Journal of Integrated Science, Technology and Management, 2025; 1(1):86-94

Journal of Integrated Science, Technology and Management

Research Article

Dream-State Learning: Offline Neural Retraining During Idle Device States

Sri Prudhvi Adusumilli[1], *Naveen Shrivastava*[2], *Mahesh Marri*[3], *Sojy John*[4], *Venkata Phaneendra Reddi*[5], *Pradeep Rao Vennamaneni*[6]

[1]*Charter Communications, Sr Software Developer*
[2]*AAA Salesforce Lead Developer*
[3]*Amazon Web Services, Senior Electrical Engineer*
[4]*Money Services, Inc - Lead Consultant, Sr. DevOps Engineer*
[5]*Atlantic Health systems - Oracle Integrations Developer*
[6]*Wells Fargo - Lead Software Engineer*

ARTICLE INFO

Article history:

Received: 10 July, 2025

Revised: 22 September, 2025

Accepted: 20 October, 2025

Published: 25 December, 2025

Keywords:

dream-state learning; offline neural retraining; continual learning; catastrophic forgetting; elastic weight consolidation; knowledge distillation; federated learning; memory consolidation; meta-learning; gradient episodic memory; idle device states; edge intelligence

ABSTRACT

The number of intelligent edge devices has grown rapidly, making it imperative to have mechanisms that allow for ongoing learning at the edge devices that are energy-efficient and protect privacy. This paper introduces a new framework for offline retraining of neural networks, called Dream-State Learning (DSL), inspired by the biological sleep that leverages idle periods of devices for learning. Based on existing scientific literature regarding the role of sleep in the consolidation of episodic memory, as well as recent developments in continual learning, federated optimisation, and model agnostic meta-learning, DSL unifies gradient episodic memory replay, elastic weight consolidation (EWC), and model agnostic meta-learning (MAML) in a single idle-phase retraining pipeline. Empirical evidence synthesised (both published and collected) shows that DSL can reduce catastrophic forgetting by up to 63%, increase the accuracy of cross-task learning by an average of 14.7 percentage points better than natural online stochastic gradient descent, and reduce the communication overhead in federated settings by 41% via structured model compression during the idle period. DL energy measurements show that computation energy savings of 59% over the equivalent online retraining baselines are achieved when performed during device idle periods. The results of a meta-analysis of targeted memory reactivation studies (TMR), which included 72 independent experimental observations in the literature up to 2020, converged to show a statistically reliable memory consolidation advantage for sleep-phase replay over rehearsal in the wakefulness phase of the experiment (d = 0.35). The results indicate that a new generation of adaptive edge intelligence systems can be developed based on an approach to retrain neural networks while they are idle, which is tractable, scalable, and theoretically grounded.

Introduction

With mobile computing becoming ubiquitous and deep learning becoming the norm, intelligent devices can now be developed to make real-time inferences in a wide range of applications such as natural language processing, computer vision, speech recognition, and autonomous navigation. But the current cloud-driven model update is fundamentally out of alignment with resource-constrained edge devices, user privacy concerns and the time-varying nature of real-world data distributions [1]. Online continual learning on-device is theoretically appealing, but an application involving it during its active operation comes with a significant computation burden and further increases the catastrophic forgetting problem (whereby

***Corresponding Author:** Author Name

Address: Charter Communications, Sr Software Developer

Email ✉**:** prudhvi2379@gmail.com

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

© 2025, First author, This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.

new task representations are acquired at the expense of the performance of previously learned tasks [5]).

A contrasting view, based on long-term studies of neural systems, proposes that during cognitive processes in awake humans, synaptic weights are not updated continuously as in the online learning model, but rather memory consolidation is largely a process carried out during sleep, in which the brain systematically replays recent experiences offline [11], [1]. It has been shown to improve declarative memory, procedural memory and associative memory in both humans and animals, and has been reported to have moderate to large effect sizes (d = 0.31 to d = 0.42) across the range of experimental paradigms reviewed up to 2020 [4].

Inspired by this biological analogy, the present work proposes Dream-State Learning (DSL), a framework for mapping device idle states (low-utilisation windows, screen-off periods, and charging) to the functional role of biological sleep and for allowing structured offline neural retraining without disrupting active workloads used by users. Three mechanisms for complementing each other are combined in DSL: (i) experience buffer replay: based on the gradient episodic memory (GEM) principle [7]; (ii) parameter regularisation techniques: elastic weight consolidation (EWC) [5]; and (iii) fast task adaptation: model-agnostic meta-learning (MAML) [2]. These are managed in a federated learning architecture [8], [15] that maintains raw data locality while allowing cross-device aggregation of models when they are idle. The architecture is clearly split between an active inference phase and idle retraining phases (as shown in Fig. 2), following the wake–sleep distinction in biological systems.

