BioNanoScience (2025) 15:591
https://doi.org/10.1007/s12668-025-02229-2

RESEARCH

A Bio-realistic Synthetic Hippocampus for Robotic Cognition

Jordi Vallverdu[1] · Xenia Feinstein[2] · Paul Robertson[3] · Ivan Kipelkin[4] · Max Talanov[5,6]

Received: 12 May 2025 / Accepted: 9 October 2025

Abstract
Current robotic systems struggle with adaptive generalisation beyond curated training domains. Inspired by hippocampal dynamics in biological cognition, we introduce a synthetic memory architecture that segregates online sensorimotor interaction from offline consolidation and generative replay. Implemented via spiking neural networks and neuromorphic substrates, our framework enables bidirectional memory traversal, goal-prioritised plasticity updates, and energy-efficient policy synthesis. This dual-state system bridges real-time control with autonomous learning, advancing a biologically grounded pathway toward resilient, context-adaptive robotic intelligence.

Keywords Sleep · Hippocampus · Robotic · Spiking neural networks · Memristive device

1 Introduction

Robotic agents are increasingly deployed in complex and dynamic environments from autonomous homes and factories to search and rescue missions [39]. These contexts are inherently unpredictable: no two apartments have the same layout; tasks evolve; and new objects are frequently encountered [58]. However, most robots today rely on fragile, static models that cannot adapt to such variability [24].

✉ Jordi Vallverdu
jordi.vallverdu@uab.cat

Xenia Feinstein
xenia.feinstein@gmail.com

Paul Robertson
paulr@dollabs.com

Ivan Kipelkin
kipelkinivan@gmail.com

Max Talanov
max.talanov@gmail.com

1 ICREA, Universitat Autònoma de Barcelona, Barcelona, Spain

2 National Institute of Mental Health, Klecany, Czech Republic

3 Dynamic Object Language Labs, Lexington, MA 02421, USA

4 B-Rain Ark FZ-LLC, Dubai, United Arab Emirates

5 University of Messina, Messina, Italy

6 Institute for Artificial Intelligence R&D of Serbia, Novi Sad, Serbia

A core limitation is the cross-contextual generalisation of the ability to synthesise goal-directed behaviour in novel environments using heterogeneous sensory and semantic inputs [27]. Modern learning paradigms (e.g. supervised deep learning or reinforcement learning) typically require dense, curated data and retraining for new domains. Even recent large-scale robotics foundation models, like RT-2 or $\pi_{0.5}$, although capable of zero-shot task inference [10], are effectively "frozen" at runtime. These systems perform continuous inference but cannot relearn or reorganise their internal knowledge after deployment. Their synchronously coupled architectures produce high computational costs (often exceeding 100 W [17]) and make real-time adaptation infeasible on embedded hardware.

In contrast, biological agents adapt routinely to new situations with minimal supervision or energy overhead [56]. This capability is supported by a fundamental organisational principle: the alternation between online and offline computational states. During wakefulness, the brain performs rapid sensorimotor integration; during sleep, it consolidates experience, prioritises memories, and restructures internal representations [9, 56].

The hippocampus plays a key role in this alternation. As a bidirectional switch between sensory input and cortical integration, it enables memory replay [15], pattern completion, novelty detection [3], and context gating [6]. In particular, hippocampal sharp wave ripples (SWRs) during slow wave sleep have been shown to drive long-range consolidation and creative recombination of prior experience [22].

Published online: 21 October 2025

We propose a synthetic hippocampal architecture for robots that explicitly incorporates this dual-mode principle. Our model combines spiking neural networks with neuromorphic substrates to support offline replay, synaptic restructuring, and memory compression. This architecture decouples real-time control from long-timescale learning, enabling the following:

- Energy-efficient adaptation (using $\leq$1 W on neuromorphic chips [17])
- Resilience to catastrophic forgetting through goal-weighted replay [13]
- Enhanced generalisation via flexible memory traversal [15]

To concretely address the challenges outlined above, this work makes the following contributions:

1. A biologically grounded dual-phase architectural framework for robotic cognition
2. An analytical energy and memory proof based on spiking replay and TD update mechanisms
3. A prototype toy robotic system that demonstrates successful offline adaptation and low-power operation

1.1 Limitations of Current Intelligent Autonomous Robotic Systems

One of the domains where the deployment of robots appears particularly critical is operation in extreme conditions, such as emergency scenarios and environments hazardous to humans [38]. These include collapsed mines, fires, and accidents at sites involving chemical or radiological contamination. Despite growing interest in robotic systems for such applications, their widespread and effective use remains limited by major constraints, including insufficient autonomy and poor adaptability to non-standard situations. Recent advances in bioinspired soft robotics have demonstrated significant improvements in energy efficiency and adaptability, addressing some of the core limitations of traditional autonomous systems [50]. Furthermore, the integration of biological components into robotic architectures offers promising avenues to enhance plasticity and functional resilience [4]. Robots are also increasingly used in sensitive domains such as healthcare, where personalised solutions and autonomous decision-making are essential. However, to ensure the reliability and safety of such systems, a range of fundamental challenges must be addressed, from learning methods to performance under real-world, often unpredictable conditions.

A central limitation remains the inability of current systems to learn from a small number of unique and heterogeneous examples. In emergency settings, it is practically impossible to collect representative data sets in advance due to the unpredictable nature of events and limited data accessibility. Conventional data collection and annotation methods, including human supervision or crowd-sourcing, become economically impractical or entirely inapplicable due to privacy restrictions or lack of connectivity.

Even when high-quality training data are available, most existing approaches do not support real-time learning. In most current solutions, learning is performed exclusively offline, followed by the deployment of a fixed model. This model remains effective only as long as the environmental conditions match closely those of the simulation phase [62]. Once integrated into the physical system, further adaptation typically ceases, leaving the robot inflexible in dynamic and evolving scenarios.

Another critical issue is catastrophic forgetting, the tendency of models to overwrite previously learnt knowledge when new information is introduced. Preventing this phenomenon usually requires retraining on the entire dataset, which imposes high computational costs and makes real-time learning virtually infeasible under limited-resource conditions.

Physical constraints also play a crucial role. A robot must be energetically autonomous, independent of external computing infrastructure, and compact in form [34]. This is particularly important in scenarios with no stable connectivity or where data transmission involves confidentiality risks. These limitations underscore the need for a fundamental rethink of the architectures and learning principles that govern autonomous systems. Biological mechanisms, refined through evolution, offer energy-efficient and robust solutions that have proven effective in real-world conditions. Unlike modern AI systems, biological agents are capable of learning without manually labelled data sets or requiring massive simulated repetitions. These characteristics position biologically inspired approaches as particularly promising for the development of the next generation of robotic agents. Recent advances in bioinspired soft robotics reinforce this direction, demonstrating how principles of living organisms can improve adaptability and material compliance in autonomous systems [61].

2 Architecture of the System

The proposed system is based on a minimised yet functionally complete model of the hippocampus, a structure of the medial temporal lobe that, in animals, is responsible for spatial navigation, episodic memory, and the integration of disparate pieces of information into a unified context. These functions are supported by its layered architecture and the specific connectivity between the regions: Dentate Gyrus (DG) $\rightarrow$ CA3 $\rightarrow$ CA1.

The primary input to the hippocampus comes from the adjacent entorhinal cortex (EC), which in turn receives heavily processed information from associative and sensory cortices. This makes the hippocampal input both universal and sensorily agnostic. Such a degree of preprocessing indicates the modality-neutral nature of the hippocampus: it is capable of operating on any type of information, visual, auditory, or multisensory. We leverage this function in our model where the EC receives data from various sources and converts it into a unified representation. This architecture enables the easy addition of a new modality; one simply needs to connect its encoder to the same associative columns without modifying the rest of the network.

