Incremental Learning Technology for Embodied Evolution of Intelligent Robots

Chunxi Yu

{000220@nkcswx.cn}

Wuxi Dipont School of Arts and Science, Wuxi, 214000, China

Abstract： With the progression of intelligent robots from controlled laboratory environments to complex and unbound real-world environments. The need for the robots to adapt to the changing environment has become of critical importance. In this study, the incremental learning methodologies are explored as the basis for achieving embodied evolution in robots. This study explores the three main incremental learning (IL) methodologies, including parameter regularization, knowledge distillation and structural expansion. The study also explores the new paradigm shift of semantic evolution using large language models and the unified world modeling approach. The study is of vital importance for the development of the next generation of autonomous robots.

Keywords: Incremental Learning, Intelligent Robots, Embodied Evolution

1 Introduction

1.1 Background

Throughout the past few decades, intelligent robots have gradually moved from the controlled environment of the laboratory to increasingly complex real-world scenarios. This has been motivated by the swift advances in the development of artificial intelligence, robotics, and sensors. Intelligent robots are being deployed more and more in such boundless scenarios as autonomous exploration, disaster response, intelligent manufacturing, space exploration, and service robots. In such unrestricted scenarios, robots are not only required to carry out assigned tasks but also have to constantly adapt to scenarios like dynamic environments, unexpected interferences, and changes in task objectives and the like. This makes the demand for robot systems that can continuously evolve and learn in unrestricted scenarios become increasingly greater. The robot system needs to keep on learning, also to develop new skills, and at the same time to retain those already learned capabilities as well. The demand for such continuous adaptation and evolution makes incremental learning and embodied revolution become two extremely important concepts in the development of the next-generation intelligent robot systems.

ICIAAI 2026, May 22-24, Singapore, Singapore

DOI 10.4108/eai.22-5-2026.2365329

Embodied evolution and incremental learning have both emerged from separate scientific traditions, but they have started to converge to achieve adaptive and autonomous robotic intelligence. The term 'Embodied Evolution' was first coined by Watson, R. A., Ficici, S. G., and Pollack, J. B. [1]. A departure from traditional optimization techniques based on simulations, the evolutionary operators are embedded directly into a population of physical robots by the researchers. Adaptability was hence achieved through real-time interaction within a physical environment to bridge the gap between simulated and real environments, thereby establishing a distributed evolutionary framework. From this pioneering work, Nolfi, S. and Floreano, D. developed a set of theoretical principles for evolutionary robotics [2]. These researchers highlighted the importance of co-evolution between morphology and control within evolutionary robotics. According to them, intelligence is a product of the dynamical interaction between body, controller, and environment. The embodied systems perspective thus provided a new framework for learning as a function of closed-loop sensorimotor interaction as opposed to algorithmic optimization. These theoretical principles hence provided a framework for understanding embodied evolution.

1.2 Recent advances of embodied evolution

In the 2010s, incremental learning was shown to play a vital role in embodied adaptation. Cully, A. presented a study on the Intelligent Trial-and-Error algorithm [3]. The robot was shown to recover from severe physical damage within minutes by rapidly searching for compensatory behaviors without retraining from scratch. In more recent years, a paradigm shift has been observed with the advent of large language models (LLMs) and vision-language models (VLMs) for embodied systems to evolve at a semantic level. Robots now use high-level reasoning components to extend learning to a level beyond low-dimensional control policies. Bärmann et al., for example, presented a study in 2024 where they developed a robot learning architecture using LLMs as cognitive cores for humanoid robots [4]. At the foundatioal model level, Wu et al. have developed LingBot - VLA which is a large scale vision - language - action foundational model for the embodied artificial intelligence and has been trained on about 20,000 hours of real world robot data [5]. Meanwhile the Robbyant team put forward LingBot-World in 2026 which is an open source world model that enables more predictive and interactive learning beyond pure imitation-based approaches [6]. In addition to large-scale fundamental models, Meng et al. developed a lifelong robot reinforcement learning framework, which this framework focuses on keeping and integrating previously acquired knowledge during continuous skill learning processes and thereby solving the problem of catastrophic forgetting in long-term specific adaptation [7]. This knowledge integration mechanism that provides a foundation for a modular and extensible embodied learning system is an important step towards general robot intelligence capable of continuously evolving in different tasks and environments. These recent developments indicate a transformation of a paradigm, from adaptation at the parameter level changing into evolution in the semantic aspect, from learning a single task changing into lifelong learning oriented toward embodied intelligence, and from simulated learning changing into gradual improvement in the real environment.