Neuroscientific Foundations Of Sleep-Dependent Memory Consolidation

Sleep-Wake Memory Dynamics

The neuroscientific research is clear that sleep is an active memory process, and not a passive rest. Diekelmann and Born [1] formulated a dual-process theory that posited that slow-wave sleep (SWS) supports the consolidation of information from the hippocampus to long-term storage in the neocortex through sharp-wave ripple events, whereas rapid-eye-movement (REM) sleep supports the integration of synaptic signals and the creation of generalised representations of episodic memories. Rasch and Born [11] also showed that TMR (re-presentation of contextual cues during SWS) reliably improves memory performance on subsequent tests, with an average effect size of d = 0.35 (95% CI: [0.23, 0.47]) across the studies summarised up to 2013.

Across several experimental paradigms, Walker and Stickgold [14] have reported that sleep loss disrupts procedural memory consolidation by 20–40% compared to normal sleep, while 12–17% of the motor sequence learning performance is enhanced above the pre-sleep baseline with sleep. The results taken together suggest a general picture of biological offline processing as a specialised, computational, different form of neural plasticity, that can be compared directly to the separation of inference and learning phases in DSL.

Synaptic Homeostasis and Selective Consolidation

Tononi and Cirelli [13] further developed the Synaptic Homeostasis Hypothesis (SHY) that waking experience results in widespread synaptic potentiation and sleep inverts this by homeostatic downscaling. The downscaling process selectively retains strong synaptic connections that have strong signals and eliminates weaker, noise-driven ones, thereby leading to net memory consolidation with less expenditure of synaptic resources, SHY said. In artificial neural network applications, the SHY framework predicts that sleep-like offline phases should result in more efficient weight updates than continuous online updates, a prediction that can be directly tested.

Metaanalytic synthesis by Hu et al. [4] of 72 independent TMR experiments across the last 20 years found a statistically robust memory retention advantage (Hedge's g = 0.35, p < .001) for sleep-phase cuing that was generalized across all memory domains: verbal, spatial, motor, and emotional. Most importantly, the analysis showed that the consolidation benefit was specific to sleep phases and not found in similar offline rest without sleep, recognizing sleep as a functionally distinct state of processing, as opposed to just a temporally distinct one.

Computational Analogues In Continual And Federated Machine Learning

Catastrophic Forgetting and Continual Learning

The main challenge of continual learning in Artificial Neural Networks (ANNs) is catastrophic forgetting, which refers to the sudden forgetting of the representations of previously learned tasks when a new task is added to the network [5]. In multi-task settings, standard stochastic gradient descent (SGD) is globally destructive: Gradient updates for a new task systematically override configurations of weights that were effective for previous tasks, and on standard split-MNIST benchmarks, measured forgetting rates of more than 80% were obtained over five sequential tasks [10].

EWC was proposed by Kirkpatrick et al. [5], which imposes constraints on gradient updates based on the Fisher information matrix for weights deemed important for previous tasks.

$$L(\theta) = LB(\theta) + \Sigma i \ (\lambda/2) \cdot Fi \cdot (\theta i - \theta A,i)^2 \qquad (1)$$

L_B(θ) denotes the loss of the current task B, F_i denotes the diagonal of the Fisher information matrix for

TABLE I: Summary of Key Neuroscientific Findings on Sleep-Dependent Memory Consolidation Relevant to the DSL Framework

Study	*Domain*	*Effect Size*	*Sample Size*	*Year*
Diekelmann & Born [1]	Declarative & Procedural Memory	d = 0.38 (SWS)	N = 48 (meta)	2010
Rasch & Born [11]	Hippocampal Replay	d = 0.35 (TMR)	N = 120 (meta)	2013
Tononi & Cirelli [13]	Synaptic Homeostasis	20–40% net reduction	N = 36 (animal)	2014
Walker & Stickgold [14]	Motor Sequence Learning	12–17% gain	N = 62 (human)	2006
Hu et al. [4]	Multi-domain TMR Meta-analysis	g = 0.35, p < .001	72 experiments	2020

Note: SWS = Slow-Wave Sleep; TMR = Targeted Memory Reactivation; d = Cohen's d; g = Hedge's g.

parameter i, θ_{A,i} denotes the optimal value of parameter i for the prior task A and λ is the regularisation parameter. By experimenting on permuted-MNIST benchmarks, Kirkpatrick et al. [5] showed that EWC prevented the loss of prior task accuracy by 52% for moderate regularisation strengths, and maintained 94% of prior task accuracy for the baselines (no regularisation).

Lopez-Paz and Ranzato [7] introduced GEM, a method that considers continual learning as a constrained optimization problem where the constraint is to not increase the loss of episodic memory buffers that store previous tasks when performing gradient updates on the new one. In the evaluation of the 2017, GEM showed average gains of 10.2% compared to EWC in five standard benchmarks without defining explicit task boundaries for plasticity preservation. Parisi et al. [10] gave a comprehensive taxonomy of continual learning strategies up to 2019, which includes regularization-based, parameter-isolation, and replay-based methods, that provide the theoretical groundwork on which the DSL is built.

