

PAPER • OPEN ACCESS

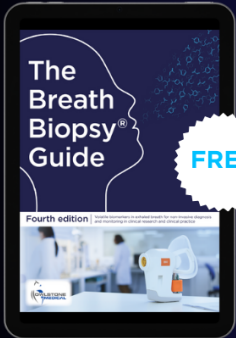
Emergence of associative learning in a neuromorphic inference network

To cite this article: Daniela Gandolfi *et al* 2022 *J. Neural Eng.* **19** 036022

View the [article online](#) for updates and enhancements.

You may also like

- [Fluid dynamic parameters of naturally derived hydroxyapatite scaffolds for *in vitro* studies of bone cells](#)
E Salerno, A d'Adamo, G Corda et al.
- [Lightweight solution for existing steel movable bridge retrofit and repair](#)
Antonella Ruzzante and Roberto Pavan
- [Assessment of existing steel bridges: codes and standard](#)
Erica Siviero and Roberto Pavan



The Breath Biopsy® Guide
Fourth edition

DOWNLOAD THE FREE E-BOOK

BREATH BIOPSY

OWLSTONE MEDICAL



PAPER

OPEN ACCESS

RECEIVED

20 December 2021

REVISED

25 March 2022

ACCEPTED FOR PUBLICATION

4 May 2022

PUBLISHED

30 May 2022

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Emergence of associative learning in a neuromorphic inference network

Daniela Gandolfi^{1,*} , Francesco M Puglisi^{2,3} , Giulia M Boiani¹, Giuseppe Pagnoni^{1,3} , Karl J Friston⁴ , Egidio D'Angelo^{5,6} and Jonathan Mapelli^{1,3,*}

¹ Department of Biomedical, Metabolic and Neural Sciences, University of Modena and Reggio Emilia, Modena, Italy

² Department of Engineering 'Enzo Ferrari', University of Modena and Reggio Emilia, Modena, Italy

³ Centre for Neuroscience and Neurotechnology, University of Modena and Reggio Emilia, Modena, Italy

⁴ Wellcome Centre for Human Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3AR, United Kingdom

⁵ Department of Brain and Behavioral Sciences, University of Pavia, Pavia, Italy

⁶ Brain Connectivity Center, IRCCS Mondino Foundation, Pavia, Italy

* Authors to whom any correspondence should be addressed.

E-mail: daniela.gandolfi@unimore.it and jonathan.mapelli@unimore.it

Keywords: predictive coding, unsupervised learning, neuromorphic electronic, brain-inspired computing

Supplementary material for this article is available [online](#)

Abstract

Objective. In the theoretical framework of predictive coding and active inference, the brain can be viewed as instantiating a rich generative model of the world that predicts incoming sensory data while continuously updating its parameters via minimization of prediction errors. While this theory has been successfully applied to cognitive processes—by modelling the activity of functional neural networks at a mesoscopic scale—the validity of the approach when modelling neurons as an ensemble of inferring agents, in a biologically plausible architecture, remained to be explored.

Approach. We modelled a simplified cerebellar circuit with individual neurons acting as Bayesian agents to simulate the classical delayed eyeblink conditioning protocol. Neurons and synapses adjusted their activity to minimize their prediction error, which was used as the network cost function. This cerebellar network was then implemented in hardware by replicating digital neuronal elements via a low-power microcontroller. **Main results.** Persistent changes of synaptic strength—that mirrored neurophysiological observations—emerged via local (neurocentric) prediction error minimization, leading to the expression of associative learning. The same paradigm was effectively emulated in low-power hardware showing remarkably efficient performance compared to conventional neuromorphic architectures. **Significance.** These findings show that: (a) an ensemble of free energy minimizing neurons—organized in a biological plausible architecture—can recapitulate functional self-organization observed in nature, such as associative plasticity, and (b) a neuromorphic network of inference units can learn unsupervised tasks without embedding predefined learning rules in the circuit, thus providing a potential avenue to a novel form of brain-inspired artificial intelligence.

1. Introduction

The predictive processing theory views brain function as a dynamical process that continuously updates a generative model of the world (i.e. internal model) to predict incoming sensory data. By acting on its environment, an organism can change its sensory inputs so to reduce—in real time—the mismatch between

the predicted and the actual incoming sensory signal through a process of Bayesian active inference [1–3]. The process is based on the minimization of a cost function, called variational free energy [4] (VFE). Although not identical, VFE can be formally related to the Helmholtz free energy, used in one of the first computational models of perceptual processing in a statistical inference engine [5]. A variety of

large-scale phenomena have been shown to conform to the free energy principle [6, 7] (FEP). However, these treatments leave a main question open: could an ensemble of free energy minimising neurons—in a biologically plausible architecture—collectively minimise their joint free energy? In other words, are large-scale phenomena emergent properties of self-organising Bayesian neurons? Recently, synaptic learning was shown to comply with Bayesian inference by tracking the probability distribution of synaptic weights and adjusting learning rates [8], but the effectiveness of single neurons organized in suitable network architectures—that perform Bayesian active inference in a behaviourally meaningful context—remained to be established.

Importantly, the FEP is thought to operate at multiple scales [9, 10]. While the minimization of free energy on a scale of milliseconds-to-seconds (as in the present work) underpins changes in synaptic activity and plasticity, free energy minimisation at evolutionary time scales has been proposed to explain the morphological changes associated with the species-specific anatomical structure of an organism [11]; namely, a natural kind of Bayesian model selection or structure learning. On this view, the cerebellum—with its peculiar internal organization and connectivity to the cerebral cortex and brainstem [12, 13]—may have been sculpted by evolution to predict the sensory consequences of motor acts and therefore minimize surprises (i.e. free energy or prediction error) about the consequences of action.

Prediction also entails novelty and error detection, temporal matching and sequence ordering [14, 15], essential functions for behaviour and thought [16]. The cerebellum needs to learn in order to adapt to the changing environment. Several forms of long-term synaptic plasticity have indeed been found in its circuits [17–19]. One of the most widely adopted experimental protocols to characterise learning mechanisms is the delayed eyeblink classical conditioning (dEBCC), a cerebellar-dependent form of associative memory that allows to predict the precise timing of contingent sensory events [20]. Interestingly, a generative model of the cerebellum implementing Pavlovian conditioning has been recently proposed [21]. However, predictive coding and active inference accounts of circuit learning have only been implemented at the level of the entire circuit, leaving ample space to investigate the role of free energy minimisation at the single cell (i.e. neurocentric) level. Clearly, implementing the FEP in a specific neuromorphic architecture, like that anticipated for the cerebellum, raises a series of questions. Does the FEP effectively apply to single neuron? What is the role of network connectivity? Will plasticity driven by neurocentric active inference emerge at the synaptic sites that mirror biology? To address these questions, we simulated an ensemble of neurons, that minimise their free energy under a very simple generative

model of their world (i.e. the ensemble of neurons to which they were connected). We used the connectivity architecture of the cerebellum to test two hypotheses: (a) the emergent behaviour of the ensemble would recapitulate associative learning of the sort seen empirically and (b) any changes to the connectivity architecture (optimised by natural Bayesian model selection) would render free energy minimisation and performance suboptimal. We tested these hypotheses using *in silico* simulations and numerical experiments.

To demonstrate the potential importance of this formulation of inference and learning for hardware implementation and edge computing, we leveraged the well-known connection between computational and thermodynamic efficiency [22, 23]. We hypothesised that the Bayes optimal computations, established in our *in silico* simulations, would translate into an efficient and expressive in-hardware implementation. To test this hypothesis, we reproduced the numerical simulations using biomimetic electronic spiking networks. A series of neuromorphic devices—that mimic neural activity and optimize electronic and robotic performance—has been proposed. Among these, field-programmable gate arrays (FPGA) have been employed to implement electronic spiking networks and therefore simulate neuronal microcircuits [24–26]. In particular, associative learning has been implemented using traditional CMOS [27–29], hybrid CMOS-memristors [30–35] or FPGA [36] and embedding bespoke learning rules in the circuit architecture. Conversely, in the current setup, learning would be free from assumptions about specific rules, speaking to the potential advantages in the design and implementation of electronic circuits [27]. This putative advance follows because synaptic plasticity follows the same (local) rules at each elementary unit, via free energy minimisation, under a simple generative model. In summary, we simulated—and subsequently implemented in hardware—a dEBCC paradigm by assembling a biologically plausible neural network architecture of the cerebellum, with distributed synaptic plasticity, where every neuron is conceived as an active inference agent that assimilates (and broadcasts) the activity impinging on its connections [37].

