INTERNATIONAL JOURNAL OF INNOVATIONS IN ENGINEERING TECHNOLOGY AND APPLIED SCIENCES (IJIETAS)

# A Comparative Study on Learning Approaches towards Artificial General Intelligence

[1*]Aditya Bhowmic, [2]Amit Sharma

[1],[2]Department of Computer Science & Engineering Sharda School of Computing Science and Engineering Sharda University, Greater Noida.
[1]adityabhawmic@gmail.com, [2]amit.sharma3@sharda.ac.in

## ABSTRACT

The concept of AGI (Artificial General Intelligence) denotes systems able to perform a diverse set of tasks in a range of environments, rather than a single narrow task. Among self-supervised prediction scaling, reward-based interaction, meta-learning, structured world models and continual learning, which strategies appear to be most promising in providing such generality? This work is a theoretical comparative study. It first puts the literature in order by presenting a taxonomy of the relationships of these paradigms to the requirements of general intelligence and its main constraint. Then it investigates a problem they all share in common: that is the inability to learn new skills sequentially while maintaining prior ones (catastrophic forgetting). The continual learning problem is formalized, standard three scenarios (task-, domain- and class-incremental) are described and main families of mitigations, regularisation (Elastic Weight Consolidation) and replay (experience replay) are analysed and compared with the upper bound of joint training. On the basis of the published evidence and the problem structure, it is argued that weight regularisation alone is insufficient for the hardest case (class-incremental) and that an explicit memory mechanism should be involved, according to the principle of complementary-learning-systems from neuroscience. The study concludes that a single learning approach is not sufficient for general intelligence, and advocates that hybridisation of large-scale learning with efficient memory and adaptation appears most justifiable.
Essentially there are two types of AIs: general AI which can learn and operate in many different contexts (General AI), and specific AI which can only work in one specific task (specific AI).

**Keywords:** Artificial General Intelligence (AGI), Continual Learning, Catastrophic Forgetting, Self-Supervised Learning, Meta-Learning, Elastic Weight Consolidation (EWC), Experience Replay.

**ARTICLE HISTORY**
**Received:** 15 July 2026, **Accepted:** 28 September 2026, **Published:** 12 October 2026

**CITATION**
Bhowmic, A., (2026). A Comparative Study on Learning Approaches towards Artificial General Intelligence, *International Journal of Innovations In Engineering Technology And Applied Sciences (IJIETAS),* 2(2), 01-08.
**DOI:** https://doi.org/10.64764/ijietas.v2.i2.01

### Problem Statement

People have associated Artificial General Intelligence (AGI) with the ability to accomplish goals in many different environments [1, 2]. The current systems, however, remain very specialised. Language models are not changing as the world changes, or a game playing agent cannot code. We don't know the most likely way of bridging that gap. Some people believe that having a very large amount of data, then scaling-up the self-supervised prediction would suffice [8, 11]. Some people believe that we only need to have interaction, and optimise for the reward, to obtain general abilities [7]. Some other people believe that you need structure (e.g. causal models, compositionality) to have generality as a human [4, 20]. And finally, some other people believe that it is not so much how many skills you already have, but how efficiently you can learn new skills [3].

**Problem addressed**

In general, the learner needs to keep learning new skills throughout the lifetime. When one NN is trained with multiple tasks sequentially, the problem of forgetting about previous tasks (catastrophic forgetting) arises [14]. First, there is a lack of short and authoritative comparisons of different learning techniques concerning AGI needs. Second, it's not clear for new (and sometimes for experts) how much simple mechanisms for continual learning could help when only one model and one output head for all the classes it saw before is required.

**Objectives and contributions**

**Survey & Taxonomy:** A comparative survey of most important learning paradigms (supervised/self-supervised, reinforcement, meta-learning, continual learning, structured, world model approaches) for AGI, with their limitations,

**Formal Statement:** A formal statement of the problem of continual learning (as a proxy of general learning requirements to learn sequentially with a single model), with standard scenarios and metrics of evaluation in continual learning

**Theory Analysis:** Analysis and discussion of two families of continual learning approaches (regularization, rehearsal) (with joint training as an upper bound), in terms of when they work or fail, and expected performance

**Implications:** Discussion of the implications for AGI and research directions

**Scope**

This work is purposefully limited to a survey and analysis of the literature. We are not building or measuring AGI, there are no new empirical results.

