

Adaptive Knowledge Consolidation for Sequential Learning in Deep Neural Networks

Bidura V M¹, Vasanthakumar S²

¹Department of Computer Science, Karpagam Academy of Higher Education, Coimbatore, Tamil Nadu, India
Email: bidura.vm[at]gmail.com

²Department of Computer Science, Karpagam Academy of Higher Education, Coimbatore, Tamil Nadu, India
Email: vasanth.chml[at]gmail.com

Abstract: *Catastrophic forgetting is a major challenge in sequential deep learning, where learning new tasks can degrade previously acquired knowledge, particularly in evolving healthcare applications. Existing approaches mainly rely on replay, regularization, or architectural strategies, while the integration of adaptive, biologically inspired sleep mechanisms in conventional deep networks remains limited. To address this gap, this paper proposes a conceptual Adaptive Knowledge Consolidation (AKC) framework integrating wake-phase learning, NREM-inspired replay, REM-inspired knowledge integration, synaptic homeostasis, and adaptive sleep scheduling. The framework is designed for sequential image classification, with planned validation using Alzheimer's disease-related brain MRI data. Evaluation will consider accuracy, forgetting measure, backward and forward transfer, sensitivity, specificity, F1-score, and AUROC. As implementation is ongoing, no numerical results are claimed. The study establishes the framework, mathematical formulation, and experimental protocol for future empirical validation.*

Keywords: continual learning; catastrophic forgetting; knowledge consolidation; sleep-inspired learning; deep neural networks; replay; medical imaging; Alzheimer's disease

1. Introduction

Deep neural networks have demonstrated strong performance in image recognition and other data-driven tasks. However, conventional training generally assumes that representative training data are available together and that the underlying data distribution remains relatively stable. In a sequential learning environment, a model instead receives tasks or data distributions over time. Updating the model for a new task can modify shared parameters and representations that were useful for earlier tasks, resulting in catastrophic forgetting [1]. Continual learning therefore seeks to balance plasticity, which allows a model to acquire new information, with stability, which protects previously learned knowledge [2].

Existing continual-learning approaches can broadly be classified into regularization-based, replay-based, architecture-based, and hybrid approaches [2]. Elastic Weight Consolidation (EWC) introduces a parameter-importance penalty that reduces changes to parameters considered important for previous tasks [1]. Synaptic Intelligence (SI) estimates parameter importance from the learning trajectory and uses this information to reduce interference during subsequent learning [3]. Learning without Forgetting (LwF) uses knowledge distillation from a previously trained model to preserve earlier capabilities when original training data are unavailable [4]. Gradient Episodic Memory (GEM) uses episodic examples and gradient constraints to reduce harmful updates to previously learned tasks [5]. These methods demonstrate that retention can be improved, but they involve trade-offs related to memory requirements, computational cost, optimization constraints, and plasticity [2].

Biological memory systems provide another perspective for addressing continual learning. Complementary Learning Systems theory describes interactions between hippocampal and neocortical learning systems, in which recent experiences

can be rapidly represented and gradually integrated into longer-term knowledge [6]. Sleep has also been investigated as a computationally relevant period for replay and memory stabilization. González *et al.* showed that computational models incorporating sleep-like mechanisms can protect memories from catastrophic interference [7]. Tadros *et al.* subsequently demonstrated that sleep-like unsupervised replay can recover previously learned information after sequential learning [8]. More recently, Sorrenti *et al.* proposed Wake-Sleep Consolidated Learning (WSCL), which organizes continual visual learning into wake, NREM, and REM stages [9].

The present work takes WSCL and the broader sleep-consolidation literature as the starting point but focuses on an adaptive knowledge-consolidation direction for sequential deep learning. The proposed framework retains the wake/NREM/REM organization while introducing explicit performance monitoring, adaptive consolidation scheduling, and synaptic-homeostasis control. The intended application is evolving medical-image classification, where retaining information from previous cohorts while incorporating new information represents an important methodological requirement.

The contribution of this paper is therefore a conceptual and experimentally testable framework, rather than a claim of completed clinical or empirical performance.