2. Method

2.1. Active inference applied to biological neural circuits

A generative model of cerebellar (figure 1(A)) synaptic inputs to a neuron can be described in terms of likelihood matrices (**A** see below, figure 1(B)) and prior beliefs over hidden states (**D**). The likelihood maps the hidden causes (**X**, figure 1(B)) to observable outcomes (**O**, figure 1(B)). The priors are over hidden states were as simple as possible; namely, that the

hidden or latent states generating (presynaptic) outcomes could be in one of two states; e.g. ‘on’ or ‘off’. Intuitively, this means that each neuron thinks the world (i.e. network) is in an ‘on’ or ‘off’ state—and adjusts its activity (i.e. Bayesian beliefs about hidden states) and connectivity (i.e. Bayesian beliefs about the likelihood and prior) to minimize VFE. This is equivalent to maximizing the evidence or marginal likelihood of its observations. Specifying a generative model of this sort allowed to characterize the activity of each neuron as encoding its expectation about the (hidden) causes of its synaptic inputs. Following the Bayes formalism in [2], the hidden state is denoted by s_t while the observable outcome is denoted by o_t (see below). The negative free energy (-FE) provides a lower bound on the log-evidence (ELBO) of the model, with equality when the approximate posterior distribution of the hidden states ($Q(s_t)$, equation (1) equals the true posterior ($P(s_t|o_t)$, equation (1) [38]:

$$Q(s_t) = \underset{Q}{\operatorname{argmin}} \text{FE} \approx P(s_t|o_t) \quad (1)$$

2.2. The cerebellar neuronal circuit model

The neuronal circuit comprised 105 neurons with the synaptic connectivity illustrated in figure 1 and table SI-1 (available online at stacks.iop.org/JNE/19/036022/mmedia). The network activity evolves with a time step accounting for synaptic integration and the action potential refractory period. The network contains both excitatory (red connections, figure 1(A)) and inhibitory (blue connections, figure 1(A)) synapses. Each neuron’s hidden state is represented by a one-dimensional quantity ($s_{t-1} \in \{0,1\}$), that can be either ‘off’ or non-firing (0) or ‘on’ or firing (1) [37]. Depending on the connectivity matrix W , if the neuron is firing, at the subsequent time step it will cause an excitatory post synaptic potential (EPSP) or an inhibitory post synaptic potential (IPSP) in its post-synaptic neurons. The EPSPs or IPSPs represent the observation or sensory input o_t for the postsynaptic neuron and are one-dimensional quantities ($o_t \in \{0,1\}$, see figure 1).

2.3. Generative model

According to the framework proposed in [37], the generative model (figure 1(B), blue left side of the box) is described by a likelihood matrix (A) mapping hidden states to sensory input, and prior beliefs over hidden states. Each synaptic connection is represented by a matrix (A_k in figure 1(B)) which encodes the probability of a specific o_t given s_{t-1} . In our setup, the A matrices for excitatory synapses are initialized as follows:

$$A_{\text{exc}} = \sigma \left(\log \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + e^{-\frac{1}{2}} \right) \quad (2)$$

while the A matrices for the inhibitory synapse are initialized as:

$$A_{\text{inh}} = \sigma \left(\log \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} + e^{-2} \right) \quad (3)$$

The σ denotes the softmax operator (normalized exponential) and the quantities of $e^{-1/2}$ and e^{-2} are weights added to differentiate the strength of inhibitory and excitatory connections. From the perspective of the generative model, these weights control the precision of the likelihood (i.e. the sensitivity to presynaptic inputs). In likelihood matrices such as the above, the element A_{ij} represents the probability that the hidden state $s_{t-1}[j]$ will induce the postsynaptic outcome $o_t[i]$. A_{inh} features higher precision compared to A_{exc} in order to mimic the higher efficacy of inhibitory synapses compared to excitatory synapses [39]. A_{inh} are responsible for decreasing the firing probability of the postsynaptic neuron.

The conditional dependence of s_t on s_{t-1} is parametrized by a probability transition matrix (B). In our work matrix B is set equal to the softmax of the identity matrix $[1,0;0,1]$, since the circuit is not equipped with specific temporal dynamics, such as oscillations, that are prescribed by a conditional dependence between states [37]. Intuitively, this means each neuron thinks that the hidden or latent states generating its observations do not change during its belief updating.

Finally, the generative model requires the initialization of the prior belief about hidden states (D). This parameter represents the firing probability of a neuron which corresponds to its ‘belief’ about the state of its presynaptic neurons at the beginning of each exchange of spikes or messages (see below). As shown in figures 1(B) and (D) at time t is updated to D at time $t + 1$, the posterior expectation of hidden state at the end of the iterative free energy minimization [37] in the variational Bayes inference process describing neuronal dynamics (see following section): figure 1(B), green right side of the neuron’s schematic soma.

2.4. Variational free-energy minimization: from Bayes rule to neuronal message passing

Variational free-energy minimization was implemented by message passing on the Forney Factor Graph representation of the generative model (see figure 2 in [40]). Analogous to the sum-product Bayesian message passing scheme [41], the variational message passing [42, 43] (VMP) involves the integration of locally computed messages. This in turn allows the reformulation of inference as an optimization process, where belief updating corresponds to the convergence of free energy to a local minimum, using an iterative gradient descent (see equation (6)). In other words, we treat neuronal dynamics as a gradient flow on VFE, following each exchange of spikes. The posterior distribution over hidden is derived by

the application of the sum-product rule [41], that is, the product of all incoming messages to the associated node. In the present framework, where each neuron acts as an inferring element, at each time point the messages update the log expectation of the hidden states via an error term (ε_t in equation (8)) [37]. The log expectation of hidden states is associated with the voltage at the neuron's soma [37] (equation (4)):

$$v_t = \ln s_t \quad (4)$$

The solution of the following equation provides posterior beliefs (boldface) about hidden states:

$$\ln \mathbf{s}_t = \ln \mathbf{B}_{t-1} \mathbf{s}_{t-1} + \ln \mathbf{B}_{t+1}^T \mathbf{s}_{t+1} + \ln \mathbf{A}^T \mathbf{o}_t \quad (5)$$

According to the VMP scheme in [37], the variational solution is expressed by introducing an auxiliary variable, the state prediction error (ε_t), which is equal to the free energy gradient:

$$\dot{v} = \varepsilon_t = - \frac{\partial FE}{\partial s_t} \quad (6)$$

$$\mathbf{s}_t = \sigma(v_t) \xrightarrow{\text{VMP}} X_t \quad (7)$$

$$\varepsilon_t = \ln \mathbf{B}_{t-1} \mathbf{s}_{t-1} + \ln \mathbf{B}_t^T \mathbf{s}_{t+1} + \ln \mathbf{A}^T \mathbf{o}_t - \ln \mathbf{s}_t \quad (8)$$

At each time step, the neuron's belief about the state of the network—in which it participates—is encoded in the probability of firing an action potential (X_t) and the softmax operator is here used to mimic a sigmoid firing-rate depolarization activation function. This accounts for first belief updating and neuronal dynamics during each time step (i.e. inference). However, we also have to consider optimization of the parameters of the generative model over multiple timesteps (i.e. learning).

By introducing a synaptic precision parameter, slower changes in synaptic efficacy also perform a gradient descent on the VFE. The synaptic precision parameter, which basically weighs the likelihood Matrix (\mathbf{A}), determines the impact of presynaptic inputs on neuronal belief updating. In this framework, the neuronal circuit is endowed with synaptic plasticity because a synapse-specific tuning of precision is implemented depending on prediction errors. In the present model, those excitatory synapses that transmit information about the state of the presynaptic neurons—that are a good match with predicted (or inferred) values—will undergo an increase in (estimated) precision, leading to more accurate predictions in the future (see section 3). Conversely, synaptic precision will decrease when predicted values conflict with sensory evidence. Inhibitory synapses will be adjusted in the opposite fashion: they will undergo an increase in precision when inferred values conflict with input evidence, and a decrease in precision when

inferred values and input are well-matched. Such configuration mimics the homeostatic role of inhibitory neurotransmission (negative feedback loops). In the present treatment (see [37]), priors P and posteriors Q densities over precision parameters (ζ^i) are chosen as gamma distributions (equations (9) and (10)):

$$P(\zeta^i) \propto \beta^i e^{-\beta^i \zeta^i} \quad (9)$$

$$Q(\zeta^i) \propto \beta^i e^{-\beta^i \zeta^i} \quad (10)$$

The posterior expectation (boldface) about the precision of sensory evidence is given by $E_{Q(\zeta^i)}[\zeta^i] = \zeta^i = (\beta^i)^{-1}$, since precision quantifies the inverse variability associated with a probability distribution. The update equation for expected sensory precision is expressed following [37].

