


Active Inference for Learning and Development in Embodied Neuromorphic Agents

Sarah Hamburg *, Alejandro Jimenez Rodriguez, Aung Htet and Alessandro Di Nuovo 

Department of Computing, Sheffield Hallam University, Sheffield S1 1WB, UK; a.jimenez-rodriguez@shu.ac.uk (A.J.R.); aung.htet@student.shu.ac.uk (A.H.); a.dinuovo@shu.ac.uk (A.D.N.)
* Correspondence: s.hamburg@shu.ac.uk

Abstract: Taking inspiration from humans can help catalyse embodied AI solutions for important real-world applications. Current human-inspired tools include neuromorphic systems and the developmental approach to learning. However, this developmental neurorobotics approach is currently lacking important frameworks for human-like computation and learning. We propose that human-like computation is inherently embodied, with its interface to the world being neuromorphic, and its learning processes operating across different timescales. These constraints necessitate a unified framework: active inference, underpinned by the free energy principle (FEP). Herein, we describe theoretical and empirical support for leveraging this framework in embodied neuromorphic agents with autonomous mental development. We additionally outline current implementation approaches (including toolboxes) and challenges, and we provide suggestions for next steps to catalyse this important field.

Keywords: active inference; neurorobotics; developmental robotics



Citation: Hamburg, S.; Jimenez Rodriguez, A.; Htet, A.; Di Nuovo, A. Active Inference for Learning and Development in Embodied Neuromorphic Agents. *Entropy* **2024**, *26*, 582. <https://doi.org/10.3390/e26070582>

Academic Editor: Alessandro Giuliani

Received: 27 May 2024

Revised: 23 June 2024

Accepted: 27 June 2024

Published: 9 July 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Embodied Agents: Inspiration from Humans

One major limitation of mainstream machine learning is the absence of a body to support self-determined learning via autonomous interaction with the environment [1]. This limitation is intrinsic to technologies like neural networks because they are inspired only by the neural components of the complex interacting subsystem that constitute a living being. Embodied agents, on the other hand, have a physical implementation (such as a body, drone, or vehicle), and exist on a continuum from weak to strong forms of embodiment (see [2] for discussion).

Neuropsychology research has extensively evidenced the embodied nature of human cognitive abilities. These are formed not only within the brain but also shaped by the body and the experiences acquired through it, both through actions, such as manipulatives, gestures, and movements [3–5], and through sensations, such as touch [6] and odours. Even purely mental phenomena like mental simulation are shaped by embodied experiences and, in turn, shape different aspects of human behaviour [7].

Importantly, the human brain is subject to energy constraints and is considered highly energy efficient (estimated power consumption is 20 Watts per day [8]—the equivalent of one energy-efficient lightbulb). For context, the current fastest supercomputer in Europe has been described as “exceptionally green” and consumes 8.5 million Watts per day [9]. While supercomputers consume multiple orders of magnitude more energy than the human brain, it is precisely the energy constraints of the latter that give rise to many of its computational properties and, ultimately, to the cognitive abilities that make us human [10].

Cognitive domains that humans excel in include learning with minimal data (aided via predictions based on mental imagery models), abstract thinking (understanding concepts that are not immediately present or tangible), creativity, problem solving, adaptability, reasoning, decision making, and collaboration (e.g., [11–14]). Importantly, these capacities

are enabled by the relatively large amount of data a human is exposed to throughout development, on one hand, combined with extensive real-time multimodal sensory input (i.e., through the skin). These constitute two novel scaling frontiers yet to be exploited by any learning algorithm to date.

Inspired by the brain, neuromorphic computing aims to create powerful energy-efficient cognitive systems and has been described as a “key enabling technology for the development of a unique generation of autonomous agents” [15]. Neuromorphic systems include plastic “synapses” (i.e., mediated by spike-timing-dependent plasticity (STDP)), and dynamic and “event-based” computing paradigms (for a review, see [16,17]). Efficiency in this context is relative to the particular energy constraints imposed by the environment and the body itself.

For instance, synapses are dynamic elements that are up- or down-scaled (i.e., pruned) depending upon said energy constraints. Extensive synaptic pruning is a crucial neurodevelopmental process whereby about 50% of the brain’s synapses are removed throughout early childhood into adulthood—enhancing energy efficiency [18]. Myelination is also a fundamental neurodevelopmental process, whereby the nervous system becomes increasingly energy efficient through the “insulation” of specific neurons at critical timepoints during development, in a coordinated pattern aligned with developmental requirements.

The dynamic aspect of computation allows for continual learning and adaptation and facilitates the agent’s coupling to environmental rhythms. In particular, the temporal architecture of the brain becomes more synchronised over neurodevelopment [19]. Finally, the event-based nature of computation allows for the agent to be acted upon by the environment (for example, it is very difficult to avoid an odour).

Tackling Developmental and Perceptual Scaling Limits

Research from developmental sciences has shown that children acquire cognitive skills—including perceptual [20], language [21], social [22], and numerical [23]—through “embodied experiences” (i.e., motor movements and interactions with objects and people). Emulating these embodied learning mechanisms in artificial agents may, therefore, promote the acquisition of human-like abilities. This developmental approach aligns with the 1950 hypothesis by Alan Turing, “Instead of trying to produce a program to simulate the adult mind, why not rather try to produce one which simulates the child’s? If this were then subjected to an appropriate course of education one would obtain the adult brain” [24].

How best to enable embodied agents to “learn like a human” remains highly conceptual. The combination of developmental robotics and neuromorphic computing is leading to the emerging field of developmental neurorobotics (see section below), which aims to inform our understanding of human-like learning mechanisms for embodied neuromorphic systems. These aspects of embodied learning need to be necessarily neuromorphic due to the importance of distributed perceptual systems like the skin and the intrinsic relationship between development and neural remodelling in the brain [18].

We propose the active inference framework (AIF) as an important tool for implementing learning mechanisms in embodied neuromorphic agents. Originating from neuroscience, the AIF offers a biologically plausible and unified explanation for how the brain processes information, learns, and generates behaviour, involving embodied perception–action loops [25]. This includes solving “hard exploration problems” and accounting for mechanisms of natural agency and behaviour [26]. Importantly, neuromorphic computing systems are thought to be well suited for implementing AIF principles, while in turn, AIF drives the development of neuromorphic architectures [27]. Even though these connections remain only theoretical, there has been considerable progress in process theories that help bridge the gap between theory and implementation of active inference [28].