1.3 Research motivation and planning

In view of these limitations of the static learning paradigm in the dynamic physical environment, this research makes an attempt to carry out a comprehensive exploration of the incremental learning method of embodied evolution in intelligent robots. Specifically, it is focused on dealing with how intelligent robots can continuously learn new skills, adapt to dynamic environments, and minimize catastrophic forgetting while maintaining past knowledge bases. To this end, the remaining part of this paper is structured as follows: Section 2 presents the theoretical background of embodied evolution and incremental learning; Section 3 analyzes the mainstream incremental learning paradigms, including parameter regularization, knowledge distillation, and structure expansion; Section 4 reviews the applications of these methods in embodied intelligent robots; Section 5 sums up the identified challenges and the potential areas that need to be explored in the future.

2 Basic principles of embodied evolution and incremental learning

2.1 Principle of embodied evolution

Embodied evolution can be regarded as a specific paradigm within evolutionary robotics. In 1999, Watson et al. defined it as an evolutionary process that occurs directly among a group of entity robots, and it is accomplished by continuously, independently, and dispersedly interacting with their environment [1]. In this regard, although traditional evolutionary algorithms for controller evolution have controllers that have evolved offline and then deployed within physical systems, the evolutionary process within embodied evolution does get directly interwoven into the robot's real-time life cycle indeed. In this situation the robot functions as a phenotype and that robot while completing tasks in the real world to evaluate its own adaptability also interacts with other robots exchanging genetic information in the form of neural network weights through local communication channels. This decentralized mode that bypasses the real disparities ensures the behaviors learned to be rooted in the physical reality of the world.

2.2 Incremental learning

Incremental learning, which is often termed as lifelong learning or continual learning, is defined as the capability of an intelligent system to sequentially learn new knowledge from a non-stationary data stream while maintaining the learned knowledge. The major challenge in incremental learning is the stability-plasticity trade-off, which requires that the system should be sufficiently plastic to incorporate new knowledge while at the same time being stable enough not to forget the knowledge learned so far due to catastrophic forgetting caused by new updates to the weights that were critical for the previous tasks. Based on the guidelines provided by Polikar and the advancements made by Kirkpatrick et al., IL methods can be categorized into the following three main methodologies [8, 9].

1）Regularization-based methods, which ensure that the knowledge learned so far is not forgotten by penalizing the weights that were critical for the previous tasks, such as EWC.

2）Replay-based methods, which store a small number of examples from previous tasks and use them to prevent forgetting, such as Incremental Classifier and Representation Learning (iCaRL).

3）Architecture-based methods, which modify the architecture of the model by increasing the number of neurons to incorporate new tasks while leaving the original architecture unchanged.

Incremental learning would thus serve as the requisite mechanism for intelligent robots to adapt to new environments through incremental evolution.

3. Research status of incremental learning techniques

Incremental learning has been extensively researched as a potential solution to address catastrophic forgetting in long-duration robotic learning systems. In the context of embodied evolution, robots need to learn new behaviors while retaining learned capabilities under changing environmental conditions. This has led researchers to propose various incremental learning techniques that can be classified into parameter regularization-based methods, knowledge distillation-based methods, network structure-based methods, among others. The following parts will bring into focus the diverse state-of-the-art technologies according to these methods, and at the same time assess their feasibility for the embodied robotic systems.

