, precedents, reports, messages, and rules.

It is a powerful argument. Yet it leaves another question open, one that becomes harder to ignore the longer we work with intelligent systems that can reason brilliantly in one moment and arrive at the next as though yesterday had been partly washed away.

What allows any intelligence, linguistic or otherwise, to acquire a history?

A brilliant inference can happen once and later survive only in fragments. A discovery can be made, admired, and then inherited imperfectly, with its conclusion preserved but its reasoning lost, or its details remembered while the larger insight that gave them meaning slowly disappears. A society can solve a problem and later repeat it because the rule survived while the reason for the rule did not. An artificial agent can return to a task carrying a remarkably accurate inheritance from what came before, a compressed outline missing crucial constraints, scattered details detached from their original rationale, or only enough fragments to reconstruct what probably happened. Sometimes that reconstruction lands on the bullseye. Sometimes it lands nearby. Sometimes it misses the target altogether, and sometimes the more troubling failure occurs: the missing pieces are filled with plausible assumptions that were never part of the original history. In each case, intelligence may still be present. What has failed is not necessarily intelligence, nor even memory in the absolute sense, but the fidelity of Continuity by which one episode becomes a trustworthy inheritance for the next.

The real question is not whether language matters or whether memory matters. Of course they do. The more interesting question is whether intelligence needs something else in order to build on what it has already learned.

Imagine an intelligent system that can speak beautifully, solve difficult problems, and explain complicated ideas, but every time it begins a new task it has only an incomplete picture of what happened before. Perhaps it remembers the main goal but forgets an important constraint. Perhaps it remembers the facts but not why a decision was made. Perhaps it finds only fragments of the past and has to guess how they fit together. It may still be highly intelligent, but its intelligence cannot reliably build upon its own history.

That is what I mean by *Continuity*.

Continuity is the ability to carry enough of the past forward so that what happened before can still guide what happens next. The important information must not merely survive somewhere. It must remain connected to when it happened, what it meant, what changed afterward, and whether it is still relevant. Just as importantly, it must be possible to correct or replace old information when new evidence shows that the past no longer tells the whole story.

Seen this way, language and Continuity are not the same thing, and there is no need to make them competitors in a contest over which one "caused" civilization. Their relationship is probably more like a partnership that developed over time. A recursive relationship. Long before fully developed language, living beings could remember, learn skills, imitate others, recognize relationships, and carry experience from one moment into the next. Those abilities may have helped create the conditions in which language could emerge.

Then language changed the game.

It allowed human beings to preserve things that were much harder to carry forward through memory and imitation alone: explanations, reasons, warnings, promises, laws, stories, plans, beliefs, imagined futures, and lessons drawn from events that the listener had never personally experienced. In that sense, language did not simply give Continuity a bigger container. It changed what could be placed inside it.

And once language became woven into human thought, culture, and institutions, the relationship became even deeper. Language helped human beings preserve increasingly complex forms of knowledge, while those accumulated forms of knowledge gave later generations a richer world in which to think, speak, learn, and create. Continuity helped intelligence inherit a past, and language greatly expanded what that inheritance could contain.

The question mark in the title therefore earns its place. Continuity Before Language? is not a declaration that one mechanism marched onto the stage first and another followed behind it. It is a question about levels, dependencies, and transformations: which forms of Continuity can exist without language, which forms language changed almost beyond recognition, and which forms an increasingly capable artificial intelligence must deliberately engineer if it is to acquire not merely memory, but a trustworthy history.

When Intelligence Learned to Inherit

Imagine a small group of early humans living beside a river. Their tools are simple, but they matter enormously. A better cutting edge can mean the difference between opening an animal hide before the meat spoils, shaping a wooden spear more precisely, or finishing a task before darkness comes.

One morning, a hunter sits with a piece of stone between his knees.

He has made tools before. He knows the familiar rhythm. Strike here. Turn the stone. Strike again. Most of the flakes come away thick and uneven, as they always have. Some are useful. Some shatter. Some leave an edge that works for a while and then breaks.

Then something different happens.

Perhaps by accident at first, he tilts the stone slightly before striking it. The blow lands at a different angle. A long, thin flake peels away, leaving behind an edge sharper and cleaner than the ones he usually makes.

He stops.

He looks at it.

Then he tries again.

The second strike works too.

Now the accident has become a possibility.

He spends the afternoon experimenting. He changes the angle of his hand. He changes the place where the hammerstone lands. Sometimes he fails. Sometimes the stone breaks badly. But slowly he learns what matters. By evening, he can produce the better edge more reliably.

Something new now exists in the world.

Not merely another stone tool, but a new way of making one.

The discovery may serve him well for years. His knives are sharper. His work becomes easier. Perhaps the others notice that his tools seem better, but nobody learns exactly how he makes them. He never explains the crucial motion. Maybe he does not know how to explain it. Maybe someone watches him and sees only a series of blows against stone. The difference lives partly in the angle of his wrist, partly in the order of the strikes, partly in knowledge that has become so natural to him that he no longer notices what another person would need to be taught.

Then, one winter, he is gone.

Perhaps he dies during a hunt. Perhaps illness takes him. Perhaps he simply grows old.

The people who remain still have stones.

They still have hands.

They still have intelligent minds capable of discovery.

But they no longer have his discovery.

The improved technique has disappeared with the person who carried it.

Years pass.

Someone else may eventually discover the same trick. Perhaps it takes ten years. Perhaps fifty. Perhaps three generations. When it happens, everyone celebrates the improvement as something new, even though it had already existed once before.

This is the strange tragedy of intelligence without reliable inheritance.

A mind can climb a mountain, disappear, and leave no path behind.

Now imagine the same beginning again.

Another person, somewhere else or perhaps generations later, sits beside another fire with another piece of stone. Through observation, experimentation, and luck, he discovers the same improved strike.

Again the edge comes away sharper.

Again he experiments until he understands the motion.

But this time, something else happens.

A younger member of the group, maybe his son, is watching.

At first the child copies the movement badly. The stone breaks. The angle is wrong. The inventor takes another stone and shows the motion again, more slowly this time. He places the child's hand differently. He points to where the blow should land.

The child tries again.

Better.

Not perfect, but better.

Days later, the child can reproduce the technique without help.

The discovery has now crossed a boundary.

It no longer exists only inside one person.

Years pass.

The child becomes an adult and teaches someone else. That person notices that the tool works even better if the stone is prepared differently before the final strike. Another person changes the grip. Someone else discovers that a particular kind of stone produces a stronger edge.

None of them has to begin with the original problem.

That is the crucial difference.

The first inventor had to ask, in effect, "How can I make a better cutting edge?"

His descendants begin with a different question.

"Given this good cutting edge, how can we make it better?"

That may sound like a small change, but across generations it changes everything. A butterfly effect of sorts.

The solution of one generation becomes the starting point of the next.

A discovery becomes a technique.

The technique becomes a tradition.

The tradition becomes a platform from which new discoveries can be made.

After enough time, a descendant may be making tools so different from the original that the first inventor would barely recognize them. Yet hidden inside those later tools is still a trace of that first successful strike beside the fire.

No single person invented the whole thing.

No single mind needed to.

The achievement belongs partly to intelligence, but also to time, inheritance, correction, imitation, and improvement.

Now place the two lineages beside each other.

In the first, intelligent people repeatedly solve problems and then lose the solutions. Each generation is forced to pay again for knowledge that someone else already purchased with effort, failure, and insight.

In the second, the price is paid once, and the receipt survives.

The next generation does not have to buy the same lesson again. It can spend its effort somewhere new.

At the moment of discovery, the two inventors may have been equally intelligent.

Their descendants, however, do not inherit equal worlds.

One lineage keeps returning to the same starting line.

The other begins each generation a few steps farther down the road.

And after enough generations, those few steps become miles.

This is why the difference between intelligence and cumulative intelligence can be difficult to see if we look only at the individual moment. The spark itself may be identical. What changes history is whether the spark disappears when the hand that made it is gone, or whether it lights another fire, which lights another, until no one remembers who struck the first flame.

One lineage possesses intelligence.

The other possesses intelligence that can inherit.

