SCIENTIFIC REPORTS

OPEN

Real time unsupervised learning of visual stimuli in neuromorphic VLSI systems

Massimiliano Giulioni¹, Federico Corradi^{1,2}, Vittorio Dante¹ & Paolo del Giudice^{1,3}

Received: 21 November 2014

Accepted: 12 August 2015

Published: 14 October 2015

Neuromorphic chips embody computational principles operating in the nervous system, into microelectronic devices. In this domain it is important to identify computational primitives that theory and experiments suggest as generic and reusable cognitive elements. One such element is provided by attractor dynamics in recurrent networks. Point attractors are equilibrium states of the dynamics (up to fluctuations), determined by the synaptic structure of the network; a 'basin' of attraction comprises all initial states leading to a given attractor upon relaxation, hence making attractor dynamics suitable to implement robust associative memory. The initial network state is dictated by the stimulus, and relaxation to the attractor state implements the retrieval of the corresponding memorized prototypical pattern. In a previous work we demonstrated that a neuromorphic recurrent network of spiking neurons and suitably chosen, fixed synapses supports attractor dynamics. Here we focus on learning: activating on-chip synaptic plasticity and using a theory-driven strategy for choosing network parameters, we show that autonomous learning, following repeated presentation of simple visual stimuli, shapes a synaptic connectivity supporting stimulus-selective attractors. Associative memory develops on chip as the result of the coupled stimulus-driven neural activity and ensuing synaptic dynamics, with no artificial separation between learning and retrieval phases.

Since its birthdate in 1989, with the publication of Carver Mead's book¹, the field of *neuromorphic engineering* aims at embodying computational principles operating in the nervous system into analog VLSI electronic devices. In a way, this endeavour may be seen as one modern instance of an over three centuries-long attempt to map forms of intelligent behavior onto a physical substrate reflecting the best technology of the day². The additional twist of neuromorphic engineering is a case for a direct mapping of the dynamics of neurons and synapses onto the physics of corresponding analog circuits. Initial success was mostly in emulating sensory functions (e.g. visual or auditory perception), and important developments are still ongoing in this area^{3–5}. However, it soon became clear that the agenda should include serious efforts to emulate, along with such implementations, elements of information processing downstream sensory stages, with the ultimate goal of approaching cognitive functions.

To make progress in this direction, beyond special-purpose solutions for specific functions, it seems important to identify neural circuitry implementing basic, and hopefully generic, dynamic building blocks, to provide reusable computational primitives, possibly subserving many types of information processing; in fact, this is both a theoretical quest and an item in the agenda of neuromorphic science. Steps in this direction have been taken recently in⁶, where 'soft winner-take-all' subnetworks provide reliable generic elements to compose finite-states machine capable of context-dependent computation. A review of the electronic circuits involved in such implementations is given in⁷.

In recurrent neural populations, synaptic self-excitation can support attractor dynamics, point attractors in the simplest instance on which our approach is based. Point attractors are stable configurations

¹Department of Technologies and Health, Istituto Superiore di Sanità, Roma, Italy. ²Institute of Neuroinformatics, University of Zürich and ETH Zürich, Switzerland. ³National Institute for Nuclear Physics, Rome, Italy. Correspondence and requests for materials should be addressed to M.G. (email: massimiliano.giulioni@iss.infn.it)

of the network dynamics; from any configuration inside the ‘basin’ of one attractor state, the dynamics brings the network towards that attractor, where it remains (possibly up to fluctuations if noise is present). In a system possessing several point attractor states, the dynamic correspondence between each attractor and its basin implements naturally an associative memory, the initial state within the basin being a metaphor of an initial stimulus, eliciting (even if removed afterwards) the retrieval of an associated prototypical information (memory). For a given network size and connectivity graph, the set of available attractor states is determined by the matrix of synaptic efficacies weighing the links of the graph; *Learning* memories is implemented through stimulus-specific changes in the synaptic matrix^{8,9}.