2. Developmental Neurorobotics

Traditional machine learning (ML) to date is not well suited to simulate continuous real-world embodied learning. In particular, traditional ML:

- Requires extensive training, computation, memory, and energy—impacting scalability.
- Does not cope well with noise, variability, and uncertainty—impacting real-world applications.
- Has difficulties generalising across tasks and environments.
- Lacks common sense—it is not able to infer, understand, or explain [29].
- Lacks sufficient transparency [30].
- Has poor performance on tasks requiring embodied intelligence [15].

This state of affairs is connected with the existing paradigms of training those systems and perhaps to the concept of training itself.

Developmental neurorobotics is an interdisciplinary research paradigm combining computational modelling, developmental psychology, and robotics to realise an embodied artificial intelligence [31]. A pervasive paradigm in this field is Piaget’s theory of cognitive development in children (1950s; [32]). In accordance with this theory, the developmental neurorobotics approach emphasises self-determined learning via interaction with the environment. Embodied agents build models (i.e., learn) based on their own interactions with the world, rather than relying on off-the-shelf pre-trained models. Robots must continuously infer, predict and adapt, and acquire new knowledge with only weak- or self-supervision (i.e., socially guided or autonomous). This helps robots to interact effectively and dynamically with their environment across a range of tasks.

It is anticipated that developmental neurorobotics systems will better serve humanlike learning—including dynamic adaptation and generalisation across tasks. In turn, such systems will also serve as new tools for investigating human development and cognition. However, the practical considerations of this paradigm remain a challenge in the field. The questions of how to set up an effective learning loop and how to exploit the different timescales of development necessitate novel approaches that extend the gradient descent algorithmic view that dominates the field.

As yet, there are no widely accepted or unified learning frameworks employed in developmental neurorobotics systems. These systems are complex, involving the integration of different modalities and learning over long timescales across different tasks. Frameworks underlying such learning may, therefore, help guide the design of these complex systems. We propose that the active inference framework (AIF) may offer such an opportunity.

3. Active Inference: A Promising Framework for Learning in Embodied Neuromorphic Agents

3.1. Active Inference

We aim here to provide a high-level overview of AIF relevant to developmental neurorobotics. For an in-depth explanation and discussion of AIF, see [33]. We additionally provide a summary box of key concepts, benefits, and challenges of AIF.

Within AIF, the brain models the world as a set of probabilities (a generative probabilistic model), which it uses to make inferences about the world and predict what is likely to happen next. The agent actively works to minimise “surprise” by bringing predicted and observed states of the world into alignment. (“Surprise” is a measure of uncertainty about the world, taking into account the quality of data. Free energy is an upper bound to surprise). Surprise minimisation is achieved by (1) adjusting models (i.e., altering perception), and/or (2) selecting actions that aim to maximise information gain and minimise prediction errors (e.g., turning your head towards an unknown noise).

AIF involves a continual loop of perception, prediction, and action. Over short timescales, perception optimises beliefs about the current state of the world (variational free energy—VFE), while over long timescales, learning optimises beliefs about the relationships between the variables that constitute the world (expected free energy—EFE) [34]. Both occur through the minimisation of “free energy” over observations in the VFE or over actions in the EFE. The ability to account for both short- and long-term learning, in addition to the hierarchical nature of the framework (i.e., deep active inference), is particularly

advantageous for neurodevelopmental frameworks, which operate over and integrate these different timescales in learning.

Compared to reinforcement learning (RL), AIF offers a more integrated view of perception and action, along with more flexible goal-setting based on prior preferences. In RL, the reward function defines an agent's goal and allows it to learn how to best act within the environment to maximise an expected reward. In AIF, any type of outcome may be more or less preferred—the implicit reward is a feature of the agent in conjunction with the environment it inhabits (for discussion, see [35]). AIF bypasses problems associated with defining reward functions (which can be difficult, particularly for real-world tasks [26]) and instead replaces these with prior beliefs about preferred outcomes. Agents are able to learn their own prior preferences and goals are flexible. In this way, AIF extends RL, encourages exploration and information seeking, and equips agents with intrinsic curiosity [36,37].

What is Active Inference?

Active inference is a framework for understanding how systems autonomously perceive, learn, and act. Rooted in Bayesian inference, it combines elements of perception, action, and learning into a unified theory. Originating from neuroscience, it is increasingly applied in machine learning and robotics.

Key Concepts:

Bayesian Inference: This is used to update beliefs about the world based on new sensory information. Beliefs are based on prior knowledge and likelihood. AIF performs approximate Bayesian inference.

Generative Models: Sensory input is predicted based on internal representations of the world. The goal of active inference is to minimise a measure of the difference between predicted and actual sensory input (prediction errors).

Free Energy Principle: This central principle posits that biological systems act to minimise a function called "free energy".

Variational Free Energy: VFE quantifies the difference between the system's internal model and the actual data observed. Minimising this ensures the model accurately reflects the observed data.

Expected Free Energy: EFE predicts free energy that will be encountered under different possible actions or action sequences. Minimising this leads to optimal decision making.

Markov Blankets: These represent conceptual boundaries (in terms of conditional independence) that separate a system from its environment. Systems update beliefs and make predictions about their environment based solely on the information contained within the blanket.

Partially Observable Markov Decision Processes: POMDPs provide a mathematical framework for modeling decision making with partial information in active inference. The free energy principle guides action selection and belief updating within this framework.

Action and Perception: These phenomena are intertwined. Actions are performed to reduce the surprise with respect to generative models (i.e., to reduce prediction error). This contrasts with traditional views that separate perception (passive) and action (active).

Learning and Adaptation: Systems continuously update their generative models to improve their predictions, allowing for adaptation to changing environments. In this case, learning is seen as inference on the parameters of the generative model.

Key Benefits:

Unified Framework: The unified framework integrates perception, action, and learning into a single model, enabling a unified approach to creating intelligent and self-organising systems.

Adaptive Learning: This continuous updating of beliefs and models allows for real-time adaptation to new information.