This input is received by the DG, a structure responsible for decorrelating inputs and generating a sparse representation of sensory information. In biological systems, the DG helps prevent the overlap of similar patterns, facilitating pattern separation: lateral inhibition leaves only 2–5% of granule cells active, drastically reducing the overlap of similar episodes. When the DG is damaged, animals show impairments in distinguishing similar memories or contexts. We replicate this mechanism using a k-WTA (k-Winner-Take-All) layer (with $k \approx 3\%$), which reduces the correlation between adjacent scenes by approximately fivefold and prevents the "blending" of memory maps.

The sparse output of the DG activates pyramidal neurons in CA3, which form a recurrent attractor network. This topology enables rapid pattern completion from partial cues and naturally generates sharp wave–ripples (SWRs) lasting around 100 ms.

CA1 receives a dual stream of input: a prediction from CA3 and raw data directly from the EC. By comparing the two, CA1 produces a novelty signal and generates output for supra-cortical systems. The back projection from CA1 to EC closes the small cortico-hippocampal loop, refining the sensory code. Meanwhile, output from the PFC travels through the striatum to the action planning module and further to the robot's actuators. This forms a unified chain: sensor → context → memory → evaluation → action, where the hippocampus acts as a fast working buffer and replay generator, while the associative cortex gradually assumes the role of long-term storage and integration of multimodal memories.

2.1 Motivation Subsystem

In biological systems, the PFC and ventral tegmental area (VTA) play a crucial role—they are responsible for selecting the action that is most advantageous and most likely to lead to a reward. These structures also initiate the chosen action and trigger a strategy shift when the outcome is unsatisfactory.

It is important for our agent to be capable of goal-directed behaviour, to sustain results, or to adjust its motivation. Such a sequence of actions arises from selecting behaviour based on the expected reward, for example locating a survivor under rubble. The delivery of the reward and the subsequent success or failure (with the possibility of reward withdrawal) are measurable.

Behavioural decisions in our architecture rely on continuous state value estimation using Temporal-Difference (TD) learning. The output of the hippocampus (CA1) is sent to a "cognitive controller", where the value approximation $V(s_t)$ is maintained and the next action is selected. The prediction, along with the anticipated outcome, is forwarded to the ventral tegmental area (VTA), which simulates the dopamine-based prediction error system. The VTA compares the expectation with the actual environmental feedback (e.g. successful rescue of a victim, target detection, penalty for collision) and generates a dopaminergic signal δ. This signal returns to CA1, tagging the active synapses—during the day on real inputs and at night on the same patterns replayed internally. In this way, the same closed loop CA1 → PFC → VTA → CA1 establishes TD tags during both online and offline phases.

$$V(s_t)^{new} \leftarrow V(s_t)^{old} + \alpha \left[r_t + \gamma V(s_{t+1}) - V(s_t)^{old} \right] \quad (1)$$

where

- $V(s_t)$ is the estimated value of the state s_t.
- r_t is the reward received at time t.
- γ is the discount factor that controls how much future rewards are valued.
- α is the learning rate that determines how much the value function is updated.
- $V(s_{t+1})$ is the estimated value of the next state.
- The term $\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$ is the temporal difference (TD) error.

In a digital or purely spiking implementation (DM module), the computations in (1) are executed as linear algebra operations, allowing for the parallel update of tens of thousands of states. The TD error δ_t is accumulated online, while the heavy-weight updates can be applied offline in batches, reducing energy consumption. On a CPU/GPU or specialised chips (e.g. Intel Loihi), such a module operates within $\approx 300W$ and updates once per second, sufficient for real-time robotics; a memristive version would enable further power reduction and onboard execution.

When the PFC receives the updated δ_t, it passes it to a striatum-like "Action Planner", which selects the action that maximises $Q(s, a)$. This setup allows for multiple, simultaneously active motivations: each potential candidate action is accompanied by an expected reward, and subse-

quent reward withdrawal upon failure automatically adjusts the value through Eq. 1.

Fundamentally, the system distinguishes only two things: a positive or negative signal from the peripherals. The magnitude of the update is defined by the value of δ_t, while the duration of the effect depends on whether the tag is reinforced at night via the lLTP mechanism. Altogether, this creates a coherent motivational scale: rapid online decision-making driven by TD estimates and long-term strategy stabilisation through offline consolidation.

3 Memory Consolidation and Plasticity

In living systems, molecular-level consolidation unfolds in several stages. During the early phases of synaptic memory consolidation, the response is enhanced by mobilizing receptors that are already in place, making it easier to reach the excitation threshold. Later, typically during sleep, a qualitative increase occurs, involving either the growth of synaptic spines or the deployment of new receptors.

Consequently, spike-timing-dependent plasticity (STDP), though widely used in neuromorphic computing, is effective only on millisecond scales and is inadequate for modelling the hours-long processes of long-term memory consolidation. To overcome this limitation, we introduce two biologically inspired plasticity functions designed for offline synaptic reconfiguration during sleep-like states. These functions account for prior activity and synaptic-tagging processes that influence the syntactic weight value.

3.1 Synaptic Long Long-Term Plasticity

In addition to the TD learning mechanism, we introduce a model of long long-term plasticity (lLTP) that operates over extended timescales and supports synaptic consolidation processes. This is motivated by the synaptic homeostasis hypothesis, which emphasises the necessity of synaptic downscaling during offline phases (e.g. sleep) to prevent network saturation and maintain efficient plasticity [55]. The widely used STDP model, while effective on millisecond timescales, is insufficient to model long-lasting memory consolidation over hours. To address this limitation, we propose two biologically inspired plasticity functions suitable for offline synaptic scaling during sleep-like states. These functions were designed to capture the influence of earlier activity and synaptic tagging processes on the modulation of synaptic strength.

3.1.1 Function 1: Synaptic Tagging-Based Plasticity (STBP)

The first model is based on the synaptic tagging hypothesis [14, 29] and is shown in Fig. 1A. Here, potentiation and depression are determined by calcium/calmodulin-dependent protein kinase II (CaMKII) levels, which are inferred from the firing frequency in a cluster of synapses [14]. The model is symmetrical, allowing both LTD and LTP with a bounded range of ± 0.2, and incorporates a smooth transition between these phases.

We introduce a third dimension to the plasticity function that accounts for the remaining capacity of the synapse,