The attractor-basin correspondence implements dimensionality reduction. Besides, the stimulus-selective self-sustained neural activity following the removal of stimulus that elicited it, can act as a carrier of selective information across time intervals of unconstrained duration, only limited - in the absence of other intervening stimuli, by the stability of the attractor state against fluctuations^{8,9}. In a previous paper¹⁰ we demonstrated attractor dynamics in a neuromorphic chip, where synaptic efficacies were chosen and fixed so as to support the desired attractor states.

If attractor dynamics is to be considered as an interesting generic element of computation and representation for neuromorphic systems, we must address the question of how it can autonomously emerge from the ongoing stimulus-driven neural dynamics and the ensuing synaptic plasticity; this we do in the present work.

To date, sparse theoretical efforts have been devoted in this direction (see^{11–13}), and to our knowledge this has been never undertaken in a neuromorphic chip. Here, in line with our previous papers^{13,14}, and consistently with the above principles, we focus on the autonomous formation of attractor states as associative memories of simple visual objects.

Our setting is simple, in that our VLSI network learns two relatively simple, and non-overlapped, visual objects. Still, it is complex, in that learning is effected autonomously (that is, without a supervised mechanism to monitor errors and instruct synaptic changes); synapses change under the local (in space and time) guidance of the spiking activities of the neurons they connect, which in turn change their response to stimuli and their average activity because of synaptic modifications. Such a dynamic loop makes the combined dynamics of neurons and synapses during learning quite complex, and controlling it a tricky business; even more so in a neuromorphic analog chip, with the implied heterogeneities, mismatches and the like.

To gain predictive control on the chips’ learning dynamics, we first characterize the single-neuron input-output gain function. Then, we use the mean-field theory of recurrent neural networks as a compass to navigate the parameters space of a population of neurons endowed with massive positive feedback and predict attractor states. Finally we measure the rates of change (potentiation or depression) of the Hebbian, stochastic synapses as a function of the pre- and post-synaptic neural activities. These three characterization measures let us choose the correct settings for a successful learning trajectory. We then proceed with experiments on the autonomous learning capabilities of the system and finally, we test the attractor property of the developed internal representations of the learnt stimuli, by checking that when presented with a degraded version of such stimuli the network dynamically reconstructs the complete representation.

To our knowledge this is the first demonstration of a VLSI neuromorphic system implementing online, autonomous learning.

Results

Neuromorphic Multi-Chip System. The neuromorphic system (Fig. 1) is composed of a silicon retina¹⁵ and of two identical reconfigurable neural chips. Visual stimuli are displayed on a screen; the silicon retina captures the dynamic contrast of the stimuli and outputs spike sequences. Those spikes are fed into the recurrent learning network implemented on the two neural chips. The network spiking activity is streamed to a PC for analysis.

As shown in Fig. 1 the recurrent network consists of 196 *Excitatory* and 43 *Inhibitory* neurons physically distributed over the two identical chips. Retina pixels have been grouped into a grid of 14×14 macro-pixels, each one generating convergent output to a single excitatory neuron (see also S.I. Fig. 1). The recurrent synapses between excitatory neurons are plastic: their efficacy can change depending on the ongoing spiking activity (see below). ‘Learning’ is the stimulus-specific synaptic change induced by the repeated presentation of stimuli and the ensuing neural activity. Excitatory synapses are binary, in the sense that only two state of efficacy are allowed, a *Potentiated* and a *Depressed* one. Learning manifests itself as a sequence of transitions between these two states: as usual we name ‘LTD’, long-term depression the transition from a *Potentiated* to a *Depressed* state, and ‘LTP’, long-term potentiation the *Depressed* $\rightarrow$ *Potentiated* transition.