Figure 1: Sequential Learning and Catastrophic Forgetting

2. Literature Survey

2.1 Continual Learning and Catastrophic Forgetting

Continual learning has been investigated under task-incremental, class-incremental, and domain-incremental settings. Reviews emphasize that the learning scenario, task boundaries, and evaluation protocol should be explicitly defined because these factors can substantially affect reported performance [2]. A recent systematic review also emphasizes that continual-learning evaluation should consider both knowledge retention and adaptation rather than relying solely on final-task accuracy [11].

Catastrophic forgetting occurs when optimization on later tasks causes substantial degradation in performance on previously learned tasks [1], [2]. Consequently, continual-learning evaluation requires measures that explicitly quantify the effect of later learning on earlier tasks.

2.2 Regularization and Parameter Protection

EWC introduces a parameter-specific regularization term based on the estimated importance of parameters for previously learned tasks [1]. Parameters with greater estimated importance are penalized more strongly when modified during subsequent learning.

SI follows a related principle but estimates parameter importance online from the contribution of parameters during learning [3]. Such methods can reduce the requirement for storing large amounts of previous training data. However, excessively strong parameter protection can restrict plasticity when new tasks require modification of shared representations [2].

2.3 Distillation and Episodic Replay

LwF uses the outputs of a previously trained network as a reference during learning of a new task [4]. This allows previously acquired behaviour to be retained without necessarily storing the complete original training set.

GEM uses episodic memory and gradient constraints to prevent updates that would increase the loss on stored examples from previous tasks [5]. Replay-based methods directly expose the model to earlier information, but memory capacity, sampling strategy, computational overhead, and data-governance requirements become important considerations [2].

These issues are especially relevant to medical imaging because patient-level data cannot simply be retained or redistributed without considering the applicable data-use and privacy requirements.

2.4 Sleep-Inspired Learning

Sleep-inspired approaches attempt to translate biological memory-consolidation principles into computational learning mechanisms. Complementary Learning Systems theory provides a theoretical basis for separating rapid learning from slower consolidation [6].

González *et al.* investigated whether computational sleep mechanisms could protect memories from catastrophic interference [7]. Tadros *et al.* demonstrated that sleep-like unsupervised replay can reduce catastrophic forgetting in artificial neural networks [8].

WSCL subsequently organized continual visual learning into wake, NREM, and REM phases [9]. The wake phase learns from current information, while NREM-style processing supports replay and consolidation. The REM-inspired stage provides an additional mechanism for representation integration and transfer [9]. These studies motivate the use of sleep-inspired computational phases, while also suggesting the need for careful experimental analysis of memory selection, consolidation timing, auxiliary information, and stability-plasticity trade-offs.

2.5 Research Gap

The existing literature establishes the usefulness of regularization, replay, and sleep-inspired consolidation, but several methodological issues remain open [1]. First, continual-learning systems require evaluation protocols that distinguish retention from adaptation and transfer [10]. Second, replay and parameter-protection methods introduce trade-offs involving memory, computational cost, and plasticity [2]. Third, WSCL demonstrates a structured wake–NREM–REM strategy, but adaptive control of when consolidation should occur remains an important design consideration [9]. Fourth, many continual-learning demonstrations rely on standard computer-vision benchmarks, whereas application to evolving medical-imaging distributions requires subject-level data separation, domain-shift analysis, and medical-classification metrics.

Therefore, this paper proposes an adaptive framework in which consolidation is activated according to measurable learning behaviour rather than being performed only at a fixed interval.

