

Winter Is Over for Continual Learning: From Continual Representation Learning to Continual Learning Agents

Luigi Quarantiello^{2,*}, Lanpei Li^{2,3}, Ehsan Tavan², Irene Testa², Giacomo Carfi², Gerlando Gramaglia², Jack Bell², Daniele Malitesta¹, Pierre Averty¹, Eric Nuerthey Coleman² and Vincenzo Lomonaco¹

¹LUISS Guido Carli University, Viale Romania 32, 00197 Rome, Italy

²Department of Computer Science, University of Pisa, 56127 Pisa, Italy

³Institute of Information Science and Technologies "A. Faedo" (ISTI), National Research Council of Italy (CNR), 56124 Pisa, Italy

Abstract

Continual Learning (CL) gained prominence with the rise of deep architectures, and then entered a winter as Large Language Models (LLMs) gave the impression that learning over time was no longer needed. We argue that this winter was not a failure but, in fact, a necessary transition. Intelligence cannot be separated from the evolving time in which an agent lives and acts, and the arrival of agentic systems makes CL return to its proper place: not the narrow question of how representations are formed, but the principle by which agents organise their knowledge as its environment changes. We call these systems *Continual Learning Agents*, designed from the start to adapt continually and consolidate knowledge over time. In this work, we follow the path from the early ideas about machines that learn over time, through the era of deep CL, towards an holistic view in which continual learning and agent design can no longer be held apart, and in which the hard problems of Artificial Intelligence (AI) and those of CL are seen to converge.

Keywords

continual learning, large language models, AI agents, adaptation, knowledge consolidation

1. Introduction

The history of machine learning is, at its core, the history of systems that learn over time. From Alan Turing's foundational vision of an educable machine [1] to modern LLMs deployed as autonomous agents, the dream of building systems that accumulate knowledge and adapt to change has never waned. Nonetheless, the path towards it has been far from linear.

Continual learning experienced a renaissance in the 2010s with the rise of deep learning, driven by seminal work on catastrophic forgetting [2] and a rich ecosystem of benchmarks, algorithms, and surveys [3, 4, 5, 6, 7]. Yet, LLMs emerged, bringing remarkable zero-shot generalization performance [8] together with web-search and retrieval capabilities [9]. Therefore, many questioned CL's relevance: if a model pre-trained on internet-scale data can handle most tasks *out of the box*, while also retrieving new information at inference time, what is left to learn continually?

We argue the answer is a great deal, but for reasons that transcend representation learning. The rise of agentic AI systems [10], which must operate continuously in dynamic environments, interact with users over extended timescales, and integrate new knowledge without compromising existing capabilities, has created precisely the conditions where CL is not merely useful but *a structural necessity*. Similar requirements arise across emerging agentic applications, including generative agent-based simulation [11]. We call this broader perspective **Continual Learning Agents**.

In this paper, we trace the intellectual lineage leading to the current scenario (Section 2), reflect on what deep CL achieved and where it fell short (Section 3), articulate the architectural vision of CL Agents (Section 4), and outline the challenges and opportunities ahead (Section 5). Our central thesis

AIxIA 2026 Discussion Papers, 24th International Conference of the Italian Association for Artificial Intelligence, Perugia, Italy, 2026

*Corresponding author.

✉ luigi.quarantiello@phd.unipi.it (L. Quarantiello); lanpei.li@isti.cnr.it (L. Li); e.tavan@studenti.unipi.it (E. Tavan); irene.testa@phd.unipi.it (I. Testa); giacomo.carfi@phd.unipi.it (G. Carfi); gerlando.gramaglia@phd.unipi.it (G. Gramaglia); jack.bell@phd.unipi.it (J. Bell); dmalitesta@luiss.it (D. Malitesta); pierre.avery@taio.fr (P. Averty); eric.coleman@phd.unipi.it (E. N. Coleman); vincenzo.lomonaco@luiss.it (V. Lomonaco)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

is: *The “winter” of CL was a transition from a method-centric, benchmark-driven discipline toward a holistic, architecture-first perspective—and this transition is now complete.* The field that has long studied how artificial systems can learn without forgetting is finally poised to build systems that genuinely, sustainably, and continuously improve.