**Literature Survey**

Following is a summary, by theme, of the literature survey of published work. It includes definitions of general intelligence, the most common learning paradigms, and analysis of continual learning methods (discussed in the section on Continual-Learning Strategies).

1. Legg & Hutter, 2007 [1]. Theme/Approach: Definition of intelligence. Main contribution: Formal Definition - Ability of an agent to accomplish goals in a wide range of environments. Limitation for AGI: Often idealised and impractical. No recipe for training.
2. Goertzel, 2014 [2]. Theme/Approach: Survey of AGI. Main contribution: Survey of Symbolic, Emergentist, Hybrid and Universalist approaches to AGI and the open challenges. Limitation for AGI: Conceptually compare not empirical.
3. Chollet, 2019 [3]. Theme/Approach: Intelligence: Measuring it. Main contribution: Definition of intelligence as efficiency of skill acquisition. ARC Benchmark introduction. Limitation for AGI: Limited scope, only a part of generality is measured.
4. Lake et al., 2017 [4]. Theme/Approach: Human-like learning. Main contribution: Justification for Causal Models, Intuitive Physics/Psychology, Compositionality, Learning to learn. Limitation for AGI: Mostly a position paper. Few systems implemented.
5. Mnih et al., 2015 [5]. Theme/Approach: Deep Reinforcement Learning. Main contribution: Single architecture that learns multiple Atari games from pixels and score. Limitation for AGI: No transfer across games. Requires large amounts of data. Each game needs separate training.
6. Silver et al. 2017 [6]. Theme/Approach: Self-play RL. Key contribution: No human game records required to learn Go from self-play. Relevant limitation to AGI: Requires known full specification simulator (board-game domain).

7. Brown et al. 2020 [9]. Theme/Approach: Large language models. Key contribution: Few-shot in-context learning at scale across many language tasks. Relevant limitation to AGI: After training, weights are fixed. Unexplained reliability gaps. Not possible to do lifelong updates.
8. Reed et al. 2022 [10]. Theme/Approach: Generalist agent. Key contribution: Multiple tasks (games, caption, robot control) using a single network. Relevant limitation to AGI: Some copying previous data. Satisfying on most tasks.
9. Finn et al. 2017 [12]. Theme/Approach: Meta-learning. Key contribution: MAML learns an initialisation that can be adapted to new task in a few gradient steps. Relevant limitation to AGI: Works for new tasks similar to training task distribution.
10. Kirkpatrick et al. 2017 [15]. Theme/Approach: Continual learning (regularisation). Key contribution: Slows down changes to important weights for previous tasks. Relevant limitation to AGI: Weak if tasks need to have same output head (penalty adds up) [18].
11. Rolnick et al., 2019 [17] Theme/Approach: Continual Learning (replay). Key Contribution: Important for mitigating forgetting in continual RL problems using experience replay. Relevant Limitation to AGI: Requires storing previous data. Privacy/storage issues.
12. van de Ven & Tolias, 2019 [18] Theme/Approach: Continual learning evaluation. Key Contribution: Class incremental learning is harder than task incremental learning (no success in regularisation). Relevant Limitation to AGI: Small benchmarks. Mostly vision-based.
13. Parisi et al., 2019 [16] Theme/Approach: Continual learning review. Key Contribution: Reviews regularisation, dynamic architecture, complementary learning approaches. Relevant Limitation to AGI: Only a review. No agreed benchmark.

Survey insights. In summary, the survey has two insights: 1) none of the approaches covers all aspects of general learner need, but all of them contribute some of them, for example: scale gives breadth [8], interaction gives goal-directed skill [5], [6], meta-learning gives fast adaptation [12], and continual learning gives retention [15], [17]; 2) one of the recurring gaps is open-ended (life-long) learning with a single model (which is the subject of this work). The mechanism of rehearsal studied here is inspired by the concept of complementary learning systems in neuroscience [19] which suggests that to consolidate new experience, the brain replays it together with the old experience.

