

Concept Paper

Not peer-reviewed version

Dreaming Machine Learning (DML): A Speculative Framework for Latent Knowledge Discovery

[Amal Jayawardena](#)*

Posted Date: 25 August 2025

doi: 10.20944/preprints202508.1716.v1

Keywords: Machine Learning (ML); Artificial Intelligence (AI); Dreaming; Creativity



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Concept Paper

Dreaming Machine Learning (DML): A Speculative Framework for Latent Knowledge Discovery

Amal Jayawardena

Independent Researcher, Australia; amaljay94@gmail.com

Abstract

Modern machine learning systems excel at data driven optimization but remain confined by training objectives and input distributions. Conventional paradigms such as supervised, unsupervised, reinforcement, and self supervised learning operate within explicit boundaries defined by task specific loss functions and curated data. In biological systems, dreaming provides essential functions such as memory consolidation, emotional regulation, and creative recombination. Inspired by this, we propose Dreaming Machine Learning, a speculative framework in which models alternate between a wake phase of conventional training and a dream phase of unconstrained exploration. During dream states, networks may replay and recombine latent embeddings, pursue alternative objectives such as novelty or entropy maximization, or generate counterfactual and hallucinatory samples that extend beyond training distributions. These internal explorations enable the discovery of latent structures, novel hypotheses, and emergent patterns that are not accessible through standard learning. We outline potential applications of Dreaming Machine Learning in scientific discovery, creativity, anomaly detection, robotics, and personalized artificial intelligence, while highlighting open challenges of evaluation, stability, ethics, and sustainability. By formalizing a wake and dream cycle, Dreaming Machine Learning aims to expand artificial intelligence beyond pattern recognition toward imagination and hypothesis generation. The future of machine intelligence may depend not only on learning by doing, but also on learning by dreaming.

Keywords: machine learning (ML); artificial intelligence (AI); dreaming; creativity

Introduction

Modern machine learning systems excel in data-driven optimization but remain constrained by their training objectives and input distributions. Conventional paradigms such as supervised, unsupervised, reinforcement, and self-supervised learning operate within explicit boundaries defined by human-provided data and task-specific loss functions. However, just as biological organisms benefit from dreaming to reorganize memories, consolidate knowledge, and generate novel associations, machine learning models may require analogous “dreaming states” to access underexplored regions of their representational space. Recent work has highlighted the need for AI systems to go beyond narrow optimization and embrace creativity, hypothesis generation, and generalization in unfamiliar domains [1].

Dreaming Machine Learning (DML) is a speculative paradigm in which models enter phases of exploration without direct input or explicit labels, guided instead by synthetic potentials, noise-injected dynamics, or counterfactual simulations. These dream states provide opportunities for networks to uncover latent structures, form novel hypotheses, and augment training with patterns not discoverable in conventional active learning. This echoes the growing recognition that learning systems may require phases of “off-task” internal computation to fully realize their potential [2].

Biological Inspiration: Human Dreaming as Data Reorganization

Human dreaming has long been hypothesized to serve critical cognitive functions, including memory consolidation, emotional regulation, and creative insight. Studies in neuroscience show that hippocampal replay during sleep plays a central role in reorganizing daily experiences into durable long-term memories [3,4]. Dreaming also supports emotional integration, reactivating emotionally salient episodes and embedding them into broader cognitive frameworks [5]. Finally, dreams foster creativity by recombining unrelated fragments of memory into novel but meaningful associations, a phenomenon that has inspired speculation about their evolutionary role [6].

These insights provide a computational analogy. Just as biological systems enter sleep cycles alternating between REM and non-REM phases, artificial networks could alternate between standard training as a wake phase and exploratory dreaming as a latent phase. Unlike passive noise, such dream cycles would be engineered for purposeful reorganization and discovery.

Defining the Dream State in ML

A dream state in machine learning can be described as a non-task-driven optimization cycle in which the objective diverges from conventional training. This could involve latent replay, where embeddings of past experiences are reactivated and stochastically recombined, paralleling hippocampal replay in sleep. It could also involve the adoption of alternative potentials, in which the model shifts from minimizing task error to maximizing novelty, entropy, or diversity. Such objectives push the system toward regions of the loss landscape normally avoided during training. Another pathway is counterfactual simulation, where assumptions are inverted and the model imagines how the world might appear if categories or dynamics were reversed. Finally, synthetic hallucination may generate inputs beyond the training distribution, re-encode them, and mine them for emergent structures.

Emerging research in unsupervised learning and generative replay already points toward the promise of such approaches. For instance, generative models that hallucinate synthetic samples have been shown to expand representational richness and improve generalization [7]. DML extends this principle, embedding such hallucination into a structured dream cycle.

Patterns Emerging from Dreams

Dream states are likely to yield outputs that appear strange, but within these distortions lie hidden structures. In vision, recombined embeddings could reveal intermediate categories that expose latent taxonomies. In language, drifting representations might capture subtle emotional or conceptual transitions. In multimodal systems, dreams may weave together text, sound, and vision in alignments not constrained by real-world data.