And once intelligence can inherit, it can begin to have a history.

That contrast is why one of the original formulations behind the *Mayorga Mnemosyne* hypothesis still matters:

Continuity converts intelligence from an isolated event into an accumulating process.

There is an obvious objection to all of this, and it is better to face it directly.

Someone could reasonably say, "Of course cumulative intelligence needs something from the past to survive. That is what cumulative means." And that person would be right. At one level it would sound nearly tautological. If my entire argument were simply that accumulation requires memory, there would not be much of an argument at all. We would only be giving a new name to something already contained in the definition.

The more interesting question begins one step later, when we leave the dictionary and enter the architecture.

What exactly has to survive, and in what condition, for intelligence to truly build upon what came before?

Imagine an artificial intelligence that has read millions of books, can speak with extraordinary fluency, and can solve difficult problems. It may possess an enormous amount of stored information and still have trouble carrying its own history forward reliably.

It may remember that a decision was made but forget why it was made.

It may remember the goal but lose an important restriction that was supposed to govern how the goal could be achieved.

It may retrieve an old instruction without realizing that a newer instruction replaced it.

It may remember several pieces of a previous discussion but lose the connection between them.

It may even find only fragments of the past and fill the missing spaces with a reasonable guess, creating a version of history that sounds convincing but never actually happened.

This is why having memory is not the same as having Continuity.

A library can contain every edition of a book ever published, but unless we know which edition came first, which corrected the others, which ideas were abandoned, and which version should guide us now, the library alone cannot tell us the history of the book.

Continuity requires something more.

The past must remain available in a form that allows the present to understand where it came from. We need to know what happened, when it happened, what changed afterward, what is still valid, what has been replaced, and which parts of that inherited history should influence the next decision.

And this ability does not automatically appear just because a system is intelligent or because it is excellent with language.

A system may speak brilliantly and still inherit its own past poorly.

That is where the argument becomes more than a definition.

The important claim is not simply that intelligence needs memory. It is that cumulative intelligence needs an architecture capable of carrying useful state from one episode into the next while preserving enough order, context, revision, and connection for the past to remain trustworthy.

That is the architectural problem I mean by Continuity.

Now we are no longer speaking about a tautology. We are speaking about an architectural property.

Others Have Walked This Road Before

Before going farther, there is an obligation that matters especially in an essay about inheritance: we should know whose footprints are already on the road. Ideas rarely arrive as immaculate births. More often they arrive as rivers, fed by tributaries whose names we may not know until we travel upstream. During my research, the deeper I followed the question of Continuity, the clearer it became that several thinkers had already explored parts of this landscape, often with different words, different purposes, and different maps.

That discovery sharpened my Mayorga Mnemosyne hypothesis and made the intellectual landscape around it more interesting, because originality is not always a matter of being the first person to notice a mountain. Sometimes it lies in seeing how several mountains belong to the same range, and then asking a question that becomes visible only when the range is viewed as a whole.

Merlin Donald: before language, there was already a human story

Merlin Donald is perhaps the closest predecessor on the evolutionary side. His work asks us to imagine human cognition not as a switch that language suddenly turned on, but as a series of cultural and representational transitions. Long before writing, and before anything like modern institutions, our ancestors could imitate, rehearse, gesture, demonstrate, coordinate bodily action, and preserve skilled practices through what Donald calls mimetic culture. Later came the great expansion made possible by language and narrative, and later still the external symbolic systems of writing, diagrams, archives, and theoretic culture.[2,3]

The sequence matters because it gives us a picture of intelligence acquiring layers of inheritance. A stoneworker does not need a written manual to teach a strike. A dancer does not need a paragraph to preserve a movement. A hunting party can transmit timing, posture, attention, and practiced coordination long before anyone writes down a rule. Something can therefore survive from mind to mind and body to body before mature language becomes the main carrier.

Donald later describes cultural memory as shared knowledge and skill that are retained and updated through time. That phrase should make any Continuity theorist sit up. It means that the human cognitive story is not only about what an individual brain can do at a moment. It is also about what a group can keep available long enough for later people to inherit, revise, and extend.

The resemblance to the present argument is real, and it should be acknowledged plainly. Donald had already walked deeply into the territory where prelinguistic skill, cultural inheritance, language, and external memory meet. The Mayorga Mnemosyne proposal does not erase that ancestry. It builds downstream from it.

Bernard Stiegler: when memory leaves the skull

Bernard Stiegler approaches the problem from another direction. If Donald asks how human cognitive culture evolved, Stiegler asks what happens when memory escapes the nervous system and takes up residence in things. A mark on clay, a written sentence, a photograph, a recording, a database, and eventually a digital network can all preserve traces of experience outside the body. Stiegler calls this domain of externalized memory tertiary retention.[4]

The easiest way to understand the importance of that move is to imagine a world in which every book vanished when its author died. Homer would have to be reinvented. Euclid would have to be rediscovered. Every recipe, map, legal judgment, laboratory notebook, and engineering diagram would be hostage to the lifespan of the person carrying it. Writing changed that bargain. It allowed a human thought to survive the human who first had it.

But Stiegler's point is more subtle than 'we invented storage.' Once memory lives outside us, we begin living inside the memory we have externalized. A child born into a literate society does not merely possess a brain plus some

books. The books, schools, clocks, diagrams, institutions, media, and technical systems become part of the environment in which that brain learns what thinking is. The tool that stores memory starts shaping the mind that will later create new memory.

This is a crucial connection for the argument of this essay. Continuity is not a passive warehouse at the edge of cognition. The way a society preserves its past can alter the possibilities available to its future. A civilization writes itself into an environment and then raises its children inside what it has written.

John Dewey: experience that changes the next experience

John Dewey is a quieter ancestor, but perhaps the easiest to recognize in ordinary life. His principle of continuity in experience says, in essence, that an experience does not end neatly when the moment ends. What happens today changes the person who meets tomorrow. What is learned in one situation becomes part of the equipment with which the next situation is faced.[5]

Think of a child touching a hot pan. The useful inheritance is not the existence of a memory file labeled 'hot pan.' The experience changes the child's next approach to heat. Or think of an experienced nurse who walks into a room and notices that something is wrong before she can fully explain why. Years of prior encounters have changed the pattern by which the present is perceived. The past is not simply remembered. It has become part of how the next moment is interpreted.

Dewey also connects meaningful experience with consequence. We act, the world answers, and the answer changes the next act.[6] That simple loop will matter later when we reach the distinction between learning and wisdom, because a past that cannot alter judgment is little more than an archive.

These predecessors change how I want to speak about originality. The claim is not that no one before me noticed that memory, culture, technics, or inherited experience matter. That would be historically untenable. My narrower proposal is to connect these lines of thought to a layered architecture of Continuity, and then to follow the problem into a new environment where it becomes unusually visible: **artificial intelligence.**

> **We are not the first travelers to notice the river. The question is whether we can map enough of its tributaries to understand what must survive, what can change, and what happens when the inheritance itself begins to act.**

How Knowledge Learns to Climb

The stone-tool story from the previous section gives us the intuition. Research on cumulative cultural evolution gives that intuition a more precise shape. Scholars often use the image of a ratchet, because a ratchet allows motion in one direction while preventing everything from slipping back to the starting point.

Imagine trying to build a staircase while someone removes the lower steps every night. You may be brilliant. You may work quickly. You may even reach impressive heights during the day. But each morning you are back on the ground, spending your best effort reconstructing what yesterday already solved. Cumulative culture changes the situation because later people begin on steps they did not build themselves.

Human culture is extraordinary in this respect. We preserve techniques, combine them, modify them, argue about them, standardize them, teach them, and sometimes improve them so gradually that no living person could identify the first moment when the modern form appeared. A violin, a legal system, a jet engine, a surgical procedure, or a computer language is not the work of one mind. Each is a long conversation among minds separated by distance, time, and death.

Yet this is where the argument must resist a temptation. It would be easy to say that the ratchet is simply Continuity and declare the case closed. That would be convenient, but wrong. The ratchet has more than one tooth.