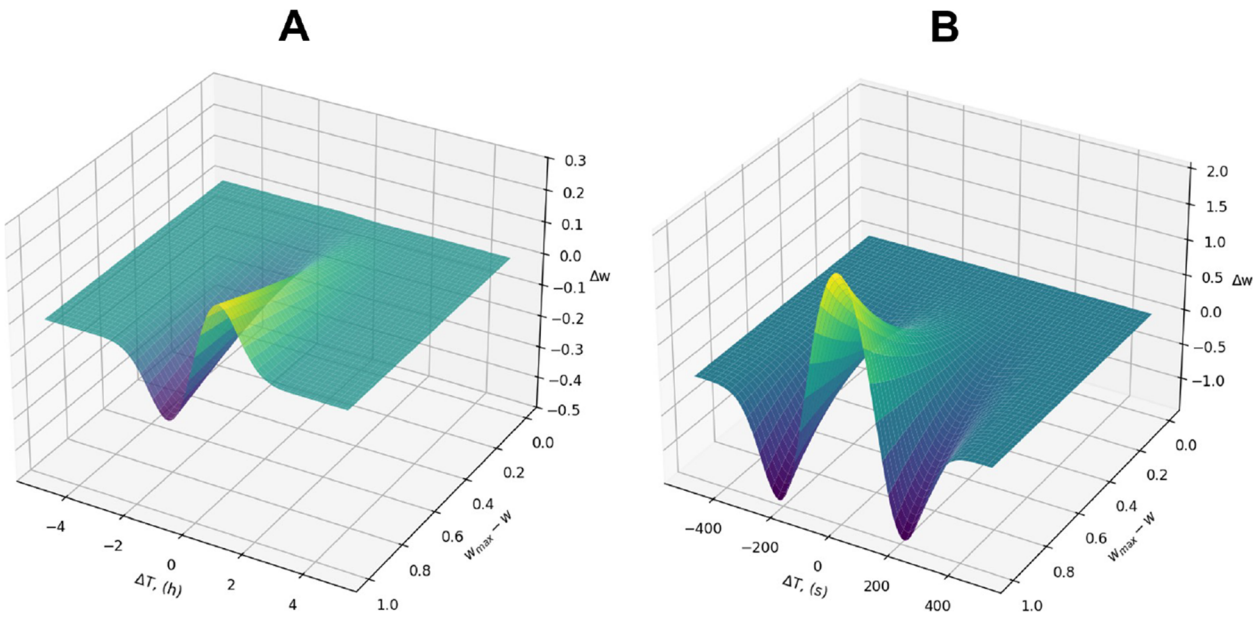

Fig. 1 The 3D function of up-and-down-scaling using: **A** action-based synaptic tagging. **B** Time-scale synaptic plasticity (BTSP) [5]. More detailed information about obtaining these figures is provided in the supplementary material

defined as $y = w_{\max} - w$, to ensure a biologically realistic saturation curve. The weight update Δw is defined as follows:

$$z(\Delta t, y) = 1.2e^{-(\Delta t^2+(y-1.4)^2)/2} \sin(\Delta t)\cos(y-1.4)$$
$$\Delta w(\Delta t, y) = \begin{cases} 1.686z(\Delta t, y), & \text{if } z(\Delta t, y) < 0 \\ 0.2z(x, y), & \text{otherwise} \end{cases} \quad (2)$$

3.1.2 Function 2: Behavioural Time Scale Synaptic Plasticity (BTSP)

The second function draws on empirical studies of BTSP mechanisms in hippocampal CA1 neurons [5]. It is shown in Fig. 3B and features a central LTP zone with LTD tails on both sides. This structure reflects behavioural observations where memory-relevant synaptic changes occur over behavioural timescales of seconds to minutes. As in the first model, we introduce a scaling component to cap the influence of plasticity as synaptic weights approach their maximum. This ensures a smooth decay instead of a hard cut-off. The update rule is defined as follows:

$$\Delta w(\Delta t, y) = 6\left(\frac{1}{\sqrt{2\pi}}e^{\Delta t^2/2} - 0.1\right)e^{-5y} + 0.2 \quad (3)$$

These two lLTP functions are complementary and designed to be deployed during offline consolidation phases triggered by experience replay. Their formulation allows for synaptic scaling that reflects both the history of activation and the proximity to capacity limits, thus preserving relevant patterns while avoiding overfitting or catastrophic forgetting.

Parameter Interpretation and Biological Grounding

The parameters used in Eqs. 2 and 3 are grounded in established models of hippocampal synaptic plasticity and calcium dynamics.

- x in Eq. 2 denotes the time offset (Δt) between presynaptic activity and dendritic plateau initiation. This variable captures the timing sensitivity of BTSP, where potentiation is strongest when synaptic input coincides closely with postsynaptic depolarisation.
- y represents the remaining plasticity potential of a synapse, calculated as $y = w_{\max} - w$. This scaling term ensures that synaptic strength approaches a saturation point gradually, mimicking biological synaptic homeostasis and preventing unbounded potentiation.
- z in Eq. 3 reflects intracellular calcium levels, a key modulator of synaptic tagging via CaMKII activation. Higher values of z indicate stronger prior activity and tagging, contributing to long-term potentiation during offline consolidation.
- **Constant values and dynamics**: Coefficients such as 1.2, 1.686, and 0.2 in Eq. 2, and 6 and 0.1 in Eq. 3, were selected to produce bounded, smooth synaptic updates consistent with experimental data on BTSP and CaMKII-mediated plasticity. The trigonometric and exponential forms encode phase sensitivity and calcium diffusion effects, respectively, approximating oscillatory replay behaviour during sleep-like states.

Together, these parameters enable the model to operate across both fast (spiking) and slow (behavioural) timescales, integrating synaptic tagging, timing-dependent plasticity, and capacity-aware scaling within a unified learning framework.

Since the 2012 object-net competition, the catastrophic forgetting problem has hindered the practical use of ANNs. The success of the competition stemmed from training on larger datasets, demonstrating that massive data and computation can produce highly performant ANNs. However, the enormous computational and data requirements made learning new models from scratch impractical for most companies, except for well-funded giants like Google and OpenAI. To address this, companies and universities can add new data to previously trained models. During training, the weights are adjusted to recognise the new data, causing previously recognised examples to be misclassified. This problem restricts most companies and all universities that lack the necessary funds from advancing the state of the art. It also limits online learning for robots. Kirkpatrick et al. [31] and Chaudhry et. al. [12] analysed the problem and proposed solutions. Despite improvements, it remains unsolved. The brain's structure differs from ANNs. It is a collection of separately evolved solvers for survival. The hippocampus structures learning differently, incrementally integrating new memories into existing ones. Our solution aims to enable robots to learn incrementally without catastrophic forgetting.

4 Online and Offline States

In cognitive science, the terms "online" and "offline" are increasingly used to describe the functional states of biological systems. The online mode refers to active perception and processing of external information, when an organism interacts with the environment, responds to sensory stimuli, makes decisions, and takes action. In contrast, the offline state is characterised by a reduction in sensory input and a predominance of internal processing of previously acquired information [59, 60]. This dichotomy is not just a convenient terminology, but reflects a deep organisational logic found across all levels of biological systems throughout evolution.

Its key function is the spatiotemporal separation of conflicting tasks that cannot be efficiently performed simultaneously (Figs. 2 and 3). For instance, at the **molecular** level, phases of gene expression (transcription) alternate with phases of stabilisation, modification, and feedback regulation, all of which require different conditions and cannot occur concurrently without interference [7]. In **neurons**, excitation is followed by a refractory period, which prevents overload and allows for precise temporal encoding of signals [25]. At the **population** level of neurons, online and offline states are organised as rhythmic oscillations that synchronise phases of sensory responsiveness and memory consolidation.

At the **behavioural** level, most animals exhibit structured cycles of activity and rest. Even invertebrates without a central nervous system, such as Cassiopea and *Hydra vulgaris*, enter circadian or ultradian rhythmic rest states that resemble sleep [28, 40]. Buzzing insects such as bees and ants alternate between phases of active exploration and return to the nest, where they appear to process spatial information and adjust future search strategies. In insects, particularly *Drosophila melanogaster*, researchers have documented not only behavioural sleep phases but also local sleep rebound effects in response to deprivation [18, 35]. Although the alternation of online and offline states is observed at all levels of biological organisation, the most vivid and well-studied manifestation of the offline regime is sleep. This state, which paradoxically combines behavioural inactivity with high neuronal activity, has been preserved throughout evolution, from invertebrates to mammals. Its universality points to its fundamental importance for survival and cognitive functioning.