2. Continual Machine Learning: A Historical Arc

The idea that intelligent systems should learn continuously dates back to the origins of the AI field itself. Turing’s 1950 paper [1] already sketched a vision of machine education: a child-like process guided by rewards and punishments, building knowledge incrementally through a curriculum of experience—continual learning *ante litteram*. Newell and Simon’s production systems [12] later offered a modular, rule-based counterpoint, where symbolic rules could be added incrementally without invalidating prior knowledge, a property that connectionist systems would struggle to replicate. Indeed, when neural networks regained momentum in the 1980s, McCloskey and Cohen [13] identified their tendency to abruptly overwrite previously learned knowledge when exposed to new data, a phenomenon they termed *catastrophic interference* that would shape the field for the following three decades.

The 1990s saw the problem formalized across several communities at once. Thrun [14] framed lifelong learning operationally for robotics, while Ring [15] introduced skill composition in reinforcement learning. In parallel, the statistical learning community addressed concept drift [16] and non-stationary environments, while the data stream mining literature developed online and incremental algorithms [17] for continuous, unbounded data. Each tradition contributed vocabulary—transfer, adaptation, plasticity, stability—that would later converge in modern CL. The most architecturally ambitious vision of this period, however, belongs to Tom Mitchell, whose canonical definition of learning as experience-driven improvement in performance [18] was later actualized in the Never-Ending Language Learner [19], a system that read the web for years, continuously refining a structured knowledge base.

The deep learning era turned this philosophy into a concrete empirical problem, and, in response, a taxonomy of solutions. Ruvolo and Eaton’s ELLA [20] offered an early demonstration of efficient lifelong learning through shared latent structure across tasks, but it was Kirkpatrick et al.’s Elastic Weight Consolidation (EWC) [4] that gave the field its first principled, Bayesian answer to catastrophic forgetting, selectively slowing learning on weights important for previous tasks. EWC anchored an entire family of *regularization-based* methods, joined by *rehearsal-based* approaches, which store exemplars of past experience and revisits them while learning from new data, and *architectural* approaches, which allocate dedicated parameters or subnetworks to new tasks. Comprehensive surveys [21, 7] and benchmarks such as CRe50 [3] unified this otherwise fragmented landscape across computer vision and robotics, bringing the field onto common empirical ground for the first time [22]. Crucially, this reframes Mitchell’s own definition: once we accept that learning is fundamentally an open-ended, experience-driven process, CL is no longer a sub-problem of machine learning. *It is machine learning, in its most essential form.*

3. The End of Continual Representation Learning?

The emergence of large foundation models [23]—pre-trained on vast corpora and adaptable via few-shot prompting or lightweight fine-tuning—appeared to many to undercut the motivation for CL. This perception was reinforced by the emergence of Retrieval-Augmented Generation (RAG) systems, which integrate web search and external retrieval directly into LLM-based pipelines [9]. If a large pre-trained model already encodes broad world knowledge and can access up-to-date information on demand, why continually update the model’s parameters to incorporate new knowledge?

This perception triggered what we term the “CL winter”: a relative loss of momentum and visibility for CL within the deep-learning mainstream, with much attention redirected toward adapting foundation models rather than studying learning dynamics per se. However, recent work has begun to re-frame deep CL within this new landscape. Zero-shot generalization and retrieval can reduce the need for immediate model updates, but they do not eliminate the need for persistent adaptation as knowledge, tasks, users,

and environments evolve. Bell et al. [24] survey three key directions for CL in the era of foundation models—namely continual pre-training, continual fine-tuning, and continual compositionality and orchestration—arguing that the core algorithmic challenges (stability-plasticity trade-off, forward and backward transfer, evaluation protocols) remain as pressing as ever, even if the architectures have changed.

We endorse this framing while contending it is necessarily incomplete. We extend it by shifting the unit of analysis from individual mechanisms of continual adaptation to the architecture of the learning agent as a whole. Deep CL, focusing on representation learning, benchmark-driven evaluation, and gradient-based solutions to catastrophic forgetting, addresses one essential dimension of learning-agent design. Yet the fundamental objectives of any intelligent agent are broader: *adaptation* (efficiently responding to novel experience and individual user needs) and *consolidation* (accumulating durable knowledge and structuring it for efficient access). Representation learning is one mechanism for pursuing these goals at the parameter level, but it is not the only one, and it is arguably not the most flexible. Continual Learning Agents, as we describe in the next section, pursue both objectives across the full architectural stack.

