What Continual Learning Actually Is

JP Vasseur · Sr Distinguished Engineer Applied AI Research Continual Learning – Sept 2026

Most conversations about Continual Learning start in the wrong place. People hear the term and think of faster model updates, shorter release cycles, or new weights shipped more often. Those can be part of it, but they are far from the whole spectrum of Continual Learning (CL) definition, which is much broader, as the rest of this document shows.

Continual Learning is the capability of a model, an agent, or a full agentic system to turn accumulated operating experience into persistent, verified improvement in future behavior, self-improvement earned through experience rather than assumed.

Today's AI reasons extremely well. AI systems retrieve information, calls tools, coordinates with other agents, and holds context within a session, but what existing AI systems such as agents do not do reliably is get better because of past experiences and thus truly self-improve. Closing that gap is, I believe, one of the defining problems of the next phase of AI.

Operating experience exposes what training never anticipated: changing conditions, tool failures, constraints specific to one customer, workflows that are technically correct but too costly to keep running, and corrections that reveal a better path nobody wrote down. When that experience turns into durable learning, it compounds through better memory, better tool selection, better routing, and reusable skills, and the system becomes more capable and more efficient simply by operating.

Here is the part people miss. Most deployed AI can adapt ***within*** a session but it does not reliably learn across sessions, because a model's weights stay static whether or not a user just corrected it, an agent just found a better tool sequence, or a workflow just succeeded after three failed attempts. Memory and retrieval-augmented generation move information from one session to the next, and memorization plays a considerable role in any learning system. And it is worth mentioning that Memory is also part of Continual Learning, because the memorization strategy itself must adapt as events unfold: what to keep, what to merge, what to forget. A memory architecture can therefore benefit from Continual Learning in its own right. Memory and RAG on their own, however, are not learning. They replay facts without ever judging whether a lesson is correct, safe, or worth generalizing, and they do not change how the system behaves. Continual Learning is about adaptation, behavior that changes over time as experience accumulates. It starts at the point where experience is selected, validated, and converted into something persistent, a better memory policy, a tool strategy learned from past failures, a new skill, a routing decision, or a model update. The ultimate goal goes one step further: learning to learn, so the system gets better at improving itself.

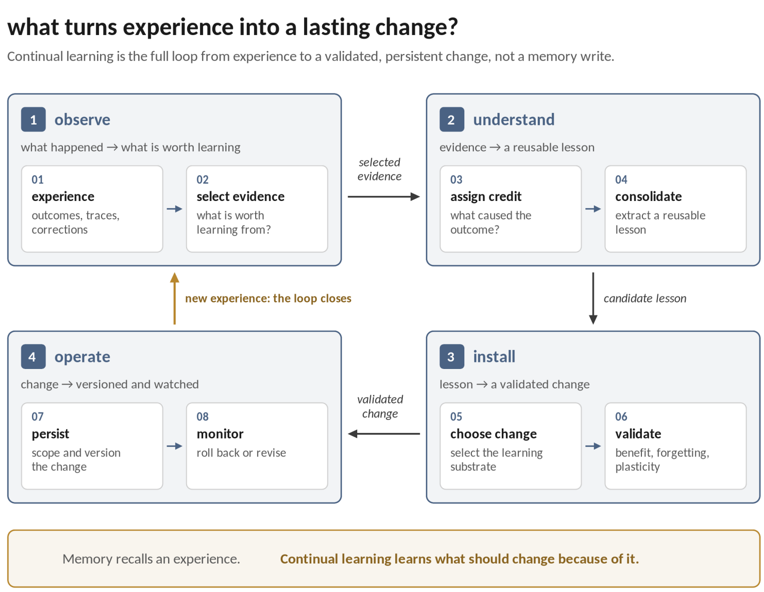

This is also why Continual Learning cannot stop at the model. At the model layer, the field has long framed the problem as learning new things (without catastrophic forgetting) of what the model already does well, the stability-plasticity trade-off leading to a number of novel approaches being developed. That work matters, but the model is only one layer. The agent that wraps it, with its memory policy, tools, prompts and skills, has to learn too, and so does the larger system of agents coordinating around it, which learns how to route, delegate, and share what one agent discovered. Continual Learning is a property of the whole system, and each layer changes in its own way, as the figure below shows.

where can a learned change live?

Continual learning is defined by the managed change, not by one place where it is stored.