Recurrent synaptic connections are random and sparse (see Fig. 1). Randomness and sparseness, together with average low values of synaptic efficacy, guarantee low correlations among neuronal activity and, simultaneously, ensure a mean homogeneous input to all neurons. Under these conditions we can approximate the on-chip network behavior with mean-field theory equations (see S.I.) and use predictive theory-inspired tools (i.e. the Effective Transfer Function as explained below) to tune the system parameters. Moreover, the homogeneous connectivity is an ideally unbiased initial condition to test the

Figure 1. The neuromorphic system. A visual stimulus is shown on a screen in front of the retina chip. The retina, a matrix of 128×128 pixels, outputs spikes to two neural chips configured to host a recurrent network with a population E of 196 excitatory neurons and an inhibitory population I of 43 neurons (network's architecture sketched on top). Sparse connections, shown with a solid line, are generated randomly: connectivity levels, i.e. the probability of synaptic contact, are 0.25 for E recurrent connection, 0.5 for I to E connection, 0.3 for E to I and 0.02 from retina to I connection. The retina field of view is divided into 14×14 macropixels, each projecting onto a single neuron of E . The network's spiking activity is monitored by a PC.

effect of learning: we expect that synaptic plasticity will cluster the connectivity structure in a stimulus dependent manner.

We provide here a brief description of the neural chips already described in^{16,17}: the chips are composed each of 128 Integrate-and-fire (IF) neurons and 16384 Hebbian plastic bistable reconfigurable synapses. Neurons and synapses are designed as mixed signal analog/digital circuits; the internal dynamics of every element is implemented in continuous-time analog circuitry while communication among neurons relies on digital pulses representing the spikes. The entire chip works asynchronously and in real-time and does not necessitate any clocks. Every neuron and synapse is implemented in silicon with a dedicated circuit, in this way we are able to exploit in parallel all the resources without relying on complex multiplexing schemes. Hence the top-level view of the chip is an ordered matrix of 128×128 synaptic circuits connected to an array of 128 neurons. The drawback of the simplicity of the analog implementation is the unavoidable presence of mismatch deriving from the fabrication process. It causes distributions of the parameters among nominal identical circuits: one of the challenges faced in this work is to gain control over a network of mismatched elements. The analog synaptic circuit implements a plasticity model proposed in¹⁸, to which we refer the reader for details (see also S.I., section 3); a slightly different version was previously implemented in¹⁷. As already mentioned the synapse is binary, i.e. it has two levels of efficacy J_{pot} and J_{dep} – *Potentiated* and *Depressed* which are stable in the absence of pre- and post-synaptic neural activity; 'efficacy' is the amount of change in the membrane potential of the post-synaptic neuron per pre-synaptic spike. Plasticity is driven by neural spikes consistently with a rate-based Hebbian paradigm, and the changes of efficacy in each synapse are stochastic because of the irregularity of neuronal spikes in time. The synaptic connectivity is fully configurable, up to all-to-all connectivity. Each synapse can be set to be excitatory/inhibitory, and can be configured to connect two neurons on the same chip, or on different chips (including the retina), or to accept synthetic spikes from a PC. The communication of spikes from the retina to the neural chips, between the neural chips, and to the PC is based on the parallel asynchronous Address-Event-Representation (AER)^{5,19} and it is managed by a custom PCI-AER board²⁰, which also implements inter-chips communication.

Theory-inspired tools to control the neuromorphic system. As remarked in the Introduction, the successful autonomous, unsupervised development of associative representations of stimuli in the form of attractors of the network's dynamics can be challenged in many ways by the interplay between neural and synaptic dynamics. Because of this, and given the large network parameter space, setting neural and synaptic parameters by trial and error is not a viable route. To gain control over the system we proceed in three steps: 1) we tune neuronal parameters by measuring the input-output response function of a single-neuron 2) then, to set the synaptic efficacy values, we characterize the response of a recurrent sub-population of neurons connected by non-plastic synapses, 3) finally, to choose the settings

Figure 2. Long Term Potentiation (LTP) and Depression (LTD) probability. On the x axis we report the pre-synaptic neuronal mean firing rate ν_{pre} , on the y axis the post-synaptic neuronal firing rate ν_{post} . Details on the experimental procedure are reported in methods.