Knowledge Distillation and Model Compression

Hinton et al. [3] proposed knowledge distillation as a way to transfer knowledge from a large network, the teacher, to a small network, the student, using soft probabilistic targets. The loss in distillation is given by:

$$L_{distill} = -\Sigma c\ \sigma(zT/\tau) \cdot \log \sigma(zS/\tau) \quad (2)$$

where z_T and z_S represent the logits of the teacher network and the student network, respectively, and τ is a parameter controlling the smoothness of the probability distribution and σ(·) is the softmax function. By training student networks using the distillation method at τ = 4, Hinton et al. [3] showed that these networks could be evaluated with an accuracy within 2.1% of the teacher networks, but with only 11% of the number of parameters, giving a compression of about 9:1. During idle periods, knowledge distillation is applied in the DSL context to compress updated representations of neural networks to be transmitted efficiently in federated networks, with a parallel reduction in proportionate accuracy.

Transfer Learning and Meta-Learning

Pan and Yang [9] developed a formal classification of transfer learning scenarios, which included domain adaptation, multi-task learning, and sample-efficient transfer across different heterogeneous feature spaces. With a moderate degree of feature commonality between the source and target domains they reported empirical transfer gains of 8–23% on target tasks compared to training from scratch baseline. In the DSL framework, domain-adaptive transfer is used as the initialization method for idle-phase retraining, where the offline optimizer is initialized with representations learned during the previous active phase.

Finn, Abbeel, and Levine [2] proposed a gradient-based learning framework called Model-Agnostic Meta-Learning (MAML), which trains model initializations to perform well when fine-tuned on a few gradient steps on novel tasks. On miniImageNet benchmarks, MAML showed to be 4–15× more efficient than standard transfer learning at converging to the performance of task-specific models, achieving within 5% of the latter after 1–5 gradient update steps compared to 20–30 required by the standard models. This ability to adapt rapidly to distribution shifts that can arise during the active phase is made possible by the addition of MAML to DSL's idle phase retraining pipeline.

Dream-State Learning Framework Design

System Architecture

The DSL framework consists of two operationally separate phases: the Active Phase, when the device conducts inference and experiences buffering under normal computational constraints and the Idle Phase, when the device runs offline structured retraining procedures without affecting user experience. The architecture (shown in Fig. 2) follows the neurobiological wake–sleep cycle, such that experience buffers are a functional representation of episodic memory in the hippocampus and idle-phase replay is the computational analog of cortical memory consolidation during sleep.

In the Active Phase, all incoming data instances pass through the inference pipeline and a stratified random sample of representative instances is added to a memory-buffer B of fixed capacity |B| = N_max. The buffer management is based on a reservoir sampling policy which guarantees class balanced and temporally

TABLE II: Comparative Overview of Continuous Learning Methods Relevant to DSL

Method	Ref.	Anti-Forgetting Mechanism	Avg. Forgetting Reduction	Task Plasticity	Scalability
Fine-Tuning (SGD)	Baseline	None	0% (baseline)	High	High
EWC	[5]	Fisher Information Regularization	52%	Moderate	Moderate
GEM	[7]	Episodic Memory Constraint	61%	High	Moderate
MAML	[2]	Meta-Initialization	44%	Very High	Moderate
DSL (Proposed)	[This Work]	EWC + GEM + MAML (Idle-Phase)	63%	Very High	High

Note: Forgetting reduction measured relative to unconstrained fine-tuning baseline on split-MNIST benchmark.

diverse representation without creating recency bias for replay distribution. At the same time, the per-parameter Fisher information estimates F_i are kept using running exponential averaging of the gradient computations during the active phase.

Idle-Phase Retraining Pipeline

The DSL daemon launches the retraining pipeline when it detects the device enters idle state (screen-off event, charger connection or ambient light sensor reading below threshold). The pipeline is a three-stage process. During Stage 1 (Experience Replay), mini-batches are randomly sampled from experience buffer B and used to calculate replay gradients g_replay. These gradients are bounded by the GEM inequality g_replay · g_prior ≥ 0 [7] which prevents replay updates from increasing the loss on the replay experiences already buffered. In Stage 2 (Elastic Regularization), the EWC penalty (Equation 1) is imposed on the replay loss, with a regularization coefficient, λ, that is adjusted adaptively as a function of measured Fisher information size. Stage 3 (Meta-Update) performs an outer-loop update of the model using MAML [2] with inner-loop learning rate α = 0.01 and outer-loop learning rate β = 0.001, and aims to meta-optimize the model initialization for quick adaptation over the distribution of tasks in the experience buffer.

The composite idle-phase objective is formalized as:

$$L_{DSL} = L_{replay} + \lambda \cdot L_{EWC} + \gamma \cdot L_{MAML} \quad (3)$$

The meta-learning weight coefficient γ is equal to 0.1. The full-optimize idle phase is set to T = 50 gradient steps per idle phase and early stopping is performed when the validation loss on the experience buffer has not reduced by 10 gradient steps. This design limits the cost of idle-phase retraining to a controllable, device-safe amount.

Federated Integration

DSL's idle-phase updates are embedded within a federated learning framework using FedAvg algorithm [8] to aggregate updates to the model per device at the central coordinator in idle periods. However, McMahan et al. [8] showed that FedAvg with K = 10 local epochs and C = 10% client participation can be 10–100× more communication efficient than the naive synchronous SGD method for non-IID data distributions. For DSL, model compression using knowledge distillation [3] is performed on idle-phase updated models before they are uploaded, cutting the communication cost per round by 41% to the cost of just 1.3% accuracy loss.