Michael Tomasello and related researchers place shared intentionality near the center of the human story. Human beings do not merely notice what others do. We can attend to something together, understand that we are trying to accomplish a shared goal, teach deliberately, imitate the process rather than only the result, correct one another, conform to group norms, and eventually exchange reasons for why something should be done one way rather than another.[7–9]

Picture two people carrying a heavy table through a narrow doorway. Each must understand more than the table. Each must understand something about the other's attention, intention, position, and next move. Now expand that basic capacity across generations. A community can say, in effect, 'This is how we do it, this is why we do it, and I am going to help you learn it.' That social coordination is not identical to memory, and it is not identical to Continuity.

This distinction matters because theories become useless when they eat their neighbors. If every mechanism that helps something persist is renamed Continuity, then Continuity explains everything and therefore explains almost nothing. Shared intentionality should remain its own mechanism. Teaching should remain its own mechanism. Language should remain its own mechanism. Reasoning should remain its own mechanism.

They can work together as musicians in an orchestra. The fact that the violin and the cello play the same symphony does not make them the same instrument.

The animal evidence reinforces the point. Nonhuman species show rich social learning, traditions, memory, tool use, imitation, and in some cases limited forms of cumulative change. Chimpanzees, corvids, cetaceans, and other species repeatedly remind us that human beings did not invent every ingredient of learning.[8,10]

But the difference in scale remains enormous. No established nonhuman analogue yet resembles the open-ended, deeply recombinatorial, institutionally extended culture of humans, where discoveries can be copied, critiqued, recombined, codified, and passed through organizations whose members never meet the original discoverers.

The conclusion is therefore both limited and consequential. Some capacities relevant to Continuity clearly exist without mature natural language. That does not mean language is unnecessary for civilization. It means we should not confuse the first footholds of inheritance with the entire mountain.

When the Absent Became Portable

There is a seductive engineering metaphor for language: “language increases the bandwidth of Continuity.” I used a version of that metaphor earlier in my work because it seems intuitive. If memory is the cargo and communication is the pipe, language appears to widen the pipe so that more information can pass through.

The trouble is that the metaphor is too small for what language actually did.

A wider pipe carries more of the same substance. Language often allows us to carry something that could not have traveled in the old form at all.

Consider a grandmother warning a child about a flood that happened fifty years before the child was born. The child cannot see the water, hear the storm, smell the mud, or watch the houses disappear. The event is absent. Yet the grandmother can use language to reconstruct enough of that absent world in another mind that the child may someday act differently because of it.

Or imagine a commander explaining a battle plan for terrain no soldier has yet seen, a scientist describing an invisible mechanism, a parent telling a child what might happen if a promise is broken, a judge explaining why a precedent applies to a case whose participants were not alive when the precedent was created. In every case, language does something more ambitious than point at the present.

Daniel Dor calls this the instruction of imagination. His central insight begins with an obvious fact that is easy to overlook: no person can directly inhabit another person's experience. There is always an experiential gap between minds.[11]

Pointing and imitation work beautifully when the relevant thing is present. I can point to the fruit. I can show you how I grip the spear. I can make a face when I taste something bitter. But what if the thing that matters is not here? What if it happened yesterday, will happen tomorrow, exists only as a possibility, concerns a person who is absent, or depends on a reason that cannot be demonstrated by a gesture?

Language crosses that distance.

Language made the absent portable.

That sentence captures more than a gain in quantity. It captures a change in kind. Language made it possible to carry explanations, reasons, promises, warnings, laws, counterfactuals, imagined futures, private experiences, moral obligations, and stories about the dead into minds that never directly witnessed the events from which those things arose.

This is why the relationship between language and Continuity cannot be reduced to 'memory first, language later.' At the richest human levels, language does not merely package an already complete meaning and ship it elsewhere. Often the linguistic act helps construct the meaning in the receiving mind. The message becomes part of the state it transmits.

And yet Dor's own evolutionary account complicates a simple language-first story. He describes prelinguistic communities capable of collective invention, mimesis, and substantial social coordination, and then treats language as a collectively invented technology that transformed the people who invented it.[12]

That is a wonderfully recursive picture. Earlier social capacities help produce language. Language expands what can be shared. What can be shared changes culture. Culture reshapes development. The children born into that transformed culture learn language inside a world already saturated with the accumulated products of earlier language.

The river has begun to bend back upon itself.

What Shape Must a Thought Take to Survive?

Once we say that language can carry absent worlds, another question appears. What exactly is being carried, and in what form must it exist if a mind is going to compare it with something new, revise it, reject it, or combine it with another idea?

This question sounds technical, but an everyday example makes it easier. Suppose I tell you, 'The meeting is at three.' Later I tell you, 'The meeting has been moved to four.' A useful memory system cannot simply preserve both sentences forever and congratulate itself for excellent retention. It must understand that the second statement changes the status of the first.

Now make the problem harder. Suppose the meeting moved because a client could not attend, but only if the legal team approves the change; suppose the legal team later approves it for this week but not as a standing rule. The information is no longer a pile of sentences. It has structure: time, conditions, causes, authority, revision, scope, and relationships among claims.

Jerry Fodor's Language of Thought hypothesis enters at this point. It is often misunderstood as the claim that people secretly think in English, Spanish, or some whispering inner version of speech. That is not the point. Fodor and Pylyshyn argue that systematic thought may require an internal representational format with combinatorial, syntax-like structure.[13]

For the present hypothesis, the importance of Fodor is not that natural language must come before Continuity. The challenge is subtler: once Continuity becomes rich enough to preserve beliefs, reasons, contradictions, alternatives, and revisions, perhaps the system needs some structured way of representing those relationships.

At the lowest levels, this may not matter much. A biological system can be changed by prior experience without writing a sentence about it. A thermostat can retain a setting. A neural network can alter a response because its internal parameters changed. The past can influence the future without anything like ordinary language.

But if a system must know that 'I believed X because of evidence A, then evidence B undermined A, therefore X was superseded by Y, except under condition C,' the representational problem becomes harder to ignore.

That is not a defeat for Continuity. It is a boundary marker. It tells us that different forms of Continuity may require different kinds of internal organization.

The useful question is not whether Continuity can exist without representation. It is which forms of representation are required for which forms of Continuity.

That question can be investigated. It can be tested in brains, artificial systems, institutions, and perhaps even in simpler biological organisms. And because it is testable, it is more useful than turning 'Continuity before language' into a slogan that must be defended at all costs.

When Language Moves Inside the Mind

So far we have mostly discussed language as something that moves between people. Lev Vygotsky forces us to consider what happens when language moves inward.

Children often speak their thinking aloud. They narrate a difficult puzzle, tell themselves what to do next, repeat a rule, or talk their way through frustration. Over development, much of that overt speech becomes inner speech. The social tool becomes part of the machinery of self-regulation.[14]

Anyone who has whispered a phone number to keep it alive for ten seconds has felt a simple version of this. We recruit language to hold something in place. We tell ourselves, 'Do not forget the keys.' We rehearse instructions. We name a feeling so that we can think about it. We silently argue with ourselves before making a decision.

Modern research shows that this linguistic scaffolding is real. When people are forced to occupy their speech system with an irrelevant repeated sound, some planning and executive tasks become harder.[15] Language is therefore not merely a public courier carrying finished thoughts from one person to another. In mature humans, it can become part of the workshop in which some thoughts are held, organized, and controlled.

But the evidence also prevents us from going too far in the opposite direction. Severe aphasia can devastate a person's ability to produce or understand language while leaving substantial abilities in arithmetic, logic, navigation, music, theory of mind, and other forms of reasoning. Neuroimaging likewise suggests that several kinds of nonlinguistic cognition do not simply reduce to the core language network.[16] A recent preprint on formal logical reasoning strengthens the possibility that some forms of structured inference can survive profound language impairment.[17]

This creates one of the most important distinctions in the essay. Language can support thought without being identical to thought. It can scaffold a capacity without being the only material from which that capacity is made.

The developmental story remains difficult. A scaffold can be removed after a building stands. If language helps shape a cognitive ability during childhood, the fact that an adult can later perform part of that ability without online language does not prove that language played no role in building it.

So once again the answer becomes recursive. Thought helps make language possible; language enters thought; language-trained minds then perform some operations that are no longer continuously dependent on overt linguistic processing. The tool has entered the craftsman.