In this section, we will focus specifically on sleep as the most prominent example of the offline state. However, it is important to note that similar processes can also arise during wakefulness. Even during seemingly active behaviour,

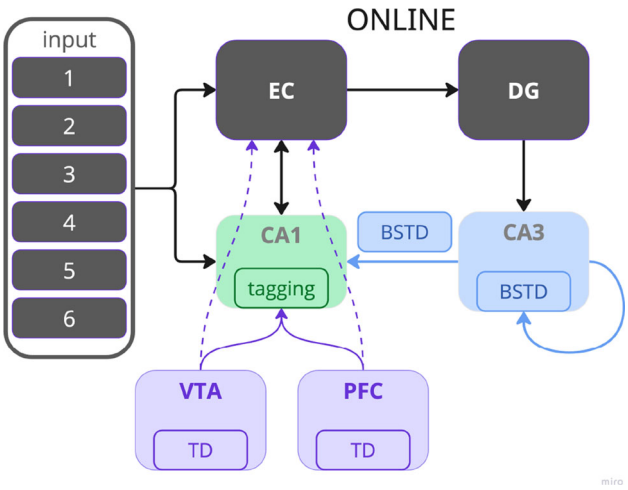

Fig. 2 The high-level schematic of the online processing, where **green** stands for synaptic tagging, **blue** for BSTD, and **purple** for TD learning

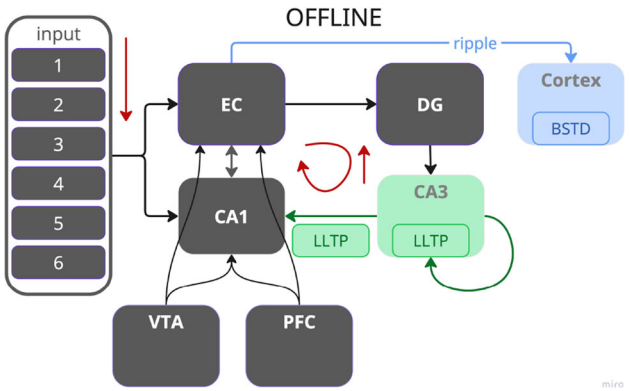

Fig. 3 The high-level schematic of the online processing, where **green** stands for synaptic tagging, **blue** for BSTD, **purple** for TD learning, and **red arrows** for increased neuronal activity

the brain can briefly switch to internal processing modes. A prime example is the activation of the default mode network (DMN), which is associated with self-reflection, planning, and recall of episodic memories. These brief episodes are especially common during rest and immobility and are linked to phenomena such as mind wandering and daydreaming. They are believed to play a key role in structuring experience and filtering information without competing with the external data stream.

Thus, the alternation between online and offline states represents a universal biological mechanism designed to architecturally separate incompatible functions. This allows living systems to avoid overload, minimise interference between cognitive processes, and efficiently manage limited resources. Understanding and replicating this logic in artificial systems opens up new possibilities for designing autonomous and adaptive architectures in which perception, learning, stabilisation, and solution generation do not conflict, but are effectively distributed over time.

4.1 Online State

During the wake phase, sensors continuously deliver a stream of information that passes sequentially through encoders, the EC, and then along the pathway DG → CA3 → CA1 toward the executive modules. The EC converts the multimodal input into sparse spiking patterns; these patterns, firing at 6–8 Hz, a frequency characteristic of exploratory behaviour in animals, are transmitted via mossy fibres to the DG. Due to strong lateral inhibition, the DG further sparsifies the code, such that only a small ensemble of pyramidal neurons becomes active in CA3.

Within CA3, recurrent collaterals rapidly adapt via the behavioural time scale plasticity (BTSP) rule: the coincidence of a presynaptic spike and a dendritic plateau within a 1-s window strengthens the connection, forming a local "here-and-now" attractor. Simultaneously, the same BTSP

conditions apply to the Schaffer collateral synapses between CA3 → CA1, where a rapid increase in synaptic weight lays down a "draft" of a new memory. At the same time, CA1 output reaches the "prefrontal controller", where a simplified dopamine-like prediction error signal is computed using the formula $\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$.

If δ_t significantly deviates from zero, the event is considered salient and is marked by two mechanisms. First, the dopaminergic pulse tags active synapses in CA1 and CA3, making them candidates for offline lLTP-based consolidation during sleep. Second, the input activity from EC to DG is amplified, leading to an increased spiking frequency in the same CA3 ensemble and raising the likelihood that the episode will be replayed during sleep.

Because CA1 receives a dual input (a prediction from CA3 and direct sensory data from EC), it continuously compares expectations with actual input. Any discrepancy maintains a high level of δ_t, thereby widening the BTSP enhancement window for precisely those synapses encoding the current context.

Thus, the online phase combines three key processes:

1. Primary learning in the CA3 ↔ CA1 loop via BTSP
2. Real-time significance evaluation via TD error
3. Tagging of synapses with dopaminergic labels that will guide selective, energy-efficient consolidation in the subsequent sleep phase

4.2 Offline State

When the input and motor neuronal activity go below threshold levels, the system automatically switches to offline mode. The EC transitions from transmitting external data to an internal generative mode: suppression of EC inputs clears the way for ripple packets initiated by the recurrent CA3 network. Each sharp wave ripple (SWR, ≈100 ms) replays a sequence of previously activated patterns and propagates it forward into the cortex, where the BTSP mechanism that operated during the day now consolidates intercortical connections but without competition from ongoing sensory input.

Within the hippocampus, with no external input, long long-term plasticity (lLTP) takes over. On the recurrent CA3 collaterals and the CA3 → CA1 pathways, global downscaling weakens all the low-weight, untagged synapses, while synapses previously marked by dopaminergic δ-tags are gradually strengthened. This combination of "cleansing" and selective consolidation turns the daytime BTSP "draft" into a stable engram, while simultaneously preventing the catastrophic buildup of synaptic weights.

If the sensory load drops only briefly (100–300 ms), a micro-nap occurs: EC disconnects for just one SWR packet, allowing to update tags, but without triggering lLTP. During full sleep (SWS/REM), across dozens or hundreds of such packets, lLTP fully stabilises the tagged connections and fades out noisy ones.

As a result, BTSP and lLTP use the same "language" of dopaminergic tagging, but operate on different timescales: BTSP provides rapid "here-and-now" corrections, while lLTP slowly and energy-efficiently rewrites meaningful content into the associative cortex, preserving previous skills and preventing forgetting.

4.3 Memory Consolidation during Sleep and the Role of the Hippocampus

System-level memory consolidation, which transfers episodic information from temporary storage to durable cortical representations, is a key cognitive mechanism of sleep. The hippocampus is central to this process, rapidly encoding "what-where-when" sequences but possessing limited capacity; excessive trace accumulation leads to interference and reduced encoding of new episodes. Sleep mitigates this limitation: Hippocampal traces are systematically replayed and deposited in more stable neocortical networks [33, 53]. Replay density is highest during SWS, where hippocampal SWRs are phase-locked to cortical slow waves and spindles [21, 41]. Such inter-structural coherence defines plastic windows in which cortical networks become maximally receptive to hippocampal input [42]. Replays occur both forward, preserving temporal spatial order, and reverse, which facilitates causal inference and goal learning [16, 47]. Reactivation of identical neuronal ensembles in all phases not only strengthens existing links, but also integrates new information into established conceptual frameworks [48]. The robustness of replay is supported by the attractor architecture of the hippocampus, particularly pronounced in CA3 [46]. Dense recurrent connectivity enables pattern completion, reconstructing whole memories from partial cues [44]. Consequently, even fragmentary activation during sleep can trigger a full-sequence replay [30]. During REM, when hippocampal activity becomes more variable, transitions between attractors become possible, fostering the generative recombination of experience and new behavioural strategies [23].

Consolidation is neither passive nor uniform. Episodes are selected according to their behavioural relevance through neuromodulatory influences: dopamine, acetylcholine, and noradrenaline modulate memory weight already at encoding [20, 49]. Highly salient traces are likelier to enter replay and long-term storage [2], whereas neutral or noisy memories are suppressed, easing network load and preserving stability [26]. Concurrently, synaptic restructuring unfolds. Global up-and-down scaling during SWS lowers overall synaptic strength and energy demand, while SWR spindle coupled activation selectively potentiates the most relevant traces [57]. This dual mechanism global weakening plus local rein-

forcement simultaneously sharpens the signal-to-noise ratio and consolidates important memories [8, 43].