Table 1: Representative Continual-Learning Strategies

Method family	Representative approach	Main mechanism	Key consideration
Regularization	EWC [1]	Protect important parameters	Excessive protection may reduce plasticity
Synaptic importance	SI [3]	Estimate parameter importance during learning	Importance estimation adds computation
Distillation	LwF [4]	Preserve previous model outputs	Requires teacher/reference outputs
Episodic replay	GEM [5]	Store examples and constrain updates	Memory and optimization overhead
Sleep-inspired replay	Tadros <i>et al.</i> [8]	Sleep-like offline replay	Replay and consolidation design
Wake-sleep consolidation	WSCL [9]	Wake + NREM + REM processing	Memory and auxiliary-data design




Figure 2: Conceptual Mapping from Consolidation to computational

3. Proposed Adaptive Knowledge Consolidation Framework

The proposed framework is referred to as **Adaptive Knowledge Consolidation (AKC)**. It is deliberately presented as a conceptual framework because its implementation and empirical validation constitute the next stage of the research.

The framework consists of five principal components:

- 1) Wake-phase learning.
- 2) Performance monitoring and adaptive sleep scheduling.
- 3) NREM-inspired replay.
- 4) REM-inspired knowledge integration.
- 5) Synaptic homeostasis.

The overall process is:

Sequential Task → Wake Learning → Performance Monitoring → Adaptive Sleep Decision → NREM Replay → REM Integration → Synaptic Homeostasis → Updated Model → Evaluation

If the monitored performance indicates that previously learned knowledge is being degraded beyond a predefined threshold, consolidation is activated. Otherwise, the model continues with sequential learning.

3.1 Problem Formulation

Let the sequential learning stream be

$$\mathcal{D} = \{D_1, D_2, \dots, D_T\},$$

where D_t denotes the dataset associated with task t , and T is the total number of tasks.

Each task is represented as

$$D_t = \{(x_i, y_i)\}_{i=1}^{N_t}, \dots \dots \dots (1)$$

where x_i represents an input image, y_i represents its corresponding label, and N_t denotes the number of samples in task t .

Let the neural network be represented by

$$f(x; \theta),$$

where θ represents the trainable parameters.

The objective of continual learning is to acquire knowledge from the current task while maintaining acceptable performance on previously learned tasks.

3.2 Wake-Phase Learning

During the wake phase, the network learns from the current task D_t . For a multi-class classification problem, the cross-entropy loss is

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{ic} \log p_{ic}, \dots \dots \dots (2)$$

where C is the number of classes, y_{ic} is the ground-truth indicator for class c , and p_{ic} is the predicted probability.

The model parameters are updated according to

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}_{CE}, \dots \dots \dots (3)$$

where η is the learning rate.

A bounded short-term memory M_S stores selected recent experiences:

$$M_S = \{(x_i, y_i)\}_{i=1}^{K_S}. \dots \dots \dots (4)$$

A long-term memory M_L stores selected information from earlier tasks:

$$M_L = \bigcup_{j < t} M_j, \quad |M_L| \leq K_L. \dots \dots \dots (5)$$

The memory capacities K_S and K_L are bounded to maintain computational feasibility.

3.3 Performance Monitoring

After learning the current task, the model is evaluated on previously observed tasks.

Let

$$A_{i,j}$$

denote the accuracy on task j after the model has completed learning through task i .

The resulting accuracy matrix is

$$A = \begin{bmatrix} A_{1,1} & - & - & \dots & - \\ A_{2,1} & A_{2,2} & - & \dots & - \\ A_{3,1} & A_{3,2} & A_{3,3} & \dots & - \\ \vdots & \vdots & \vdots & \ddots & - \\ A_{T,1} & A_{T,2} & A_{T,3} & \dots & A_{T,T} \end{bmatrix}.$$

The diagonal represents performance immediately after learning each task, while later entries for earlier tasks indicate the effect of subsequent learning.

Forgetting for task j after the final task can be calculated as

$$F_j = \max_{i \in \{j, \dots, T-1\}} A_{i,j} - A_{T,j}. \dots \dots \dots (6)$$

The final average forgetting measure is

$$FM = \frac{1}{T-1} \sum_{j=1}^{T-1} F_j. \dots \dots \dots (7)$$

A larger value of FM represents greater degradation of previously acquired performance.

3.4 Adaptive Sleep Scheduling

Unlike a fixed sleep interval, AKC uses measured learning behaviour to determine when consolidation should occur.