2.5. Eye-blink conditioning protocol and data analysis

The temporal evolution of the system was simulated in cycles (i.e. time-steps), where each time step can be thought as equivalent to a period of two milliseconds, in accordance with known neuronal dynamics. Rate calculation and stimulations are therefore to be interpreted in this temporal frame of reference. Associative learning was induced by pairing a generic sensory stimulus (the conditioned stimulus (CS)) with an eyeblink-eliciting unconditioned stimulus (US). The former was produced by the repetitive stimulation (10 stimuli at 2 Hz) of (simulated) PN with 120 ms bursts at 400 Hz, while the latter was conveyed as a repetitive stimulation (10 stimuli at 2 Hz) of (simulated) trigeminal nuclei (TN) with 30 ms bursts at 400 Hz. TN stimulation was always delivered within the pontine nuclei (PN) burst stimulation window: more precisely, the ends of the two stimulation windows coincided. Each simulation comprised the subsequent and independent repetition of three conditions: US, CS and combined stimulation (USCS).

To quantify the emergence of conditioning, the spiking probability of the readout neurons (RNs) was assessed in terms of total number of spikes (spike area) generated in the temporal window corresponding to the CS. The number of spikes associated with a specific time step was obtained by introducing an upper threshold of 90% for the firing probability of RNs output. The spike histogram representation of the RN was also divided in two time-windows for each CS stimulus window, one in the first half of the window and the other in the second half. This allowed us to identify early RN activation in the USCS condition compared to CS and US alone.

2.6. Hardware implementation

The circuit architecture was realized by connecting 105 identical digital neurons, each implemented by means of a commercial, off-the-shelf, low-cost

and low-power microcontroller, the STM32L475. This device is equipped with an ARM® Cortex®-M4 core, with digital signal processing and floating-point unit that can run up to 80 MHz, as well as with several integrated peripherals like analog-to-digital and digital-to-analog converters, fast communication buses, controllers, operational amplifiers, comparators, timers, and cryptography units. However, only the core was powered and exploited, disabling all the peripherals to optimize the energy consumption profile. All digital neurons were set to run at a core frequency of 16 MHz, were supplied by a 3 V DC power source, and received an externally generated low-frequency (i.e. 80 Hz) clock signal to synchronize their activity, with each clock cycle lasting 12.5 ms; corresponding to an individual time-step. Each digital neuron is characterized by a given number of input terminals (in addition to the one dedicated to the clock signal) that depends on the specific neuron class (e.g. GrC, GoC, etc), and by one output terminal, which is then connected to a specific input terminal of one or more neurons, according to the network connectivity scheme. Those neurons whose inputs were not connected to the output terminal of other neurons, were stimulated by programmable voltage sources according to the eye-blink conditioning protocol—described in section 2.3—to reproduce the possible presence of spikes at the input of the neural network. Therefore, digital voltage signals (3 V rectangular pulses) are present on the input and output terminals of each neuron, closely mimicking the spiking activity of biological neurons. At the rising edge of the low-frequency clock signal, each digital neuron reads the information at each of its input terminals (i.e. whether a spike is present or not) and combines this information with the stored values of the synaptic precision for each synapse (each input terminal is associated with either an excitatory or an inhibitory synapse). All digital neurons were identically programmed to perform the VFE minimization process described in section 2.2 and evaluate their firing probability, which was used to probabilistically determine whether they should output a spike at the next rising edge of the clock signal. The only difference between different neuron classes (e.g. GrC, GoC, etc) was the number of inputs and the associated kind of synaptic connection (excitatory or inhibitory). Calculations were all performed by the digital core unit, using the same code employed in the simulation of the delayed eye-blink conditioning protocol described in section 2.3 and, in detail, in section 3.2.

The core frequency of the microcontroller was chosen to be 16 MHz, i.e. the minimum value that allowed enough time for the execution of the VFE minimization, while keeping the overall duration of the time-step in the ≈ 10 ms range (close to the typical spike timing in biological systems), which in turn allowed achieving low-power real-time unsupervised learning in commercial hardware. Notably,

each digital neuron dissipated just 4.8 mW, as quantified by measuring the current drawn at the voltage supply terminals, and the whole network power consumption was about 0.5 W. Furthermore, the network started to exhibit conditioned learning after only five stimulations (i.e. about 1200 time-steps), which implies that the system was fully trained in just 15 s, with an overall energy consumption < 8 J. Notably, the frequency-normalized power dissipation (i.e. $4.8 \text{ mW}/16 \text{ MHz} = 0.3 \text{ mW MHz}^{-1}$) and the frequency-normalized number of instructions per second (i.e. $20 \text{ MIPS}/16 \text{ MHz} = 1.25 \text{ MIPS MHz}^{-1}$) are equivalent to those of the ARM processor constituting the core of the SpiNNaker chip, although implementing the same task with the supported neuron models in SpiNNaker would require a few thousands neurons, as opposed to only 105 for the current approach.

Compared to existing computational simulations of the delayed eyeblink conditioning (software-implemented neural network running on an Intel i7 2600 CPU [29]), the network described in the present work reduces the number of requisite neurons by a factor > 60 , the power consumption by a factor ≈ 148 , the overall energy consumption for training by a factor > 434 , and the corresponding energy-delay product figure of merit (for full training) by a factor > 1272 . For further comparison against existing solutions, the proposed model was benchmarked against a recently proposed simulative model for associative memory in the form of a spiking neural network that was tested both on a low-power CPU and on Intel's Loihi artificial intelligence accelerator—highly optimized for spiking neural algorithms [27]. The best energy efficiency (1.1 W) was achieved during the execution of the association task on the pre-trained model running on the Loihi hardware. However, the energy and the time required for the model training was not disclosed, and this is typically a time- and energy-hungry process. Conversely, the present implementation dissipates just 0.5 W to perform fully unsupervised online training and inference, which speaks to the deployment of autonomous learning, smart agents at the edge. Such an approach has the unprecedented advantage of not requiring any embedded learning rules, since these emerge from the network connectivity that, furthermore, can be easily re-programmed in-field by adopting programmable neuron-to-neuron connections (as occurs in FPGA devices).

3. Results

A neural network with adapting activity according to Bayesian principles was recently proposed in a fully connected circuit comprising purely excitatory elements [37]. We have extended this approach to a biologically inspired architecture, while also introducing inhibitory connections.




Figure 1. FE minimization in the mossy fiber pathway activation of a cerebellar network of inferring neurons. (A) Cerebellar cortical circuit. Scheme of the circuit used to mimic the activity of single cerebellar inferring neurons. Neurons are represented as colored circles, excitatory connections are depicted with red lines and inhibitory connections with blue lines. Abbreviations: Pontine nuclei (PN), mossy fibers (MFs), granule cells (GrCs), Golgi cells (GoCs), Purkinje cells (PCs), deep cerebellar nuclei (DCNexc excitatory, DCNinh inhibitory). Shaded blocks are not included in the circuit configuration simulated in this figure (Trigeminal nuclei—TN, inferior olive—IO, readout neurons—RNs). In the subsequent figures, shaded blocks are included in circuit configuration. (B) Schematic representation of a single neuron acting as a free energy minimizing element. Each neuron (blue and green box) encodes, in terms of firing probability (X), an inference about the hidden state generating its inputs through an iterative minimization of VFE. At time t , the i th neuron receives inputs from different presynaptic neurons ($o_i^j(t)$) indicated by the connectivity matrix (W). The VFE minimization is implemented by message passing (green shaded area) within the generative model (blue shaded area) based on likelihood (A), transition probability (B), prior belief (D) and presynaptic outcomes (o) matrices. The firing probability X encodes the update of the prior beliefs about the presynaptic hidden states (D). Every neuron in the network can be presynaptic for other neurons at a subsequent time step ($o_k^i(t+1)$). (C) The image shows the probability of firing for every neuron in the network (y -axis) at each time step (x -axis). Each row represents the time course of activity of a specific neuron. Separation between neuronal layers is indicated by white horizontal lines. Network activity is elicited by the stimulation of PN (Conditioned Stimulus, black arrows) as explained in section 2. (D) The plot shows the average FE time course for each class of neurons. The black line shows the mean FE obtained by averaging the contribution of all neurons at a given time in the network.

3.1. Single inferential neurons in a cerebellar circuit

Single synthetic neurons—actively influencing the signals presented at their input—were assembled to mimic a simplified version of the cerebellar cortex (figure 1(A)). The neurons were connected through either excitatory or inhibitory synapses that were initialized with the same weights (except for the sign). The neurons adjusted their activity and their connection strength following the minimization of VFE (for convenience, this was measured as a change in negative free energy). The inferred firing probability (X , see section 2), which can be viewed as the ‘expectation’ of a neuron about the hidden state generating its presynaptic inputs (figure 1(B)), was monitored for each neuron in the network

(figure 1(C)). The system was initialized with random states (D in figure 1(B); see section 2) and left free to adjust its activity for 250 time-steps. The circuit was activated with trains of pulses delivered to mossy fibers (MFs). This caused a synchronization in the probability of firing in granule cells (GrCs), thereby modifying Purkinje cell (PC) and deep cerebellar nuclei (DCN) activities downstream. Interestingly, in a few cycles of burst stimulation, feedforward and feedback inhibitory loops passing through Golgi cells (GoCs) evinced an inhomogeneous activity in GrCs. The GrCs, followed by the DCN cells, initiated and sustained the FE increase (figure 1(D)).