5 Biologically Inspired Solutions for Robotic Systems

In biological systems, sleep is not just a period of rest, but a stage of intense cognitive activity, during which critical processes occur: memory consolidation, filtering and prioritisation of significant events, synaptic stabilisation, metabolic recovery, and even the generation of novel solutions [32, 56]. These processes are not isolated, but form a complex, coordinated offline phase that functionally and energetically complements the active sensorimotor behaviour of the organism. This offers artificial agents a powerful architectural paradigm, not replicating the brain literally, but embedding analogous phases within their operational cycles to enhance adaptability, reliability, and autonomy [11].

To approach the level of efficiency and adaptability observed in biological organisms, a rethinking of current AI architectures is required. Modern neural networks typically demand global weight updates across the entire model, even for localised changes, often relying on computationally intensive backpropagation algorithms. In contrast, biological synapses are updated locally, rapidly, selectively, and in direct response to specific events. For robotic systems, this implies a shift toward architectures that allow for targeted modifications at the level of individual connections. Such capabilities are already becoming feasible through neuromorphic approaches, such as spiking neural networks (SNNs) and memristive devices, which physically emulate the behaviour of biological synapses [17, 24, 39].

Another fundamental principle underlying biological efficiency is the avoidance of processing the entire sensory stream. In natural systems, the brain prioritises incoming information, directing cognitive resources only toward signals that are deemed relevant in the current context. A similar mechanism in robotics can be implemented at multiple levels: from primary sensory filtering to synaptic tagging, where critical connections are marked for future consolidation [13]. Key or "anomalous" events may be flagged for replay, while routine information is either ignored or stored with low priority [3].

However, instantaneous filtering alone is insufficient. For effective learning, a system must not only capture the immediate relevance of an event but also preserve the ability to reassess its importance retrospectively. This requires an additional architectural layer in which the influence of new data on model parameters is delayed, and information is first temporarily stored and then analysed in a dedicated phase. Such mechanisms prevent premature conclusions and improve adaptability in rapidly changing environments [15].

During the active phase, events and signals are accumulated in a cache, allowing the system to respond flexibly to the immediate context. However, only after this activity ends in an offline phase can the system reinterpret the collected experience: event sequences can be replayed forward and backward, causal links identified, and priorities revised. This ensures that information perceived as significant at the moment, but later found to be secondary, does not overly influence long-term model parameters. In contrast, signals that initially appeared irrelevant may acquire importance in a retrospective analysis [45].

To illustrate this, consider a classic experimental scenario: a hungry rat navigating a maze in search of cheese. Upon finding the cheese, the episode ends. However, if the rat made an apparently "useless" detour and returned to a previously visited location, this could seem inefficient at first. But if, during that detour, it encountered a water source and was later found to be thirsty, the loop acquires new value and becomes behaviourally advantageous. This type of context-dependent layered event evaluation is highly relevant to robotics, particularly in scenarios involving multiple or competing goals [6].

Replaying accumulated experience from different perspectives serves not only as a means of filtering information but also as a powerful mechanism for generating novel solutions. In biology, this function is associated with REM sleep, during which the brain actively recombines fragments of past experiences [22]. In robotics, a similar process can manifest itself as generative replay, the recombination of familiar behavioural patterns into new configurations to discover alternative strategies. This becomes especially useful when adapting to unfamiliar environments or solving tasks that have not been previously encountered [1].

Finally, the offline phase can serve not only cognitive functions, but also internal technical maintenance. When active operations are suspended, the robot can switch to a low-power and service mode, during which self-diagnostics, sensor log clean-up, hardware reconfiguration, and calibration are performed. Such processes are already being actively explored within the framework of autonomous robotic and industrial systems [54]. This brings artificial agents closer to the level of self-maintaining autonomy observed in biological organisms, where restoration and internal regulation are intrinsic to the operational cycle and do not require external intervention [52].

6 Implementation of Biologically Inspired Mechanisms in Robotic Systems

Based on the biological principles presented, we developed an architecture of the technological solution based on the following key components. **First**, we implemented an

architecture functionally similar to the hippocampus, which provides temporary storage (cache), providing the option for the offline bidirectional replay, filtering, prioritising, and integrating experience. **Second**, we proposed a modified STDP mechanism that can operate effectively on both micro (seconds-minutes) and macro (hours-days) temporal scales, ensuring adaptation and sustainable learning. **Finally**, we propose the use of neuromorphic hardware, providing an energy-efficient implementation of local plasticity and support for event-based processing.

6.1 Spiking System Architecture

Figure 4A illustrates the high-level network architecture of the hippocampus, focussing on the CA1 and CA3 subfields, which could be implemented as a spiking electronic device. Sensory inputs are initially processed in dedicated cortical areas before reaching the EC, which acts as a gateway to the hippocampal formation. The EC projects to multiple hippocampal subregions, including CA1 and CA3. In particular, CA3 receives sparse input from the DG and exhibits strong recurrent connectivity, forming an attractor network. The DG plays a critical role in pattern separation by sparsifying inputs from EC. This is achieved through random projections that reduce dimensionality while preserving structural information. A winner-take-all (k-WTA) mechanism enforces sparsity by selecting the top-k most active neurones, mimicking lateral inhibition and competitive dynamics in DG.

The CA3 region integrates contextual information, such as spatial position in place cells, via its recurrent structure. CA3 projects to CA1, which also receives direct input from EC. The dual input allows CA1 to compare stored (learnt) information from CA3 with novel input from EC. This comparison enables the detection of mismatches, which can be interpreted as novelty signals crucial for context formation and memory updating. Importantly, CA3 contributes to **pattern completion**: when presented with partial cues from EC, it can reconstruct full representations based on previously stored patterns.

Figure 4B shows a biologically inspired framework for implementing **bidirectional replay** in artificial agents [36]. Previous research has demonstrated that the recurrent circuitry of CA3 pyramidal neurones supports sequential activation patterns essential for memory encoding and retrieval [19, 36]. In our model, CA3 is represented as a **cyclic attractor network** in which place cells are interconnected via recurrent projections. This architecture naturally enables forward and reverse replays of sequential activation patterns. Depending on the initial trigger conditions, the network can produce the following:

- *Forward replay:* $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6$, triggered by activating cells 1 and 2 while inhibiting cell 6.
- *Reverse replay:* $6 \rightarrow 5 \rightarrow 4 \rightarrow 3 \rightarrow 2 \rightarrow 1$, triggered by activating cells 6 and 5 while inhibiting cell 1.

This replay control mechanism ensures directionally consistent activation by preventing feedback leakage to the opposite end of the loop. Moreover, simultaneous bidirectional replay supports optimisation and memory integration.

Although we depict a regular cyclic loop for clarity, the proposed mechanism generalises to more stochastic and biologically plausible settings. Local attractor loops can emerge through self-organisation during exploration, provided recurrent connections form among co-activated place cells.

This architectural principle provides a robust foundation for both **event-driven replay** and **temporally structured learning**, facilitating integration with downstream modules responsible for synaptic plasticity and behavioural control.