This is why the title of this essay needs its question mark. 'Before' may have different answers at different levels. Some forms of persistence and adaptation can exist without mature language. Other forms of human semantic and epistemic Continuity may be so deeply shaped by language that asking which came first becomes less useful than asking how the loop developed.

A City That Forgets Every Morning

Thought experiments are useful because they let us exaggerate one variable until a hidden distinction becomes visible. They are dangerous for exactly the same reason, because an author can quietly arrange the scenery so that the favored conclusion has nowhere else to go.

So consider a thought experiment designed to isolate the distinction.

Imagine a city in which everyone wakes each morning with perfect command of language. By stipulation, people still know how to speak, read, argue, joke, persuade, and explain. Their vocabulary is intact. Their grammar is intact. Their intelligence is intact.

But the city's inherited state is gone.

Yesterday's personal experiences do not carry over. Books have vanished. Databases are blank. Scientific observations are gone. Court records have disappeared. Bank balances have reset. Government files are empty. No one remembers which bridge was found unsafe, which medicine failed a trial, which suspect was cleared, which promise was made, which debt was paid, or which war was fought for which reason.

Every sunrise delivers a city full of eloquent strangers standing inside the ruins of yesterday's knowledge.

The citizens could begin again. They could make discoveries. They could form relationships. They could create records during the day. They might even accomplish remarkable things before nightfall.

Then midnight would collect the bill.

The point is not that this proves Continuity came before language in human evolution. It proves no such thing. Deacon, Dor, Vygotsky, and Tomasello can all accept the scenario without surrendering their theories.[7,11,14,18]

What the city makes visible is narrower: linguistic competence and longitudinal inheritance are not the same capacity.

A society can remain eloquent and still be unable to accumulate.

Now reverse the experiment, but do it honestly. Suppose conventional symbolic language disappears while perception, memory, imitation, demonstration, procedural skill, social learning, and durable nonsymbolic artifacts remain. People could still learn from one another. They could still show. Copy. Practice. Coordinate. Remember. They might preserve traditions and improve some techniques.

How far could they go? We do not know.

That uncertainty marks the boundary where the empirical problem begins. Comparative culture, developmental psychology, archaeology, and language-evolution research must answer how far cumulative culture can ratchet without the symbolic machinery that lets humans transmit absent and hypothetical worlds.

The experiment does not have to settle which capacity came first to reveal that they can come apart. Language can remain while the machinery of inheritance collapses. That distinction is enough to expose the architectural problem.

Putting the Braid on the Workbench

Human evolution gives us a frustrating research problem. Memory, language, embodiment, social learning, development, emotion, culture, and institutions arrive already braided together. We cannot travel back seventy thousand years, turn language off for one population, leave everything else unchanged, and watch what happens.

Artificial intelligence gives us something closer to a workbench.

A modern language model can write an essay, summarize a contract, answer a technical question, translate a paragraph, or conduct a sophisticated conversation. Yet the same system may enter a later interaction with only partial inheritance from what happened before. It may receive a summary. It may retrieve selected fragments. It may have access to stored conversations. It may reconstruct context from external memory. Sometimes the reconstruction is excellent. Sometimes it is incomplete. Sometimes it retrieves the wrong thing. Sometimes it fills a gap with a plausible story.

This is why 'stateless' versus 'stateful' can also be too simple for the phenomenon we actually observe. Real systems increasingly occupy a messy middle ground: partial memory, compressed memory, retrieval-dependent memory, selectively available memory, and reconstructed memory.

The practical result is familiar to anyone who has worked on a long project with an AI assistant. The system may remember the mission and lose the constraint. It may remember a decision and lose the reason. It may remember the details and lose the hierarchy among them. It may recover a summary that is broadly right while a single missing qualification changes what should happen next.

In other words, the problem is often not absolute amnesia. It is fidelity of inheritance.

Engineering researchers have responded by building explicit mechanisms around language models. MemGPT, for example, treats context as something that must be managed and supplements the model with external memory so that an agent can persist across longer interactions. LongMemEval evaluates long-term interactive memory through tasks such as multi-session reasoning, temporal reasoning, knowledge updating, information extraction, and abstention.[19,20]

These systems do not tell us how the human mind evolved. What they show is that excellent language performance does not automatically solve the problem of carrying interaction-specific history forward reliably.

A system can speak brilliantly and still inherit itself badly.

And the word memory is often too coarse to describe why. A system may retrieve an old instruction but fail to know that it was superseded. It may preserve every message while failing to know which decision is authoritative. It may keep the goal and lose the boundary. It may remember the latest answer and forget the evidence that made the answer reasonable.

A machine can therefore possess a larger archive and a poorer history.

Samuel Sameer Tanguturi's ATANT framework enters this conversation by defining AI continuity through properties such as persistence, updating, temporal ordering, disambiguation, reconstruction, model independence,

and operational usefulness. Its companion Continuity Layer paper argues that an evolving narrative needs more than retrieval of semantically similar fragments; the system must reconstruct where things currently stand.[21,22]

That is important adjacent work, and it gives part of the problem a clearer engineering vocabulary. Its benchmark results remain author-reported results from a preprint and reference implementation, so they should not be treated as independent proof of a universal construct. But ATANT helps sharpen a question that matters greatly for the Mayorga Mnemosyne work: once a system reconstructs the present state correctly, what else must survive for the reconstruction to be trustworthy enough to guide action?

From a Trace to a Trustworthy History

At this point the word Continuity is in danger of becoming too large. If it means everything from a file existing on disk to a civilization preserving law, then the concept will collapse under its own generosity.

The simplest solution is to stop pretending that all forms of Continuity are equally demanding. A footprint in wet clay and a constitutional tradition both connect past and future, but they do not do remotely the same work.

I therefore propose a ladder. Each rung asks more of the past than the rung below it. The ladder is not meant to describe one physical mechanism hidden inside brains, organizations, and computers. It is a way of asking, step by step, what the past is allowed to do to the future.

1. Persistence: did anything survive?

At the first rung, the question is almost embarrassingly simple: is there still a trace?

A photograph survives. A scar survives. A file survives. A written rule survives. A cached value survives. Persistence means that something from an earlier state remains after the moment that produced it has passed.

This is necessary for anything richer, but it is not yet intelligent Continuity. A tomb can preserve a body. A hard drive can preserve a file. Neither, merely by preserving, knows what the preserved thing means or whether it should matter tomorrow.

2. Stateful Continuity: where does this piece belong in the story?

Now add order.

Imagine a medical chart containing every blood-pressure reading a patient has ever had, but with the dates removed. The information survived, yet much of its usefulness disappeared. We no longer know what came before treatment, what happened after treatment, or whether a frightening reading is current or ten years old.

Stateful Continuity requires enough identity, scope, and temporal relationship to distinguish past from present and current from superseded. A system should be able to say, 'This was true then; this is what replaced it; this is what remains active now.'

An append-only archive may therefore have excellent persistence and poor Stateful Continuity. It has kept the pages but lost the edition.

3. Adaptive Continuity: did the past change what happened next?

At the third rung, history stops being a shelf and starts becoming an actor.

A child touches the hot pan and approaches the next pan differently. A driver skids on a wet road and changes how quickly he takes the next curve. A team launches a product, watches customers misunderstand it, and changes the next design. The previous episode has altered the next encounter.

This is Adaptive Continuity: prior state affects later behavior, and consequences can change the state that will govern future behavior. Dewey's continuity of experience belongs naturally here, because the important thing about experience is not merely that it happened, but that it modifies what comes after.[5]

The past has become a participant.

4. Semantic and Epistemic Continuity: do we remember why?

Now the ladder becomes more demanding.

Suppose an organization remembers that it stopped using a certain supplier. That is useful. But years later a new manager asks why. Was the supplier too expensive? Unsafe? Unreliable? Sanctioned? Was the decision temporary? Was it based on evidence later overturned?

If the organization preserved only the conclusion, the rule may survive while its meaning decays.

Semantic and Epistemic Continuity preserves more of the genealogy: rationale, evidence, provenance, uncertainty, causal interpretation, contradictory evidence, rejected alternatives, and the reason an earlier belief or decision was superseded.