By monitoring the dynamical evolution of synaptic efficacies, we observed that the emergence of a patterned network discharge was paralleled by a

widespread increase in synaptic efficacy (i.e. precision) at the MFs-GrCs connections (see section 2 and figure S1), coherently with findings on long-term plasticity observed both *in vitro* and *in vivo* [37]. In other words, basic functional properties of the granular layer network emerged simply from a local optimization of synaptic weights as a consequence of VFE minimization.

3.2. Delayed eyeblink conditioning

We then extended the cerebellar network by including the ‘extracerebellar pathways’ (figure 1(A); shaded neurons and connections) passing through the inferior olivary neurons (IOs) in order to investigate the dBCC, a typical paradigm of associative learning. In this form of associative learning, a neutral stimulus (conditioned stimulus, CS) overlaps and co-terminates with an eyeblink-eliciting stimulus (unconditioned stimulus, US), resulting—over time—in the emergence of an eyeblink conditioned response to the CS alone. The connectivity pattern was therefore extended to include IO inputs from TN and a layer of RNs, which was used to monitor network activity. In order to properly simulate the CS and the US protocols, two distinct pathways activating RNs were created (figure 1(A)). On the one hand, the CS pathway, which includes the PN, activating via MFs both GrCs and GoCs, with the PN also exciting DCN neurons through MF collaterals. On the other hand, the US pathway, which includes TN and activates RN both via a direct connection and via the IO-PC-DCN pathway.

Both the US and CS pathways eventually activate PCs, which exert their inhibition on DCN neurons. The excitatory subpopulation of DCN activates RNs as well as the inhibitory neurons of DCN, which in turn inhibit IO. The activity of these cells was monitored and compared in different conditions. As for the cerebellar network (cf figure 1), synaptic precisions of all connections were automatically adjusted on the sole basis of FE minimization.

In order to simulate dBCC, the CS stimulation protocol evoked firing activity in MFs, whereas RNs were left mostly silent (figure 2(A)). Conversely, the US stimulation protocol evoked an intense firing in RNs (figures 2(B) and (D)). The CS and US were suitably timed in order to generate the overlapping and co-terminations typical of dBCC (15 time-steps of US overlapped with the ending part of the 60 time-steps of CS; see section 2). In figure 1(C), trains of stimuli induced a strong increase of synchronized firing probability in the layers closest to the stimulation sites (i.e. in GoCs and GrCs). Here, the probability of firing was reduced after the first layers (figures 2(A) and (B)) both with the CS and US. The CS and US responses were markedly different, reflecting circuit connectivity. At visual inspection, the combined application of CS and US (USCS) increased and anticipated RNs responses (figure 2(D)). Notably, FE

started to increase after the 4th burst of stimuli during CS and USCS, whereas it remained almost constant during US (figure 2(E)). To test the hypothesis that any departure from the proposed biomimetic structure should result in suboptimal behavior, we explored different network configurations (see supplementary material, SI3–SI8). We found that the joint or collective free energy—our measure of the network performance—was mainly dependent on the chosen connectivity pattern and, to a lesser extent, on the number of neurons in the network (figure 2(F)). In brief, as predicted, any departure from the biomimetic cerebellar architecture degraded network performance; suggesting that our approximation to the cerebellar connectivity was itself a free energy minimizing structure. In other words, the kind of connectivity structure that would have emerged had used the joint free energy as the basis of Bayesian model selection (i.e. structure learning).

Following common dBCC protocols, we analyzed and measured the response activity of the RNs. The firing probability was converted into action potentials via a thresholding procedure and the number of spikes collectively generated by RNs neurons was counted in the three conditions and compared (see section 2). During the US, the number of RNs spikes was roughly constant throughout the stimulation protocol (green line in figures 3(D)–(F)), whereas during the CS (blue line), RNs neurons were mostly silent. Only a small increase in the RNs activity could be observed in the last two CS stimulation bursts, likely reflecting changes in the MFs-GrCs synaptic efficacy following the repeated stimulation (figures 3(A) and (C)). During the combined USCS protocol, the collective number of spikes markedly increased starting from the fifth stimulus burst (figure 3(F)). The cumulative spikes count generated in the last five bursts was in fact larger than in the first five bursts (figure 3(D), +53%). Moreover, when spike count was performed only considering the first half of the CS stimulation window, an even larger increase during the USCS protocol was observed during the last five bursts compared to the first five bursts (figure 3(E); +255%). The increase and anticipation of responses during USCS confirmed the emergence of a form of associative learning. This effect is consistent with experimental findings *in vivo*, which report an anticipation of eyelid closure and of spikes generation in DCN and pre-motor neurons [20]. Finally, it should be noted that the network became silent after the end of the dBCC protocol. Then a further presentation of CS, US, and USCS quickly re-established associative learning, suggesting the presence of consolidation mechanisms (figure SI-2). Different from other dBCC simulations, like those using spiking [25] or biologically realistic neurons, we have adapted the classical protocol to a network where learning rules were not embedded into synaptic sites. Although the dBCC task is the same,

