Agential AI for Integrated Continual Learning, Deliberative Behavior, and Comprehensible Models

Extended Abstract

Zeki Doruk Erden
École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
zeki.erden@epfl.ch

Boi Faltings
École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

Contemporary machine learning faces key limitations, such as a lack of integration with planning, incomprehensible structure, and inability to learn continually. We present initial design for system, Agential AI (AAI), that overcomes these issues. AAI's core is a learning method that models temporal dynamics with guarantees of completeness, minimality, and continual learning. It integrates this with a behavior algorithm that plans on a learned model and encapsulates high-level behavior patterns. Preliminary experiments on a simple environment show AAI's effectiveness and potential.[1]

KEYWORDS

Continual learning; Planning; Behavior encapsulation

ACM Reference Format:
Zeki Doruk Erden and Boi Faltings. 2025. Agential AI for Integrated Continual Learning, Deliberative Behavior, and Comprehensible Models: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Machine learning currently relies on continuous approximations of environmental relations through fixed architectures like neural networks (NNs). While this has solved many challenges [7, 9, 15], core limitations remain [3, 8, 10, 14], including the inability to learn continually without destroying past knowledge or constraining assumptions like task boundaries [6, 12] or replay data [1], incomprehensible structure [13], difficulty integrating learned information into deliberative behavior [2], and non-decomposability of learned behaviors [11]. These limitations stem from relying on fixed, unstructured models rather than learning the environment in a structured manner. To address these, we introduce the early design of a framework called Agential AI (AAI), composed of (1) Modelleyen – a gradient-free structural learning mechanism that enables continual learning without destructive adaptation, (2) a planner that uses learned models to deduce goal-directed actions, (3) a behavior encapsulation mechanism that can decompose behaviors into subunits. We outline AAI's design principles and demonstrate proof-of-concept results in a simple environment.

[1]A full version of this work can be found in [4].

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

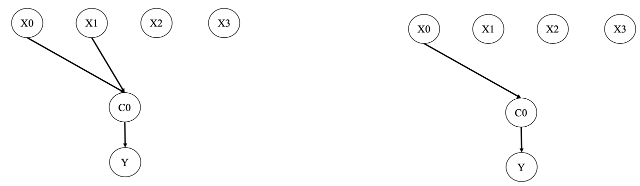

(a) (b)

Figure 1: Illustration of connection formation and refinement. The target relationship is $X0 \rightarrow Y$. (a) $X0, X1 \rightarrow Y$ observed. Connections from both $X0$ & $X1$ are formed. (b) $X0 \rightarrow Y$ observed. $X1$ is deduced unnecessary.

2 MODELLEYEN

Modelleyen is an algorithm for learning a model for an environment, enabling continual learning and information reuse while staying consistent with past experiences. Its core mechanism is a local variation and selection process, playing key role in achieving continual learning and structured environment modeling, forming the foundation for all other capabilities. Modelleyen can be applied to any prediction task, but this work specifically focuses on modeling the dynamics of sequential environmental observations.

The core organizational units in Modelleyen are *state variables (SVs),* which represent environmental observations, their dynamics, and internal computational units capturing relationships between SVs. A specific type, *conditioning SVs (CSVs),* drives learning by modifying its source and target composition. This occurs in a two-step process (Fig. 1): first, when an active SV lacks an explanation, a new CSV is created, initially linking all active observational SVs as positive (activatory) sources. Then, as the CSV state stabilizes, inactive positive sources are refined. An analogous mechanism applies to negative (inhibitory) sources. It can be shown (Theorem 1 in [4]) that these steps ensure past CSV responses remain intact after modifications, except during the initial formation of negative sources (a one-time event per CSV), offering a potent continual learning guarantee from the lowest level of organization.

Fundamentally different from methods like NNs, Modelleyen updates its model instantly with new information, leading to initial "overfitting" that allows for precise generalization afterwards via refinement to a minimal structure. This approach underlies the continual learning capabilities of our design.

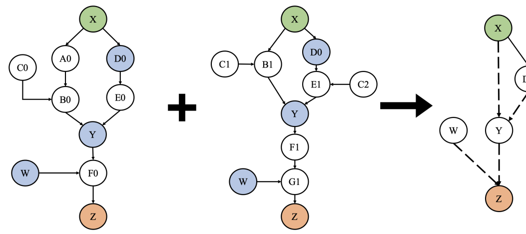

Figure 2: Illustration of the behavior encapsulation process, where two ANs representing alternative pathways are derived from a unified AN generated by the planner, and the pathways between X and Z are encapsulated by identifying reliable pathways in both networks, forming a new AN (right) where each dashed encapsulated edge contains copies of the corresponding subnetworks from the original AN variants.

3 PLANNER & BEHAVIOR ENCAPSULATION

Our planner design operates by constructing an action network (AN) -a dependency graph linking current states to goal states with alternative paths- using the model learned by Modelleyen, by tracing back from desired goals to current states. This contrasts with methods like model-based RL, which samples forward from initial states. The behavior encapsulation mechanism restructures the planner's exhaustive action networks into a hierarchical, comprehensible format, improving both interpretability and reusability. It does so by isolating alternative pathways into separate action networks, refining them by extracting common components, and encapsulating the subnetworks along these paths in different ANs into behavioral units (Fig. 2). Beyond enhancing action network interpretability, this method enables modular behavior reuse; though it is currently demonstrated independently, with full integration into agent operations left for future work.