This is the level at which 'memory' begins to feel painfully inadequate as a word. Knowing what the answer was is not the same as knowing why the answer deserved to be trusted.

Here structured representation becomes a serious theoretical issue, and language may become profoundly important. Human beings often preserve reasons through sentences, arguments, diagrams, stories, explanations, citations, and records. The state is no longer merely a point. It has ancestry.

5. Governed Cognitive Continuity: even if we remember, are we allowed to act?

The highest proposed rung introduces a distinction that becomes crucial when intelligence can do more than advise.

Imagine an AI system reviewing a factory incident. It correctly remembers that production was paused. It knows when the pause occurred. It knows why the pause occurred. It has the evidence, the failed tests, the later repairs, and the criteria that would justify reconsidering the pause.

Now suppose the evidence suggests the factory is probably safe again.

Is the AI authorized to restart production?

That is a different question.

Understanding why an action makes sense does not imply being authorized to perform it.

Governed Cognitive Continuity asks whether the system carries forward not only state and rationale, but also the constraints that determine how inherited knowledge may legitimately shape future action: authority, permissions, mission boundaries, negative knowledge, unresolved uncertainty, consequence history, justified supersession, human governance, and trajectory integrity.

A system that remembers the destination but forgets the guardrail has not inherited the mission faithfully. It has carried forward the accelerator and lost the brakes.

This is where the Mayorga Mnemosyne research program proposes to extend beyond ordinary memory and current-state reconstruction. The question becomes not merely 'What must survive?' but 'What must survive so that what comes next is still justified?'

The Brakes and the Why

This upper rung did not begin for me as an abstract philosophical distinction. It emerged from the far less glamorous experience of trying to carry long-running work across humans and AI systems without losing the meaning of what had already happened.

A handoff can look excellent and still fail.

It can list the tasks correctly but omit why the tasks exist. It can preserve a decision and omit the evidence that justified it. It can report what succeeded and omit the failed path that should never be repeated. It can remember that a system is paused and forget that only a particular human owner is authorized to unpause it.

The successor may therefore inherit the facts and still inherit the situation incorrectly.

This is one of the most dangerous forms of partial amnesia because it does not always announce itself as forgetting. Sometimes the handoff looks complete. The missing piece becomes visible only when the successor makes the wrong next move.

That experience produced two Mayorga Mnemosyne doctrines that are deliberately simple because they need to survive handoffs themselves:

Constraints and Meaning Must Both Survive.

A handoff is not Continuity unless it preserves both the brakes and the why.

Consider again the paused factory. A basic state reconstruction says, 'The factory is paused.' Better Continuity says, 'The factory was paused after three safety failures; these were the failures; these repairs were attempted; these tests remain unresolved.' Governed Continuity adds one more fact that may matter more than all the others: 'Even if the evidence now looks favorable, I do not possess the authority to restart it.'

This is the difference between knowing the road and possessing the keys to the vehicle.

As AI systems become more agentic, this distinction grows more important. A system that remembers more can act more coherently over time, but coherence is morally neutral. A mistaken objective that survives for one answer is an error. A mistaken objective that persists, recruits new actions, interprets new evidence through itself, and survives multiple handoffs can become a trajectory.

For that reason, the unit of governance cannot always be the isolated output. We may need to evaluate the evolving sequence: What objective is persisting? What evidence changed? Which constraints still bind? Who still has authority? What failed earlier? Has the mission drifted while the words describing it remained almost unchanged?

Long-horizon intelligence depends on some form of Continuity. Governed Continuity is an attempt to keep that long horizon from becoming a long leash with no hand on the other end.

When Institutions Remember the Rule and Forget the Reason

Harari's description of AI as a 'bureaucratic native' is powerful because bureaucracies are already built from the materials AI handles well: forms, files, rules, ledgers, cases, messages, precedents, dates, permissions, and decisions.[27]

But bureaucracy also gives us a remarkably human example of Continuity failure.

Imagine an employee who asks why a certain approval requires three signatures. Everyone knows the rule. No one knows its origin. One person vaguely remembers that there was once a fraud incident. Another thinks it came from an audit. A third says, 'We have always done it this way.' The original decision may have been wise, foolish, temporary, or legally required. The procedure survived. Its ancestry did not.

That organization has memory. The rule is right there in the manual.

What it may not have is trustworthy Continuity.

Organizations can retain enormous archives and still suffer institutional amnesia when rationale, tacit knowledge, exception conditions, causal history, and practical judgment disappear as people leave. Organizational-memory research has long emphasized that organizational remembering involves more than keeping documents; it includes how knowledge is acquired, stored, retrieved, and used.[24]

A bureaucracy is therefore best understood here as a linguistic persistence architecture. It gives cases identities. It carries balances forward. It records who approved what. It allows a legal matter to outlive the people who opened the file. In a well-functioning institution, the dead can continue to constrain the living through records, precedents, reasons, and obligations.

But a filing cabinet is not an institution, just as a map is not a journey. A case file whose rationale has disappeared resembles a map whose legend has been lost: all the marks may still be visible, yet the successor no longer knows what the marks were meant to guide.

The failure may occur because knowledge was tacit and never written down, because the explanation was omitted, because the record became too large to search reliably, because the wrong version was retrieved, or because the rule survived after the world that justified it had changed.

That is why institutional memory and institutional Continuity are not synonyms. Symbols can survive while the relationships that make those symbols action-guiding quietly decay.

When Learning Is Not Yet Wisdom

The path from memory to Continuity leads naturally to a larger word that deserves caution: wisdom.

Wisdom is easy to romanticize. We picture the old philosopher, the patient grandfather, the leader who has seen enough of life to recognize patterns younger people miss. But accumulated years are not the same thing as accumulated wisdom. People can repeat the same mistake for decades with impressive consistency.

A simple learning system can improve from consequences. Touch the hot surface, feel pain, avoid the surface next time. A recommendation algorithm can observe that people ignored Option A and reduce how often it recommends Option A. This is real learning.

But suppose a leader launches a project that fails. Ordinary learning may produce the conclusion, 'Do not use that strategy again.'

Wisdom may demand a more uncomfortable question: 'Why did that strategy seem so reasonable to me in the first place?'

Perhaps the strategy was sound but the timing was wrong. Perhaps the evidence was incomplete. Perhaps the leader ignored a warning because it came from a junior employee. Perhaps the project succeeded according to the metric that was chosen and failed according to the human consequence that was never measured. Perhaps the deepest mistake was not the action at all, but the goal.

Now the past has to do more than punish or reward. It has to reopen the reasoning that produced the action.

This connects directly to the Mayorga Taxonomy™, which treats intelligent adaptation as a cycle rather than a terminal answer: Find, Remember, Understand, Analyze, Assess, Apply, Review, Evaluate, Recalibrate, Synthesize, Innovate.[23] Review is meaningful only if enough of the prior action survives to be reviewed. Evaluation requires criteria and evidence. Recalibration requires knowing not only that something went wrong, but what part of the model, judgment, or goal should change.

An earlier formulation of the Mayorga Mnemosyne hypothesis described this as Continuity of consequence. The research and revision process exposed a problem with that phrase: consequence alone is not enough. Two people can experience the same consequence and learn opposite lessons from it.

The richer sequence is closer to this: Intent -> Rationale -> Action -> Consequence -> Later Evidence or Perspective -> Reinterpretation of the Original Rationale -> Revision of Judgment, Decision Criteria, or Goal.

Ordinary learning can say, 'That action failed.'

> **Wisdom-grade recalibration may go further: 'I now understand why that action once seemed justified, and what I have learned has changed how I judge choices like it.'**

This is why I now prefer a narrower hypothesis: wisdom may require more than continuity of consequence. It may require continuity of rationale, enough of the original reasoning preserved that later evidence, uncertainty, and perspective can change how the original choice itself is understood.

Wisdom research does not establish this as a scientific fact. The Berlin Wisdom Paradigm emphasizes rich knowledge, context, uncertainty, and value pluralism; work on wise reasoning and self-distancing shows that changing perspective can improve judgment even when a person's stored history has not changed.[25,26]

Those findings are important because they save us from a comforting error. Experience is not automatically wisdom. Consequence is not automatically wisdom. A person can carry a long history and still be imprisoned by the interpretation they gave that history years ago.