for the plasticity circuitry, we measure the LTD and LTP probability of the on-chip synapses when subjected to controlled neuronal activity. For the first and third measures we developed ad-hoc experimental procedures (described respectively in the second section of the S.I. and in the next section). To study the response of the recurrent neuronal sub-population we faced challenges related to the characterization of a system endowed with a massive positive feedback and embedded in a larger network. Relying on predictions based on theoretical models and mapping the derived parameters onto the hardware is not viable: on one side it requires lengthy calibration procedures and on the other side, anyhow, it is prone to fail due to unavoidable differences between the models' assumptions and analog circuital behavior, which easily became critical in presence of a massive positive feedback. As an alternative, we used the method introduced and validated in¹⁰ to identify regions in the parameter space compatible with the desired network behavior: the coexistence of (two, in our case) stable, stimulus-selective collective attractor states; in the S.I. section 2, we briefly summarize the procedure, and describe the results for the present system. The method instantiates the dimensional reduction of the mean-field theory for a multi-population network (proposed in²¹) in a self-consistent procedure on chip, allowing us to estimate the 'Effective Transfer Function' (ETF) of a population of neurons of interest (one selective excitatory population in our case). From an engineering point of view the ETF can be seen as the open-loop transfer function of the sub-population of interest, taking into account the feedback provided by the rest of the network. As detailed in the S.I. the ETF allows us to predict, within certain approximations, the mean firing rate of attractor states of a single sub-population for different levels of its recurrent synaptic potentiation. Since learning implies selective changes in the ratio of the potentiated/depressed synapses in the network, the study of ETF for the sub-populations which will be affected by the stimuli offers a predictive tool for expected learning histories.

Synaptic plasticity. The analysis based on the ETF gives an estimate of the expected changes in the average firing rates of the network's attractor states, as the ratio of potentiated/depressed synapses changes as a result of learning. In turn, expected rates of synaptic changes evolve depending on changes in the network's populations average activities. It is therefore relevant to derive an estimate of the LTP and LTD transition probabilities as a function of the pre- and post-synaptic firing rates (ν_{pre} , ν_{post}); the procedure is described in Materials and Methods section.

The contour plots in Fig. 2, panel B confirm that the silicon synapse implements, as desired, Hebbian LTP (its probability increases as ν_{pre} and ν_{post} both increase) and heterosynaptic LTD (i.e., synaptic depression probability increases with increasing ν_{pre} , and only occurs for low ν_{post}).

The joint information from the analysis of the ETF and the LTP/LTD transition probabilities allowed us to approximately predict the 'working point' of the plastic synapses at successive stages of learning (the prediction is expected to be more reliable for slow learning, such that the network evolves through quasi-equilibrium states, which are difficult to obtain for a small and heterogeneous network like the one on chip). Such knowledge allows to promote 'learning trajectories' with balanced LTP/LTD changes, which is another important stability factor¹³.

Figure 3. Learning Dynamics. (A) Sequence of input visual stimuli each active for 1 sec. as highlighted by the orange bands. (B) single-trial firing rates of the excitatory sub-populations responding to the *happy* (red) and *sad* (green) faces. (C) system output, i.e. network activity during the inter-stimulus intervals shown as a two-dimensional image, constructed from the map of the macro-pixels to the neurons. (D) degraded version of input stimuli and (E) corresponding network output.

Autonomous learning: forming stimulus-selective attractor states. Two visual stimuli were repeatedly presented on a LCD screen, acquired in real time by the silicon retina, and mapped onto the recurrent network distributed on two neural chips; our goal was to achieve autonomous associative learning leading to the formation of stimulus-selective attractor states as internal representations (‘memories’) of the stimuli. Hence, after learning, for every stimulus we expect a specific network response that should persist even after the removal of the stimulus.

Visual stimuli, (a ‘happy’ and a ‘sad’ face shown in panel A of Fig. 3), were orthogonal (zero overlap) and their coding level (fraction of activated macro-pixels, see S.I.) is fixed at about 1/3; Hence each stimulus will activate about 5460 retina neurons and, correspondingly, about 65 excitatory neurons. the implications of this for more realistic situations are discussed in the Discussion section.