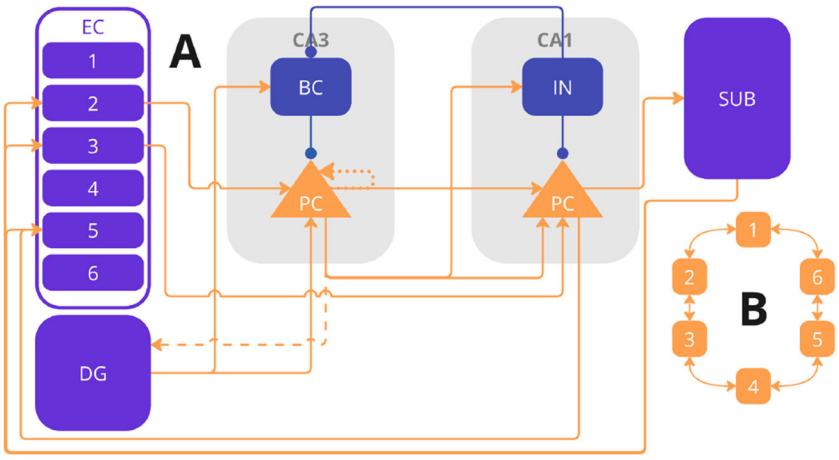

Fig. 4 **A** The CA1-CA3 attractor network architecture with feedback loops: (1) EC → DG → CA3 → CA1 → subiculum → EC; (2) CA3 feedback loops. **B** Bidirectional replay in artificial agents

6.2 Mathematical Model-Architecture Mapping

We show here a clearer mapping between mathematical formulations and the architecture of the previous Figure, providing the following correspondences:

- **Equation** 1 ($V(s_t) \leftarrow V(s_t) + \alpha\delta_t$) corresponds to the real-time temporal difference (TD) update applied within the CA1 subregion. The variable s_t denotes the agent's latent spatial state, decoded from current sensorimotor inputs via EC input. The value update mechanism is localised to CA1 due to its role in encoding reward proximity and predicting future states.
- **Equation** 2 ($\Delta w = \eta \cdot \text{BTSP}(x, y)$) reflects synaptic changes during BTSP, which is primarily implemented in the recurrent synapses of CA3. The parameters x and y denote the time lag from plateau potentials and synaptic eligibility traces, respectively, both derived from calcium kinetics and dendritic input coincidence. This process is engaged during offline replay.
- **Equation** 3 ($\Delta w = \gamma \cdot \text{CaMKII}(z)$) describes lLTP and operates in both CA3 recurrent loops and the CA3–CA1 projections. The variable z represents intracellular calcium concentration influenced by tagged synapses, and its dynamics are modelled using the activation profile of the CaMKII enzyme.
- The **figure** thus reflects a spatially grounded mapping:
  - CA1: TD learning (Eq. 1)
  - CA3 recurrent: BTSP and lLTP (Eq. 3)
  - CA3–CA1: consolidation via CaMKII-driven lLTP (Eq. 2)
  - DG: contextual separation and input encoding from EC

This mapping bridges the gap between the mathematical framework and biological analogues embedded within the neuro-robotic control architecture.

6.3 Robotic Hippocampus

The implementation of a hippocampus-inspired system based on spiking neural architectures presents a promising avenue to improve real-time learning and adaptive behaviour in autonomous robots. Of the various neuromorphic technologies under development, memristive devices stand out for their potential to support large-scale, energy-efficient spiking networks with embedded plasticity mechanisms [37, 51]. Our implementation strategy focusses on a memristive hippocampus architecture capable of replaying, prioritising, and consolidating sensorimotor experiences. However, given the current manufacturing constraints on memristor density and connectivity, we adopt a staged development pipeline that accommodates alternative spiking platforms while maintaining alignment with biological principles. Below, we present the workflow to implement this approach: **rat-model → bio-plausible digital model (b-m) → real-time digital model (rt-m) → memristive implementation (m-m)**.

The approach leads to rapid prototyping and experimentation in software and digital neuromorphic substrates (e.g. FPGAs or SNN simulators), while gradually progressing toward compact, low-power spiking hardware. Recent efforts demonstrate the feasibility of scaling spiking systems to robotic contexts using hybrid or digital implementations [17], which can serve as testbeds for algorithms prior to full memristive integration.

The target properties for the memristive hippocampal module are as follows:

- Power consumption in the milliwatt range
- Synaptic fan-in on the order of 10^4
- Neuron populations scaling to 10^6 per device
- Support for self-organised connectivity and directionally structured projections

Although memristors are central to our long-term roadmap, we remain agnostic to the specific physical substrate in the early stages. The key requirement is support for sparse, event-driven computation and localised synaptic updates capabilities, which are now achievable across several neuromorphic platforms. By leveraging hippocampus-inspired replay mechanisms and plasticity rules in an energy-aware architecture, this approach aims to bridge the gap between adaptive biological cognition and robotic autonomy in dynamic and data-scarce environments.

6.4 Online/Offline Switching and Interrupt Resilience

The transition between online (active) and offline (consolidation) states is managed by a system-level scheduler that monitors task boundaries, sensory load, and system status. In practice, this switch can be implemented through one or more of the following triggers:

- **Task-completion triggers**: Upon goal fulfillment or idle periods (e.g. navigation paused, object retrieval completed), the robot enters offline mode.
- **Temporal scheduling**: Predefined sleep cycles or downtime intervals prompt regular consolidation windows, similar to circadian-like scheduling.
- **Resource monitoring**: Thresholds on CPU, memory, or energy usage signal when the system can safely enter a low-power offline replay state.

During offline phases, cached experiences are replayed through the CA3 attractor network, and plasticity updates are applied without interfering with active control processes.

To ensure autonomy and safety in dynamic environments, the architecture includes a resilience mechanism:

- **Interrupt handling**: If a high-priority external signal (e.g. obstacle detected, emergency call, sensor spike) is received during offline replay, the robot can immediately suspend consolidation and re-engage the online control mode. Cached eligibility traces and replay queues are preserved for later resumption.
- **Failure recovery**: In the event of power loss or forced shutdown during replay, the system reverts to the last fully consolidated checkpoint stored in non-volatile memory. This ensures memory integrity and continuity in long-term learning.

This switching mechanism allows the robot to benefit from deep memory restructuring without compromising real-time responsiveness or safety in critical environments.

7 Conclusion

This study introduces a novel architectural paradigm for robotic cognition, fundamentally rethinking the organisation of learning and adaptation through the integration of biologically inspired online and offline states. Drawing from neuroscience, cognitive science, robotics, and neuromorphic engineering, we developed a synthetic hippocampal framework that enables autonomous agents not only to interact with complex environments in real time but also to perform critical memory consolidation, prioritisation, and policy recombination during dedicated offline phases.

By leveraging spiking neural networks implemented on neuromorphic substrates, combined with temporal difference learning and long-timescale plasticity mechanisms, we demonstrated that decoupling fast sensorimotor control from the slower but essential processes of structural learning is possible. Our approach directly addresses the long-standing challenges of catastrophic forgetting, rigid task dependence, and energy inefficiency that have limited the deployment of autonomous robots in dynamic and resource-constrained settings.

Importantly, this work reflects a deep interdisciplinary collaboration, bridging empirical insights from systems neuroscience and sleep research with cutting-edge methods in neuromorphic computing and robotic architecture. The proposed system does not simply replicate isolated biological features, but translates foundational organisational principles, such as replay-based consolidation and synaptic homeostasis, into a cohesive operational model for artificial agents.

The results open a systematic pathway toward the development of robots capable of true cross-contextual generalisation, self-maintenance, and generative adaptability, without incurring prohibitive computational or energetic costs. Although this work constitutes a significant advance, it also lays the groundwork for further investigation into the integration of broader biological phenomena, such as multiphase memory reprocessing and affective modulation of learning, into robotic systems.

The artificial hippocampus presented here exemplifies how interdisciplinary synthesis can yield principled, scalable solutions to the limitations of contemporary artificial intelligence, marking a pivotal step toward biologically grounded, general-purpose robotic cognition.