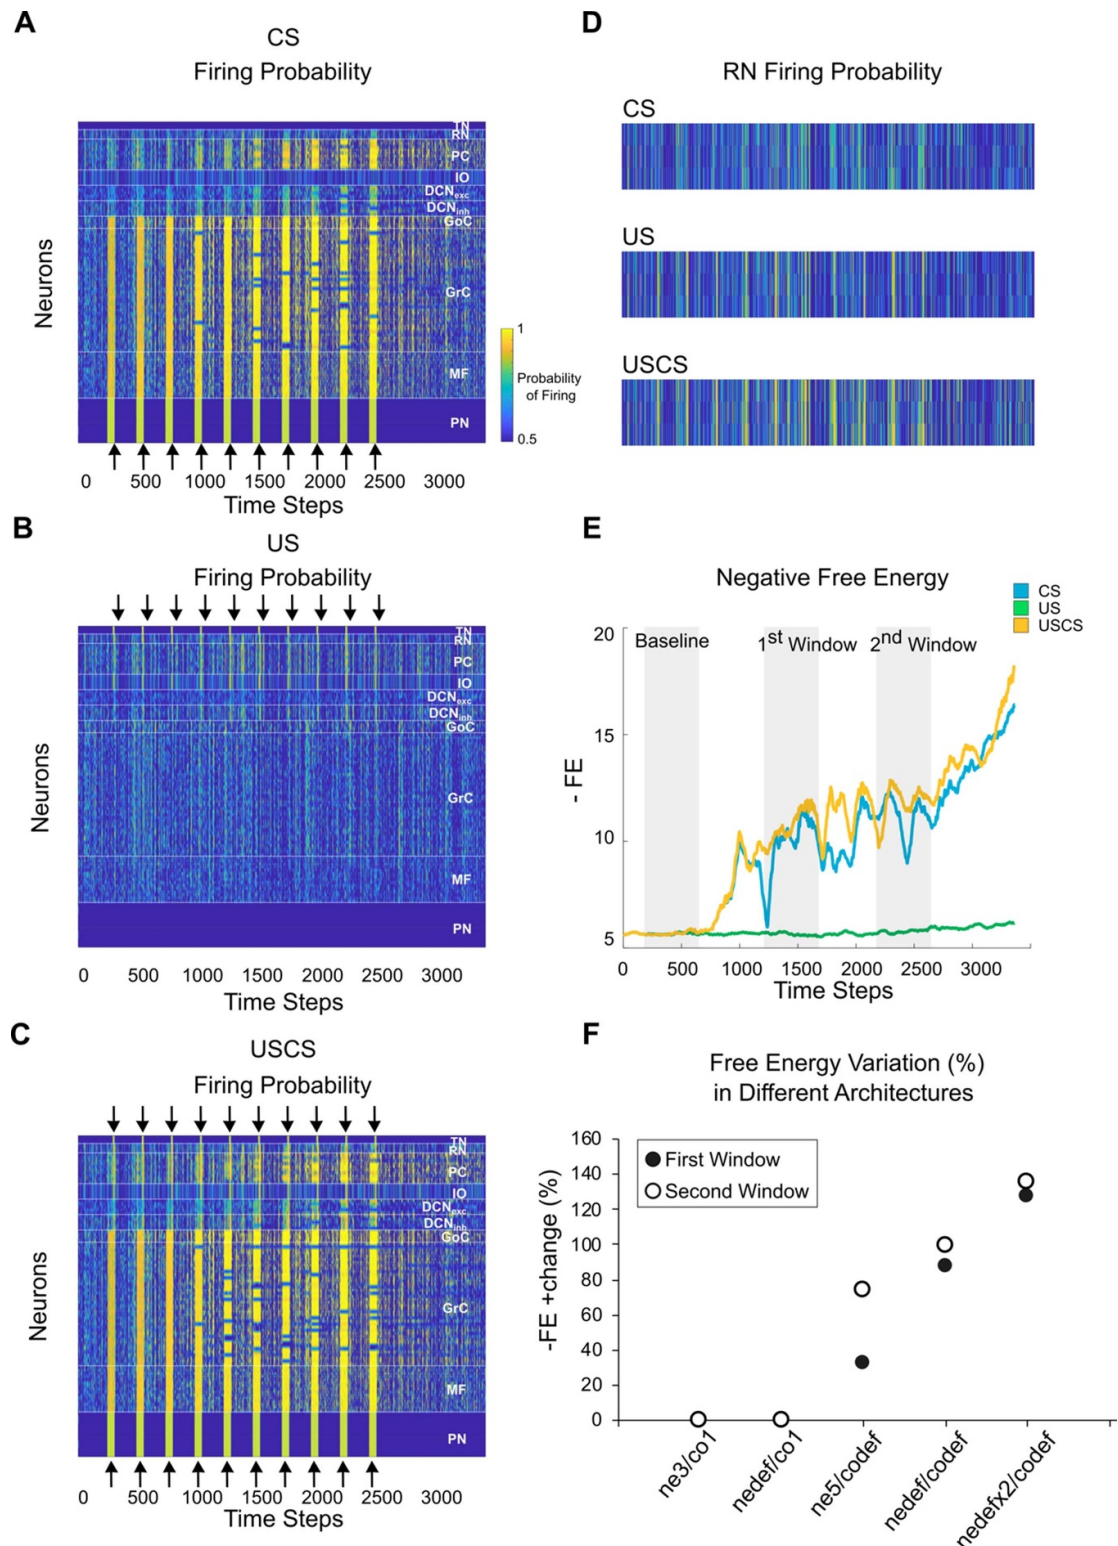
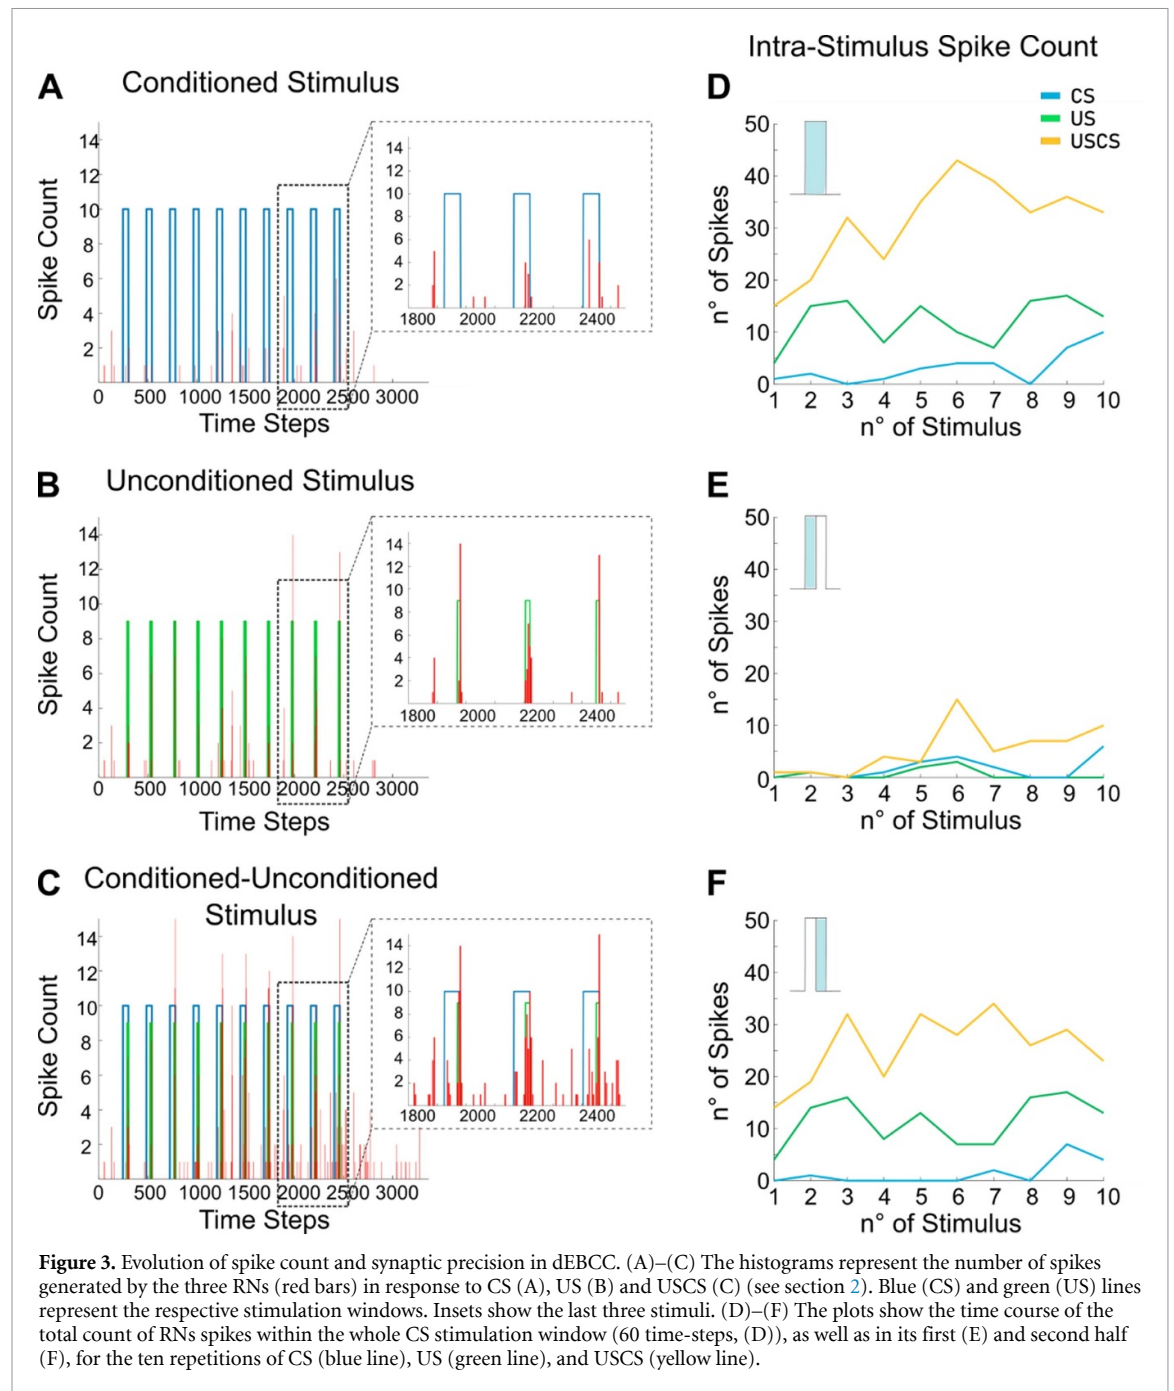


Figure 2. Associative learning through FE minimization in dBCC. (A)–(C) Firing probability of the network obtained as in figure 1(C) in response to the CS (A), US (B) and USCS (C) protocols. (D) Readout neurons firing probability in response to the CS, US and USCS (C). (E) The plot shows the time course of the FE during the different protocols (US—green; CS—blue; USCS—yellow). (F) The plot shows changes in the mean FE taken in separate time-windows (shaded areas in panel (E)) during the USCS protocol with different network configurations: ne3/co1 (minimal network), nedef/co1 (standard network with minimal connections), ne5/co1 (homogeneous and standard connected network), nedef/co1 (standard network, same as in (A)–(C)), nedefx2/co1 (double-sized network) (see section 2 and supplementary materials for details). The changes in the 1st window (black) and 2nd window are calculated with respect to baseline and highlight the importance of the network architecture for efficient applications of the FEP.




our modelled system is unique in terms of (a) the number of repetitions [27–29, 36], (b) the complexity of the neuronal network, (c) the absence of explicit learning rules.

The emergence of associative learning—through long-term plasticity in conventional neural networks—typically takes tens of cycles of stimulus presentation. In our case, the system starts to evince stable learning after only five repetitions. In short, a network of inferring neurons, that minimize a single (VFE) cost function, exhibit the functional properties and synaptic plasticity consistent with known experimental data, in the absence of bespoke learning rules.