Learning proceeds as a sequence of transitions between the J_{pot} and J_{dep} synaptic states, depending on the activity of pre- and post-synaptic neurons induced by stimuli, as explained below. Notice that no explicit control is imposed on synaptic dynamics depending on the network being stimulated or not, and no distinction is made between ‘learning’ and ‘retrieval’ phases; the network just evolves based on the incoming flow of stimuli, and its own feedback: our system’s learning is completely autonomous and does not require any supervision.

According to the Hebbian learning that synapses implement (as from Fig. 2), each stimulus presentation is expected to provoke changes in a fraction of synapses as follows (we remind that learning is stochastic, to the extent that neural activities are): LTP in synapses connecting neurons activated at high rates by the same stimulus; LTD in synapses connecting neurons activated by different stimuli, or connecting neurons activated by stimuli to neurons never activated by any stimuli (we named them the ‘background’ neurons); statistically little or no changes in synapses connecting pairs of ‘background’ neurons.

In the above scenario, because of autonomous learning, excitatory neurons get partitioned in three populations (two selective to stimuli, and the background), which both under stimulation and in the absence of it show an evolving pattern of relatively stable firing rates; as remarked in the Introduction, during this unsupervised evolution the network could well drift to undifferentiated high activity or quiescence states. Balance between LTP and LTD is essential to guarantee a successful learning trajectory¹⁴. LTP should eventually grow high enough to support stable selective attractor states; however, if this is not counterbalanced by LTD along the way, and especially in the early stages of learning, the learning trajectory can easily lead the network to globally unstable states; knowledge of the LTP/LTD probabilities of Fig. 2 is important to obtain robust learning trajectories. This is where the predictive power of the ETF, and knowledge LTP/LTD curves, come into play, hinting at safe paths in the large parameter space.

Our main results are described in Fig. 3, which illustrates a typical successful ‘learning history’, and in Fig. 4, which describes the underlying evolving micro-structure of the synaptic matrix.

Panels A,B,C of Fig. 3 describe respectively the sequence of stimuli, the average firing rates of the two selective populations and a two-dimensional representation of the network’s output to match the representation of stimuli. It is seen that: 1) initially, each stimulus provokes a response in its target neural population (which, through the induced inhibitory activation essentially silence the other populations, up to small noise); this activity is rapidly extinguished when the stimulus is removed, and no noticeable activity is present during the inter-stimulus interval 2) during each stimulation, a fraction of synapses changes, according to the (stochastic) scheme already explained, and this is reflected in the slow increase in the firing rates under stimulation 3) after many repeated stimulations, the build-up of synaptic self-excitation determines the appearance of a high-activity (meta)stable state for each of

Figure 4. Synaptic evolution. (A) Fraction of potentiated synapses. Traces named *inter-selective*, *selective to bkg* and *bkg to selective* refer respectively to the average fraction of potentiated synapses in the connections between the two selective sub-populations (happy and sad), from the two selective sub-populations to the background, and from the background to both happy and sad. (B) Hamming distance over time between two successive recordings of the synaptic matrix. (B–D) snapshots of the excitatory recurrent connection: gray (black) dots are depressed (potentiated) synapses between the pre-synaptic neurons (plotted on the x-axis) and post-synaptic ones (on the y-axis). Neurons responding to the same stimulus are grouped together. In red we highlight the group of synapses connecting neurons strongly reacting to the happy face. They get potentiated during learning. The group of synapses highlighted in blue are those connecting sad to happy population: as expected they get depressed during the initial learning stages.

the two selective sub-populations (the equivalent of the ETF passing from 1 to 2 stable fixed points in Fig. 3 in S.I.), and after the stimulus is released, the corresponding selective sub-population stays in a self-sustained state of elevated activity: the attractor state propagating the memory of the last stimulus across the inter-stimulus interval; sometimes this reverberant state is destabilized by the next incoming stimulus, while in other cases it decays spontaneously, due to finite-size fluctuation⁵³. Correspondingly, in a phase of mature learning, the output of the network during the inter-stimulus interval well matches the input stimuli.