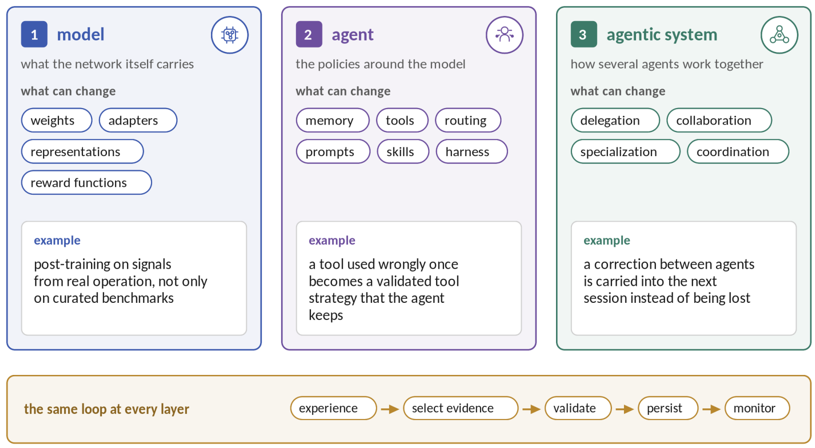

Continual Learning has a long biological history

Physarum polycephalum, a single cell with no neurons, habituates to repeated stimuli and stores spatial memory in its network morphology and in the chemical trails it leaves behind, a state that survives dormancy. C. elegans runs habituation and associative learning on a fully mapped connectome of 302 neurons, with memories that last about a day. Songbirds close an explicit training loop: a juvenile memorizes a tutor song, compares its own practice against that template through auditory feedback, and keeps correcting drift into adulthood. Humans combine fast hippocampal encoding of single episodes with slow neocortical consolidation during sleep, so one experience is captured in seconds and turned into general knowledge over days to years. From one cell to the human brain the substrate changes, but the principle holds, experience is retained and reused to improve future behavior.

None of this is free, because evidence does not travel cleanly between systems built for different contexts, proving that a change actually helped is a causal-attribution problem rather than a correlation one, and once a lesson can move between deployments nobody owns its lifecycle by default. Real governance has to include security controls such as provenance, signing, and resistance to poisoned evidences, safety controls such as watching for distribution shift, negative transfer, and evaluator contamination, and one rule above all of it, that promotion is eligibility rather than deployment, so every consuming model, agent, or system keeps the authority to reject, constrain, or roll back what it is given.

what can learn, how fast, and where?

Above the model, no learning surface has a durable learning loop today.

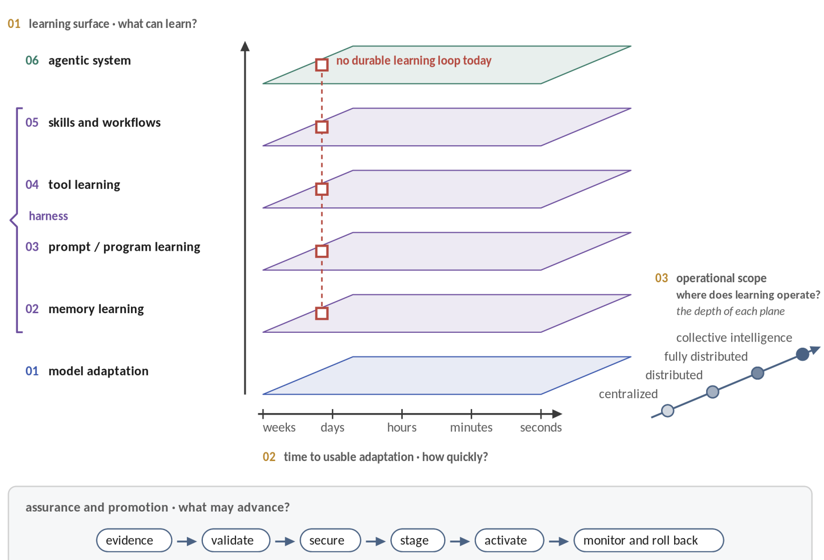

Challenges

The promise is extraordinary, and the work is far from done. The figure above maps the space. Learning can happen at multiple surfaces, from model adaptation to the agentic system, at speeds from weeks down to seconds, and at scopes from a single deployment to collective intelligence. Above the model, none of these surfaces has a durable learning loop today and building one raises hard problems. The first is generalization: deciding when one experience, or a handful of them, says something general enough to change future behavior, and when it is noise, an exception, or a local artifact. The second is timing. Adapt too slowly and the system wastes what it sees. Adapt too fast and it overreacts to a few events and chases noise. Forgetting cuts both ways, since a model must not lose what it already does well, while a memory must drop what has become stale or wrong. And proving that a change actually helped is a causal question, which correlation alone cannot answer.

At scale the difficulty moves to the system. Every deployment that improves itself also drifts away from the others, so fleets of agents become more and more heterogeneous, and sharing a lesson between self-improving agents that no longer share the same tools memory strategy or context is still an open problem. Safety and security follow directly: provenance and signing of evidence, resistance to poisoned experience, detection of distribution shift and negative transfer and the right of every consumer to reject or roll back what it receives. None of these problems is trivial. In my view, all of them can be overcome, and each already has a credible engineering path.

Conclusion

The reward is worth the effort. Today a model improves in discrete steps. Experience is collected, filtered, aggregated into a batch, and folded into the next training run weeks or months later, and almost everything else is lost: the interactions, tool calls, failures, recoveries, and corrections that never make it into a dataset. Continual Learning changes the unit of learning. A system that learns from its own experience learns from every interaction it has while it operates instead of waiting for the next release. The amount of experience at stake is gigantic far larger than the curated data that reaches any training run, and every deployed model, agent, and agentic system produces more of it every second. Once that experience can be turned safely into validated improvement, and shared between systems under the right governance, AI stops being a product that is shipped and becomes a capability that compounds. I believe this is one of the largest untapped sources of progress in AI today.

This is the first post in a series on Continual Learning, with deeper analysis, research outcomes and results.