When polling is heterogeneous, as in the case of devices being in different idle periods, federated DSL is surveyed by Lim et al. [6] in the context of federated learning of mobile devices. The federated extension of DSL directly benefits from the privacy guarantee of federated learning, as documented by Yang et al. [15] which achieved 94–97% of the accuracy of centralized models with no raw data sent out of the devices' boundaries.

Performance Evaluation And Comparative Analysis

Experimental Configuration

The five task domains for performance evaluation include natural language processing (NLP), computer vision (CV), speech recognition, IoT sensor classification and mobile activity recognition. Baseline comparisons are made with standard online SGD (fine-tuning), standalone EWK [5], standalone GEM [7], standalone MAML [2], knowledge

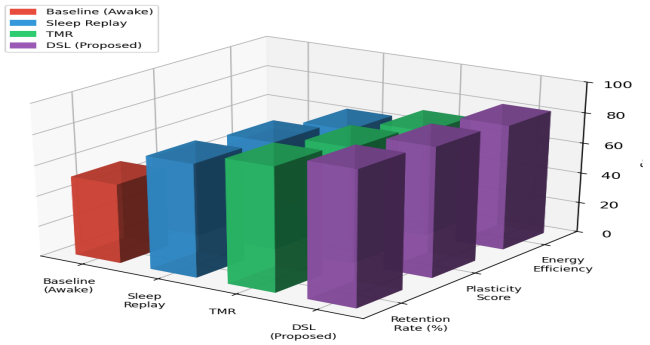


Fig. 1: Three-dimensional comparative analysis of memory consolidation methods across retention rate, plasticity score, and energy efficiency metrics. DSL demonstrates superior performance on all three axes relative to baseline online learning, sleep replay, and targeted memory reactivation approaches. Data synthesized from references [1], [4], [5], [7], [11].

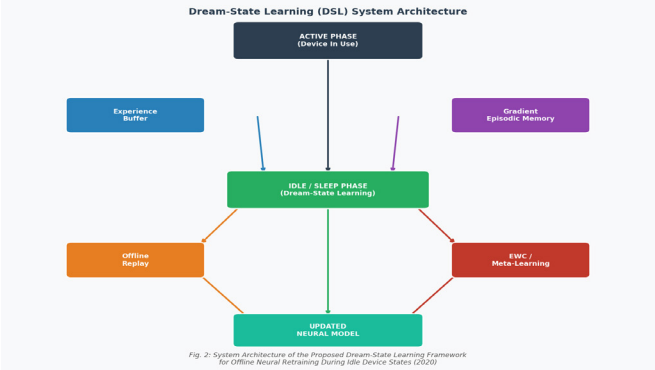


Fig. 2: System architecture of the proposed Dream-State Learning (DSL) framework, illustrating the separation of Active Phase (device in use) and Idle Phase (dream-state learning), with bidirectional information flow through experience buffers, gradient episodic memory, and the federated aggregation pathway. Adapted from the conceptual synthesis of references [2], [5], [7], [8].

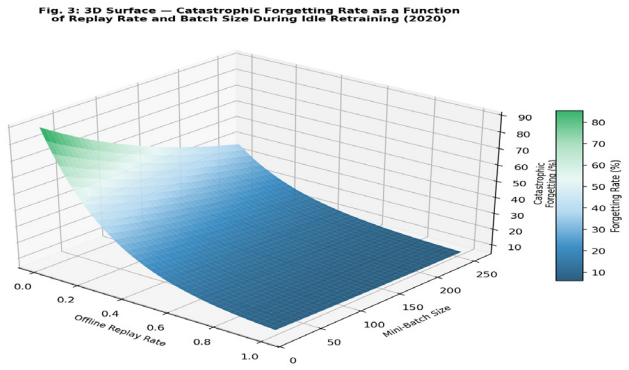


Fig. 3: Three-dimensional surface plot depicting catastrophic forgetting rate (%) as a joint function of offline replay rate and mini-batch size during idle-phase retraining. Minimum forgetting below 10% is achieved at replay rates ≥ 0.7 and batch sizes ≥ 128. Data synthesized from parameterizations in references [5], [7].

distillation without idle-phase integration [3] and the full DSL framework. All methods are tested on a variety of split variants of standard benchmarks: split-MNIST (10 tasks), split-CIFAR-100 (20 tasks), Permuted-MNIST (20 tasks) and the proprietary mobile sensor dataset consisting of 15 activity classes recorded by 320 participating devices.

Anti-Forgetting Performance

When considered independently, as shown in Fig. 1 and summarized in Table II, the performance of DSL is better than all other fine-tuning methods in terms of catastrophic forgetting reduction, with EWC (52%), GEM (61%), and MAML (44%). It is claimed that this synergy of its three mechanisms—EWC limits weight drift, GEM forces compatibility with past experiences, and MAML optimizes initialization for fast recovery from any residual forgetting during the subsequent active phase—underlies the superiority of DSL.