**Agi Learning Approaches Taxonomy**

We organise the different approaches by their main mechanism, strength w.r.t. AGI and main limitation.
Supervised/self-supervised Learning. Main mechanism: Predict (a label/next token) from a very large dataset [8], [9]. Main strength for AGI: Much knowledge. Few-shot adaptation. Main limitation: Static after training. High compute. Low grounding.

**Reinforcement Learning. Main mechanism:** Maximise reward from interacting with environment [5], [7], [21]. Main strength for AGI: Goal-oriented sequential decision making. Main limitation: Low sample efficiency. Requires reward design and simulators.

**Meta-learning. Mechanism:** learn how to learn in a distribution of tasks [12], [13]. AGI strength: rapid adaption to new tasks. AGI weakness: tasks similar to training distribution.

**Continual / lifelong learning. Mechanism**: retain old skills while learning new ones [15]-[17]. AGI strength: unbounded learning of skills. AGI weakness: forgetting. Memory/regularisation scalability.

**Structured and world-model approaches. Mechanism:** learn cause and effect models, use them for planning [4]. AGI strength: efficient reasoning. Transfer [20]. AGI weakness: relatively few mature large-scale implementations.

**Supervised and self-supervised learning**

Large models trained on predicting the next token or a human-provided label learn general knowledge and are then able to adapt to new tasks with few examples in context [8], [9]. Some authors interpret this breadth as an indication of early signs of general capability [11]. But since we fix the parameters at the end of the learning process, there is no learning of new skills from the model's own experience and the connection to real world is weak.

**Reinforcement learning**

Reinforcement learning (RL) is a learning paradigm where an agent learns to take actions that maximize its cumulative reward by interacting with an environment [21]. RL has been successful in video games [5] and in self-play for Go [6]. It has been proposed that the reward signal can be sufficient to grow the skills of intelligence [7]. However, RL is sample inefficient and highly depends on the design of reward function. RL has been typically used in well-specified simulators and for each task in isolation, where the task is trained individually.

**Meta-learning**

Learning to learn, or meta-learning, is about training the learner to quickly adjust to a new task from a distribution of tasks [12], [13]. This is a direct response to Chollet's efficiency criterion [3]. However, it assumes new tasks are similar to the meta-training ones. This is not guaranteed in an open-ended environment.

**Structured/world-model approaches**

For learning from minimal data, Lake et al. [4] propose that we learn by constructing the causal model of the world and compose known concepts to build new concepts. LeCun [20] proposes architectures for learning and utilizing predictive world models for planning. These approaches offer efficient reasoning and transfer. However, only few of them are implemented at large scale.

**Continual and lifelong learning**

Continual learning is the process of learning new skills while not forgetting old skills [15][16][17]. It is not an alternative to the above but an additional property on them, since the foundation model, the RL agent and the meta-learner all need to be continuously learning post-deployment, which is problematic. Thus, this report focuses on continual learning for the rest of the report.

**The Continual-Learning Problem**

**Catastrophic forgetting**

Catastrophic forgetting (also known as catastrophic interference) was identified by McCloskey and Cohen [14] in connectionist networks's sequential learning. When a new task is learned, the common weights are updated so that the network can learn for the new data. But the same weights are used for the previous tasks. The performance may collapse for previous tasks. This difficulty is due to the distributed representation, since the capability of generalisation may make the knowledge suffer from interference. This problem is often referred to as the stability-plasticity dilemma, since the learner needs to be enough plastic to absorb new knowledge and be enough stable to hold on to previous knowledge.

**Formal setting**

Given a sequence of T tasks (1, 2, …, T). Data for task t (D_t) is shown to the learner before data for task t+1 (D_(t+1)) (that's what makes it different than regular multi-task learning). i.e. there is some restriction or removal of access to previous data. The learner is tested on the test data for all tasks j after training on task t.