Embodied and Situated: Perception and cognition are intertwined with the agent's body (embodied) within a particular environment (situated).

Key Challenges:

Complexity: The mathematical and computational complexity of implementing active inference can be high.

Scalability: Scaling active inference models to large and complex applications remains an ongoing challenge.

3.2. Key Features of Embodied Learning and Development in the Light of AIF

There is both theoretical and empirical support for AIF as a promising tool for learning and development in embodied neuromorphic agents. First, we highlight key features of embodied learning and development [32]:

- Learning is cumulative and progresses in complexity.
- Concrete and abstract concepts are a continuum; both are learned by linking concepts to embodied perceptions [38].
- Learning results from self-exploration with the world, often in combination with social interaction.
- The importance of sensorimotor skills (including the discovery of one's own body), communication skills, and social skills.

We have additionally summarised the current literature outlining the requirements for embodied "intelligent" agents (Table 1).

Table 1. Summary of requirements for embodied "intelligent" agents (Sources: [15,29,39–43]).

Domain	Requirements
Cognition	Predictive, flexible, brain–body–environment model, incorporating value systems.
Computation	Sparse representations, efficient, morphological, incorporating information-processing biases.
Learning	Continual, open-ended, hierarchical multi-task, generalisable (including zero-shot), leveraging plasticity mechanisms.

Krichmar [40] stated that neurorobotics is the ideal methodology to address brain-like features and called for the community to focus on general cognition rather than particular behaviours. Additional emphasis is placed on the ability to continuously infer, predict, and adapt (e.g., [15]).

3.3. Theoretical Support for Active Inference in Embodied Neuromorphic Agents

Embodiment is a key feature of AIF [44,45]—perception and cognition are deeply situated and intertwined in the embedded context of the agent and its environment [46]. In AIF, the brain has even been described as "taking a back seat to the body" [46]. This is similar to the concept of morphological computation (whereby certain processes are performed by a robot's body with minimal neural control [47]). The importance of morphological computation in neurorobotics was emphasised by Krichmar [40].

There is also an emphasis on brain–body–environment interactions. In AIF, classic physical world distinctions between "agent" and "environment" do not exist. Instead, the statistical tool of Markov blankets is employed in models to represent boundaries between systems with an external and internal state. Sometimes blankets correspond to a physical boundary (e.g., a cell exchanging energy in a tissue), while other times they do not (e.g., external force in a moving pendulum) [48]. This particular conceptual aspect fits the requirements of the computational needs required to exploit the extent of distributed sensory systems like the skin.

Furthermore, the drive to reduce uncertainty, which underpins the AIF, has been described as comparable to curiosity [34]. The emphasis on curiosity-driven embodied activity in AIF suggests it may offer a useful framework for self-supervised learning. Curiosity is widely accepted as a powerful driver of cognitive development in humans.

Da Costa et al. [49] described AIF as an interesting framework for robotic applications because it unifies state-estimation, control, and world model learning as inference processes that are solved by optimising a single objective function—free energy. Furthermore, it endows robots with adaptive capabilities, which are central to real-world applications. According to [49], key features of AIF which may address current challenges in robotics include:

- Accurate and robust state tracking, including the integration of multiple sensory streams.
- Adaptive model-based and shared control.
- Learning and grounding, including learning from sparse and noisy observations and organisation of knowledge into hierarchical modular representations.
- Operational specification, safety, transparency, and explainability.
- Energy efficiency, qualified by the considerations mentioned above.

Based on these properties, AIF may lead to the development of robots that are context adaptive, safe, social, and collaborative [49]. The outlined properties of AIF also confer benefits for regulated use cases (e.g., medical applications), wearables and exoskeletons with a degree of autonomy, and neurotechnologies integrated with the human nervous system [49]. Additionally, emphasising the role of principles (like the FEP) beyond individual algorithms appears to be an attractive solution for the satisfaction of abstract norms of behaviour like social norms.

Efficiency is also emphasised by the AIF. According to [25], optimising a generative model is comparable to optimisation under the efficient coding principle. Furthermore, the FEP has been shown to encompass other principles like the efficient coding hypothesis [25]. Efficiency, in this case, is a contingent aspect because even though the FEP is theoretically efficient from the mentioned points of view, current implementations still rely on gradient descent optimization and are, nevertheless, inefficient computationally. This is bound to change with emerging neuromorphic approaches and implementation.

Conceptually, it was claimed by Lanillos et al. [36] that AIF could potentially generate high-order cognitive and metacognitive capabilities, such as monitoring, self explainability, and to some degree, “awareness”. The authors claimed this is because the model does not just predict sensations, but also their predictability, enabling robots to monitor their precision about prediction, and thus, become self-attentive.

Finally, both AIF and neuromorphic computing originate from neuroscience. It is claimed that the emphasis on neuroscience in AIF reduces the gap between engineering and life sciences [49] and makes it possible to bridge connections between neuroscience, robotics, AI, and psychology [34]. Thus far, much of the work in neuromorphic computing has focused on hardware development [17]. With regard to neuromorphic computing algorithms, [17] outline the clear benefits for neural network-style computation. Together, this suggests that AIF may be particularly suited to neuromorphic agents.

3.4. Empirical Support for Active Inference in Embodied Neuromorphic Agents

AIF has been shown to perform as well as traditional ML methods in some simple environments, and research indicates it may perform better in environments featuring volatility, ambiguity, and context sensitivity [49]. Learning has included the ability to generalise prior knowledge to new stimuli, resulting in a “one-shot learning” capacity qualitatively similar to that observed in humans [34].

Implementations of AIF in embodied systems have included simulated robot arms for searching, reaching, and manipulating [50–52]; a model for estimation and control in a humanoid robot [53,54]; and multisensory body perception and action in a humanoid robot [55].

The latter of these studies used proprioceptive and visual input and demonstrated natural behaviours for upper body reaching and head object tracking. The AIF algorithm was validated in terms of noise robustness and multisensory integration and was also applied for reaching and grasping a moving object [55]. The results show that closed-

loop adaptation and multisensory integration are two important characteristics of the AIF approach [55].

Together, studies have demonstrated the benefits of AIF for robot body estimation [53,55], navigation [56,57], fault-tolerant behaviour [58], planning [36], self/other distinction [54], and complex social cognition [59]. AIF appears particularly useful for applications where the dynamics of the robot and/or the task are uncertain [36].