Let FM_t represent the current forgetting estimate and let ΔA_{old} represent the change in average accuracy over previously learned tasks.

A simple trigger is defined as

$$S_t = \begin{cases} 1, & FM_t > \tau_F \text{ or } \Delta A_{old} < -\tau_A, \\ 0, & \text{otherwise,} \end{cases}$$

where

- $S_t = 1$ indicates that consolidation is activated,
- τ_F is the forgetting threshold, and
- τ_A is the allowable decrease in previous-task accuracy.

The thresholds must be determined using validation data or a predefined calibration procedure. The final test set must not be used to tune the sleep trigger.

3.5 NREM-Inspired Replay

When consolidation is activated, the NREM phase replays information stored in short-term and long-term memory.

A generic replay objective is

$$\mathcal{L}_{NREM} = \mathbb{E}_{(x,y) \sim M_S} [\mathcal{L}(y, f(x; \theta))] + \alpha \mathbb{E}_{(x,y) \sim M_L} [\mathcal{L}(y, f(x; \theta))], \quad \dots(8)$$

where α controls the contribution of long-term memories.

Knowledge distillation can additionally be introduced:

$$\mathcal{L}_{KD} = T^2 D_{KL} (p_{\theta_{old}}^{(T)} \| p_{\theta}^{(T)}), \quad \dots(9)$$

where T is the distillation temperature and D_{KL} represents the Kullback–Leibler divergence.

The NREM objective can therefore be expressed as

$$\mathcal{L}_{NREM}^* = \mathcal{L}_{NREM} + \lambda_{KD} \mathcal{L}_{KD}. \quad \dots(10)$$

3.6 REM-Inspired Knowledge Integration

The REM-inspired phase is intended to support representation integration and future transfer.

Let D_{aux} represent an auxiliary or “dream” dataset that is separate from the current task and test data.

The auxiliary learning objective is

$$\mathcal{L}_{REM} = \mathbb{E}_{(x,y) \sim D_{aux}} [\mathcal{L}(y, f(x; \theta))]. \quad \dots(11)$$

If unlabelled auxiliary data are used, a consistency objective can be considered:

$$\mathcal{L}_{cons} = \|f(x; \theta) - f(\tilde{x}; \theta)\|^2, \quad \dots(12)$$

where x and $\tilde{x}$ represent two augmented versions of the same auxiliary sample.

The auxiliary dataset must be selected carefully because a substantial mismatch between the auxiliary and target domains may produce negative transfer.

3.7 Synaptic Homeostasis

Synaptic homeostasis is treated as a computational control mechanism rather than an exact biological simulation.

Let θ_k^{old} represent the parameter value after the previous consolidation stage and let Ω_k represent the importance of parameter k .

A stability loss can be defined as

$$\mathcal{L}_{stab} = \sum_k \Omega_k (\theta_k - \theta_k^{old})^2. \quad \dots(13)$$

The parameter importance may be estimated using a Fisher-information-based approach or another continual-learning importance estimator.

The overall objective becomes

$$\mathcal{L}_{AKC} = \mathcal{L}_{new} + \lambda_{NREM} \mathcal{L}_{NREM}^* + \lambda_{REM} \mathcal{L}_{REM} + \lambda_{stab} \mathcal{L}_{stab}. \quad \dots(14)$$

The coefficients λ_{NREM} , λ_{REM} , and λ_{stab} must be determined experimentally using the validation protocol.




Figure 3: Proposed adaptive knowledge-consolidation framework

Fig. 3. Proposed adaptive knowledge-consolidation architecture. Wake learning is followed by performance monitoring and, when triggered, NREM replay, REM integration, and synaptic homeostasis.