4. Toward Continual Learning Agents

The rise of agentic AI creates conditions for a richer realization of CL than representation learning alone can provide. We use the term *Continual Learning Agents* to describe AI agents architecturally designed to pursue adaptation and consolidation as first-class system objectives, not merely as training-time algorithmic constraints. **Adaptation** denotes the capacity to efficiently update behavior in response to new situations, user preferences, task distributions, and environmental change, without requiring costly full retraining. **Consolidation** denotes the capacity to accumulate knowledge and capability layer upon layer: adding what is new, discarding what is obsolete, organizing retained knowledge for efficient access, and continuously refining the agent’s model of a changing world. Viewed in this way, CL becomes an emergent property of the overall system architecture rather than the responsibility of any single algorithmic component.

Several architectural trends converge to make CL Agents practically realizable. One direction is represented by Mixture-of-Experts architectures [25], which route computation across specialized sub-networks and support incremental capability expansion, although adding new experts without disrupting existing ones or degrading the router remains challenging. Building on this principle, the continual model routing framework by Bell et al. [26], a concrete instantiation of the continual compositionality direction, proposes that agents should not be tied to a single model but should navigate an evolving ecosystem of specialized models, routing inputs to the most competent expert and integrating their outputs. This architectural perspective dissolves the boundary between “the model” and “the environment of models”, enabling a form of distributed and compute-aware, collective CL at societal scale. It also raises a central question for the next generation of CL systems: under which conditions can a continually evolving set of specialized experts match or surpass the performance of a monolithic model, while offering better adaptability and computational efficiency?

A complementary direction is represented by memory architectures [27], which support knowledge consolidation through persistent storage, organization, and retrieval. RAG [28] allows agents to query external knowledge stores at inference time, extending working memory beyond model weights. Structured knowledge bases, such as wikis, knowledge graphs, relational stores, serve as persistent, editable representations of accumulated world knowledge. Agent scaffolding and memory-management frameworks, such as LangChain [29] and MemGPT [30], provide mechanisms for orchestrating tools, retrieval, context management, and persistent interaction state across long horizons. Together, these mechanisms approximate Mitchell’s vision: a system that continuously reads, stores, and refines what it knows.

Beyond parameters, memories, and external knowledge bases, Continual Learning Agents also need to be able to continuously learn at the workflow level. Modern agentic architectures rarely rely on a single model invocation. Instead, such systems operate through dynamic combinations of models, tools, APIs,

retrieval mechanisms, memory components, and human-in-the-loop interactions [31, 32, 33]. From that point of view, CL should not only concern what an agent knows but also how it operates through time. It implies choosing the right tools or model for each case, performing task decomposition, retrieving external information and asking for human assistance [34, 35, 26]. Each interaction generates an execution record that includes information about the initial goal, intermediate decisions, selected tools, failures, retries, costs, latency, human interventions, and the outcome of the task. Such execution records can later be consolidated into reusable procedural knowledge. In this sense, open model hubs turn continual compositionality and continual model routing into more than isolated architectural choices: they become mechanisms through which Continual Learning Agents can reuse, select, and combine evolving capabilities over time, approaching monolithic-model performance through collections of compute-aware experts [36, 26].

5. What's Next?

The convergence of agentic AI and CL opens an era of genuine opportunity. We now have the computational infrastructure, architectural building blocks, and societal deployment context to build agents that learn and adapt sustainably, effectively, and efficiently. After years focused on narrow objectives (e.g., catastrophic forgetting mitigation, benchmark accuracy, parameter efficiency), we can now pursue what the field was always intended to address: artificial systems that continuously improve in the real world.

However, this opportunity carries serious responsibilities. Agents that improve autonomously raise novel safety challenges absent from static models. A Continual Learning Agent can drift in unpredictable directions, acquire biases through adversarial experience injection, or consolidate harmful beliefs through uncured web interactions. Ensuring that adaptation remains aligned with human values, and that consolidation does not entrench errors or amplify misalignment, requires safety frameworks that treat the learning process itself, not just the initial model, as a safety-critical component.