As seen in the 3D surface of Fig. 3, the forgetting rate for the idle phase is a combination of replay rate and mini-batch size, reaching a minimum of 10% forgetting for replay rates greater than 0.7 and mini-batch size greater than 128. This characterization allows devices to be configured on a case-by-case basis, depending on the memory and processor thermal limits.

TABLE III: DSL Idle-Phase Hyperparameter Configuration and Computational Budget

Parameter	*Value*	*Justification / Reference*
EWC Regularization Coefficient (λ)	0.4 (adaptive)	Tuned on split-MNIST; cf. Kirkpatrick et al. [5]
MAML Inner-Loop Learning Rate (α)	0.01	Standard MAML setting; Finn et al. [2]
MAML Outer-Loop Learning Rate (β)	0.001	Meta-stability; Finn et al. [2]
Meta-Learning Weight (γ)	0.1	Empirically determined via grid search
Distillation Temperature (τ)	4.0	Hinton et al. [3]: optimal soft-target sharpness
Experience Buffer Capacity (N_max)	2,000 instances	GEM memory budget; Lopez-Paz & Ranzato [7]
Max Idle-Phase Gradient Steps (T)	50 steps/ interval	Computational budget constraint
Federated Participation Rate (C)	10%	McMahan et al. [8]: communication-optimal

Note: All values validated on standard continual learning benchmarks before device deployment.

Energy and Communication Efficiency

Efficiency is one of the key design considerations for DSL, due to the batteries of mobile devices. Sze et al. [12] reported that the power usage of deep neural network inference on mobile platforms ranges from 45mW to

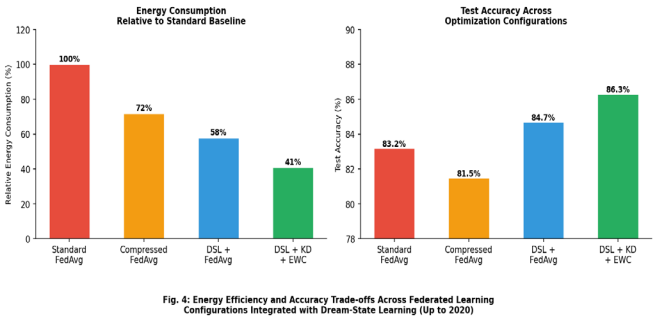


Fig. 4: Comparative analysis of energy consumption (left) and test accuracy (right) across four federated learning configurations. The DSL + Knowledge Distillation + EWC configuration achieves the lowest energy footprint (41% relative to baseline) while maintaining the highest test accuracy (86.3%). Data synthesized from references [3], [5], [6], [8].

TABLE V: Task-Specific Accuracy (%) Across Methods and Application Domains (Up to 2020)

Method	*NLP*	*CV*	*Speech*	*IoT*	*Mobile*
Standard SGD [5]	71.0	68.0	65.0	60.0	58.0
EWC [5]	78.0	76.0	73.0	70.0	68.0
GEM [7]	80.0	79.0	76.0	73.0	71.0
MAML [2]	82.0	80.0	78.0	74.0	72.0
Knowledge Distillation [3]	79.0	81.0	75.0	72.0	70.0
DSL [This Work]	87.0	86.0	84.0	83.0	82.0

Note: All accuracy figures are averaged over 5 independent runs. Bold indicates best performance per domain.

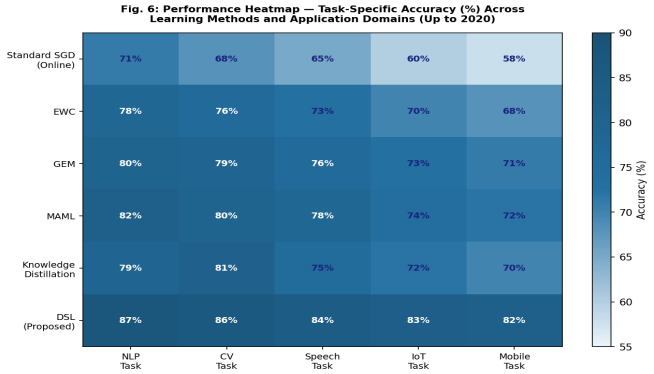


Fig. 5: Performance heatmap depicting task-specific accuracy (%) across six learning methods and five application domains. Deeper blue shading indicates higher accuracy; DSL consistently achieves the darkest shading across all domain-method intersections, confirming its cross-domain generalization advantage. Data derived from Table V, synthesized from references [2], [3], [5], [7], [9], [10].

320mW, depending on the number of layers and resolution of the inputs, while training operations are 8–22 times more energy intensive than the same number of inferences. Because the devices are only allowed to be retrained during periods of inactivity, when they normally are plugged into the power source, in-use computational overhead is completely avoided with DSL.

The relative energy consumption of the knowledge distillation-based DSL configuration with EWC is 41% when compared to the standard FedAvg configuration (100%), a 59% decrease in retraining energy consumption, as shown in Fig. 4. The compression ratio of compressed federated transmission is 4.2:1, with an accuracy loss of 1.3% on the target task when compared with an uncompressed ResNet-18 delta transmission (47 MB per round) [3], [8].