3.1 Parameter regularization-based methods

Parameter regularization-based approaches have been proposed as methods that aim to retain the learned knowledge while limiting the updates of the most important neural network parameters during the process of learning new tasks. The most prominent contribution in this area has been made, where the authors proposed the EWC method for preventing catastrophic forgetting in neural networks [9]. The importance of the parameters was calculated using the Fisher Information Matrix, and the updates of the most important parameters were penalized during the subsequent learning process. This way, the neural networks were able to retain the performance for the learned tasks while adapting to the new tasks in a sequential manner. An algorithm framework diagram of EWC is shown in Fig. 1, which demonstrates how robots protect critical weights during sequential learning, allowing skill refinement without overwriting neural connections essential for existing robotic capabilities.

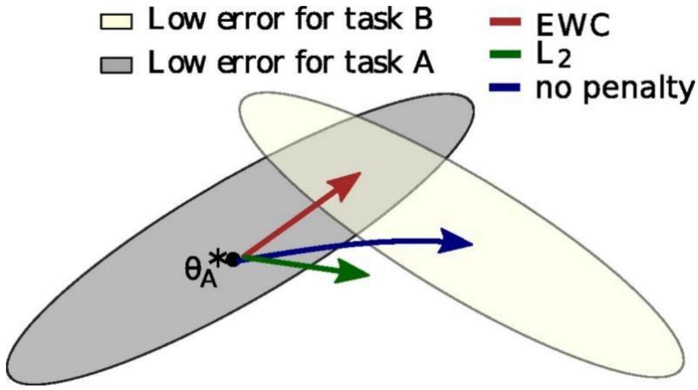

Fig. 1. An Algorithm Framework Diagram of EWC [9].

Zenke et al. proposed the Synaptic Intelligence (SI), which improved the Elastic Weight Consolidation method by taking into consideration the importance of the parameters over the entire training process using the Fisher Information Matrix approximation [10]. This helped the neural networks in the process of stable incremental learning in dynamic environments where the tasks were fed in a sequential manner. Another regularization-based method for preventing catastrophic forgetting in neural networks was proposed, where the importance of the parameters was calculated based on the output function and the parameter updates were made in the Memory Aware Synapses (MAS) [11].

These regularization-based approaches for incremental learning have been found to be very useful as they do not require the storage of the previous training data and the modification of the neural network architecture. This property has been found to be very useful for the implementation of the incremental learning paradigm in the development of embodied robotic systems that have limited computational and memory capabilities. However, the parameter regularization-based approaches have been found to have many shortcomings as the number of learned tasks increases, resulting in the degradation of the ability of the neural networks in the process of learning new tasks. Therefore, the parameter regularization-based approaches can be employed in scenarios where the changes in the tasks and the sequence length are relatively gradual, such as in the development of autonomous robotic systems for refining the learned skills in the long term.

3. 2 Knowledge distillation-based methods

The approaches based on the knowledge distillation concept avoid the phenomenon of catastrophic forgetting by transferring the information from the preceding model to the new model during training, based on the alignment of the output distributions of the two models. The concept of knowledge distillation was first proposed by Hinton et al. [12]. They showed that the output distributions represented by a soft probability distribution are more informative than hard labels.
A prominent example of a knowledge distillation-based incremental learning method is the Learning without Forgetting (LwF) concept, proposed by Li and Hoiem [13]. This method involves training the model on new tasks while simultaneously preserving the responses to the preceding tasks through the distillation loss, where the output distributions of the current model are aligned with the output distributions of the preceding model. This allows the model to learn new information while retaining the representations learned from the preceding tasks. Another prominent incremental learning method is iCaRL, proposed by Rebuffi et al. [14]. This method uses the concept of knowledge distillation, combined with the concept of exemplars, to incrementally learn new object categories while retaining a consistent feature representation. This method has shown promising results in the class incremental learning benchmark, serving as the basis for many incremental learning methods used for robotic vision tasks. An algorithm framework diagram of LwF is shown in Fig. 2. By maintaining output consistency, LwF ensures vision-based representations remain robust as the robot adapts incrementally to new environmental tasks.