At the same time, the CL community must undergo a methodological reorientation. For too long, the field has been method-centric: defined by overspecialized algorithms and narrow benchmarks that capture only a fragment of the real challenge. With the emergence of CL Agents, the right questions are not “which regularization technique best prevents forgetting on a fixed task sequence?” but rather: what architectural principles enable sustained adaptation and consolidation at the system level? How do we evaluate genuine, open-ended lifelong improvement? How do we decompose the problem of building CL agents into modules that can be engineered, tested, and made safe?

This transition toward Continual Learning Agents also demands new evaluation mechanisms. Rather than measuring performance on fixed task sequences alone, future benchmarks should quantify an agent’s ability to accumulate knowledge, personalize interactions, recover from failures, and adapt over long deployment horizons while remaining safe and efficient. Developing standardized evaluation protocols for open-ended continual improvement may become as important as developing new CL algorithms themselves.

The LLM ecosystem—with its momentum, infrastructure, and industrial investment—provides the wind at our backs for this reorientation. We can, at last, abandon the vision of CL as a corrective patch on top of gradient descent, and embrace it as the architectural soul of intelligent agency. The winter was not a defeat. It was the necessary clearing that makes a new spring possible.

Acknowledgments

This research was partially supported by the FIS2 Grant from Italian Ministry of University and Research (Grant ID: FIS2023-03382).

Declaration on Generative AI

During the preparation of this work, the authors used Gemini and Claude in order to: Paraphrase and reword, Grammar and spelling check. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] A. M. Turing, Computing machinery and intelligence, *Mind* 59 (1950) 433–460.
- [2] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, Y. Bengio, An empirical investigation of catastrophic forgetting in gradient-based neural networks, *arXiv preprint arXiv:1312.6211* (2013).
- [3] V. Lomonaco, D. Maltoni, CORE50: a new dataset and benchmark for continuous object recognition, in: *Proceedings of the 1st Annual Conference on Robot Learning (CoRL)*, PMLR, 2017, pp. 17–26.
- [4] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al., Overcoming catastrophic forgetting in neural networks, *Proceedings of the National Academy of Sciences* 114 (2017) 3521–3526.
- [5] F. Zenke, B. Poole, S. Ganguli, Continual learning through synaptic intelligence, in: *International conference on machine learning*, Pmlr, 2017, pp. 3987–3995.
- [6] Z. Li, D. Hoiem, Learning without forgetting, *IEEE transactions on pattern analysis and machine intelligence* 40 (2017) 2935–2947.
- [7] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, T. Tuytelaars, A continual learning survey: Defying forgetting in classification tasks, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44 (2021) 3366–3385.
- [8] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, Y. Iwasawa, Large language models are zero-shot reasoners, *Advances in neural information processing systems* 35 (2022) 22199–22213.
- [9] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, M. Wang, H. Wang, Retrieval-augmented generation for large language models: A survey, *arXiv preprint arXiv:2312.10997* (2023).
- [10] D. B. Acharya, K. Kuppan, B. Divya, Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey, *IEEE Access* 13 (2025) 18912–18936.
- [11] C. Adornetto, A. Mora, K. Hu, L. I. Garcia, P. Atchade-Adelomou, G. Greco, L. A. A. Pastor, K. Larson, Generative agents in agent-based modeling: Overview, validation, and emerging challenges, *IEEE Transactions on Artificial Intelligence* 6 (2025) 3165–3183.
- [12] A. Newell, H. A. Simon, *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ, 1972.
- [13] M. McCloskey, N. J. Cohen, Catastrophic interference in connectionist networks: The sequential learning problem, in: *Psychology of Learning and Motivation*, volume 24, Academic Press, 1989, pp. 109–165.
- [14] S. Thrun, A lifelong learning perspective for mobile robot control, in: *Intelligent robots and systems*, Elsevier, 1995, pp. 201–214.
- [15] M. B. Ring, *Continual learning in reinforcement environments*, The University of Texas at Austin, 1994.
- [16] G. Widmer, M. Kubat, Learning in the presence of concept drift and hidden contexts, *Machine Learning* 23 (1996) 69–101.
- [17] V. Losing, B. Hammer, H. Wersing, Incremental on-line learning: A review and comparison of state of the art algorithms, *Neurocomputing* 275 (2018) 1261–1274.
- [18] T. M. Mitchell, *Machine Learning*, McGraw-Hill, New York, NY, 1997.
- [19] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, B. Yang, J. Betteridge, A. Carlson, B. M. Dalvi, M. Gardner, B. Kisiel, et al., Never-ending learning, *Communications of the ACM* 61 (2018) 103–115.
- [20] P. Ruvolo, E. Eaton, ELLA: An efficient lifelong learning algorithm, in: *Proceedings of the 30th International Conference on Machine Learning*, PMLR, 2013, pp. 507–515.
- [21] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, S. Wermter, Continual lifelong learning with neural networks: A review, *Neural Networks* 113 (2019) 54–71.
- [22] T. Tuytelaars, B. Liu, V. Lomonaco, G. van de Ven, A. Cossu, Deep Continual Learning (Dagstuhl Seminar 23122), *Dagstuhl Reports* 13 (2023) 74–91. URL: <https://drops.dagstuhl.de/entities/document/10.4230/DagRep.13.3.74>. doi:10.4230/DagRep.13.3.74.
- [23] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al., On the opportunities and risks of foundation models, *arXiv preprint arXiv:2108.07258* (2021).