Within neuromorphic research, Isomura et al. [60] recently used a biologically plausible “canonical neural network” architecture to explore AIF. The authors cast the cost functions of these networks as variational free energy under an implicit generative model. Results suggested that the delayed modulation of Hebbian plasticity (a process integral to forming associations between neurons in delayed reward tasks) is a realisation of active inference. The authors suggest that casting cost functions in this way may dramatically reduce the complexity of designing self-learning neuromorphic hardware by offering a simple architecture with low computational cost. Such human-like plasticity is also likely advantageous for neurodevelopmental approaches.

Gandolfi et al. [61] recently explored AIF principles in a biologically plausible neuromorphic system built with conventional off-the-shelf hardware (low-powered microcontrollers). The authors created a simplified cerebellar circuit simulating a delayed eyeblink classical conditioning paradigm. In this model, neurons and synapses adjusted their activity to minimise their prediction error. This led to associative plasticity and unsupervised learning. Importantly, inference capabilities emerged entirely from connectivity and from the “Bayesian engine” in neurons (whereby individual neurons acted to minimise their prediction error) rather than from hardwired learning rules. Learning was also particularly rapid and required considerably fewer units and less power than traditional bioinspired networks.

Gandolfi et al. [61] highlight the remarkable simplicity of the generative model entailed by each element—each assumes one hidden state that can be in one of two levels. This is likened to spin-glass models in physics, in which the states can be “up” or “down”. Each element infers whether its world is in an “up” or “down” state by assimilating sensory (presynaptic) input. The authors suggest their experiments could be adopted to implement brain-like predictive capabilities into neuromorphic computers and robotic controllers.

AIF was also recently explored in a novel biosynthetic neuromorphic system called “DishBrain”. In this system, *in vitro* neural networks were integrated with *in silico* computing and embedded in a simulated-game world mimicking “Pong” [62]. The system was closed-loop in that it provided feedback on the causal effect of its behaviour, which afforded the *in vitro* culture “embodiment”. Note, the authors use the term “embodiment” here to refer to the “separation of internal versus external states where feedback of the effect of an action on a given environment is available” [62] (this aligns with functional perspectives of embodiment as discussed in [2]).

Applying AIF principles resulted in apparent learning within 5 min of real-time gameplay, and the system was said to exhibit “synthetic biological intelligence”.

As yet, we are not aware of any studies implementing AIF in neurorobotic systems (robots using neuromorphic hardware). We have, however, identified two recent studies using Bayes frameworks with neuromorphic hardware for robotic navigation and vision. Bayesian and AIF models share some similarities in terms of their probabilistic modelling approach. However, Bayesian models typically treat action as a separate stage of processing (decisions are made based on inferred probabilities), while AIF models treat action as an integral part of perception (the brain actively seeks out sensory information that reduces uncertainty and prediction error).

Tang et al. [63] proposed a spiking neural network (SNN) architecture with Bayes inference to solve a crucial problem in robotics—simultaneous localization and mapping (SLAM). The authors integrated this neuromorphic algorithm into neuromorphic hardware (Intel’s Loihi processor) and demonstrated 100 times less energy and comparable accuracy to the widely used GMapping algorithm. Ayyad et al. [64] presented a neuromorphic

vision-based approach to robot perception and control. Their novel approach for the event-based detection and tracking of circular objects in the scene leveraged a Bayesian framework to retain the asynchronous nature of neuromorphic cameras.

Finally, it is noteworthy that AIF has been linked to development, including morphogenesis (a form of collective cellular intelligence) [65,66]. Pio-Lopez et al. [65] demonstrated that AIF could be used to simulate disorders of morphogenesis, potentially leading to new treatment approaches for developmental disorders.

We envisage the role of AIF in development may be conceptually extended through framing cognitive neurodevelopment as: (1) adjusting models (i.e., incorporating new rules), and/or (2) selecting actions that aim to maximise information gain and minimise prediction errors (i.e., neurodevelopmental processes, such as myelination and synaptic pruning). In this way, agents act on their “internal” environment to progress developmentally. This has implications for the design of developmental neurobotic systems.

3.5. Summary

There is considerable theoretical support for utilising AIF in embodied neuromorphic agents. Empirical support is also emerging. The features and mechanisms discussed above demonstrate how embodied neuromorphic AIF agents have the potential to meet all the requirements for embodied “intelligence” outlined previously (see summary Table 1). Additional benefits identified include self-generated learning (through mechanisms akin to innate curiosity), meta-cognitive abilities, including self-monitoring and explainability (conferring safety benefits for socially collaborative and regulated use cases), multisensory, and low-power consumption.

According to Bartolozzi et al. [15], “A real breakthrough in the [neurorobotics] field will happen if the whole system design is based on biological computational principles, with a tight interplay between the estimation of the surroundings and the robot’s own state, and decision making, planning and action”. We argue that AIF may provide this, and potentially, in turn lead to the development of new tools for investigating cognitive neurodevelopmental processes.

4. Current AIF Implementation Approaches in Embodied Agents

According to de Vries [67], to implement AIF, at a minimum, engineers must express an agent’s prior preferences or constraints that underwrite behaviour. Design should focus on two tasks:

1. Specification of the agent’s model and inference constraints.
2. A recipe to continually minimise the free energy in that model under situated conditions driven by environmental interactions.

Current open-source tools for creating AIF agents include Statistical Parametric Mapping (SPM), PyMDP [68], ForneyLab [69] and RxInfer [67]. The latter of these uses a reactive message-passing framework whereby the inference process consists entirely of a (parallelizable) series of small steps (messages) that individually and independently contribute to free-energy minimisation [67].

As an example, the standard implementation approach in AIF research at present, for the discrete case, is the partially observable Markov decision processes (POMDP). A step-by-step guide to this is provided by [34]. In short, the POMDP formulation “describes beliefs about abstract states of the world, how they are expected to change over time, and how actions are selected to seek out preferred outcomes or rewards based on beliefs about states” [70]. Due to the concept of partial observability, an agent can be uncertain about its beliefs (states are “hidden”). It infers how likely it is to be in one hidden state or another based on observations (i.e., sensory input), which provide probabilistic information. The Markov property ensures that decision making implicitly includes all relevant knowledge about past states within current state beliefs [70]. Despite POMDP being a popular implementation approach in AIF research, robotics studies have used a range of different methods (see Table 2 for summary).