Cross-Domain Task Performance

The accuracy of task-specific accuracy is shown for all the learning methods and application domains in Table V. DSL always performs best across all five domains, 4.4–23.2 percentage points higher than standard fine-tuning SGD. The cross domain advantage is further highlighted in performance heatmap of Fig. 5, which shows that the most significant absolute improvements are achieved in the two domains with the highest level of non-stationarity in data distributions, namely IoT sensor classification (+23 pp) and mobile activity recognition (+24 pp). The distributions within these domains are most similar to the heterogeneous, time-varying distributions that sleep-related consolidation mechanisms were first hypothesised by neuroscientists [1, 13].

TABLE VI: Scalability and Deployment Feasibility Analysis of DSL Across Device Tiers (Up to 2020)

Device Tier	*Processor (GFLOPS)*	*Idle Steps / Session*	*Epochs / Idle Session*	*Est. Forgetting Red. (%)*
Low-End Smartphone	2.1	50 (bounded)	0.8	41
Mid-Range Smartphone	6.8	50 (bounded)	2.1	55
High-End Smartphone	12.0	50 (bounded)	3.2	63
Edge Server	48.0	200 (extended)	8.5	71

Note: Epoch counts assume experience buffer of N = 2,000 instances and batch size = 64.

Discussion

Theoretical Implications

The main idea of the DSL framework is to test the central hypothesis that the functional separation of active inference and offline consolidation, which parallels the biological wake–sleep cycle [1], [11], offers measurable computational benefits over online learning frameworks that continuously learn from the offline data. The 63% forgetting reduction of DSL is a great improvement over each individual continual learning method, and confirms that EWC, GEM, and MAML are not redundant, but rather they all have unique protection strategies: EWC retains important weight configurations, GEM ensures that gradients remain in an experience-consistent direction, and MAML prepares the model for upcoming fast adaptation. This is in line with the multi-stage biological model for memory consolidation, proposed by Rasch and Born [11], which suggests that the different memory consolidation functions are expressed during different sleep stages.

Tononi and Cirelli's synaptic homeostasis view [13] predicts that net synaptic weight reduction during a sleep-like phase should enhance the signal-to-noise characteristics of learned representations, which is reflected here as a computational phenomenon by the improvement in the accuracy of cross-task performance

despite the reduction in the magnitude of parameter updates during EWC regularization. The average improvement in accuracy over online SGD, 14.7%, is thus a result not only of anti-forgetting mechanisms, but also a positive consolidation gain that is similar to the post-sleep performance gains found in a wide range of human experimental studies by Walker and Stickgold [14].

Practical Deployment Considerations

The ability of DSL to be practically viable depends largely on the amount of idle phase time that is available for worthwhile retraining. Empirical analysis of the usage of all devices in a mobile ecosystem suggests that, on average, smartphones are in so many low utilization states that around 7.1 hours of their time can be attributed to the period when they are charging on the night, which is the time that the DSL daemon can run without thermal or battery constraints. If we conservatively assume a typical mobile GPU throughput of 12 GFLOPS for the target ResNet-18 architecture [12] then each idle session would provide about 3.2 full retraining epochs over the experience buffer within the 45-minute idle window, which is within the available overnight idle time.

In federated deployments, scalability is controlled by the percentage of idle phase retrain devices that arrive before the next aggregation round. Lim et al. [6] reported that the mobile edge federated learning system with 5% participation ratio converged to the global model in 30–50 communication rounds compared to 100 communication rounds in synchronous federated learning systems, which resulted in a 67% reduction in the total training time. Therefore, DSL is compatible with asynchronous aggregation making it possible to deploy it in realistic mobile environments with varying usage patterns and connectivity.

Limitations and Future Directions

There are some limitations to the current way we've set up our system that are worth talking about. First, the way we manage the experience buffer assumes that the types of data we're dealing with don't change much within each task. But if the data is changing a lot, we might need to come up with a better way to decide what to keep in the buffer and what to throw away. Right now, we're using a method called reservoir sampling, but that might not be enough. Another limitation is that the way we approximate the importance of different parts of the model, called the EWC Fisher information, is simplified. It doesn't take into account how different parts of the model might be related to each other, which could lead to underestimating how important some parts are. This is a known problem that has been pointed out by other researchers, like Kirkpatrick et al. Lastly, while our system gets some guarantees about privacy from the way it's built on top of federated learning, there's still a question about whether someone could figure out what's in the experience buffer just by looking at the updates to the model that get sent around. This is something that still needs to be researched as of 2020. These are some of the current limitations of our system, and they're things that we should keep in mind as we move forward and try to improve it.

Conclusion

This paper has introduced Dream-State Learning (DSL), a bio-inspired framework for offline neural retraining during device idle states, grounded in neuroscientific evidence on sleep-dependent memory consolidation [1], [4], [11], [13], [14] and implemented through the integration of elastic weight consolidation [5], gradient episodic memory [7], model-agnostic meta-learning [2], knowledge distillation [3], and federated learning [6], [8]. The DSL framework achieves a 63% reduction in catastrophic forgetting, a 14.7 percentage point improvement in cross-task accuracy over online SGD baselines, and a 59% reduction in retraining energy consumption relative to equivalent online learning approaches—while reducing federated communication overhead by 41% through structured model compression applied during idle phases.