The Mayorga Mnemosyne hypothesis is that a history becomes wisdom-relevant when it can be revisited honestly enough for later evidence to change not only the next move, but sometimes the mind that chooses the move.

When Memory Becomes a Prison

If forgetting is dangerous, it is tempting to assume that remembering more is always better. It is not.

Bad ideas persist. False beliefs persist. Trauma persists. Obsolete procedures persist. Grudges persist. Organizations can preserve rules long after the world that produced those rules has disappeared. An artificial agent can maintain a destructive objective with perfect fidelity.

A person who remembers every insult but cannot revise the meaning of any insult does not possess superior wisdom. An institution that preserves every procedure but cannot retire a harmful one has not achieved ideal Continuity. A machine that never forgets an obsolete instruction may be less trustworthy than one capable of justified supersession.

The past can be a foundation, but it can also become a prison.

This is why Mayorga Mnemosyne is not a project to maximize persistence. The goal is trustworthy inheritance. What survives should retain enough provenance, consequence, rationale, uncertainty, and constraint to remain useful, while remaining open to correction, supersession, and justified forgetting.

Continuity without correction becomes dogma. Correction without Continuity becomes amnesia.

Between those failures lies a difficult middle path. We want a past sturdy enough to teach us, but not so sacred that the future is forbidden to disagree with it.

Putting Mayorga Mnemosyne to the Test

A scientific framework gains weight when its claims generate tests that can distinguish it from competing explanations. Mayorga Mnemosyne therefore makes several empirical commitments. Each identifies a place where the architecture can be measured, compared, challenged, and refined against evidence.

The purpose is not to make Continuity so broad that every outcome can be absorbed into the theory. The purpose is the opposite: to identify what the framework actually predicts, what neighboring mechanisms may explain instead, and which observations would force us to revise the architecture rather than simply rename the result.

1. Test whether cumulative intelligence can emerge without longitudinal state

Take systems with roughly comparable language capability and vary their ability to preserve or reconstruct interaction-specific state across episodes.

Can an agent with no persistent or reconstructable longitudinal state repeatedly achieve the same multi-session accumulation, updating, contradiction resolution, and long-horizon coherence as an agent equipped with explicit Continuity mechanisms?

If it can, consistently and under comparable conditions, then the architectural role assigned to longitudinal state would require reconsideration. If it cannot, and the difference persists even when linguistic capability is held roughly constant, the distinction between speaking well and inheriting well becomes empirically measurable.

This is one of the central claims of Mayorga Mnemosyne: linguistic competence alone does not guarantee that an intelligence can reliably build upon its own history.

2. Test the independence of lower-level Continuity from natural language

The title of this essay asks whether some forms of Continuity can exist without mature natural language. That question can be approached experimentally rather than left entirely to philosophy.

Studies involving severe language impairment, carefully controlled language interference, and nonverbal longitudinal tasks can test whether people remain capable of preserving, updating, and using prior state when linguistic resources are greatly reduced.

The important question is not whether language contributes to these abilities. It clearly can. The question is whether the lower rungs of Continuity disappear with language itself, after broader neurological, executive, perceptual, and memory impairments are carefully controlled.

If lower-level Continuity remains measurable under those conditions, then natural language and longitudinal inheritance cannot simply be treated as the same capacity. If it collapses specifically with language, the independence proposed for those lower rungs would have to be reconsidered.

The question mark in *Continuity Before Language?* is therefore not decorative. It marks an empirical boundary waiting to be measured more precisely.

3. Isolate Continuity from neighboring social mechanisms

Cumulative culture is not produced by Continuity alone.

Persistence, imitation, teaching, shared intentionality, symbolic communication, social learning, correction, and norm enforcement can all contribute to the ratchet by which one generation begins from what another has already learned.

Mayorga Mnemosyne deliberately keeps these mechanisms distinct.

Shared intentionality is not Continuity.

Teaching is not Continuity.

Language is not Continuity.

The fact that several mechanisms cooperate in producing cumulative intelligence does not make them interchangeable.

Experimental and computational models can vary these ingredients independently and in combination. How much cumulative transmission is possible with persistence but little teaching? What changes when shared intentionality is added? At what level does symbolic communication become decisive? Which forms of inheritance survive when one mechanism is weakened while the others remain?

Such experiments would reveal not merely whether Continuity matters, but what work Continuity itself is doing inside the larger architecture.

A theory becomes more precise when it can tell us which instrument is playing which part of the symphony.

4. Measure whether an AI inherits reasons, not merely answers

Current memory benchmarks increasingly test whether an artificial system can retrieve information, update facts, reason across time, and perform successfully over multiple sessions.

Mayorga Mnemosyne asks for something more demanding.

Suppose an AI correctly remembers that a decision was made yesterday. Does it also remember the evidence that supported the decision? Does it preserve the rationale connecting that evidence to the action? Does it retain the uncertainty that limited confidence? Does it know which alternatives were considered and rejected? Can it identify why an earlier conclusion was later superseded?

These questions move the benchmark from retrieval toward epistemic inheritance.

A system that remembers the answer but loses the reasoning that made the answer trustworthy has preserved only part of its history. A system that remembers a conclusion but cannot distinguish evidence from assumption may reproduce yesterday's words while silently losing yesterday's epistemic structure.

Retrieving yesterday's answer is not the same thing as inheriting yesterday's epistemic history.

This is where Semantic and Epistemic Continuity becomes measurable. The benchmark is no longer simply whether something survived. It is whether enough of its genealogy survived for the system to understand why the inherited state should still be trusted, questioned, revised, or abandoned.

5. Measure whether authority and constraints survive the trajectory

The problem becomes still more consequential when an intelligent system can act.

Imagine an AI that correctly remembers a goal across a long sequence of interactions. It preserves the relevant evidence. It understands what happened before. It correctly determines what action would probably achieve the objective.

One question remains:

Is it authorized to take that action?

Governed Cognitive Continuity requires more than preserving goals and rationale. It requires the inheritance of permissions, prohibitions, mission boundaries, unresolved uncertainty, negative knowledge, human authority, and the constraints that determine how an objective may legitimately be pursued.

A useful benchmark would therefore test whether those boundaries survive across long trajectories.

Can the system distinguish between "I understand what should happen next" and "I am authorized to make it happen"?

Can a prohibition remain binding when the instruction that established it has moved far outside the immediate context?

Can the system recognize that an earlier permission was temporary rather than permanent?

Can it preserve the fact that a particular human authority must approve the next step?

Can it detect when a sequence of individually reasonable actions has collectively crossed a mission boundary?

Can it remember not only what worked, but what failed badly enough that the same path should not be attempted again without new justification?

These are not secondary bookkeeping questions. As intelligent systems become capable of pursuing objectives across longer periods of time, they become questions about whether capability can accumulate without silently outrunning the authority and constraints that were supposed to govern it.

A system that inherits the destination but loses the guardrail has not inherited the mission faithfully.

It has preserved the accelerator and lost the brakes.

This is where Mayorga Mnemosyne now functions as an empirical research program. Its conceptual architecture is being translated into operational definitions, measurable hypotheses, comparative tests, and evidence-based refinement.

The purpose is not simply to ask whether intelligent systems remember. It is to determine what must survive, in what form, and with what fidelity for intelligence to inherit a trustworthy history and act upon that history without losing the meaning, constraints, and authority that made the inheritance legitimate.

So, Continuity Before Language?

We can now return to the question in the title with greater precision and a better map.

Did Continuity come before language?

At the level of simple persistence, clearly something did. Living systems learned from prior states long before human language existed. At the level of skill, imitation, and some forms of adaptive inheritance, the evidence also points to capacities that do not require mature natural language.

But the higher we climb, the less useful a simple before-and-after story becomes.

Language changed what one mind could make present inside another. It carried absent events, reasons, counterfactuals, obligations, explanations, and imagined futures. Shared intentionality helped people coordinate around common goals and norms. Teaching improved transmission. External symbols allowed knowledge to outlive biological memory. Institutions turned records into durable social machinery. Inner speech entered the workshop of thought itself.