**Three scenarios**

As per [18] by van de Ven and Tolias there are 3 (increasingly hard) scenarios:

Task-incremental: Task is known at test time, thus allowing you to use a distinct head (output) for each task.
Domain-incremental: No or need for task knowledge as the output space remains constant across tasks, but the input distribution changes.

**Class-incremental:** No task knowledge. Must learn to discriminate all classes seen so far with a single output layer.

Class-incremental is the hardest (most similar to general intelligence) since the task is implicit. And hence most instructive on 'what does a continual-learning method actually do' in this setting.

**Evaluation metrics**
Let A [t,j] be the test accuracy for task j after training on task t. Usual metrics
Final avg. accuracy = (1/T) \sum_j A[T,j] the average accuracy across all T tasks after the last task is learnt
Forgetting = (1/(T-1)) \sum_(j<T) ( max_t<T A[t,j] - A[T,j] ). The average drop from best accuracy on each earlier task to its final accuracy.

The final average accuracy is the overall competence at the end of the sequence, and forgetting is the sole measure of loss in previously learned skills, so it's not possible to use good performance on the most recent task to obscure forgetting.

**Continual Learning Strategies**
We consider four reference strategies. These encompass both of the major classes of continual learning techniques, regularisation and rehearsal [16], as well as both standard bounding strategies.

**Naive Sequential Learning (lower bound)**
This is the simplest approach, which simply trains one task at a time without any form of protection for prior knowledge. It is the performance of an unprotected network. It is the baseline against which the other methods are measured. Standard back-propagation with stochastic gradient descent is used for training [22].

**Regularisation: Elastic Weight Consolidation**
Elastic Weight Consolidation (EWC) [15] retains knowledge by constraining the variation of important weights from previous tasks. After a task has been learnt, the sensitivity of the output distribution of the model to each parameter is estimated as its importance by the Fisher information F_i at the time that parameter was learnt. A quadratic penalty is added to the loss for learning a subsequent task:

$L(\theta) = L_t(\theta) + (1/2\ \lambda)\Sigma_i F_i (\theta_i - \theta^*_i)\ 2$ where the loss of the current task is L t , the previous task's weights are $\theta *$ and $\lambda$ controls how much we protect old knowledge. A Bayesian interpretation of the penalty is that it represents the posterior distribution of the weights after the outcome of the previous tasks, which we approximate by a Gaussian around the value $\theta *$. This means that we keep "elastic" the values of weights which are important for previous tasks around their previous values. Two points are worth noting. Firstly, the gradient of the penalty will be stiff with a large $\lambda$ which can lead to an unstable optimisation. A small $\lambda$ gives little protection. Secondly, the penalties grow with the number of tasks, which will make it more difficult for the network to learn new tasks as the number of tasks increases [15].

**Rehearsal:** Experience replay
Replay of a few examples from previous tasks along with the current task data during the training process [17]. By doing this, the learner sees the example of old tasks in training, and thus the gradient followed is a mix of old and new tasks. The gradient is not purely driven by the new task. This replay approach resembles the memory replay in the complementary learning system theory [19]. It proposes that the hippocampus encodes new experience quickly and then replays them to the neocortex. It enables a slow interleaved learning

process without overwriting the previous memories. There are some disadvantages of keeping the previous task's data. The storage is possibly limited. Privacy may be concerned. Also, it is not clear what examples to keep.

**Joint training (upper bound)**

Joint training means that we train all the classes in one time with the combined data. This is not a sequential constraint-satisfying approach, therefore, it is not a continual learning method. But this is a showing of the same architecture's ability without forgetting. This is the usual upper bound in the continual learning evaluation [16], [18].

**Theoretical Analysis**

This section states what the expected behaviour of strategies in the class-incremental scenario would be based on problem structure and published evidence. Note that they are expected behaviours analytically, not new measurements.