- [24] J. Bell, L. Quarantiello, E. N. Coleman, L. Li, M. Li, M. Madeddu, E. Piccoli, V. Lomonaco, The future of continual learning in the era of foundation models: Three key directions, arXiv preprint arXiv:2506.03320 (2025).
- [25] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, J. Dean, Outrageously large neural networks: The sparsely-gated mixture-of-experts layer, in: Proceedings of the 5th International Conference on Learning Representations (ICLR), 2017.
- [26] J. Bell, G. Carfi, G. Gramaglia, V. Lomonaco, Continual model routing in evolving model hubs, arXiv preprint arXiv:2605.28577 (2026).
- [27] V. Dorovatas, M. Schwerin, A. D. Bagdanov, L. Caccia, A. Carta, L. Charlin, B. Hammer, T. L. Hayes, T. Hess, C. Kanan, et al., Modular memory is the key to continual learning agents, arXiv preprint arXiv:2603.01761 (2026).
- [28] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive NLP tasks, in: Advances in Neural Information Processing Systems (NeurIPS), volume 33, 2020, pp. 9459–9474.
- [29] H. Chase, Langchain, 2022. URL: <https://github.com/langchain-ai/langchain>, software framework for building agents and LLM-powered applications.
- [30] C. Packer, V. Fang, S. Patil, K. Lin, S. Wooders, J. Gonzalez, Memgpt: towards llms as operating systems. (2023).
- [31] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, Y. Cao, React: Synergizing reasoning and acting in language models, 2023. URL: <https://arxiv.org/abs/2210.03629>. arXiv:2210.03629.
- [32] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu, A. H. Awadallah, R. W. White, D. Burger, C. Wang, Autogen: Enabling next-gen LLM applications via multi-agent conversations, in: First Conference on Language Modeling, 2024. URL: <https://openreview.net/forum?id=BAakY1hNKS>.
- [33] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, Y. Lu, Y. Lin, X. Cong, X. Tang, B. Qian, S. Zhao, L. Hong, R. Tian, R. Xie, J. Zhou, M. Gerstein, dahai li, Z. Liu, M. Sun, ToolLLM: Facilitating large language models to master 16000+ real-world APIs, in: The Twelfth International Conference on Learning Representations, 2024. URL: <https://openreview.net/forum?id=dHng2O0Jjr>.
- [34] X. Huang, W. Liu, X. Chen, X. Wang, H. Wang, D. Lian, Y. Wang, R. Tang, E. Chen, Understanding the planning of llm agents: A survey, 2024. URL: <https://arxiv.org/abs/2402.02716>. arXiv:2402.02716.
- [35] S. Fan, X. Cong, Y. Fu, Z. Zhang, S. Zhang, Y. Liu, Y. Wu, Y. Lin, Z. Liu, M. Sun, Workflowllm: Enhancing workflow orchestration capability of large language models, in: International Conference on Learning Representations, volume 2025, 2025, pp. 24498–24525.
- [36] L. Gong, J. Zhang, M. Wei, H. Zhang, Z. Huang, What Is the Intended Usage Context of This Model? An Exploratory Study of Pre-Trained Models on Various Model Repositories, ACM Transactions on Software Engineering and Methodology 32 (2023) 1–57. URL: <https://dl.acm.org/doi/10.1145/3569934>. doi:10.1145/3569934.