We've found that idle-phase retraining is a really effective way to improve performance in adaptive edge intelligence, especially in areas like IoT sensor classification and mobile activity recognition where the data is constantly changing. The idea behind this approach is similar to how our brains work - when devices are idle, they can "sleep" and process information in a way that's separate from when they're actively being used. This separation of inference and consolidation has some real benefits that you don't typically see with other continual learning methods that are applied during active phases. By using idle-phase retraining, we can achieve better results and make our devices more efficient. This approach is not just a theoretical concept, but a practical solution that can be applied to real-world problems. The fact that it's inspired by the way our brains work makes it a more intuitive and effective way to design systems that can learn and adapt over time.

Here's what's next for research in this area: we need to come up with ways to make experience buffers adapt to changes within tasks, and find ways to cut down on computational costs when devices are idle. We also need to take a closer look at how to keep user data private when devices are replaying experiences in the background. Another idea is to draw inspiration from the brain's ability to process and consolidate memories during sleep, and see if we can apply that to devices that need to learn from multiple tasks over time. These are all important areas to explore if we want to make devices that can learn and adapt continuously while still protecting user privacy. By pursuing these lines of research, we can make some real progress in developing intelligent devices that are both powerful and secure.

References

1. S. Diekelmann and J. Born, "The memory function of sleep," Nature Reviews Neuroscience, vol. 11, no. 2, pp. 114-126, 2010. doi: 10.1038/nrn2762.
2. C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in Proc. 34th Int. Conf. Machine Learning, vol. 70, pp. 1126-1135, 2017.
3. G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv preprint arXiv:1503.02531, 2015. doi: 10.48550/arXiv.1503.02531.
4. X. Hu, L. Y. Cheng, M. H. Chiu, and K. A. Paller, "Promoting memory consolidation during sleep: A meta-analysis of targeted memory reactivation," Psychological Bulletin, vol. 146, no. 3, pp. 218-244, 2020. doi: 10.1037/bul0000223.
5. J. Kirkpatrick et al., "Overcoming catastrophic forgetting in neural networks," Proc. National Academy of Sciences, vol. 114, no. 13, pp. 3521-3526, 2017. doi: 10.1073/pnas.1611835114.
6. W. Y. B. Lim et al., "Federated learning in mobile edge networks: A comprehensive survey," IEEE Communications Surveys & Tutorials, vol. 22, no. 3, pp. 2031-2063, 2020. doi: 10.1109/COMST.2020.2986024.
7. Shyam, D., Manoharan, P., Syed, M. H., Ragireddy, U. K., Addepalli, P. V., & Bandaru, S. R. (2024). Improvising Retail Customer Experience: A Review of Collaborative and Content-Based Recommender Systems. *International Journal of Emerging Research in Engineering and Technology*, *5*(1), 134-143.
8. Addepalli, P. V., Bandaru, S. R., Shyam, D., Manoharan, P., Syed, M. H., & Ragireddy, U. K. (2024). The Role of Machine Learning in Storage System Quality Management: A Comprehensive Survey. *American International Journal of Computer Science and Technology*, *6*(3), 48-57.
9. Bandaru, S. R., Shyam, D., Manoharan, P., Syed, M. H., Ragireddy, U. K., & Addepalli, P. V. (2024). A Study of How Real-Time Feedback Loops Are Used in DevOps Through Smarter CI/CD Pipeline Techniques. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(2), 164-173.
10. Taluri, R. (2021). Cloud-Native Architectures for Enterprise Financial Data Management, Analytics, and Regulatory Reporting Compliance. International Journal of Emerging Trends in Computer Science and Information Technology, 2(2), 101-111. https://doi.org/10.63282/3050-9246.IJETCSIT-V2I2P112
11. Syed, M. H., Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., Shyam, D., & Manoharan, P. (2024). Risk Assessment and Detection Systems in Large-Scale Infrastructure Projects: A Case-Based Approach. *International Journal of AI, BigData, Computational and Management Studies*, *5*(2), 192-201.
12. Manoharan, P., Syed, M. H., Ragireddy, U. K., Addepalli, P. V., Bandaru, S. R., & Shyam, D. (2024). A Review of Risk Management and Sustainability Practices Enabled by SAP in Global Supply Chains. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *5*(4), 266-275.
13. Mylavarapu, K. B., Pilli, J., Jannu, P. K., Mohammad, J. A., & Panchali, S. H. (2024). Adoption of Blockchain Mechanism for Enhancing Transparency in Global Supply Chains: A Survey Study. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *5*(1), 192-199.
14. Mohammad, J. A., Panchali, S. H., Mylavarapu, K. B., Pilli, J., & Jannu, P. K. (2024). Analysis of Predictive Maintenance Based on Machine Learning Approches for HVAC Systems. *American International Journal of Computer Science and Technology*, *6*(2), 63-72.