**Why naive training fails**

The class-incremental scenario, when training a new task, only involves the new classes. The cross-entropy loss that includes the softmax shared output encourages the network to increase scores of the current class and decrease all other classes including the ones learned previously. But there is no signal for the opposite direction. Hence the network learns only the latest classes. It forgets almost all previous classes but it works well for the latest classes. This is the classic version of catastrophic forgetting [14].

**Why regularisation is expected to have trouble**

EWC restricts the update of important weights for the previous tasks, and it is effective for cases where the tasks are separated by the task specific outputs (e.g. the case of task incremental [15]). But in the case of class-incremental, there is also a need to learn the boundaries between different task-classes. Without previous data for the previous tasks, the penalty doesn't provide the contrast information between old classes and new classes. The only effect is to make the change of weights slower. Which is exactly what van de Ven & Tolias [18] show: that the methods based solely on regularisation don't work well in the case of class-incremental. And the analysis shows there is a trade-off when tuning $\lambda$. Too large may lead to instability in the optimisation, and too small may turn the method to a naive training. So this is not that EWC is ineffective in general, but that alone it's poorly suited to the worst case scenario.

**This will enable rehearsal to help. this is because rehearsal should help.**

Replay provides examples of old classes, intermingled in the same mini-batches as new classes. The model is directly benefited in regard to the missing inter-task contrast and is still punished for the lack of knowledge of the classes whose examples it stores. The forgetting effect of experience replay is very significant as demonstrated by Rolnick et al. [17] and the rehearsal family is the most effective in class-incremental benchmarks [18]. The buffer is, however, limited in size so that the examples in it may not be the most representative; hence the old tasks are still expected to be learnt less well than the most recent one, and less well than under joint training. Rehearsal helps to minimize forgetting but does not prevent it, and works best when memory is extensive and varied.

**Expected ordering**

The qualitative order in the class-incremental setting after combining analysis is as follows: naive sequential training and EWC are at the bottom, replay is at a much higher level and the joint training is at the highest level. The order in which the methods are presented here is not an indication of any specific accuracy values. Empirical studies should be carried out to confirm the magnitudes as described in Future Scope.

**Discussion:** Implication for Agi

Let's not rely on weights for lifelong learning. The main point of the survey is that a learner needs to keep learning skills through one model and has to have an explicit method for retaining the previous skills, which is consistent with the complementary-learning-systems theory [19]. The various tasks are not distinguishable by means of regularisation.

Scalable and selective memory is the largest bottle neck. Learning throughout a person's life is an open-ended process, a system that cannot store everything. The big hurdle for achieving AGI will be making memory scalable and selective, meaning what to store, how to compress it, or how to create it rather than store it [16].

Continual learning is a complement to the other paradigms. They provide breadth in the form of foundation models [8] and goal-driven behaviour through RL [5]–[7] and rapid adaptation through meta-learning [12] and efficient reasoning through structured models [4] and [20]. None of them keeps learning after deployment without forgetting, and that's where continual learning comes in handy as a limitation for each of them.

Evaluation matters. The apparent merit of a method varies based upon the context within which it is assessed [18]. A method that is successful only in the task-incremental should only be credited for the task-incremental. Any discussion of progress towards general intelligence should, therefore, mention the situation and acknowledge that forgetting has occurred as well as accuracy.

**Limitations**

This study is a review and analysis without any new experiments; it depends on the results of other published studies and the theory behind them. The treatment of continual learning focuses on one regularisation technique, one rehearsal technique and joint training, while excluding detailed treatment of dynamic-architecture and parameter-isolation and generative-replay approaches [16]. The taxonomy is not exhaustive and the literature selected covers a spectrum of taxonomies, paradigms and methods of continual learning that simplify methods which are related in practice. The discussion focuses on one aspect of general intelligence (continual skill acquisition), but not on what it is and how it's going to be done.