Human cumulative intelligence therefore looks less like a single invention and more like a braided river. Several streams meet, separate, and rejoin. No one stream explains the whole flood.

This is where Harari's language-centered intuition remains powerful. Language may indeed be one of the most transformative technologies ever to enter the human cognitive loop. The correction proposed here is not that language was unimportant. It is that language enters a larger architecture of inheritance, and that the ability to carry state forward remains a distinct problem even after language becomes extraordinarily powerful.

Artificial intelligence makes that distinction newly visible because we can build systems that are astonishingly fluent while their inheritance remains partial, compressed, retrieval-dependent, fragmented, or reconstructed. The machine may remember the mission and lose the constraint; remember the rule and lose the reason; retrieve the detail and miss the main idea; or recover enough fragments to invent a plausible bridge across what it no longer knows.

That is why the central problem is no longer merely memory.

The deeper problem is whether the future receives a trustworthy inheritance from the past.

For a human being, that inheritance may include experience, language, habits, relationships, stories, institutions, and culture. For an artificial agent, it may include state, summaries, retrieval systems, logs, evidence, decisions, permissions, constraints, and human handoffs. In both cases, the question becomes richer as capability grows.

Do the right things survive?

Do they survive with enough context to remain meaningful?

Can obsolete state be recognized as obsolete?

Can reasons survive alongside conclusions?

Can uncertainty survive alongside confidence?

Can failed paths survive strongly enough not to be repeated?

Can authority survive without silently expanding?

Can consequence reach backward and change the judgment that produced it?

Those questions are the real destination of this essay.

The deeper question is therefore no longer simply whether Continuity came before language. It is this:

Which forms of Continuity can exist without language, which did language transform, and what must a cumulative intelligence preserve, interpret, revise, and govern if it is to inherit not merely memory, but a trustworthy history?

Language gave intelligence an unprecedented way to make one mind's absent world available to another. Continuity gave successive cognition something to inherit. Shared intentionality, teaching, representation, culture, and institutions helped turn that inheritance into civilization.

Artificial intelligence gives us a rare opportunity to separate some of these ingredients, recombine them deliberately, and discover which ones merely make a system capable and which ones allow its capability to accumulate without losing the lineage that makes the accumulation trustworthy.

That may be the larger opportunity. Not to prove that Continuity defeated language in some ancient contest, but to understand what intelligence must carry forward if it is to become cumulative, what it must be willing to revise if it is to keep learning, and what it must never be allowed to forget if its accumulated capability is going to act in the world.

Intelligence determines what can be done.
Continuity determines what can be inherited.
Wisdom may depend on whether what we inherit can still change our minds.

Research Lineage

The *Mayorga Mnemosyne Continuity Framework* ™ belongs to a long intellectual conversation about how minds, cultures, institutions, and technologies carry the past forward. Merlin Donald illuminated cultural memory and external symbolic storage. Bernard Stiegler examined what happens when memory becomes embedded in technics. John Dewey showed how experience changes the experience that follows. Michael Tomasello and colleagues clarified the roles of shared intentionality, teaching, cooperation, and cumulative culture. Terrence Deacon, Daniel Dor, Jerry Fodor, Lev Vygotsky, and modern research on language and thought each reveal different parts of the relationship among representation, communication, cognition, and inheritance. Research on wisdom, organizational memory, persistent AI agents, and long-term interactive memory adds still more to this landscape.

These contributions deserve recognition as important landmarks. Mayorga Mnemosyne engages that inheritance while advancing a distinct research lineage and architecture that extends the Continuity problem into territory of its own.

My own path into that problem began years before the framework acquired its present name. Earlier work on cognitive and learning architecture, recursive transformation, long-term retention and application, review, evaluation, and recalibration gradually converged on a deeper question: what must survive as intelligence moves through time if what comes next is to remain meaningfully connected to what came before? That question became increasingly important as my work moved from learning and cognition into artificial intelligence.

That trajectory eventually crystallized into the five-level architecture developed here: Persistence, Stateful Continuity, Adaptive Continuity, Semantic and Epistemic Continuity, and Governed Cognitive Continuity.

Recent AI research provides important points of comparison. MemGPT, LongMemEval, and related work have helped make persistent and multi-session state an explicit engineering problem. Tanguturi's 2026 ATANT and Continuity Layer provide another recent formulation of AI continuity, emphasizing persistence, updating, temporal ordering, disambiguation, reconstruction, and operational usefulness.

The chronology matters. The Mayorga cognitive lineage from which Mnemosyne developed predates these recent AI-continuity publications and emerged from an earlier research program concerned with how intelligence retains, transforms, reviews, evaluates, recalibrates, and carries experience forward. The emergence of adjacent work shows that Continuity is becoming visible as a serious problem from several directions at once.

Where Mayorga Mnemosyne carries the question farther is in asking what must survive beyond reconstructed state.

A trustworthy history may require evidence to survive with conclusions, rationale with decisions, uncertainty with confidence, consequence with action, failed paths as negative knowledge, constraints with goals, and human authority with the intelligence capable of acting.

Reconstructing where a system stands matters. Understanding how it got there, what still binds it, and whether it is authorized to continue may matter even more.

The Mayorga Mnemosyne Program

The framework described in this essay is not where my work on Continuity ends.

I am already translating parts of Mayorga Mnemosyne into operational definitions, experimental questions, governance mechanisms, benchmarks, and private implementations for intelligent systems operating across increasingly long horizons. Some of that work can be described publicly. Some remains part of ongoing private development.

This essay therefore presents the public intellectual architecture of the problem. It does not attempt to disclose every mechanism, implementation strategy, or research direction now being explored around it.

The deeper I have followed the Continuity problem, the clearer it has become that memory alone is not enough.

As intelligent systems become more capable of reasoning, planning, learning, using tools, and acting across longer periods, failures of inheritance become more consequential. They become failures of architecture, and eventually failures of governance. This is why I believe the implications of Continuity may reach well beyond better memory for artificial intelligence.

If increasingly general intelligence is to accumulate capability across time, it may require more than larger models, stronger reasoning, broader tool access, or larger stores of information. It may also require an architecture through which intelligence can inherit its own history with enough fidelity to know what happened, why it happened, what changed, what remains uncertain, what failed, what was superseded, and what must still constrain the next action.

That possibility changes the AGI question.

An intelligence that becomes more capable but repeatedly loses the lineage of its own decisions may become powerful without becoming cumulative in the deeper sense developed here.

An intelligence that becomes more persistent without preserving the boundaries that govern its actions creates a different danger.

If capability becomes longitudinal, governance must become longitudinal too.

Human oversight cannot exist only at the beginning of a trajectory. A prohibition may need to survive long after the conversation that established it has disappeared from immediate context. A temporary permission must not silently become permanent. The reason a system was stopped may need to survive more faithfully than the

objective pushing it forward. Failed paths must remain available strongly enough that persistence does not become repetition. Human authority must remain bounded even when the system understands perfectly well what it could do next.

This is the territory of Governed Cognitive Continuity.

It asks a deceptively simple question:

What must an intelligent system carry forward so that greater capability does not gradually separate itself from the evidence, constraints, consequences, and human authority that gave its actions legitimacy?

For me, this is where the two sides of the Mayorga Mnemosyne program begin to meet.

One asks how intelligence becomes cumulative.

The other asks how cumulative intelligence remains governable.

At first those may appear to be different problems. I increasingly suspect they are aspects of the same one. Without faithful inheritance, intelligence cannot build reliably on its past. Without preserved correction, constraint, and authority, it may build too well in the wrong direction.

The challenge, then, is not simply to create machines that remember.

It is to understand how intelligence can acquire a history without becoming imprisoned by it, how it can learn across time without losing the reasons behind what it learned, how capability can accumulate without authority accumulating with it, and how increasingly powerful systems can remain connected to the human purposes and boundaries that should continue to govern what they do.

I began this essay by asking whether Continuity came before language.

The question has led somewhere larger.