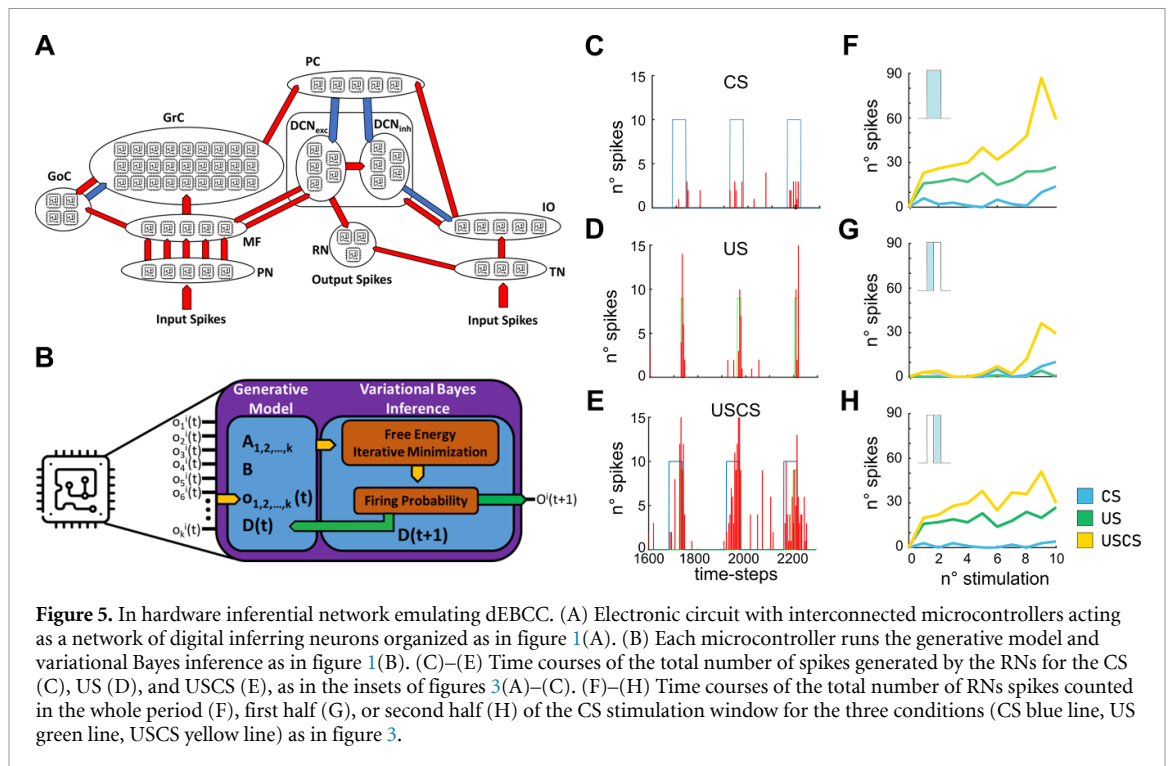
The most credited hypothesis—regarding the mechanisms that underwrite delayed eyeblink conditioning—is long-term synaptic plasticity in the cerebellar cortical circuit. In particular, experimental data have shown that long-term depression (LTD) at the GrCs-PCs synapses and the concomitant long-term potentiation (LTP) at MF-DCN synapses are required for learning in dBCC [20, 21] (a more complex scenario includes other forms of plasticity). To establish a parallel mechanism in our simulations, we monitored the evolution of synaptic precision, which is equivalent to synaptic efficacy.

At all circuit synapses, precision was updated according to FE minimization (section 2) but



only some synapses changed their precision or efficacy significantly. All GrC-PC connections increased precision during CS, whereas 15% of

GrC-PC synapses decreased precision during USCS (figure 4(A)). This result closely matches experimental observations of LTD at the parallel fibers to



PCs synapses, driven by the ‘teaching signal’ conveyed by IOs activity (figures 4(A) and (B)). Additionally, during USCS, precision increased at MF-DCN synapses (more markedly at MF-DCN_{exc} than MF-DCN_{inh}) (figure 4(C)). This analysis suggests that the selective expression of synaptic plasticity—at certain synapses in the network—is required for the dBCC to emerge (see [26, 34]). We explored different network configurations (see supplementary material, SI-3, SI-4, SI-5, SI-6, SI-7 and SI-8) to assess the robustness of these dBCC results. Associative learning was impaired when connections were either homogeneously or randomly distributed across the network (rather than being organized hierarchically as in biology), suggesting that learning is not a general property exhibited by any set of interconnected neurons but rather that it depends sensitively on network architecture (see supplementary materials). Notably, the selective deletion of either synaptic connections or groups of neurons reduced associative learning task and changed the time course of free energy minimization (see figures SI-7 and SI-8).

3.3. Active inference in digital electronic neurons

The approach described here lends itself to an efficient and expressive electronics in-hardware implementation. To showcase the advantages for an in-hardware implementation of the proposed paradigm—in which associative memory emerges spontaneously from network connectivity through active inference—we assembled a network of digital electronic neurons, with the same connectivity modelled in silico above (section 2 and supplemental information figures SI-9 and SI-10). The circuit

architecture is reported in figure 5(A). Notably, all digital neurons have the same structure, with the only difference between neuron classes (i.e. PCs, GrCs, GoCs) being the number of input lines and the corresponding number and class of synapses (excitatory and inhibitory). All digital neurons were implemented by means of identical low-power ARM processors (section 2), each one encoding the generative model and the variational Bayes inference engine that underwrites iterative VFE minimization, as reported in figure 5(B). The same stimulation protocols used for software simulation (i.e. CS, US, and combined USCS) were also supplied to the hardware network, by stimulating the input terminals of the digital neurons representing the PN and the TN (depending on the specific protocol) using voltage pulses that reproduced the spikes employed in the in silico simulations (see section 2).

As for the case of in silico simulations, each neuron was initialized with random activity and network activity was monitored by recording firing in the three RN digital neurons, using the same thresholding procedure adopted in silico (see section 2). Consistent with the results of software simulations, the number of spikes observed during the US protocol was almost constant in the stimulation period, whereas almost no response was detected at RNs neurons during the CS protocol. As expected, during the combined USCS protocol, the collective number of spikes produced by the RNs significantly increased from the fifth presentation of stimulus train onward (+75%; figure 5(F)). The hardware implementation also showed an increase and anticipation of responses during the combined

USCS (figures 5(F)–(H)), confirming that the electronic network was able to perform associative learning. Interestingly, the in-hardware implementation with digital neurons showed superior performance (full unsupervised training in about 15 s by dissipating just 0.5 W) and re-programmability, when compared to other existing solutions (see section 2).

4. Discussion

We have shown that a network of inferring neurons, connected within a brain-inspired architecture, can implement functional properties observed in nature through the local minimization of a general cost function. The FEP was applied to a cerebellar-like network subject to a classic associative learning paradigm, the dEBCC, which is perhaps the most investigated example of a predictive learning process. In this protocol, the statistical dependency of a conditioned stimulus, upon an unconditioned stimulus, is learnt by the neuronal microcircuit, setting the basis for correct anticipatory responses. The cerebellum is at the core of predictive systems in mammals and operates by comparing predicted states (e.g. the expected sensory consequences of a motor action) with actual states (e.g. the actual sensory feedback), learning from errors and updating its internal representations (through long-term synaptic plasticity) in order to generate predictions that minimize future errors. In nature, cerebellar learning involves several diverse forms of plasticity at multiple synaptic sites. The fact that, in our simulations, plasticity was expressed selectively to enable associative learning, supports the tenet that specific biochemical and cellular mechanisms at circuit synapses have evolved in order to minimize systemic free energy. Characteristic patterns emerged for synaptic weights, in particular those at the parallel fibers—PCs whose prominent role in dEBCC is supported by several experimental findings [45], and those at MFs—granule cell synapses, which varied from contact to contact, generating the rich set of transmission dynamics (figure S1) anticipated by a recent study [46]. The wiring architecture also proved crucial, as the predictive capabilities of the circuit changed with the neuronal population size and was lost by perturbing circuit hierarchy and connectivity (figures S3–S6). Thus, the specific architecture of the cerebellar network, which was early recognized to play a key role in the Motor Learning Theory [47], also appears to comply with the FEP, in the sense of model or structure learning [48–50].

The present approach lends itself well to generalization. When modelling a specific situation, first a plausible hierarchical neural network architecture for the given task can be selected, and then various alternatives are tried and compared using the (joint or collective) free energy as a performance metric. Furthermore, this procedure could be automatized so that it effectively performs structure learning

based on free-energy minimization [48–50]. In this (Bayesian model selection) scheme, the best architecture for the chosen task is automatically identified by the joint requirement of maximizing accuracy, while minimizing model complexity. This follows because model evidence or marginal likelihood can always be expressed as accuracy minus complexity [51]. The generality of such framework makes networks of inferring neurons promising candidates for building general-purpose devices. This scenario is supported by the results obtained following the selective disruption of specific groups of neurons and synapses that impaired associative learning (see figure 1). It would be interesting to also apply this approach to other brain circuits, for example the isocortex or the hippocampal formation.