Table 2: Functional specification of the proposed framework

Component	Input	Function	Expected effect
Wake learning	Current task data	Acquire new representations	High plasticity
Performance monitor	Task-wise validation accuracy	Estimate retention/forgetting	Trigger information
NREM replay	Short-/long-term memory	Reinforce prior knowledge	Lower forgetting
REM integration	Auxiliary/dream data	Promote broader representations	Potentially improve transfer
Synaptic homeostasis	Parameter/importance signals	Control excessive change	Stability–plasticity balance
Adaptive scheduler	Forgetting/performance signals	Choose consolidation time	Avoid unnecessary sleep cycles

3.8 Algorithmic Description

Algorithm 1. Adaptive Knowledge Consolidation

Input: Sequential datasets $D_1, \dots, D_T$, neural network $f(\cdot; \theta)$, short-term memory M_S , long-term memory M_L , thresholds τ_F, τ_A .

Output: Final model θ_T , accuracy matrix A .

- 1) Initialize model parameters θ , M_S , and M_L .
- 2) For $t = 1, \dots, T$:
 - Train the model on current task D_t .
 - Store selected current-task samples in M_S .
- 3) Update M_L using a bounded memory-selection strategy.
- 4) Evaluate the model on all previously observed tasks.
- 5) Update the accuracy matrix A .
- 6) Calculate the current forgetting measure FM_t .
- 7) If $FM_t > \tau_F$ or $\Delta A_{old} < -\tau_A$:
 - Perform NREM replay using M_S and M_L .
 - Perform REM-inspired integration using D_{aux} .
 - Apply synaptic-homeostasis regularization.
 - Update model parameters.
- 8) Evaluate the consolidated model.
- 9) Calculate AA, FM, BWT, FWT, sensitivity, specificity, F1-score, and AUROC.
- 10) Repeat experiments using multiple random seeds.
- 11) Report mean and standard deviation across runs.

4. Experimental Methodology and Evaluation Plan

The experimental study should be conducted in stages so that the contribution of each component can be independently evaluated. First, a strong baseline and the published WSCL configuration should be reproduced as closely as practical. Second, adaptive scheduling and synaptic homeostasis should be introduced individually. Third, the complete framework should be evaluated on standard continual-learning benchmarks and subsequently on medical-imaging data.

This staged design avoids attributing improvements to the complete framework without determining which individual component contributes to the observed change.

4.1 Datasets and Task Construction

4.1.1 CIFAR-10

CIFAR-10 contains 60,000 color images belonging to 10 classes, including 50,000 training images and 10,000 test images [13].

For the proposed class-incremental experiment, the ten classes can be divided into five sequential tasks:

$$D_1 = \{C_1, C_2\}, D_2 = \{C_3, C_4\}, \dots, D_5 = \{C_9, C_{10}\}.$$

The same class order must be used for every compared method.

Recommended preprocessing includes normalization using training-set statistics, random horizontal flipping and random cropping during training, with deterministic preprocessing during validation and testing.

4.1.2 Tiny-ImageNet

Tiny-ImageNet provides a larger and more challenging visual benchmark than CIFAR-10. A possible class-incremental protocol is to divide its 200 classes into ten sequential tasks containing 20 classes each.

The same task order and evaluation protocol must be applied across all compared methods.

4.1.3 Alzheimer's Disease Neuroimaging Initiative

The Alzheimer's Disease Neuroimaging Initiative (ADNI) provides longitudinal and multicentre research data for Alzheimer's disease and related research. Access to ADNI data is subject to applicable data-use requirements [14].

Potential diagnostic categories include:

- Cognitively normal (CN)
- Mild cognitive impairment (MCI)
- Alzheimer's disease (AD)

However, the continual-learning sequence should not be created merely by arbitrarily separating these diagnostic labels. The task sequence should represent a meaningful distributional or cohort change, such as acquisition period, scanner/site group, or longitudinal visit, depending on the selected ADNI subset.

All images from a participant must remain within the same train, validation, or test partition to avoid subject leakage.

The medical-imaging experiment should therefore satisfy the following requirements:

- No participant overlap between training, validation, and testing.
- Test images must never be used for replay.
- Test data must not be used to tune sleep thresholds.
- MRI preprocessing must be documented and consistently applied.
- Scanner/site and acquisition information should be recorded when relevant.
- Class imbalance should be reported.
- ADNI data-use requirements must be followed.
- Results should be described as research classification results rather than clinical diagnosis.