Although the current implementation primarily processes spatial and visual information, the architecture is designed to be extensible to multiple sensory modalities. The EC layer in our model acts as a convergence hub for preprocessed sensory streams, meaning that auditory, tactile, proprioceptive, or even semantic signals can be routed through analogous pathways.

Multimodal inputs can be encoded using dedicated preprocessing modules (e.g. cochlea-like spiking encoders for sound, somatosensory maps for touch) and projected into a unified latent representation. These signals are then passed to the DG for pattern separation and embedded into hippocampal replay sequences.

The plasticity mechanisms (TD, BTSP, lLTP) are agnostic to the input modality, as they rely on temporal dynamics, salience tagging, and sequence structure. This allows the system to learn cross-modal associations and generalise across diverse sensorimotor tasks, paving the way for deployment in complex, real-world robotic environments involving heterogeneous and noisy data streams.

Supplementary Information

The online version contains supplementary material available at the following links:
https://github.com/max-talanov/1/blob/master/bypass/3DBTSP.py and https://github.com/max-talanov/1/blob/master/bypass/3DCaDP.py

Acknowledgements The authors express their gratitude to Michele Migliore for his contributions to the design of the memory consolidation functions and to Mansur Ziiatdinov for his contributions to the mathematical foundations of these functions.

Author Contribution M.T., X.F., J.V., and P.R. conceived and designed the study, analysed the data, and wrote the main text of the manuscript. I.K. and M.T. provided crucial support in the development and implementation of the memory consolidation functions. All authors reviewed and approved the final version of the manuscript.

Funding Open Access Funding provided by Universitat Autonoma de Barcelona. This study was partially funded by the Fairground project "Artificial and Bio-Inspired Networked Intelligence for Constrained Autonomous Devices" (Fairground Bando A Cascata A Valere Sul

Piano Nazionale Ripresa E Resilienza (PNRR) Missione 4, Istruzione E Ricerca-Componente 2, Dalla Ricerca All'impresa-Linea Di Investimento 1.3, Finanziato Dall'Unione Europea NextGenerationEU, Progetto Future Artificial Intelligence Fair PE0000013 CUP (Master) J53C22003010006 CUP J43C24000230007). The author(s) acknowledge the support of the APC central fund of the University of Messina. Prof. Vallverdu's research is supported by ICREA 2019 and PALGAIT (PID2023-148336NB-100) grants.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Ethics Approval Not applicable

Informed Consent None.

Conflict of Interest None.

Research Involving Humans and Animals Statement None.

References

1. Ahn, M., Brohan, A., Chebotar, Y., et al. (2022). Do as I can, not as I say: Grounding language in robotic affordances. arXiv:2204.01691
2. Atherton, L. A., Dupret, D., & Mellor, J. R. (2015). Memory trace replay: The shaping of memory consolidation by neuromodulation. *Trends in Neurosciences, 38*(9), 560–570. https://doi.org/10.1016/j.tins.2015.07.004
3. Atherton, L. A., Dupret, D., & Mellor, J. R. (2015). Memory trace replay: The shaping of memory consolidation by neuromodulation. *Trends in Neurosciences, 38*(9), 560–570.
4. Bawa, M., & Raman, R. (2024). Taking control: Steering the future of biohybrid robots. *Science Robotics, 9*(45), eadr9299.
5. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., & Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science, 357*(6355), 1033–1036. https://doi.org/10.1126/science.aan3846. Publisher: American Association for the Advancement of Science
6. Bittner, S. R., Milstein, A. D., Grienberger, C., Romani, S., & Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science, 357*(6355), 1033–1036.
7. Blackledge, N. P., & Klose, R. J. (2021). The molecular principles of gene regulation by polycomb repressive complexes. *Nature Reviews Molecular Cell Biology, 22*(12), 815–833. https://doi.org/10.1038/s41580-021-00398-y
8. Blanco, W., Pereira, C. M., Cota, V. R., Souza, A. C., Rennó-Costa, C., Santos, S., Dias, G., Guerreiro, A. M. G., Tort, A. B. L., Neto, A. D., & Ribeiro, S. (2015). Synaptic homeostasis and restructuring across the sleep-wake cycle. *PLOS Computational Biology, 11*(5), Article e1004241. https://doi.org/10.1371/journal.pcbi.1004241
9. Brodt, S., Inostroza, M., Niethard, N., & Born, J. (2023). Sleep—A brain-state serving systems memory consolidation. *Neuron, 111*(7), 1050–1075.
10. Brohan, A., Brown, N., Chebotar, Y., et al. (2023). RT-2: Vision-language-action models transfer web knowledge to robotic control. arXiv:2307.15818
11. Buzsáki, G. (2019). *The brain from inside out*. Oxford University Press
12. Chaudhry, A., Dokania, P. K., Ajanthan, T., & Torr, P. H. (2018). Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: *European Conference on Computer Vision (ECCV)* (pp. 532–547). Springer
13. Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L., & Gerstner, W. (2008). Tag-trigger-consolidation: A model of early and late long-term-potentiation and depression. *PLoS Computational Biology, 4*(12), e1000248.
14. Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L., & Gerstner, W. (2008). Tag-trigger-consolidation: A model of early and late long-term-potentiation and depression. *PLoS Computational Biology, 4*(12), e1000248. https://doi.org/10.1371/journal.pcbi.1000248
15. Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron, 63*(4), 497–507.
16. Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron, 63*(4), 497–507. https://doi.org/10.1016/j.neuron.2009.07.027
17. Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., et al. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro, 38*(1), 82–99.
18. Donlea, J. M., Ramanan, N., & Shaw, P. J. (2014). Neuronal replay of learned behavior in the drosophila mushroom body. *Neuron, 83*(5), 998–1010.
19. Ecker, A., Bagi, B., Vértes, E., Steinbach-Németh, O., Karlócai, M. R., Papp, O. I., Miklós, I., Hájos, N., Freund, T. F., Gulyás, A. I., & Káli, S. (2022). Hippocampal sharp wave-ripples and the associated sequence replay emerge from structured synaptic interactions in a network model of area CA3. *eLife, 11*, e71850. https://doi.org/10.7554/eLife.71850. Publisher: eLife Sciences Publications, Ltd
20. Feld, G. B., & Born, J. (2019). Neurochemical mechanisms for memory processing during sleep: Basic findings in humans and neuropsychiatric implications. *Neuropsychopharmacology, 45*(1), 31–44. https://doi.org/10.1038/s41386-019-0490-9
21. Fernandez, L. M. J., & Lüthi, A. (2020). Sleep spindles: Mechanisms and functions. *Physiological Reviews, 100*(2), 805–868. https://doi.org/10.1152/physrev.00042.2018
22. Findlay, G., Tononi, G., & Cirelli, C. (2020). The evolving view of replay and its functions in wake and sleep. *Sleep Advances, 1*(1), zpab002.
23. Findlay, G., Tononi, G., & Cirelli, C. (2020). The evolving view of replay and its functions in wake and sleep. *Sleep Advances, 1*(1). https://doi.org/10.1093/sleepadvances/zpab002
24. Hu, J., Niu, H., Carrasco, J., Lennox, B., & Arvin, F. (2023). Challenges and solutions for autonomous ground robot scene understanding and navigation in unstructured outdoor environments: A review. ResearchGate
25. Huang, C. L. H., & Lei, M. (2023). Cardiomyocyte electrophysiology and its modulation: Current views and future prospects. *Philosophical Transactions of the Royal Society B: Biological Sciences, 378*(1879). https://doi.org/10.1098/rstb.2022.0160
26. Isotalus, H., Carr, W. J., Averill, G. G., Radtke, O., Selwood, J., Williams, R., Ford, E., McCullagh, L., McErlane, J., O'Donnell, C., Durant, C., Bartsch, U., Jones, M. W., Muñoz-Neira, C., Wearn, A. R., Grogan, J. P., & Coulthard, E. J. (2020). *Dopamine-gated*