**Conclusion**

The study compared and classified the core learning approaches that are relevant to AGI and explored the sequential learning problem and catastrophic forgetting as a problem common across all the approaches. It formalized the class-incremental scenario and studied regularization, rehearsal and joint training, concluding that rehearsal is the most effective of the simple mechanisms studied and that any small explicit memory is best suited for the hardest scenario, regularization is least effective. This is consistent with the published evidence [15]–[18] and with complementary-learning-systems theory [19]. The general conclusion is that there is no single method of learning that can be used to achieve general intelligence. The most defendable route is large-scale learning, efficient memory and rapid adaptation, and making the memory scalable and selective is a fundamental open problem.

**Future Scope**

Empirically validate the analytical expectations from the Theoretical Analysis on deep networks, larger sets of data, and longer sequences of tasks to see if they are true at scale.

Include additional techniques (generative replay, dynamic architectures, parameter isolation, and combinations of meta-learning and continual learning).

**Efficient replay:** buffer size, sampling strategies and memory compression.

Apply to reinforcement learning problems, and include forward and backward transfer as measures in addition to accuracy and forgetting.

Investigate the possibility of rehearsal and regularisation with very large pre-trained models.

**References**


1. S. Legg and M. Hutter, "Universal intelligence: A definition of machine intelligence," Minds and Machines, vol. 17, no. 4, pp. 391-444, 2007.
2. B. Goertzel, "Artificial general intelligence: Concept, state of the art, and future prospects," Journal of Artificial General Intelligence, vol. 5, no. 1, pp. 1-48, 2014.
3. F. Chollet, "On the measure of intelligence," arXiv:1911.01547, 2019.
4. B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," Behavioral and Brain Sciences, vol. 40, e253, 2017.
5. V. Mnih et al., "Human-level control through deep reinforcement learning", Nature, vol. 518, pp. 529-533, 2015.
6. D. Silver et al., "Mastering the game of Go without human knowledge," Nature, vol. 550, pp. 354-359, 2017, doi: 10.1038/nature24270.
7. D. Silver, S. Singh, D. Precup, and R. S. Sutton, "Reward is enough," Artificial Intelligence, vol. 299, art. 103535, 2021.
8. T. Brown et al., "Language models are few-shot learners," Advances in Neural Information Processing Systems, vol. 33, pp. 1877-1901, 2020.
9. L. Ouyang et al., "Training language models to follow instructions with human feedback," in Advances in Neural Information Processing Systems, vol. 35, pp. 27730-27744, 2022.
10. S. Reed et al., "A generalist agent," Transactions on Machine Learning Research, 2022, arXiv: 2205.06175.
11. S. Bubeck et al., "Sparks of artificial general intelligence: Early experiments with GPT-4," arXiv, 2023.
12. C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in Proc. 34th Int. Conf. Machine Learning (ICML), vol. 70 of PMLR, pp. 1126-1135, 2017.
13. T. Hospedales, A. Antoniou, P. Micaelli and A. Storkey, "Meta-learning in neural networks: A survey," IEEE Trans. on Pattern Analysis and Machine Intelligence, vol. 44, no. 9, pp. 5149-5169, 2022.
14. M. McCloskey, and N. J. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," Psychology of Learning and Motivation, vol. 24, pp. 109-165, 1989.
15. J. Kirkpatrick et al., "Overcoming catastrophic forgetting in neural networks," Proc. National Academy of Sciences, vol. 114, no. 13, pp. 3521-3526, 2017.
16. G. I. Parisi, R. Kemker, J. L. Part, C. Kanan and S. Wermter, "Continual lifelong learning with neural networks: A review," Neural Networks, vol. 113, pp. 54-71, 2019.
17. D. Rolnick, A. Ahuja, J. Schwarz, T. Lillicrap and G. Wayne, "Experience replay for continual learning", Advances in Neural Information Processing Systems, vol. 32, 2019.
18. G. M. van de Ven and A. S. Tolias, "Three scenarios for continual learning," arXiv:1904.07734, 2019.
19. J. L. McClelland, B. L. McNaughton and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex," Psychological Review, vol. 102, no. 3, pp. 419-457, 1995.
20. Y. LeCun (2022), “A path towards autonomous machine intelligence”, OpenReview v.0.9.2.
21. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
22. S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.