The inference network introduced here shows properties that may have implications for computational technologies. First, it was more than an order of magnitude faster than classical artificial neural networks, which require hundreds of cycles for learning to occur [52, 53]. Secondly, learning was from one [54] to four orders [36] of magnitude smaller than other biologically inspired networks (e.g. [22–24, 34]), which typically involve thousands of units implementing specific learning rules at each synaptic stage. Thirdly, learning was implemented using conventional off-the-shelf hardware components, demonstrating low power consumption and high capability of dynamic reconfiguration. These unique features could result in unparalleled advantages for electronic implementation and artificial intelligence applications. For example, since the inference network capabilities emerge entirely from its connectivity and from the Bayesian engine in neurons (rather than from hardwired learning rules), the connections between neurons could simply be changed through programmable switch matrices, like those adopted in FPGA devices. Moreover, inference networks use the same Bayes inference engine and identical hardware for different classes of neuron. Thus, the ease of implementation and reprogramming would make the inferential networks quite attractive in terms of maintenance costs and smart re-use of resources.

A notable feature of this approach is that since every element in the network is trying to minimize its free energy, the joint free energy, namely the free energy of the ensemble, is minimized. The joint free energy is minimized because it is an expression of the mutual predictability of the system, and we also expect that a system in which everything is mutually predictable would also manifest a quantifiable thermodynamic efficiency. In this setting, it would be interesting to investigate how the free energy minimization time-course correlates with declining power consumption.

In summary, the inferential network approach presented here has implications extending beyond

the specificity of the dEBCC. The remarkable similarity between our biomimetic simulations and the kind of sentient behavior seen in real animals speaks to the importance of network structure: our results show that the emergent behavior of an ensemble of free energy minimizing elements depend sensitively on their composition and connectivity. This structure determines the ability of each element to predict the ensemble of which it is a constituent—and exogenous inputs (e.g. the US and CS). A key contribution of this work (i.e. proof of principle) is that a structured composition of isomorphic computing elements can form the basis of biomimetic computation. What is remarkable here is the simplicity of the generative model entailed by each element of the ensemble. Effectively, this generative model assumes one hidden state that can be in one of two levels. This is reminiscent of spin-glass models in physics [55] in which the states can be ‘up’ or ‘down’. In effect, each element is trying to infer whether its world is in an ‘up’ or ‘down’ state by assimilating sensory (presynaptic) inputs and acting upon its world by emitting action potentials; i.e. pinging other members of the ensemble. We refer to spin-glass models—and the binary nature of inference—to draw a connection with quantum mechanical formulations of the FEP [56]. In this setting, one could regard our simulations of learning as aligning quantum reference frames, so that the belief state of each atomic agent comes to ‘sense’ or ‘measure’ the states of others. On a quantum mechanical reading, the emergent synchronization—that underwrites the minimization of joint VFE—might be understood as entanglement under the Principle of Unitarity. Please see [56] for further discussion along these lines. This perspective on our experiments may be usefully revisited in the context of quantum computation and could be adopted to implement brain-like predictive capabilities into neuromorphic computers and robotic controllers through convenient software and hardware implementations. Active inference devices may turn out to represent a competitive alternative for several computational applications supporting energy-efficient [57] autonomous functions in artificial intelligence and robotics. Indeed, the importance of reward-free, self-organization—to maximize mutual predictability and generalized synchrony—is emerging in many fields; ranging from theoretical neuroscience [58] through to *in vitro* cell cultures [59]. The work reported here, shows that this fundamental kind of behavior can be instantiated both in silico—and in-hardware.

5. Conclusion

We have presented an innovative approach for modelling artificial neural networks, where each neuron behaves as if it is inferring (predicting) the causes of its own inputs, according to the principle of

free energy minimization. We have applied this approach to a biologically realistic cerebellar circuit in a classical conditioning paradigm and have shown that the network adjusts its activity and synaptic weights in agreement with known neurophysiology. Because of its peculiar features—very simple, general-purpose neuronal units, rapidity of learning, energetic efficiency—this approach could pave the way to the development of a new generation of biomimetic artificial intelligence systems. The selection of specific network architectures for given tasks could be performed through structure learning, based on the minimization of joint free energy of the network or ensemble. Furthermore, we believe that the implementation of such active inference units into tailored neuromorphic chips, may have the potential to revolutionize information technology, because of its remarkable combination of performance and energetic efficiency.

Data availability statement

Data and codes are available on a dedicated Github repository upon request to the corresponding authors.

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgment

We thank Dr Casper Hesp for helpful discussion during the revision phase of the manuscript.

Funding

This work has been funded by the following grants:

- D G is supported by Italian Ministry of university and Research (MUR): DM n° 1062 del 10/08/2021 PON—‘Ricerca e Innovazione 2014–2020’, ‘Azione IV.6 Contratti di ricerca su tematiche Green’
- M G B is supported by Italian Ministry of University and Research (MUR): ‘Dipartimenti di Eccellenza’ Program (2018–2022)
- K J F is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z) and a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1)
- SMART-BRAIN; partnering project to the European Union’s Horizon 2020 Framework Program for Research and Innovation under the specific GA N° 785907 (Human Brain Project SGA2) to J M.
- European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement N° 945539 (Human Brain Project SGA3) to E D.

Authors contributions

Conceptualization: D G, J M; Methodology: D G, F M P, J M; Simulation: D G, G M B; Hardware implementation: F M P; Investigation: D G, F M P, G M B, G P, E D A, J M; Writing—original draft: D G, F M P, G P, E D A, J M; Writing—review & editing: D G, F M P, G M B, G P, K J F, E D A, J M.

ORCID iDs

Daniela Gandolfi  <https://orcid.org/0000-0003-2315-2309>
 Francesco M Puglisi  <https://orcid.org/0000-0001-6178-2614>
 Giuseppe Pagnoni  <https://orcid.org/0000-0002-8272-8091>
 Karl J Friston  <https://orcid.org/0000-0001-7984-8909>
 Egidio D'Angelo  <https://orcid.org/0000-0002-6007-7187>
 Jonathan Mapelli  <https://orcid.org/0000-0002-0381-1576>