These patterns provide candidates for knowledge expansion. Just as human dreams can inspire artistic breakthroughs or scientific hypotheses, machine dreams could uncover structural relationships that would otherwise remain invisible. This aligns with arguments in cognitive neuroscience that the brain actively constructs “virtual realities” during dreaming to facilitate adaptive learning [8].

Applications of Dreaming ML

The potential applications of DML extend across science and society. In scientific discovery, dreaming models could generate molecular structures or biological pathways that go beyond established databases, echoing calls for AI to move from data-fitting to hypothesis generation [9]. In creativity, dream-inspired generative systems may yield music, art, or literature unbounded by cultural priors. For anomaly detection, divergences between dream distributions and real-world inputs could reveal “unknown unknowns,” providing new approaches to risk identification. In

robotics, embodied agents might evolve locomotion or manipulation strategies overnight through internal simulation, without requiring constant physical trial-and-error.

Most provocatively, personalized AI companions could develop continuity of memory and emotion by dreaming. By revisiting prior interactions in dream states, agents may generate something akin to an inner life, reshaping future behaviors in ways that feel more adaptive and human-like. This would represent a profound shift, extending AI from pattern recognition into imagination.

Challenges and Open Questions

DML raises as many questions as it answers. The problem of evaluation is fundamental: how can we measure the utility of dreams if their value lies in exploring beyond ground truth? Stability is another open concern, since dream states may risk destabilizing network weights, leading to catastrophic forgetting. Ethical concerns loom large, as unbounded dreamlike recombinations may yield harmful outputs or reinforce undesirable biases. Finally, the energy cost of prolonged or parallel dream phases is significant, raising questions of sustainability at a time when the computational cost of deep learning is already under scrutiny [10].

Toward a Dream–Wake Cycle in ML

A promising way forward is to formalize a dual-objective cycle. During the wake phase, the system optimizes against explicit tasks with external data. During the dream phase, it engages in exploration guided by novelty, counterfactual recombination, or synthetic hallucination. Integration occurs when useful dream artifacts are identified and reinforced into the task model.

This cycle mirrors biology, where waking life grounds the organism in reality and dreaming reorganizes and expands that reality. Embedding such cycles into machine intelligence could transform AI into systems capable not only of learning from data, but also of imagining beyond it.

Conclusions

Dreaming Machine Learning is a speculative but provocative hypothesis: that artificial systems, like biological organisms, require phases of dreaming to achieve full cognitive potential. By alternating between wake and dream phases, models could uncover hidden structures, generate creative hypotheses, and achieve generalization that extends beyond current limits. Although technical, ethical, and philosophical challenges remain, recognizing dreaming as a missing pillar of machine intelligence could open a new frontier in AI research. The future of intelligence may depend not only on learning by doing, but also on learning by dreaming.

References

1. Kejriwal, M. *et al.* Can AI have common sense? Finding out will be key to achieving machine intelligence. *Nature* **634**, 291-294 (2024). <https://doi.org/10.1038/d41586-024-03262-z>
2. Kucyi, A., Kam, J. W. Y., Andrews-Hanna, J. R., Christoff, K. & Whitfield-Gabrieli, S. Recent advances in the neuroscience of spontaneous and off-task thought: implications for mental health. *Nature Mental Health* **1**, 827-840 (2023). <https://doi.org/10.1038/s44220-023-00133-w>
3. Ji, D. & Wilson, M. A. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience* **10**, 100-107 (2007). <https://doi.org/10.1038/nn1825>
4. Klinzing, J. G., Niethard, N. & Born, J. Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience* **22**, 1598-1610 (2019). <https://doi.org/10.1038/s41593-019-0467-3>
5. Goldstein, A. N. & Walker, M. P. The Role of Sleep in Emotional Brain Function. *Annual Review of Clinical Psychology* **10**, 679-708 (2014). <https://doi.org/https://doi.org/10.1146/annurev-clinpsy-032813-153716>
6. Cai, D. J., Mednick, S. A., Harrison, E. M., Kanady, J. C. & Mednick, S. C. REM, not incubation, improves creativity by priming associative networks. *Proceedings of the National Academy of Sciences* **106**, 10130-10134 (2009). <https://doi.org/doi:10.1073/pnas.0900271106>

7. Shin, H., Lee, J. K., Kim, J. & Kim, J. Continual learning with deep generative replay. *Advances in neural information processing systems* **30** (2017).
8. Hobson, J. A. & Friston, K. J. Waking and dreaming consciousness: Neurobiological and functional considerations. *Progress in Neurobiology* **98**, 82-98 (2012).
<https://doi.org/https://doi.org/10.1016/j.pneurobio.2012.05.003>
9. Stokes, J. M. *et al.* A Deep Learning Approach to Antibiotic Discovery. *Cell* **180**, 688-702.e613 (2020).
<https://doi.org/10.1016/j.cell.2020.01.021>
10. Guan, L. Reaching carbon neutrality requires energy-efficient training of AI. *Nature* **626**, 33 (2024).
<https://doi.org/10.1038/d41586-024-00200-x>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.