The method founded on the concept of knowledge distillation is a flexible way of transferring information among models, especially when the system has a need to learn new information without having to retrain the model from the very beginning. This concept, which is very promising to be applied to the incremental learning methods of robot vision, especially when the robot system needs to learn new information, for instance, when interacting with humans.

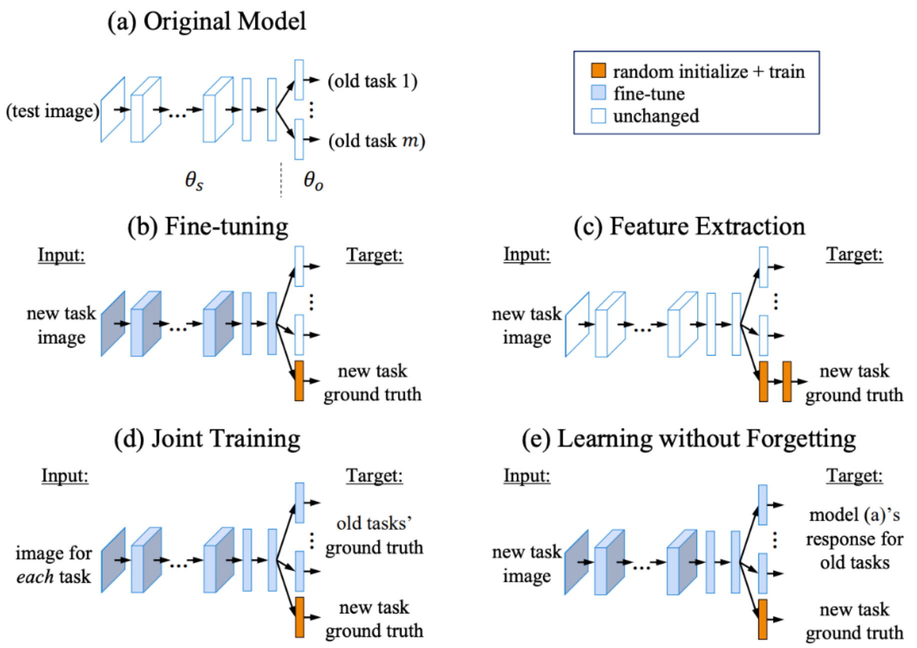

Fig. 2. An Algorithm Framework Diagram of LwF. (a) Original Model; (b) Fine-tuning; (c) Feature Exraction; (d) Joint Training; (e) Learning without Forgetting [13].

3.3 Network structure-based methods

The way of structure expansion in incremental learning is to update the neural network architecture in order to add new tasks and also keep the existing knowledge. These methods do not control the parameter updates but rather add more network capacity to the model so as to integrate new tasks without interfering with the existing knowledge.

One of the most important and impact papers on the subject is the Progressive Neural Networks (PNN) proposed by Rusu et al. [15]. In the PNN model, the addition of a new task results in the addition of more columns to the network, with the existing columns remaining constant. The lateral connections allow the model to leverage the knowledge learned from the previous tasks without changing the parameters learned earlier, thus reducing the forgetting of existing knowledge. Another important paper on the subject is the Dynamically Expandable Networks (DEN), proposed by Yoon et al. [16]. In the model, the network has the ability to automatically expand the architecture to incorporate new tasks. The model selectively adds neurons to the network and retrains the network on the relevant part of the network, thus enhancing the efficiency of the model. An algorithm framework diagram of PNN is shown in Fig. 3, which highlights open-ended learning.

Under the background of embodied evolution, the structure expansion method is of great significance because robots in the long-term operation process are very likely to come across entirely brand-new tasks or complex situations in the environment. The robot thus would be required to assimilate new skills without bungling the existing ones. But the major issue with this approach is the network structure that is continuously increasing, and this could lead to an increase in computing cost. So, the approach based on architecture is rather important in the context of open learning because a new skill has to be learned each time the robot operates.