Table 2: Proposed Dataset and Task Construction

Dataset	Modality	Classes	Proposed task sequence	Main purpose
CIFAR-10 [13]	RGB images	10	5 tasks $\times$ 2 classes	Initial class-incremental benchmark
Tiny-ImageNet	RGB images	200	10 tasks $\times$ 20 classes	Larger-scale visual benchmark
ADNI [14]	Structural MRI	CN/MCI/AD	Cohort/time/site-based sequence	Medical-imaging domain-shift evaluation

4.2 Metrics for Catastrophic Forgetting

Let $A_{i,j}$ represent the accuracy on task j after learning through task i .

Average Accuracy

Final average accuracy is

$$AA = \frac{1}{T} \sum_{j=1}^T A_{T,j}. \quad \dots \dots \dots (15)$$

Higher AA indicates stronger overall final performance.

Forgetting Measure

For task j ,

$$F_j = \max_{i \in \{j, \dots, T-1\}} A_{i,j} - A_{T,j}. \quad \dots \dots \dots (16)$$

The final forgetting measure is

$$FM = \frac{1}{T-1} \sum_{j=1}^{T-1} F_j. \quad \dots \dots \dots (17)$$

Lower FM indicates better retention.

Backward Transfer

Backward transfer is defined as

$$BWT = \frac{1}{T-1} \sum_{j=1}^{T-1} (A_{T,j} - A_{j,j}). \quad \dots \dots \dots (17)$$

Negative BWT indicates degradation of earlier-task performance, while positive BWT indicates improvement.

Forward Transfer

Forward transfer can be expressed as

$$FWT = \frac{1}{T-1} \sum_{j=2}^T (\tilde{A}_j - b_j), \quad \dots \dots \dots (18)$$

where $\tilde{A}_j$ represents performance on task j before direct training on that task and b_j represents the corresponding baseline performance of a model without prior-task learning.

4.3 Medical-Image Evaluation

For the medical-imaging stage, accuracy alone should not be used as the sole evaluation measure. Sensitivity, specificity, F1-score, and AUROC should be reported together with continual-learning metrics.

Sensitivity

$$Sensitivity = \frac{TP}{TP+FN}. \quad \dots \dots \dots (19)$$

Specificity

$$Specificity = \frac{TN}{TN+FP}. \quad \dots \dots \dots (20)$$

F1-Score

$$F1 = 2 \frac{Precision \times Recall}{Precision + Recall}. \quad \dots \dots \dots (21)$$

For multi-class classification, macro-F1 should be reported.

AUROC

For multi-class classification, AUROC should be calculated using a one-vs-rest strategy and macro averaging.

Repeated experimental runs should be reported using mean $\pm$ standard deviation, and confidence intervals should be provided where appropriate.

Table 3: Evaluation Metrics

Metric	Purpose	Desired interpretation
Average Accuracy (AA)	Overall final performance	Higher
Forgetting Measure (FM)	Quantify degradation of earlier tasks	Lower
Backward Transfer (BWT)	Measure effect of later learning on earlier tasks	Less negative/ more positive
Forward Transfer (FWT)	Measure benefit of previous knowledge on future tasks	Higher
Sensitivity	Detect positive medical cases	Higher
Specificity	Correctly identify negative cases	Higher
F1-score	Balance precision and recall	Higher
AUROC	Measure discrimination across thresholds	Higher

4.4 Experimental Configurations

Table 4: Staged Experimental Plan

Phase	Purpose	Model/Data	Main comparison
Phase I	Baseline reproduction	ResNet-18 + standard benchmark	Naive sequential learning vs established CL baseline
Phase II	Consolidation analysis	Same backbone/task sequence	Wake only vs Wake+NREM vs Wake+REM vs combined
Phase III	Adaptive framework	Standard benchmark + medical imaging	Fixed scheduling vs adaptive scheduling
Phase IV	Component ablation	Complete framework	Remove one component at a time

All experiments are planned experiments and must not be interpreted as completed results.