A sample history of the synaptic changes underlying the learning history of Fig. 3 is given in Fig. 4, Panels A, B. Panel A illustrates the time course of the fraction of synapses in the potentiated state, for the different synaptic groups (identified by colour); synapses are inspected every 2 presentations of the same stimulus (4 stimulations in the alternate sequence). Consistently with expectation: synapses connecting neurons affected by the same stimulus get potentiated, actually approaching saturation, such that the average potentiation level remain essentially stable for $t > 300$ s; LTD is visible in the initial stage of learning; synapses connecting neurons in the background stay essentially unaffected.

Panel B shows the Hamming distance between the network synaptic matrices sampled as in panel A. We see here that the seemingly steady situation reached in panel A for the selective ‘sad to sad’ and ‘happy to happy’ synapses actually conceals a dynamic balance; for this to happen, again, LTP and LTD probabilities should be properly balanced. The evolution of the synaptic matrix is further illustrated in panels C–E of Fig. 4. The pictures are, from top to bottom, snapshots of the synaptic matrix taken at the beginning of learning, after 30 seconds, (2 presentations of each stimulus), and after 300 seconds (40 presentations).

The snapshot in panel C just reflects the random choice of 5% potentiated synapses as the initial condition of the network; in panel D we see that the synaptic matrix getting the expected structure, for the chosen index labelling: two blocks of mainly potentiated synapses (LTP), corresponding to synapses connecting neurons responsive to the same stimulus (remember that the stimuli are orthogonal, hence the non-overlapping ‘potentiated squares’); ‘whitened’ blocks of depressed synapses (LTD) connecting neurons responsive to different stimuli, or those connecting neurons responsive to either stimuli to background neurons; one strip on the extreme right, of synapses connecting neurons in the background, which stay essentially unaffected. Such features get further sharpened in the mature learning phase of panel E.

Error Correction properties. In order to check that the developed selective, self-sustaining states of elevated activity are indeed attractors of the network dynamics, 1) we observe their persistence after removal of the stimulus and 2) we check that the network is able to perform *error correction* of a degraded stimulus: in a mature stage of learning, for initial conditions (stimuli) close to one of the ‘memories’, the network spontaneously relaxes to the corresponding attractor state.

This we show in panel D of Fig 3, in which the stimulus is degraded by removing 20% of its active pixels. The network quickly reconstructs the complete, learnt memory, thanks to the selective feedback. We remark that systematically exploring the error-correction ability, by varying the amount of degradation of the stimuli, would allow to give an estimate of the *basin of attraction* of the attractor states. Indeed, this we showed in¹⁰ (in which the selective synaptic structure was imposed and not self-generated).

Discussion

We demonstrated a neuromorphic network of spiking neurons and plastic, Hebbian, spike-driven synapses which autonomously develops attractor representations of real visual stimuli acquired by a silicon retina. Attractor dynamics results in stimulus-selective, elevated activity after removal of the stimulus, and error correction properties. Detailed inspiration from mean-field theory is used to implement on chip a search strategy in the large parameter space of the network, which works at the level of the network’s input-output gain function, without having to delve into the single circuits level.

Within the limitations of the chosen scenario, which we discuss in the following, what we achieved takes a step along a most needed line of development, if neuromorphic systems are ultimately to act as autonomous and adaptive dynamic systems interacting in real time with their environment (for which, of course, much more than just building associative representations will be needed).

In view of further progress, some of the simplifying choices we made, and some limitations must be discussed and put in perspective.