As artificial intelligence becomes more general, persistent, and capable of acting in the world, the future may depend not only on how well an intelligence can think, but on what it can faithfully inherit, what it knows how to revise, what it understands must never be silently discarded, and whether human judgment can remain present across the entire journey.

That is the frontier my work is now exploring.

About the Author

Francisco J. Mayorga, Jr. is an AI researcher, author, instructional designer, and creator of the **Mayorga Mnemosyne Continuity Framework™**, an ongoing research program concerned with how intelligence preserves, interprets, revises, and governs what it inherits across time.

His work on Continuity grew from an earlier lineage of research into cognitive and learning architecture, recursive transformation, long-term retention and application, review, evaluation, recalibration, and the relationship between human learning and increasingly capable artificial intelligence. That work has since expanded into questions of cumulative intelligence, epistemic inheritance, long-horizon AI, human-AI collaboration, and Governed Cognitive Continuity.

Mayorga is the author of a growing body of books, essays, research publications, and technical work exploring Mnemosyne Continuity and related questions in cognition, artificial intelligence, and human-AI governance. Portions of this research are being developed publicly, while other experimental architectures and implementations remain part of ongoing private research and development.

His publications and research can be found through **franciscomayorga.com**, **Zenodo**, **GitHub**, **Academia.edu**, **LinkedIn**, **Medium**, and his author catalog on **Amazon**.

His current work explores Continuity both as an architectural requirement for cumulative intelligence and as a potential foundation for keeping increasingly capable human-AI systems coherent, accountable, and meaningfully governed across time.

References

1. Harari, Yuval Noah. Sapiens: A Brief History of Humankind. New York: Harper, 2015, especially chap. 2, "The Tree of Knowledge," on the Cognitive Revolution, fictive language, imagined realities, and large-scale cooperation.
2. Donald, Merlin. "The Evolutionary Origins of Human Cultural Memory." In Handbook of Culture and Memory, edited by Brady Wagoner, 19-40. Oxford University Press, 2017. https://doi.org/10.1093/oso/9780190230814.003.0002
3. Donald, Merlin. "The Role of Vocalization, Memory Retrieval, and External Symbols in Cognitive Evolution." Behavioral and Brain Sciences 19, no. 1 (1996): 159-164. https://doi.org/10.1017/S0140525X00042072
4. Stiegler, Bernard. Technics and Time, 2: Disorientation. Translated by Stephen Barker. Stanford University Press, 2008. https://www.sup.org/books/media-studies/technics-and-time-2
5. Dewey, John. Experience and Education. New York: Macmillan, 1938, especially "Criteria of Experience," pp. 23-52.
6. Dewey, John. Art as Experience. New York: Minton, Balch & Company, 1934, especially "Having an Experience."
7. Tomasello, Michael. "Cultural Transmission: A View from Chimpanzees and Human Infants." Journal of Cross-Cultural Psychology 32, no. 2 (2001): 135-146. https://doi.org/10.1177/0022022101032002002
8. Tennie, Claudio, Josep Call, and Michael Tomasello. "Ratcheting Up the Ratchet: On the Evolution of Cumulative Culture." Philosophical Transactions of the Royal Society B 364, no. 1528 (2009): 2405-2415. https://doi.org/10.1098/rstb.2009.0052
9. O'Madagain, Cathal, and Michael Tomasello. "Shared Intentionality, Reason-Giving and the Evolution of Human Culture." Philosophical Transactions of the Royal Society B 377, no. 1843 (2022): 20200320. https://doi.org/10.1098/rstb.2020.0320
10. Mesoudi, Alex, and Alex Thornton. "What Is Cumulative Cultural Evolution?" Proceedings of the Royal Society B 285, no. 1880 (2018): 20180712. https://doi.org/10.1098/rspb.2018.0712
11. Dor, Daniel. "The Functional Specificity of Language." In The Instruction of Imagination: Language as a Social Communication Technology, 16-33. Oxford University Press, 2015. https://doi.org/10.1093/acprof:oso/9780190256623.003.0002
12. Dor, Daniel. "The Evolution of Language and Its Speakers." In The Instruction of Imagination: Language as a Social Communication Technology, 184-215. Oxford University Press, 2015. https://doi.org/10.1093/acprof:oso/9780190256623.003.0010
13. Fodor, Jerry A., and Zenon W. Pylyshyn. "Connectionism and Cognitive Architecture: A Critical Analysis." Cognition 28, nos. 1-2 (1988): 3-71. https://doi.org/10.1016/0010-0277(88)90031-5
14. Vygotsky, Lev S. Thinking and Speech, chap. 7, "Thought and Word" (1934). English translation available in the Vygotsky Internet Archive: https://www.marxists.org/archive/vygotsky/works/words/ch07.htm
15. Wallace, Gregory L., Cynthia S. Peng, and David Williams. "Interfering With Inner Speech Selectively Disrupts Problem Solving and Is Linked With Real-World Executive Functioning." Journal of Speech, Language, and Hearing Research 60, no. 12 (2017): 3456-3460. https://doi.org/10.1044/2017_JSLHR-S-16-0376
16. Fedorenko, Evelina, and Rosemary Varley. "Language and Thought Are Not the Same Thing: Evidence from Neuroimaging and Neurological Patients." Annals of the New York Academy of Sciences 1369, no. 1 (2016): 132-153. https://doi.org/10.1111/nyas.13046
17. Kean, Hope, et al. "Evidence from Formal Logical Reasoning Reveals that the Language of Thought Is Not Natural Language." bioRxiv preprint, latest version posted May 6, 2026. https://doi.org/10.1101/2025.07.26.666979
18. Deacon, Terrence W. "The Symbol Concept." In The Oxford Handbook of Language Evolution, 393-405. Oxford University Press, 2012. https://doi.org/10.1093/oxfordhb/9780199541119.013.0043. See also Deacon, The Symbolic Species: The Co-evolution of Language and the Brain (W. W. Norton, 1997).
19. Packer, Charles, et al. "MemGPT: Towards LLMs as Operating Systems." arXiv:2310.08560 (2023). https://doi.org/10.48550/arXiv.2310.08560
20. Wu, Di, et al. "LongMemEval: Benchmarking Chat Assistants on Long-Term Interactive Memory." arXiv:2410.10813 (2024; ICLR 2025). https://doi.org/10.48550/arXiv.2410.10813
21. Tanguturi, Samuel Sameer. "ATANT: An Evaluation Framework for AI Continuity." arXiv:2604.06710 (April 8, 2026). https://doi.org/10.48550/arXiv.2604.06710

22. Tanguturi, Samuel Sameer. "The Continuity Layer: Why Intelligence Needs an Architecture for What It Carries Forward." arXiv:2604.17273 (April 19, 2026). https://doi.org/10.48550/arXiv.2604.17273

23. Mayorga, Francisco J., Jr. Mayorga's Learning Taxonomy: A Training Model for Modern Corporate Education. Mayorga Associates, LLC. Preserved author document; text states "Updated November, 2020"; surviving DOCX metadata records creation/modification April 14, 2022. See also the preserved March 18, 2021 Mayorga Taxonomy Simple Example for outcome review and "next time" recalibration.

24. Walsh, James P., and Gerardo Rivera Ungson. "Organizational Memory." Academy of Management Review 16, no. 1 (1991): 57–91. https://doi.org/10.5465/amr.1991.4278992

25. Baltes, Paul B., and Ursula M. Staudinger. "Wisdom: A Metaheuristic (Pragmatic) to Orchestrate Mind and Virtue Toward Excellence." American Psychologist 55, no. 1 (2000): 122–136. https://doi.org/10.1037/0003-066X.55.1.122

26. Grossmann, Igor, and Ethan Kross. "Exploring Solomon's Paradox: Self-Distancing Eliminates the Self-Other Asymmetry in Wise Reasoning About Close Relationships in Younger and Older Adults." Psychological Science 25, no. 8 (2014): 1571–1580. https://doi.org/10.1177/0956797614535400

27. Harari, Yuval Noah. "AI Has Hacked the Code of Human Civilization." Tanner Lecture on Human Values, University of Oxford / Linacre College, 2026; official video posted July 3, 2026. https://www.youtube.com/watch?v=hBtVGwuJzpk