4 EVALUATION

Setup: We validate AAI in a simple finite-state machine (FSM) environment that models various temporal succession patterns and is divided into three subtypes to evaluate continual learning performance. The agent learns actions step by step with Modelleyen, selects actions based on the planner operating on the current learned model state. We assess the average steps required to reach the goal as a metric for effectively modeling the environment and evaluating planner performance, while also presenting an example of behavior encapsulation applied to planner-generated action networks.

Planning with a learned model and continual learning: Our experiments show that Modelleyen can continuously learn a functional model of the environment without compromising existing knowledge when the environment subtype changes, allowing the planner to effectively operate on this learned model to achieve a specified goal state. This is illustrated in Figure 3, which demonstrates a decrease in the time it takes for the agent to reach the goal as learning progresses, while also retaining previous performance levels in specific environment subtypes even after exposure to and learning from multiple other subtypes.

Behavior encapsulation: Figure 4 presents a sample action network alongside a demonstration of the resulting encapsulated

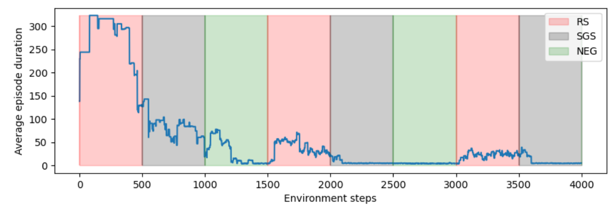

Figure 3: Average (across 5 trials) episode durations throughout learning with changing environment subtypes. Vertical limits show the environment changes.

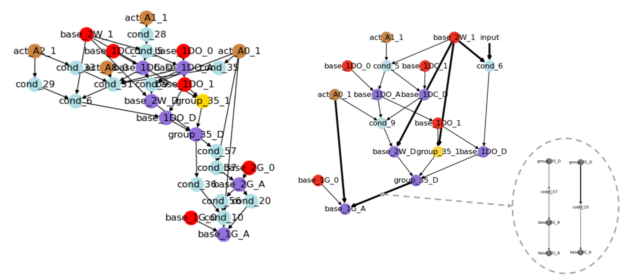

(a) Full action network. **(b) Encapsulated action network.**

Figure 4: Example of action networks on test environment. Bold edges are encapsulated. Each node represents a different state variable, and each edge represents conditioning and succession relations between them.

action network (AN). Despite the inherent complexity of the full action network, even in this simple environment, encapsulation effectively transforms it into a comprehensible, structured, and minimal format. As previously mentioned, the identified subgoals, pathways, and encapsulated components can serve as foundational subpolicies for future behaviors, although we have yet to fully integrate this mechanism into the agent's ongoing operations.

5 CONCLUSION

Agential AI, comprising Modelleyen, Planlayan, and a behavior encapsulator, addresses key challenges in classical machine learning by improving continual learning, enhancing comprehensibility, integrating learning and planning effectively, and enabling behavior decomposition into hierarchical structures. Its primary strength lies in constructing a structured model of the environment while preserving past information using a local variation and selection method. Future work will focus on extending Modelleyen to move beyond the current Markovian assumption, which relies on immediate temporal successions, and to adapt it for high-dimensional structured observation spaces, such as visual data. Both of these can be addressed by redefining the observation space as a network and modifying the foundational operation of Modelleyen to function on networks rather than lists of SVs; see [5] for an extension that applies this approach to visual processing.

REFERENCES

[1] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems* 33 (2020), 15920–15930.

[2] Sinan Çalışır and Meltem Kurt Pehlivanoğlu. 2019. Model-free reinforcement learning algorithms: A survey. In *2019 27th signal processing and communications applications conference (SIU)*. IEEE, 1–4.

[3] Jeff Clune. 2019. AI-GAs: AI-generating algorithms, an alternate paradigm for producing general artificial intelligence. *arXiv preprint arXiv:1905.10985* (2019).

[4] Zeki Doruk Erden and Boi Faltings. 2025. Agential AI for Integrated Continual Learning, Deliberative Behavior, and Comprehensible Models. arXiv:2501.16922 [cs.AI] https://arxiv.org/abs/2501.16922

[5] Zeki Doruk Erden and Boi Faltings. 2025. Continually Learning Structured Visual Representations via Network Refinement with Rerelation. arXiv:2502.13935 [cs.CV] https://arxiv.org/abs/2502.13935

[6] Maxwell J Jacobson, Case Q Wright, Nan Jiang, Gustavo Rodriguez-Rivera, and Yexiang Xue. 2022. Task Detection in Continual Learning via Familiarity Autoencoders. In *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 1–8.

[7] Abdullah Ayub Khan, Asif Ali Laghari, and Shafique Ahmed Awan. 2021. Machine learning in computer vision: a review. *EAI Endorsed Transactions on Scalable Information Systems* 8, 32 (2021), e4–e4.

[8] Yann LeCun. 2022. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review* 62, 1 (2022).

[9] Yuxi Li. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274* (2017).

[10] Gary Marcus. 2018. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631* (2018).

[11] Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek. 2021. Hierarchical reinforcement learning: A comprehensive survey. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–35.

[12] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671* (2016).

[13] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. 2019. Explainable AI: A brief survey on history, research areas, approaches and challenges. In *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II 8*. Springer, 563–574.

[14] Anthony M Zador. 2019. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature communications* 10, 1 (2019), 3770.

[15] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).