References

- [1] Kawato M 1999 Internal models for motor control and trajectory planning *Curr. Opin. Neurobiol.* **9** 718–27
- [2] Friston K J, Kilner J and Harrison L 2006 A free energy principle for the brain *J. Physiol.* **100** 70–87
- [3] Rao R P N and Seinowski T J 2002 Predictive coding, cortical feedback, and spike-timing dependent plasticity (Cambridge: MIT Press) pp 297–315
- [4] Friston K J 2010 The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11** 127–38
- [5] Dayan P, Hinton G E, Neal R M and Zemel R S 1995 The helmholtz machine *Neural Comput.* **7** 889–904
- [6] Moran R J, Campo P, Symmonds M, Stephan K E, Dolan R J and Friston K J 2013 Free energy, precision and learning: the role of cholinergic neuromodulation *J. Neurosci.* **33** 8227–36
- [7] Sales A C, Friston K J, Jones M W, Pickering A E and Moran R J 2019 Coeruleus tracking of prediction errors minimize cognitive flexibility: an active inference model *PLoS Comput. Biol.* **15** e1006267
- [8] Aitchinson L, Jegminat J, Menendez J A, Pfister J P, Pouget A and Latham P E 2021 Synaptic plasticity as Bayesian inference *Nat. Neurosci.* **24** 565–71
- [9] Parr T and Friston K J 2017 Working memory, attention, and salience in active inference *Sci. Rep.* **7** 14678
- [10] Friston K J 2010 Some free-energy puzzles resolved: response to Thornton *Trends Cogn. Sci.* **14** 54–55
- [11] Sella G and Hirsh A E 2005 The application of statistical physics to evolutionary biology *Proc. Natl Acad. Sci. USA* **102** 9541–6
- [12] Ivry R B, Spencer R M, Zelaznik H N and Diedrichsen J 2002 The cerebellum and event timing *Ann. New York Acad. Sci.* **978** 302–17
- [13] D'Angelo E and Casali S 2013 Seeking a unified framework for cerebellar function and dysfunction: from circuit operations to cognition *Front. Neural Circuits* **6** 116
- [14] Ivry R B and Baldo J V 1992 Is the cerebellum involved in learning and cognition? *Curr. Opin. Neurobiol.* **2** 212–6
- [15] Ghajar J and Ivry R B 2009 The predictive brain state: asynchrony in disorders of attention? *Neuroscientist* **15** 232–42
- [16] Ito M 2008 Control of mental activities by internal models in the cerebellum *Nat. Rev. Neurosci.* **9** 304–13
- [17] Mapelli J and D'Angelo E 2007 The spatial organization of long-term synaptic plasticity at the input stage of cerebellum *J. Neurosci.* **27** 1285–96
- [18] Gao Z, van Beugen B J and De Zeeuw C I 2012 Distributed synergistic plasticity and cerebellar learning *Nat. Rev. Neurosci.* **13** 619–35
- [19] Mapelli J, Gandolfi D, Vilella A, Zoli M and Bigiani A 2016 Heterosynaptic GABAergic plasticity bidirectionally driven by the activity of pre- and postsynaptic NMDA receptors *Proc. Natl Acad. Sci. USA* **113** 9898–903
- [20] Koekkoek S K E, Hulscher H C, Dortland B R, Hensbroek R A, Elgersma Y, Ruigrok T J H and De Zeeuw C I 2003 Cerebellar LTD and learning-dependent timing of conditioned eyelid responses *Science* **301** 1736–9
- [21] Friston K J and Herreros I 2016 Active inference and learning in the cerebellum *Neural Comput.* **28** 1812–39
- [22] Jarzynski C 1997 Nonequilibrium equality for free energy differences *Phys. Rev. Lett.* **78** 2690–3
- [23] Sengupta B, Stemmler M B and Friston K J 2013 Information and efficiency in the nervous system—a synthesis *PLoS Comput. Biol.* **9** e1003157
- [24] Yang S, Wang J, Li S, Deng B, Wei X, Yu X and Li H 2015 Cost-efficient FPGA implementation of basal ganglia and their Parkinsonian analysis *Neural Netw.* **71** 62–75
- [25] Yang S, Wang J, Li S, Deng B, Wei X, Yu X and Li H 2018 FPGA implementation of hippocampal spiking network and its real-time simulation on dynamical neuromodulation of oscillations *Neurocomputing* **282** 262–76
- [26] Yang S, Deng B, Wang J, Li H, Lu M, Che Y, Wei X and Loparo K A 2019 Scalable digital neuromorphic architecture for large-scale biophysically meaningful neural network with multi-compartment neurons *IEEE Trans. Neural Networks Learn. Syst.* **31** 148–62
- [27] Hampo M et al 2020 Associative memory in spiking neural network form implemented on neuromorphic hardware *ICONS 2020: Int. Conf. on Neuromorphic Systems 2020* vol 5 pp 1–8
- [28] Antonietti A, Martina D, Casellato C, D'Angelo E and Pedrocchi A 2019 Control of a humanoid NAO robot by an adaptive bioinspired cerebellar module in 3D motion tasks *Computational Intelligence and Neuroscience* vol 2019 pp 4862157
- [29] Antonietti A, Casellato C, Garrido J A, Luque N R, Naveros F, Ros E, D'Angelo E and Pedrocchi A 2016 Spiking neural network with distributed plasticity reproduces cerebellar learning in eye blink conditioning paradigms *IEEE Trans. Biomed. Eng.* **63** 210–9
- [30] An H, Zhou Z and Yi Y 2017 Memristor-based 3D neuromorphic computing system and its application to associative memory learning *Proc. 2017 IEEE 17th Int. Conf. on Nanotechnology (IEEE-NANO) (July 2017)* vol 2017 pp 555–60
- [31] Yang J, Wang L, Wang Y and Guo T 2017 A novel memristive hopfield neural network with application in associative memory *Neurocomputing* **227** 142–8
- [32] Liu X, Zeng Z and Wen S 2016 Implementation of memristive neural network with full-function pavlov associative memory *IEEE Trans. Circuits Syst.* **63** 1454–63
- [33] Eryilmaz S B 2014 Brain-like associative learning using a nanoscale non-volatile phase change synaptic device array *Front. Neurosci.* **8** 205
- [34] Moon K et al 2014 Hardware implementation of associative memory characteristics with analogue-type resistive-switching device *Nanotechnology* **25** 495204
- [35] Pershin Y and Di Ventra M 2009 Experimental demonstration of associative memory with memristive neural networks *Nat. Prec* **2009** 1
- [36] Yang S, Wang J, Zhang N, Deng B, Pang Y and Azghadi M R 2021 CerebellumMorphic: large-scale neuromorphic model and architecture for supervised motor learning *IEEE Trans. Neural Networks Learn. Syst.* **Feb 23** 1–15

- [37] Palacios E R, Isomura T, Parr T and Friston K J 2019 The emergence of synchrony in networks of mutually inferring neurons *Sci. Rep.* **9** 6412
- [38] Friston K J, FitzGerald T, Rigoli F, Schwartenbeck P and Pezzulo G 2017 Active inference: a process theory *Neural Comput.* **29** 1–49
- [39] Zhou S and Yu Y 2018 Synaptic E-I balance underlies efficient neural coding *Front. Neurosci.* **12**
- [40] Van de Laar T W and De Vries B 2019 Simulating active inference processes by message passing *Front. Robot. AI* **6** 20
- [41] Kschischang F R, Frey B J and Loeliger H A 2001 Factor graphs and the sum-product algorithm *IEEE Trans. Inf. Theory* **47** 498–519
- [42] Winn J, Bishop C M and Jaakkola T 2005 Variational message passing *J. Mach. Learn. Res.* **6** 661/694
- [43] Dauwels J 2007 On variational message passing on factor graphs 2007. *ISIT 2007. IEEE Int. Symp. on Information Theory (IEEE)* **2007** 2546–50
- [44] Painkras E, Plana L A, Garside J, Temple S, Galluppi F, Patterson C, Lester D R, Brown A D and Furber S B 2013 SpiNNaker: a 1-W 18-core system-on-chip for massively-parallel neural network simulation *IEEE J. Solid-State Circuits* **48** 1943–1953
- [45] Boele H J, Peter S, Ten Brinke M M, Verdonchot L, Jpelaar A C H I, Rizopoulos D, Gao Z, Koekkoek S K E and De Zeeuw C I 2018 Impact of parallel fiber to Purkinje cell long-term depression is unmasked in absence of inhibitory input *Sci. Adv.* **4** eaas9426
- [46] Casali S, Tognolina M L, Gandolfi D, Mapelli J and D'Angelo E 2020 Cellular-resolution mapping uncovers spatial adaptive filtering at the rat cerebellum input stage *Commun. Biol.* **3** 1–15
- [47] Eccles J C, Ito M and Szentagothai J 1967 *The Cerebellum as a Neuronal Machine* (Berlin: Springer)
- [48] Tervo D G, Tenenbaum J B and Gershman S J 2016 Toward the neural implementation of structure learning *Curr. Opin. Neurobiol.* **37** 99–105
- [49] Gershman S J 2017 Predicting the past, remembering the future *Curr. Opin. Behav. Sci.* **17** 7–13
- [50] Smith R, Schwartenbeck P, Parr T and Friston K J 2020 An active inference approach to modeling structure learning: concept learning as an example case *Front. Comput. Neurosci.* **14** 41
- [51] Penny W D, Stephan K E, Mechelli A and Friston K J 2004 Comparing dynamic causal models *Neuroimage* **22** 1157–72
- [52] Sanger T D and Kawato M 2020 A cerebellar computational mechanism for delay conditioning at precise time intervals *Neural Comput.* **32** 2069–84
- [53] Xu T, Xiao N, Zhai X, Chan P K and Tin C 2018 Real-time cerebellar neuroprosthetic system based on a spiking neural network model of motor learning *J. Neural Eng.* **15** 016021
- [54] Luque N R, Garrido J A, Carrillo R R, Tolu S and Ros E 2011 Adaptive cerebellar spiking model embedded in the control loop: context switching and robustness against noise *Int. J. Neural Syst.* **21** 385–401
- [55] Yedidia J S, Freeman W T and Weiss Y 2005 Constructing free-energy approximations and generalized belief propagation algorithms *IEEE Trans. Inf. Theory* **51** 2282–312
- [56] Fields C, Friston K, Glazebrook J F and Levin M 2021 A free energy principle for generic quantum systems (arXiv:2112.15242)
- [57] Obringer R, Rachunok B, Maia-Silva D, Arbabzadeh M, Nateghi R and Madani K 2021 The overlooked environmental footprint of increasing internet use *Resour. Conserv. Recycl.* **167** 105389
- [58] Isomura T, Parr T and Friston K 2019 Bayesian filtering with multiple internal models: toward a theory of social intelligence *Neural Comput.* **2013** 2390–431
- [59] Isomura T and Friston K J 2018 *In vitro* neural networks minimize variational free energy *Sci. Rep.* **8** 16926