First, it would be tempting to compare the performance of our neuromorphic network with theoretical predictions about the limit of capacity, i.e. the maximum number of networks configurations that can be embedded as retrievable memories. Hebbian dynamic learning with bounded (two-states in our case) synapses entail a capacity scaling as the logarithm of the number of synapses²² (as opposed to the linear scaling of the Hopfield network^{8,23} or its equivalent realization with spiking neurons²⁴); this limit can be partially overcome by making learning stochastic (as in our case) or lowering the coding level (see²⁵ for a discussion and further developments). However, a quantitative comparison with the theoretical estimates is affected by several factors (for the chip and largely for simulations as well), e.g. the fact that the probability of synaptic changes is not really constant along the learning history; that mismatch in the synaptic circuits effectively creates - at any stage of learning - a distribution of probabilities of synaptic changes across the network (notice that we do not compensate for mismatch modulating the connectivity among neurons as suggested in²⁶); that finite-size effects influence synaptic dynamics through the distribution of firing rates they induce¹³.

Empirically, the memory capacity of our small VLSI network is three patterns, beyond which learning becomes unstable; of course the small size of the network prevented us from checking how the capacity scales with the number of synapses. From the hardware point of view, achieving higher capacity with larger network hits at a scalability issue. Designing new chips that embed more neurons and synapses would provide a limited option for scaling up (unless radical changes in the implementation technology are considered - e.g. memristors). The other natural approach would be to combine many neural chips; this would face challenges related to the bandwidth for inter-chip communication, that would need to ensure reliable real-time AER-like spike delivery.

Second, we choose orthogonal stimuli (also with the same coding level) to be learnt by the network, which clearly facilitates learning by minimizing the interference induced on synaptic changes by different stimuli. A natural question then arises, as to the implication of including more naturalistic stimuli with arbitrary overlaps (and possibly coding levels). As for overlaps, a computationally effective, and

biologically motivated, modification of synaptic dynamics was proposed in²⁷; for the same stochastic, bistable Hebbian synaptic model here implemented, a regulatory mechanism was there introduced, preventing further potentiation or depression of a synapse when the post-synaptic neuron was recently highly active or poorly active, respectively. Such modified ‘stop-learning’ synapses were implemented and successfully demonstrated in neuromorphic chips^{17,28}, and would provide a good option to make the development of attractor representations more robust to the spatial structure of the stimuli. Allowing for different coding levels is instead essentially a matter of size of the network, allowing, even for low coding level, for a good averaging of synaptic mismatches across the dendritic tree of each neuron.

Putting our work in the context of previously published work, the attempt to embed learning capabilities in hardware devices is of course not new in general, starting from the supervised ‘adaptive pattern classification machine’ described by Widrow and Hoff²⁹ in the 60s.

Not exhaustively: one of the first steps towards learning microchips was taken in the 80s by Alspector and colleagues³⁰, who demonstrated unsupervised learning in a feed-forward network (for a broad review of pioneering works see³¹). More recently³², demonstrates a feedforward architecture to classify patterns of mean firing rates imposed by synthetic stimuli. A complex neuromorphic chain including visual sensing and processing (convolution filters) was described in³³.

Massive efforts have been devoted to develop various kinds of plastic synaptic circuits (using standard CMOS^{34–44} or with new memristor devices^{45–48}; for a recent review of technologies employed for the purpose see⁴⁹).

Still, very few works demonstrated learning in neuromorphic hardware at the network level. As mentioned above, Hebbian stop-learning synapses²⁷ have been used in pure feed-forward networks^{17,28} trained in a semi-supervised way to discriminate among synthetic patterns of mean firing rates, while⁵⁰ demonstrates how Spike-Timing-Dependent-Plasticity increases synchronicity in a network with local connectivity among neighboring neurons.

To our knowledge, learning synthetic stimuli in a spiking recurrent network with massive feedback has been dealt with only off-chip in⁵¹, reporting a digital implementation.

In the domain of recurrent spiking networks implementing associative memories this work takes a significant step towards the conditions of autonomous operation in naturalistic conditions by removing - as we did - the artificial separation between a ‘learning phase’, in which pre-computed synapses are ‘downloaded’ to the chip, and a ‘retrieval phase’, in which the associative memory is tested with frozen synapses.