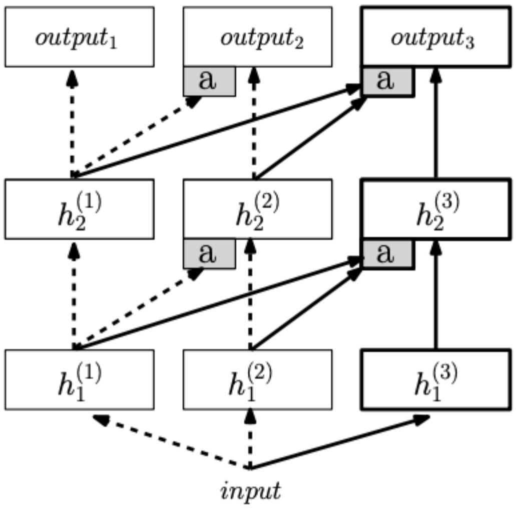

Fig. 3. An Algorithm Framework Diagram of PNN [15].

4. Current limitations and future prospects

Including incremental learning into the framework of embodied evolution has brought about a revolutionary path for developing robots which can learn and adapt to the surrounding environment in the same way as living organisms. However, as research progresses from basic behavioral controllers to more complex cognitive learning, quite a few systematic limitations are revealed, and these limitations now set the boundary between controlled research and truly autonomous, lifelong learning robots.

4.1 Stability-plasticity dilemma and semantic generalization

Generalization is one of the core limitations, which is defined by the stability-plasticity dilemma inherent to incremental learning systems. A learning robot requires to have a certain degree of malleability so as to learn new cognitive skills but also requires stability in order to avoid catastrophic forgetting, in which situation new learning data will overwrite the neural connections of previously learned skills. Recent breakthroughs in large language models and vision-language-action models make it possible for robots to learn complex cognitive skills through human-machine interaction, yet they cannot be generalized to environments outside the definition of the training data, in which the sensory inputs are quite different from what the model originally encountered. In order to make further progress in this regard, it is necessary to come up with a number of knowledge evolution cycles. These cycles that can achieve the generalization of semantics are enabling robots to learn a set of underlying principles for specific tasks, ensuring that their behavior libraries are quite reliable in various situations.

4.2 Data utilization efficiency and world modeling

The efficiency of data utilization is of great importance to the evolution of the real world, especially considering the persistent phenomenon of the gap that exists between simulation and actual manifestation. Traditionally, embodied evolution mainly depends on real-time interactions to make sure that behaviors are based on physical realities; however, this is often not practical in terms of time. With recent developments like LingBot-World, there has been a focus on using a large quantity of data to construct a unified world model. However, the future of evolution is based on reducing these amounts significantly using high efficiency learning paradigms. With the application of the intelligent trial and error paradigm, future robots that are able to recover from damage or adapt to new tasks within a few minutes can be foreseen. By making use of causal reasoning within a world model, the simulation of outcomes in a robot's mind can be accomplished; in this way, the efficiency is increased significantly, and the evolution can be accelerated correspondingly.

4.3 Physical constraints and hardware safety awareness

The physical constraints uniquely bind the embodied evolution with the physical laws and the hardware constraints. The controller is closely connected with the form of the robot, and the form deteriorates over time. Although basic research has emphasized the importance of co-morphology, the simultaneous evolution of the robot's body and brain, the current progressive embodied evolution methods are increasingly regarded as a problem centered on software. A very promising path for future research is integrating physical security awareness into the progressive controller updates. This demands the design of autonomous algorithms to monitor the mechanical limits or hardware degradation conditions of robots, thereby adjusting the learning process in a corresponding way. In this way, it can be ensured that the discovery of new motor skills does not cause structural problems for the robot, thus maximizing the service life of the robot in harsh environments.