Table 2. Summary of AIF implementation methods used in robotic studies.

Robot Skill (Relevant Studies)	Summary of Method for Implementing AIF
Body estimation [53,55]	Lanillos & Cheng [53] present body estimation and control as a free-model active inference problem combining state-of-the-art regressors with free energy lower bound minimization. Oliver et al. [55] approximated the robot's body through variational inference (i.e., minimising the variational free energy (VFE) bound using the error between the expected and the observed sensory information). Free energy optimization was performed using gradient descent.
Body perception and action [71]	Used "PixelAI", a novel pixel-based deep active inference algorithm, which provides model-free learning and unifies perception and action into a single variational inference formulation. It combines the FEP with deep convolutional decoders.
Navigation [56,57,72]	Çatal et al. [56] cast the simultaneous localization and mapping (SLAM) problem in terms of a hierarchical Bayesian generative model. The agent reasons on two different levels: on a higher level for long-term navigation and on a lower level for short-term perception. At the lower level, the model builds upon earlier work [26] (build and learn the generative model using deep artificial neural networks, which are trained on sequences of action–observation pairs), in order to learn and infer belief states from sequences of observations and actions. Burghardt & Lanillos [57] used laser sensors. They defined the true state (position) of the robot, and the position belief of the robot. Estimation was solved by computing the posterior distribution by optimizing VFE. Taniguchi et al. [72] combined sequential Bayesian inference using particle filters and information gain-based destination determination in a probabilistic generative model ("SpCoAE" method). The method achieves AIF by selecting candidate points with the maximum information gain and performing online learning from observations obtained at the destination.
Planning [73]	This study specified nominal behaviour offline through behaviour trees and used a "leaf node" to specify the desired state to be achieved, rather than an action to execute. The decision of which action to execute to reach the desired state was performed online through active inference, resulting in "continual online planning and hierarchical deliberation".
Goal-directed behaviour [74]	This study brings together AIF-inspired behaviour generation and biologically plausible SNNs. It demonstrated that "goal-directed, anticipatory behaviour can emerge from projecting intentions through continuously unfolding spike dynamics onto motor input".
Complex social cognition [59]	This study used a multi-layered "PV-RNN" model (this is an RNN-type model that can instantiate predictive coding and active inference in a continuous spatio-temporal domain) with two branches (vision and proprioception) connected through an associative module. The model predicts incoming visual sensation and proprioception simultaneously; prediction error is back-propagated through time, and each module is modulated to maximize the evidence lower bound.

A key challenge for the AIF approach is the design of meaningful generative models and prior beliefs (preferences). Lanillos et al. [36] suggest that the required set of priors could be obtained through a proper curriculum, within a framework by which humans teach robots. This developmental approach is similar to Turing's hypothesis outlined above.

5. Suggestions to Catalyse AIF for Developmental Neuromorphic Agents

5.1. Approaches

There are a number of potential avenues to integrate AIF for learning and development in embodied neuromorphic agents. When considering future directions, it is noteworthy that AIF is a multi-scale framework, and therefore potentially applicable across the neuromorphic "stack"—from hardware and sensors to modality integration, computations, and global cognitive orchestration.

One potential approach is to demonstrate a previous AIF robotic study in neuromorphic hardware. For example, the study by Traub et al. [74] may be particularly applicable for translation into neuromorphic systems as it utilized recurrent spiking neural networks (SNNs) and the eligibility propagation algorithm ("e-prop"; a method for training SNNs which relies on local learning rules at each synapse). This algorithm has recently been

implemented in SpiNNaker [75] and SpiNNaker 2 [76]. Traub et al. [74] noted that e-prop is ideal for continuous online learning scenarios. An alternative approach would be to implement and expand a previous AIF neuromorphic study in a robotic system. For example, recreate Bayesian neurons [61] or Bayesian circuits [62] in neuromorphic hardware situated within a humanoid robot.

“One-shot” learning is also of particular interest. [34] used AIF to model concept learning and demonstrated “one-shot” generalisation to new stimuli. Although this study did not incorporate either robotics or neuromorphic hardware, it may be possible to implement a similar state–space expansion and reduction approach in a neurobotic system (i.e., through exploring the role of “open slots”). This could also be combined with “Bayesian One-Shot learning” techniques to optimise multimodality training (e.g., [77]).

5.2. Benchmarking

The field would likely benefit from standardised benchmarking tasks. The OpenAI Gym platform (gymnasium.farama.org) has been valuable in providing this for the AI community, including AIF studies [78]. However, developmental neurorobotics (the focus of this perspective) requires embodied tasks due to the emphasis on self-determined learning via interaction with the environment. This field would additionally benefit from tasks capable of identifying trajectories and “stages” of learning over time, in addition to “mental age” mapping, in order to enhance the validity of comparisons with human development.

We, therefore, suggest that it may be valuable to adapt a battery of cognitive tests commonly used in children for use in embodied humanoid agents. For example, theory of mind tasks (such as the “Sally-Anne” test [79]), the Bayley scale for infant and toddler development [80], and specific tasks associated with each stage of Piaget’s stages of cognitive development (including tasks related to object permanence, conservation, classification, and abstract reasoning).

5.3. Open-Source Resources

The AI field has greatly benefited from open-source resources. We advocate for the development of a resource offering a range of different open-source biologically plausible neuromorphic circuits—such as cortical column, ring attractor, hippocampal, and central pattern generator. An existing open neuromorphic platform could be leveraged to create this—such as BrainScales2 [81], which offers spiking information transmission and plastic synapses. The development of such a resource would provide AI communities with modular composable blocks for creating novel neuromorphic solutions.

Author Contributions: Conceptualization, S.H., A.J.R., and A.D.N.; methodology, S.H., A.J.R., and A.D.N.; analysis, S.H., and A.J.R.; resources, A.D.N.; writing—original draft preparation, S.H. and A.J.R.; writing—review and editing, S.H., A.J.R., A.H., and A.D.N.; supervision, A.D.N.; funding acquisition, A.D.N. All authors have read and agreed to the published version of the manuscript.