5. Findings, Results and Discussion

5.1 Findings from Existing Literature

The existing literature provides several observations that motivate the proposed framework. First, continual-learning evaluation requires consideration of both retention and adaptation rather than current-task accuracy alone [2]. Second, regularization and replay mechanisms can reduce interference but introduce memory, computational, or plasticity trade-offs [1]. Third, computational sleep mechanisms provide evidence that replay and consolidation can help preserve previously learned information [7]. Finally, WSCL demonstrates a structured wake–NREM–REM approach to continual visual learning [9].

These published observations provide the motivation for investigating adaptive knowledge consolidation.

5.2 Published Results as Baseline Context

The results reported for WSCL belong to the published WSCL study and should not be interpreted as results of the proposed AKC framework [9]. The proposed paper does not reproduce those numerical results as if they were obtained from the present framework.

5.3 Expected Experimental Findings

The proposed framework has not yet been experimentally validated. Therefore, numerical results are intentionally not reported.

The following are hypotheses to be tested:

- 1) NREM replay may improve retention of earlier tasks.
- 2) REM-inspired auxiliary learning may influence forward transfer and representation generalization.
- 3) Adaptive scheduling may reduce unnecessary consolidation when performance remains stable.
- 4) Adaptive scheduling may activate consolidation when measurable forgetting increases.
- 5) Synaptic homeostasis may regulate the stability–plasticity trade-off.

- 6) The complete framework may involve additional computational cost because consolidation introduces additional processing stages.

These statements are hypotheses rather than established findings.

5.4 Ablation Strategy

Ablation is necessary because the proposed framework contains several interacting components. The following configurations should be compared:

- Wake only**
Measures forgetting without consolidation.
- Wake + NREM**
Tests whether replay improves retention.
- Wake + REM**
Tests whether auxiliary information affects transfer.
- Wake + NREM + REM**
Tests the combined effect of the two sleep-inspired stages.
- Full framework without adaptive scheduling**
Tests whether adaptive timing provides an additional effect.
- Complete AKC framework**
Tests the complete proposed configuration.

The comparison should jointly consider AA, FM, BWT, FWT, and computational cost.

Table 5: Ablation and Reporting Structure

Configuration	Research question	Main metrics
Wake only	How much forgetting occurs without consolidation?	AA, FM, BWT, FWT
Wake + NREM	Does replay improve retention?	FM, BWT
Wake + REM	Does auxiliary learning affect transfer?	FWT, AA
Wake + NREM + REM	Do the two stages complement each other?	AA, FM, FWT
Full without adaptive scheduling	Does scheduling influence performance?	AA, FM, computation
Complete AKC	What is the combined effect?	All metrics

6. Discussion

The proposed framework addresses the stability–plasticity trade-off in continual learning, where models must acquire new information while preserving previously learned representations [2]. Sleep-inspired consolidation separates rapid learning from subsequent memory processing, with previous studies showing that replay and sleep-like mechanisms can reduce forgetting [7]. WSCL further demonstrates the use of wake, NREM, and REM stages for continual visual learning [9].

The main methodological extension is adaptive consolidation scheduling. Rather than using fixed intervals, the framework monitors task-wise performance and triggers consolidation when forgetting or instability exceeds predefined thresholds. Thresholds should be selected using validation data or a predefined calibration procedure to avoid test-set bias.

For medical imaging, evaluation should jointly consider knowledge retention, adaptation to new cohorts, transfer, and

classification performance. ADNI offers a longitudinal resource for Alzheimer’s research, but the dataset subset and task sequence must be clearly defined, with participant-level separation to prevent data leakage [14].

The framework has several limitations. Synaptic homeostasis is a computational abstraction, REM performance depends on appropriate auxiliary information, and replay requires additional storage and computation. Most importantly, the framework has not yet been empirically validated. Alzheimer’s MRI classification should therefore be considered a research task rather than clinical diagnosis without appropriate clinical validation.