4.4 Decentralized transfer learning and collective evolution

The key for the practical application of embodied intelligence will be the learning transfer among different platforms. At present, contemporary incremental learning benchmarks like iCaRL are mainly for benchmark testing in relation to incremental single-agent task updates. However, the real embodied evolution occurs in a group of agents in a decentralized process. Therefore, future research should be focused on promoting that kind of decentralized transfer learning model in which a group of robots will communicate with each other locally and share genetic knowledge. This would make it possible for researchers to design unified knowledge architectures, which would then relieve the catastrophic forgetting situation of the entire robot swarm, and thus promote a collective evolution process in which the experiences of each robot can contribute to the shared foundation of robot intelligence.

5 Conclusion

In conclusion, this study provides a comprehensive analysis of how incremental learning methods can be effectively integrated as an integral component of the embodied evolution of intelligent robotic systems. The results show that parameter regularization techniques such as synaptic intelligence can help in enabling efficient and memory-saving optimization of robot skills in systems having limited computing power. Also, knowledge distillation and extension techniques provide a reliable framework for integrating new skills without affecting the existing knowledge base. Moreover, using large language models as the cognitive core and unified visual-language-action models like LingBot-World to achieve semantic evolution represents a significant shift of the evolution mechanism from parameter evolution towards causal evolution and natural human-computer interaction. In the future the research within this field will center on developing decentralized transfer learning models and content that is related to hardware security awareness making it possible for the robot groups to share genetic knowledge in a manner that is in line with the physical limitations and wear of their mechanical bodies. By dealing with the problems of stability and plasticity and by using intelligent trial and error to enhance data utilization efficiency, the progress of the general embodied foundation model can then be promoted. The contribution of this research consists in filling the gap between simulated intelligence and real-world intelligence and having put forward a path that leads to collective evolution, on which the whole robot swarm can evolve as a compact unit as if relying on its own life through sharing knowledge frameworks.

References

[1] Watson, R. A., Ficici, S. G., Pollack, J. B.: Embodied evolution: Embodying an evolutionary algorithm in a population of robots. Proc. IEEE Congress on Evolutionary Computation (CEC), Washington, DC, USA, pp. 335–342 (1999)

[2] Nolfi, S., Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press (2000)

[3] Cully, A.: Robots that can adapt like animals. Nature. Vol. 521, pp. 503–507 (2015)

[4] Bärmann, L., Kartmann, R., Peller-Konrad, F., Niehues, J., Waibel, A., Asfour, T.: Incremental learning of humanoid robot behavior from natural interaction and large language models. Frontiers in Robotics and AI, 11, Article 1455375 (2024)

[5] Wu, W., Lu, F., Wang, Y., et al.: A Pragmatic VLA Foundation Model. arXiv preprint arXiv:2601.18692 (2026)

[6] Robbyant Team, Gao, Z., Wang, Q., et al.: Advancing Open-source World Models. arXiv preprint arXiv:2601.20540 (2026)

[7] Meng, Y., et al.: Preserving and Combining Knowledge in Robotic Lifelong Reinforcement Learning. Nature Machine Intelligence, Vol. 7, No. 2, pp. 256-269, (2025)

[8] Polikar, R., Upda, L., Upda, S.S., Honavar, V.: Learn++: An incremental learning algorithm for supervised neural networks. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), Vol. 31, No. 4, 497–508 (2001)

[9] Kirkpatrick, J., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences. Vol. 114, No. 13, pp. 3521–3526 (2017)

[10] Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. Proc. International Conference on Machine Learning (ICML), Australia, Sydney, pp. 3987-3995 (2017)

[11] Aljundi, R., et al.: Memory aware synapses: Learning what (not) to forget. arXiv preprint arXiv:1711.09601 (2018)

[12] Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)

[13] Li, Z., Hoiem, D.: Learning without forgetting. arXiv preprint arXiv:1606.09282 (2016)
[14] Rebuffi, S.-A., et al.: iCaRL: Incremental classifier and representation learning. arXiv preprint arXiv:1611.07725 (2017)
[15] Rusu, A. A., et al.: Progressive neural networks. arXiv preprint, arXiv:1606.04671 (2016)
[16] Yoon, J., et al.: Lifelong learning with dynamically expandable networks. arXiv preprint arXiv:1708.01547 (2018)