Funding: This work is funded by the Engineering and Physical Sciences Research Council (EP/X018733/14).

Institutional Review Board Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Oudeyer, P.Y. Autonomous development and learning in artificial intelligence and robotics: Scaling up deep learning to human-like learning. *Behav. Brain Sci.* **2017**, *40*, 275. [CrossRef] [PubMed]
2. Manzotti, R. Embodied AI beyond Embodied Cognition and Enactivism. *Philosophies* **2019**, *4*, 39. [CrossRef]
3. Glenberg, A. Embodiment as a unifying perspective for psychology. *WIREs Cogn. Sci.* **2010**, *1*, 586–596. [CrossRef] [PubMed]
4. Pfeifer, R.; Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence*; MIT Press: Cambridge, MA, USA, 2006.
5. Wilson, M. Six views of embodied cognition. *Psychon. Bull. Rev.* **2002**, *9*, 625–636. [CrossRef] [PubMed]
6. Montagu, A. *Touching: The human significance of the skin*; William Morrow Paperbacks: New York, NY, USA, 1986.

7. Frank, C.; Schack, T. The Representation of Motor (Inter)action, States of Action, and Learning: Three Perspectives on Motor Learning by Way of Imagery and Execution. *Front. Psychol.* **2017**, *8*, 678. [CrossRef]
8. Balasubramanian, V. Brain power. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, 2107022118. [CrossRef] [PubMed]
9. LUMI Ranked Third on Top500 List of World's Fastest Supercomputers—The Fastest Supercomputer in Europe. Available online: <https://www.lumi-supercomputer.eu/lumi-ranked-third-on-top500-list-of-worlds-fastest-supercomputers-the-fastest-supercomputer-in-europe/> (accessed on 2 July 2024).
10. Barlow, H. Possible principles underlying the transformation of sensory messages. *Sens. Commun.* **1961**, *1*, 217–233. [CrossRef]
11. Kosslyn, S.; Thompson, W.; Ganis, G. *The Case for Mental Imagery*; Oxford University Press: Oxford, UK, 2006; Volume VI. [CrossRef]
12. Borghi, A.; Barca, L.; Binkofski, F.; Tummolini, L. Varieties of abstract concepts: Development, use and representation in the brain. *Philos. Trans. R. Soc. B Biol. Sci.* **2018**, *373*, 20170121. [CrossRef]
13. Zaidel, D. Creativity, brain, and art: Biological and neurological considerations. *Front. Hum. Neurosci.* **2014**, *8*, 389. [CrossRef]
14. Bobrowicz, K.; Thibaut, J.P. The Development of Flexible Problem Solving: An Integrative Approach. *J. Intell.* **2023**, *11*, 119. [CrossRef]
15. Bartolozzi, C.; Indiveri, G.; Donati, E. Embodied neuromorphic intelligence. *Nat. Commun.* **2022**, *13*, 1024. [CrossRef] [PubMed]
16. Aitsam, M.; Davies, S.; Nuovo, A. Neuromorphic Computing for Interactive Robotics: A Systematic Review. *IEEE Access* **2022**, *10*, 122261–122279. [CrossRef]
17. Schuman, C.; Kulkarni, S.; Parsa, M.; Mitchell, J.; Date, P.; Kay, B. Opportunities for neuromorphic computing algorithms and applications. *Nat. Comput. Sci.* **2022**, *2*, 10–19. [CrossRef]
18. Huttenlocher, P. Synaptic density in human frontal cortex—developmental changes and effects of aging. *Brain Res.* **1979**, *163*, 195–205. [CrossRef] [PubMed]
19. Uhlhaas, P.H.; Roux, F.; Rodriguez, E.; Rotarska-Jagiela, A.; Singer, W. Neural synchrony and the development of cortical networks. *Trends Cogn. Sci.* **2010**, *14*, 72–80. [CrossRef] [PubMed]
20. Vernon, D. Cognitive vision: The case for embodied perception. *Image Vis. Comput.* **2008**, *26*, 127–140. [CrossRef]
21. Cangelosi, A.; Riga, T. An Embodied Model for Sensorimotor Grounding and Grounding Transfer: Experiments with Epigenetic Robots. *Cogn. Sci.* **2006**, *30*, 673–689. [CrossRef] [PubMed]
22. Cangelosi, A.; Schlesinger, M. *Developmental Robotics: From Babies to Robots*; MIT Press: Cambridge, MA, USA, 2015.
23. Pecyna, L.; Cangelosi, A.; Nuovo, A. A robot that counts like a child: A developmental model of counting and pointing. *Psychol. Res.* **2020**, *86*, 2495–2511. [CrossRef]
24. Proudfoot, D. Child machines. In *The Turing Guide*; Copeland, J., Bowen, J., Sprevak, M., Wilson, R., Eds.; Oxford University Press: Oxford, UK, 2017. [CrossRef]
25. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. [CrossRef]
26. Çatal, O.; Wauthier, S.; Boom, C.; Verbelen, T.; Dhoedt, B. Learning Generative State Space Models for Active Inference. *Front. Comput. Neurosci.* **2020**, *14*, 574372. [CrossRef]
27. Fields, C.; Friston, K.; Glazebrook, J.; Levin, M.; Marciànò, A. The Free Energy Principle drives neuromorphic development. *Neuromorphic Comput. Eng.* **2022**, *2*, 042002. [CrossRef]
28. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; Pezzulo, G. Active inference: A process theory. *Neural Comput.* **2017**, *29*, 1–49. [CrossRef] [PubMed]
29. Butz, M. Towards Strong AI. *KI—Künstliche Intell.* **2021**, *35*, 91–101. [CrossRef]
30. Marcus, G. Deep Learning: A Critical Appraisal. *arXiv* **2018**, arXiv:1801.00631. [CrossRef]
31. Asada, M.; Hosoda, K.; Kuniyoshi, Y.; Ishiguro, H.; Inui, T.; Yoshikawa, Y.; Ogino, M.; Yoshida, C. Cognitive Developmental Robotics: A Survey. *IEEE Trans. Auton. Ment. Dev.* **2009**, *1*, 12–34. [CrossRef]
32. Piaget, J. *The Psychology of Intelligence*; Routledge: London, UK, 2003. [CrossRef]
33. Parr, T.; Pezzulo, G.; Friston, K. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*; The MIT Press: Cambridge, MA, USA, 2022. [CrossRef]
34. Smith, R.; Schwartenbeck, P.; Parr, T.; Friston, K. An Active Inference Approach to Modeling Structure Learning: Concept Learning as an Example Case. *Front. Comput. Neurosci.* **2020**, *14*, 41. [CrossRef]
35. Sajid, N.; Ball, P.; Parr, T.; Friston, K. Active inference: Demystified and compared. *Neural Comput.* **2021**, *33*, 674–712. [CrossRef]
36. Lanillos, P.; Meo, C.; Pezzato, C.; Meera, A.; Baioumy, M.; Ohata, W.; Tschantz, A.; Millidge, B.; Wisse, M.; Buckley, C.; et al. Active Inference in Robotics and Artificial Agents. Survey and Challenges. *arXiv* **2021**, arXiv:2112.01871.
37. Kawahara, D.; Ozeki, S.; Mizuuchi, I. A Curiosity Algorithm for Robots Based on the Free Energy Principle. In Proceedings of the 2022 IEEE/SICE International Symposium on System Integration (SII), Narvik, Norway, 9–12 January 2022; pp. 53–59. [CrossRef]
38. Di Nuovo, A.; Cangelosi, A. Abstract Concept Learning in Cognitive Robots. *Curr. Robot. Rep.* **2021**, *2*, 1–8. [CrossRef]
39. Hawkins, J. Special report: Can we copy the brain?—What intelligent machines need to learn from the Neocortex. *IEEE Spectr.* **2017**, *54*, 34–71. [CrossRef]
40. Krichmar, J. Neurorobotics—A Thriving Community and a Promising Pathway Toward Intelligent Cognitive Robots. *Front. Neurobotics* **2018**, *12*, 42. [CrossRef]
41. Eppe, M.; Wermter, S.; Hafner, V.; Nagai, Y. Developmental Robotics and its Role Towards Artificial General Intelligence. *KI—Künstliche Intell.* **2021**, *35*, 5–7. [CrossRef]