7. Conclusion

This paper presented Adaptive Knowledge Consolidation (AKC), a conceptual framework for sequential deep learning inspired by biological memory consolidation. The framework combines wake-phase learning, NREM-inspired replay, REM-inspired knowledge integration, synaptic homeostasis, and adaptive sleep scheduling to address the stability–plasticity challenge associated with catastrophic forgetting.

The framework provides a mathematical formulation, an adaptive consolidation mechanism, an algorithmic procedure, and a structured experimental protocol. Evaluation is based on task-wise accuracy, average accuracy, forgetting measure, backward transfer, forward transfer, sensitivity, specificity, F1-score, and AUROC.

The planned evaluation begins with standard continual-learning benchmarks and progresses toward medical-imaging experiments using appropriately constructed Alzheimer’s disease-related MRI data. The proposed framework is currently conceptual and has not yet been empirically validated. Consequently, no numerical performance claims are made in this paper. Future work will implement the framework, conduct controlled ablation studies, compare it against established continual-learning baselines, and evaluate its behaviour on sequential medical-imaging data.

References

- [1] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, J. Quan, M. Ramalho, A. Grabska-Barwinska, C. K. Grefenstette, K. Teh, A. D. T. B. H. H. et al., “Overcoming catastrophic forgetting in neural networks,” *Proc. Natl. Acad. Sci. USA*, vol. 114, no. 13, pp. 3521–3526, Mar. 2017, doi: 10.1073/pnas.1611835114.
- [2] 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, May 2019, doi: 10.1016/j.neunet.2019.01.012.
- [3] F. Zenke, B. Poole, and S. Ganguli, “Continual learning through synaptic intelligence,” in *Proc. 34th Int. Conf. Mach. Learn. (ICML)*, vol. 70, 2017, pp. 3987–3995.
- [4] Z. Li and D. Hoiem, “Learning without forgetting,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham, Switzerland: Springer, 2016, pp. 614–629, doi: 10.1007/978-3-319-46493-0_37.

- [5] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 6467–6476.
- [6] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory," *Psychological Review*, vol. 102, no. 3, pp. 419–457, Jul. 1995, doi: 10.1037/0033-295X.102.3.419.
- [7] O. C. González, Y. Sokolov, G. P. Krishnan, J. E. Delanois, and M. Bazhenov, "Can sleep protect memories from catastrophic forgetting?" *eLife*, vol. 9, Art. no. e51005, 2020, doi: 10.7554/eLife.51005.
- [8] T. Tadros, G. P. Krishnan, R. Ramyaa, and M. Bazhenov, "Sleep-like unsupervised replay reduces catastrophic forgetting in artificial neural networks," *Nature Communications*, vol. 13, Art. no. 7742, 2022, doi: 10.1038/s41467-022-34938-7.
- [9] A. Sorrenti, G. Bellitto, F. Proietto Salanitri, M. Pennisi, S. Palazzo, and C. Spampinato, "Wake-sleep consolidated learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 36, no. 7, pp. 12668–12679, Jul. 2025, doi: 10.1109/TNNLS.2024.3458440.
- [10] Y.-C. Hsu, Y.-C. Liu, A. Ramasamy, and Z. Kira, "Re-evaluating continual learning scenarios: A categorization and case for strong baselines," arXiv:1810.12488, 2018.
- [11] Q. Yang, L. Wang, J. Wicker, and G. Dobbie, "Continual learning: A systematic literature review," *Neural Networks*, vol. 195, Art. no. 108226, 2026, doi: 10.1016/j.neunet.2025.108226.
- [12] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. G. Slabaugh, and T. Tuytelaars, "A continual learning survey: Defying forgetting in classification tasks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 7, pp. 3366–3385, Jul. 2022, doi: 10.1109/TPAMI.2021.3057446.
- [13] A. Krizhevsky, "Learning multiple layers of features from tiny images," University of Toronto, Toronto, ON, Canada, Tech. Rep., 2009.
- [14] Alzheimer's Disease Neuroimaging Initiative (ADNI), "Alzheimer's Disease Neuroimaging Initiative," ADNI Data Resource, accessed Sep. 2026.