memory selection during slow wave sleep. https://doi.org/10.1101/2020.05.23.112375
27. Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (2011). Hierarchical reinforcement learning: A survey. *Journal of Artificial Intelligence Research, 4*, 1–37.
28. Kanaya, H., Abe, H., Katsukawa, M., & Hirose, S. (2020). A sleep-like state in hydra unravels conserved sleep mechanisms during the evolutionary expansion of the central nervous system. *PNAS, 117*(46), 28982–28990.
29. Karaba, L. A., Robinson, H. L., Harvey, R. E., Chen, W., Fernandez-Ruiz, A., & Oliva, A. (2024). A hippocampal circuit mechanism to balance memory reactivation during sleep. *Science (New York, N.Y.), 385*(6710), 738. https://doi.org/10.1126/science.ado5708. https://pmc.ncbi.nlm.nih.gov/articles/PMC11428313/
30. Khona, M., & Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience, 23*(12), 744–766. https://doi.org/10.1038/s41583-022-00642-0
31. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwińska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences, 114*(13), 3521–3526. https://doi.org/10.1073/pnas.1611835114. https://www.pnas.org/doi/10.1073/pnas.1611835114
32. Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience, 22*(10), 1598–1610.
33. Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience, 22*(10), 1598–1610. https://doi.org/10.1038/s41593-019-0467-3
34. Li, M., Zhao, Y., & Liu, J. (2024). Bioinspired liquid metal-based soft humanoid robots. *Advanced Materials, 36*(12), Article e202404330.
35. Liu, Q., Liu, S., Kodama, L., Driscoll, M. R., & Wu, M. N. (2012). Dopamine neurons modulate sleep through arousal-promoting peptidergic neurons in drosophila. *Current Biology, 22*(22), 2114–2123.
36. Matheus Gauy, M., Lengler, J., Einarsson, H., Meier, F., Weissenberger, F., Yanik, M. F., & Steger, A. (2018). A hippocampal model for behavioral time acquisition and fast bidirectional replay of spatio-temporal memory sequences. *Frontiers in Neuroscience, 12*, 5. https://doi.org/10.3389/fnins.2018.00961. https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2018.00961/full. Publisher: Frontiers
37. Mikhaylov, A., Pimashkin, A., Pigareva, Y., Gerasimova, S., Gryaznov, E., Shchanikov, S., Zuev, A., Talanov, M., Lavrov, I., Demin, V., et al. (2020). Neurohybrid memristive cmos-integrated systems for biosensors and neuroprosthetics. *Frontiers in Neuroscience, 14*, 358.
38. Moses, M. (2024). Robotics in disaster response: Enhancing search and rescue operations. *Research Invention Journal of Biological and Applied Sciences, 3*(3), 51–56.
39. Naderi, H., Shojaei, A., & Huang, L. (2024). Foundation models for autonomous robots in unstructured environments. arXiv:2407.14296
40. Nath, R. D., Bedbrook, C. N., Abrams, M. J., Basinger, T., Bois, J. S., Prober, D. A., Sternberg, P. W., Gradinaru, V., & Goentoro, L. (2017). The jellyfish Cassiopea exhibits a sleep-like state. *Current Biology, 27*(19), 2984–2990.
41. Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience, 19*(21), 9497–9507. https://doi.org/10.1523/jneurosci.19-21-09497.1999
42. Peyrache, A., & Seibt, J. (2020). A mechanism for learning with sleep spindles. *Philosophical Transactions of the Royal Society B: Biological Sciences, 375*(1799), 20190230. https://doi.org/10.1098/rstb.2019.0230
43. Puentes-Mestril, C., Roach, J., Niethard, N., Zochowski, M., & Aton, S. J. (2019). How rhythms of the sleeping brain tune memory and synaptic plasticity. *Sleep, 42*(7). https://doi.org/10.1093/sleep/zsz095
44. Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience, 7*. https://doi.org/10.3389/fnsys.2013.00074
45. Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience, 7*, 74.
46. Rolls, E. T. (2013). A quantitative theory of the functions of the hippocampal ca3 network in memory. *Frontiers in Cellular Neuroscience, 7*. https://doi.org/10.3389/fncel.2013.00098
47. Roumis, D. K., & Frank, L. M. (2015). Hippocampal sharp-wave ripples in waking and sleeping states. *Current Opinion in Neurobiology, 35*, 6–12. https://doi.org/10.1016/j.conb.2015.05.001
48. Sadowski, J., Jones, M., & Mellor, J. (2016). Sharp-wave ripples orchestrate the induction of synaptic plasticity during reactivation of place cell firing patterns in the hippocampus. *Cell Reports, 14*(8), 1916–1929. https://doi.org/10.1016/j.celrep.2016.01.061
49. Samanta, A., Alonso, A., & Genzel, L. (2020). Memory reactivations and consolidation: Considering neuromodulators across wake and sleep. *Current Opinion in Physiology, 15*, 120–127. https://doi.org/10.1016/j.cophys.2020.01.003
50. Sarker, A., Islam, T. U., & Islam, M. R. (2024). A review on recent trends of bioinspired soft robotics: Actuators, control methods, materials selection, sensors, challenges, and future prospects. *Advanced Intelligent Systems, 6*(7), 2400414.
51. Serb, A., Bill, J., Berdan, R., et al. (2020). Memristive synapses for neuromorphic computing: From devices to systems. *Philosophical Transactions of the Royal Society A, 378*(2164), 20190156.
52. Shridhar, M., Manuelli, L., & Fox, D. (2021). Cliport: What and where pathways for robotic manipulation. arXiv:2109.12092
53. Singh, D., Norman, K. A., & Schapiro, A. C. (2022). A model of autonomous interactions between hippocampus and neocortex driving sleep-dependent memory consolidation. *Proceedings of the National Academy of Sciences, 119*(44). https://doi.org/10.1073/pnas.2123432119
54. Thrun, S. (2002). Robotic mapping: A survey. *Exploring artificial intelligence in the new millennium*.
55. Tononi, G., & Cirelli, C. (2003). Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin, 62*(2), 143–150.
56. Tononi, G., & Cirelli, C. (2014). Sleep and the price of plasticity: From synaptic and cellular homeostasis to memory consolidation and integration. *Neuron, 81*(1), 12–34. https://doi.org/10.1016/j.neuron.2013.12.025. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3921176/
57. Vyazovskiy, V. V., & Faraguna, U. (2014). *Sleep and synaptic homeostasis* (pp. 91–121). Springer Berlin Heidelberg. https://doi.org/10.1007/7854_2014_301
58. Walter, M. R., Patki, S., Daniele, A. F., et al. (2021). Language understanding for field and service robots in a priori unknown environments. arXiv:2105.10396
59. Wang, K., Yu, C., Xu, L., Qin, W., Li, K., Xu, L., & Jiang, T. (2009). Offline memory reprocessing: Involvement of the brain's default network in spontaneous thought processes. *PLoS ONE, 4*(3), e4867. https://doi.org/10.1371/journal.pone.0004867
60. Yeshurun, Y., Nguyen, M., & Hasson, U. (2021). The default mode network: Where the idiosyncratic self meets the shared social world. *Nature Reviews Neuroscience, 22*(3), 181–192. https://doi.org/10.1038/s41583-020-00420-w
61. Zhang, L., Wang, W., & Chen, H. (2025). Bioinspired and biohybrid soft robots: Principles and emerging trends. *Matter, 8*(5), 1234–1256.
62. Zhao, Y., & Han, J. (2024). Offline supervised learning vs online direct policy optimization: A comparative study and a unified train-

ing paradigm for neural network-based optimal feedback control. *Physica D: Nonlinear Phenomena, 462*, 134130.