42. Nguyen, P.H.; Kovač, M. Adopting Physical Artificial Intelligence in Soft Aerial Robots. *IOP Conf. Ser. Mater. Sci. Eng.* **2022**, *1261*, 012006. [[CrossRef](#)]
43. Ijspeert, A.J.; Crespi, A.; Ryczko, D.; Cabelguen, J.-M. From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model. *Science* **2007**, *315*, 1416–1420. [[CrossRef](#)] [[PubMed](#)]
44. Friston, K. Embodied inference: Or ‘I think therefore I am, if I am what I think’. In *The Implications of Embodiment: Cognition and Communication*; Imprint Academic: Exeter, UK, 2011; pp. 89–125.
45. Ramstead, M.; Kirchhoff, M.; Friston, K. A tale of two densities: Active inference is enactive inference. *Adapt. Behav.* **2020**, *28*, 225–239. [[CrossRef](#)] [[PubMed](#)]
46. Linson, A.; Clark, A.; Ramamoorthy, S.; Friston, K. The Active Inference Approach to Ecological Perception: General Information Dynamics for Natural and Artificial Embodied Cognition. *Front. Robot. AI* **2018**, *5*, 21. [[CrossRef](#)]
47. Pfeifer, R. Morphological Computation: Connecting Brain, Body, and Environment. In *Biologically Inspired Approaches to Advanced Information Technology*; Ijspeert, A., Masuzawa, T., Kusumoto, S., Eds.; Springer: Berlin/Heidelberg, Germany, 2006; pp. 2–3. [[CrossRef](#)]
48. Hipólito, I.; Es, T. Enactive-Dynamic Social Cognition and Active Inference. *Front. Psychol.* **2022**, *13*, 855074. [[CrossRef](#)]
49. Da Costa, L.; Lanillos, P.; Sajid, N.; Friston, K.; Khan, S. How Active Inference Could Help Revolutionise Robotics. *Entropy* **2022**, *24*, 361. [[CrossRef](#)]
50. Pio-Lopez, L.; Nizard, A.; Friston, K.; Pezzulo, G. Active inference and robot control: A case study. *J. R. Soc. Interface* **2016**, *13*, 20160616. [[CrossRef](#)] [[PubMed](#)]
51. Pezzato, C.; Ferrari, R.; Corbato, C. A Novel Adaptive Controller for Robot Manipulators Based on Active Inference. *IEEE Robot. Autom. Lett.* **2020**, *5*, 2973–2980. [[CrossRef](#)]
52. Van de Maele, T.; Verbelen, T.; Catal, O.; De Boom, C.; Dhoedt, B. Active Vision for Robot Manipulators Using The Free Energy Principle. *Front. Neurobotics.* **2021**, *15*, 642780. [[CrossRef](#)]
53. Lanillos, P.; Cheng, G. *Active Inference with Function Learning for Robot Body Perception*; Technical University of Munich: Munich, Germany, 2018.
54. Lanillos, P.; Gerven, M. Neuroscience-inspired perception-action in robotics: Applying active inference for state estimation control and self-perception. *arXiv* **2021**, arXiv:2105.04261.
55. Oliver, G.; Lanillos, P.; Cheng, G. An Empirical Study of Active Inference on a Humanoid Robot. *IEEE Trans. Cogn. Dev. Syst.* **2021**, *14*, 462–471. [[CrossRef](#)]
56. Catal, O.; Verbelen, T.; Maele, T.; Dhoedt, B.; Safron, A. Robot navigation as hierarchical active inference. *Neural Netw.* **2021**, *142*, 192–204. [[CrossRef](#)]
57. Burghardt, D.; Lanillos, P. Robot Localization and Navigation through Predictive Processing using LiDAR. *arXiv* **2021**, arXiv:2109.04139.
58. Baioumy, M.; Pezzato, C.; Corbato, C.; Hawes, N.; Ferrari, R.; Kamp, I.; Bibal, A.; Bouadi, T.; Frénay, B.; Galárraga, L.; et al. Towards Stochastic Fault-Tolerant Control Using Precision Learning and Active Inference. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*; Kamp, M., Gallicchio, C., Schiele, G., Pernkopf, F., Graça, G., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 681–691. [[CrossRef](#)]
59. Ohata, W.; Tani, J. Investigation of the Sense of Agency in Social Cognition, Based on Frameworks of Predictive Coding and Active Inference: A Simulation Study on Multimodal Imitative Interaction. *Front. Neurobot.* **2020**, *14*, 61. [[CrossRef](#)] [[PubMed](#)]
60. Isomura, T.; Shimazaki, H.; Friston, K. Canonical neural networks perform active inference. *Commun. Biol.* **2022**, *5*, 55. [[CrossRef](#)] [[PubMed](#)]
61. Gandolfi, D.; Puglisi, F.; Boiani, G.; Pagnoni, G.; Friston, K.; D’Angelo, E.; Mapelli, J. Emergence of associative learning in a neuromorphic inference network. *J. Neural Eng.* **2022**, *19*, 036022. [[CrossRef](#)]
62. Kagan, B.; Kitchen, A.; Tran, N.; Habibollahi, F.; Khajehnejad, M.; Parker, B.; Bhat, A.; Rollo, B.; Razi, A.; Friston, K. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron* **2022**, *110*, 3952–3969. [[CrossRef](#)]
63. Tang, G.; Shah, A.; Michmizos, K. Spiking Neural Network on Neuromorphic Hardware for Energy-Efficient Unidimensional SLAM. *arXiv* **2019**, arXiv:1903.02504.
64. Ayyad, A.; Halwani, M.; Swart, D.; Muthusamy, R.; Almaskari, F.; Zweiri, Y. Neuromorphic vision based control for the precise positioning of robotic drilling systems. *Robot. Comput.-Integr. Manuf.* **2023**, *79*, 102419. [[CrossRef](#)]
65. Pio-Lopez, L.; Kuchling, F.; Tung, A.; Pezzulo, G.; Levin, M. Active inference, morphogenesis, and computational psychiatry. *Front. Comput. Neurosci.* **2022**, *16*, 988977. [[CrossRef](#)] [[PubMed](#)]
66. Friston, K.; Levin, M.; Sengupta, B.; Pezzulo, G. Knowing one’s place: A free-energy approach to pattern regulation. *J. R. Soc. Interface* **2015**, *12*, 20141383. [[CrossRef](#)] [[PubMed](#)]
67. de Vries, B. Toward Design of Synthetic Active Inference Agents by Mere Mortals. *arXiv* **2023**, arXiv:2307.14145.
68. Heins, C.; Millidge, B.; Demekas, D.; Klein, B.; Friston, K.; Couzin, I.; Tschantz, A. pymdp: A Python library for active inference in discrete state spaces. *J. Open Source Softw.* **2022**, *7*, 4098. [[CrossRef](#)]
69. Cox, M.; Laar, T.; Vries, B. A Factor Graph Approach to Automated Design of Bayesian Signal Processing Algorithms. *Int. J. Approx. Reason.* **2019**, *104*, 185–204. [[CrossRef](#)]
70. Smith, R.; Friston, K.; Whyte, C. A step-by-step tutorial on active inference and its application to empirical data. *J. Math. Psychol.* **2021**, *107*, 102632. [[CrossRef](#)] [[PubMed](#)]