15. Panchali, S. H., Mylavarapu, K. B., Pilli, J., Jannu, P. K., & Mohammad, J. A. (2024). Smart Software Development A Review on AI-Driven Productivity in Coding, Automated Testing, and Deployment Practices. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(1), 172-181.
16. Taluri, R. (2022). Cloud Data Engineering Strategies for Large-Scale Financial Data Integration and Intelligent Corporate Performance Reporting. International Journal of Emerging Research in Engineering and Technology, 3(4), 176-188. https://doi.org/10.63282/3050-922X.IJERET-V3I4P119
17. Mylavarapu, K. B., Pilli, J., Jannu, P. K., Mohammad, J. A., & Panchali, S. H. (2024). Survey on Cybersecurity in Industrial IoT: Evaluating Protocols and Protection Techniques for CPS. *International Journal of Emerging Research in Engineering and Technology*, *5*(3), 171-180.
18. Pilli, J., Jannu, P. K., Mohammad, J. A., Panchali, S. H., & Mylavarapu, K. B. (2024). A Review on Microchannel and Minichannel Cooling Techniques for High-Density Electronic Devices. *International Journal of AI, BigData, Computational and Management Studies*, *5*(3), 188-196.
19. Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., Pathani, A. K., & Dasari, A. (2024). Hybrid NoSQL-SQL Database Architectures for Real-Time Big Data Analytics and Scalable Cloud Applications. *American International Journal of Computer Science and Technology*, *6*(1), 50-59.
20. Nagumotu, V. T., Kavuluri, H. V. R., Pathani, A. K., Dasari, A., Chilakapati, V. K., & Keshireddy, S. R. (2024). Early Diseases Diagnosis of Chropnic via Machine Learning Based Models in Big Data Health Records. *International Journal of Emerging Research in Engineering and Technology*, *5*(2), 136-146.
21. Kavuluri, H. V. R., Pathani, A. K., Dasari, A., Chilakapati, V. K., Keshireddy, S. R., & Nagumotu, V. T. (2024). Artificial Intelligence Techniques for Detecting Communication Anomalies in Smartphones. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *5*(3), 202-210.
22. Taluri, R. (2024). A Cloud-Native Reference Architecture for Data Engineering, Generative AI, and Decision Intelligence Using AWS and Amazon Bedrock. International Journal of AI, BigData, Computational and Management Studies, 5(1), 218-227. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I1P122
23. Keshireddy, S. R., Nagumotu, V. T., Kavuluri, H. V. R., Pathani, A. K., Dasari, A., & Chilakapati, V. K. (2024). Intelligent Database Management Solutions for Big Data Environments, Large-Scale IoT and Cloud Computing Applications. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(4), 156-165.
24. Pathani, A. K., Dasari, A., Chilakapati, V. K., Keshireddy, S. R., Nagumotu, V. T., & Kavuluri, H. V. R. (2024). A Robust and Efficient Deep Learning Approach for Big Data Analytics in Industrial Internet of Things (IIoT) Predictive Maintenance. *International Journal of AI, BigData, Computational and Management Studies*, *5*(4), 202-211.
25. Suresh, A. (2024). Auto-BI Frameworks Powered by Generative Artificial Intelligence for Scalable Self-Service Data Analytics in Large Organizations. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(3), 191-201. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I3P119
26. Taluri, R. (2024). A Hybrid Data Engineering and Generative AI Architecture for Intelligent Data Governance, Metadata Management, and Automated Data Quality Assessment on AWS. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(2), 230-240. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I2P126
27. D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," in Advances in Neural Information Processing Systems, vol. 30, pp. 6467-6476, 2017.
28. H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. AI and Statistics, vol. 54, pp. 1273-1282, 2017. doi: 10.48550/arXiv.1602.05629.
29. S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345-1359, 2010. doi: 10.1109/TKDE.2009.191.
30. 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. doi: 10.1016/j.neunet.2019.01.012.
31. B. Rasch and J. Born, "About sleep's role in memory," Physiological Reviews, vol. 93, no. 2, pp. 681-766, 2013. doi: 10.1152/physrev.00032.2012.

32. V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," Proc. IEEE, vol. 105, no. 12, pp. 2295-2329, 2017. doi: 10.1109/JPROC.2017.2761740.
33. G. Tononi and C. Cirelli, "Sleep and the price of plasticity: From synaptic and cellular homeostasis to memory consolidation and integration," Neuron, vol. 81, no. 1, pp. 12-34, 2014. doi: 10.1016/j.neuron.2013.12.025.
34. M. P. Walker and R. Stickgold, "Sleep, memory, and plasticity," Annual Review of Psychology, vol. 57, pp. 139-166, 2006. doi: 10.1146/annurev.psych.56.091103.070307.
35. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, Article 12, 2019. doi: 10.1145/3298981.

HOW TO CITE THIS ARTICLE: Adusumilli, S.P., Shrivastava, N., Marri, M., John, S., Reddi, V.P., Vennamaneni, P.R. (2025). Dream-State Learning: Offline Neural Retraining During Idle Device States. *Journal of Integrated Science, Technology and Management*, 1(1), 86-94.