71. Sancaktar, C.; Gerven, M.; Lanillos, P. End-to-End Pixel-Based Deep Active Inference for Body Perception and Action. In Proceedings of the 2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob, Valparaiso, Chile, 26–30 October 2020; pp. 1–8. [[CrossRef](#)]
72. Taniguchi, A.; Tabuchi, Y.; Ishikawa, T.; Hafi, L.; Hagiwara, Y.; Taniguchi, T. Active Exploration based on Information Gain by Particle Filter for Efficient Spatial Concept Formation. *Adv. Robot.* **2023**, *37*, 840–870. [[CrossRef](#)]
73. Pezzato, C.; Corbato, C.; Bonhof, S.; Wisse, M. Active Inference and Behavior Trees for Reactive Action Planning and Execution in Robotics. *IEEE Trans. Robot.* **2022**, *39*, 1050–1069. [[CrossRef](#)]
74. Traub, M.; Legenstein, R.; Otte, S. Many-Joint Robot Arm Control with Recurrent Spiking Neural Networks. In Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Prague, Czech Republic, 27 September–1 October 2021; pp. 4918–4925. [[CrossRef](#)]
75. Perrett, A.; Summerton, S.; Gait, A.; Rhodes, O. Online learning in SNNs with e-prop and Neuromorphic Hardware. In Proceedings of the 2022 Annual Neuro-Inspired Computational Elements Conference, Virtual, 28 March–1 April 2022; pp. 32–39. [[CrossRef](#)]
76. Rostami, A.; Vogginger, B.; Yan, Y.; Mayr, C. E-prop on SpiNNaker 2: Exploring online learning in spiking RNNs on neuromorphic hardware. *Front. Neurosci.* **2022**, *16*, 1018006. [[CrossRef](#)]
77. Fei-Fei, L.; Fergus, R.; Perona, P. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.* **2006**, *28*, 594–611. [[CrossRef](#)] [[PubMed](#)]
78. Cullen, M.; Davey, B.; Friston, K.; Moran, R. Active Inference in OpenAI Gym: A Paradigm for Computational Investigations Into Psychiatric Illness. *Biol. Psychiatry. Cogn. Neurosci. Neuroimaging* **2018**, *3*, 809–818. [[CrossRef](#)] [[PubMed](#)]
79. Baron-Cohen, S.; Leslie, A.; Frith, U. Does the autistic child have a “theory of mind”? *Cognition* **1985**, *21*, 37–46. [[CrossRef](#)]
80. Balasundaram, P.; Avulakunta, I. Bayley Scales of Infant and Toddler Development. In *StatPearls*; StatPearls Publishing: Treasure Island, FL, USA, 2023.
81. Pehle, C.; Billaudelle, S.; Cramer, B.; Kaiser, J.; Schreiber, K.; Stradmann, Y.; Weis, J.; Leibfried, A.; Müller, E.; Schemmel, J. The BrainScaleS-2 Accelerated Neuromorphic System with Hybrid Plasticity. *Front. Neurosci.* **2022**, *16*, 795876. [[CrossRef](#)] [[PubMed](#)]

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.