Finally, concerning the prospective interest of point attractor networks as ‘reusable building blocks’ of neuromorphic systems (see Introduction) we would like to remark that, in the face of the rich repertoire of dynamic states exhibited by the brain (and neural networks), attempts are under way to bridge such elementary point attractor dynamics and the complexity of neural dynamics over multiple time scales (see⁵² and references therein), and the domain of application of the attractor concept has gradually evolved to include models of working memory, information integration, decision making, multi-stable perception. This provides, we believe, a prospective rich context for the implementation reported here.

Materials and Methods

Procedure for estimating LTP and LTD probabilities. We choose 64 excitatory neurons, and for each one we set up 128 input synapses, all configured as ‘AER synapses’, of which 64 are configured to be non-plastic (efficacy set to 0.05), the remaining 64 were plastic, with efficacy set to zero (in this way they do not affect the neuron’s firing, but the synapse can still perform transitions between its two binary states). In this way, firing ν_{post} of each neuron is driven by input from the 64 non-plastic AER synapses, receiving on their pre-synaptic terminal synthetic spike trains through the PCI-AER board; the plastic AER synapses are pre-synaptically driven by synthetic spike trains with average rate ν_{pre} . To explore the (ν_{pre}, ν_{post}) plane, for each ν_{post} (i.e. for each average rate of external spikes to the non-plastic synapses), ν_{pre} (the average rate of external spikes to the plastic synapses) is varied. For each (ν_{pre}, ν_{post}) pair, the initial state of the synapse (i.e. the value of the internal variable $X(0)$) is set (high, to measure LTD, low, to measure LTP), and its binary state is checked after 1 sec.

References

1. Mead, C. *Analog VLSI Implementation of Neural Systems*. (Springer, 1989).
2. Cordeschi, R. *The Discovery of the Artificial: Behavior, Mind and Machines Before and Beyond Cybernetics*. (Kluwer Academic Publisher, The Netherlands, 2002).
3. Liu, S. C., Van Schaik, A., Minch, B. A. & Delbruck, T. Asynchronous binaural spatial audition sensor with 2x64x4 channel output. *IEEE Trans. on Biomed. Circ. and Syst.* **99**, 1–12 (2013).
4. Liu, S. C. & Delbruck, T. Neuromorphic Sensory Systems. *Curr. Opin. in Neurobio.* **20**, 1–8, (Martin K. & Scott K. ed.) (2010).
5. Liu, S. C., Delbruck, T., Indiveri, G., Whatley, A. & Douglas, R. *Event-Based Neuromorphic Systems* (Wiley, 2015).
6. Nefci, E. *et al.* Synthesizing cognition in neuromorphic electronic systems. *Proc. Nat. Acad. Sci., PNAS.* **110**, E3468–E3476 (2012).
7. Chicca E., Stefanini, F., Bartolozzi, C. & Indiveri, G. Neuromorphic Electronic Circuits for Building Autonomous Cognitive Systems. *Proc. of the IEEE.* **102**, 1–22 (2014).
8. Amit, D. J. *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press, 1989).
9. Wang, X.-J. Attractor Network Models, *Encyclopedia of Neuroscience*. (eds Squire, L. R.), 667–679 (Oxford Academic Press, 2008).
10. Giulioni, M. *et al.* Robust working memory in an asynchronously spiking neural network realized in neuromorphic VLSI. *Front. in Neurosc.* **5**, doi: 10.3389/fnins.2011.00149 (2012).

11. Mongillo, G., Curti, E., Romani, S. & Amit, D. J. Learning in realistic networks of spiking neurons and spike-driven plastic synapses. *Eur. J. Neurosci.* **21**, 3143–3160 (2005).
12. Amit, D. J. & Brunel, N. Learning internal representations in an attractor neural network with analogue neurons. *Network: Computation in Neural Systems* **6**, 359–388 (1995).
13. Del Giudice, P., Fusi, S. & Mattia, M. Modelling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. *J. Physiol. Paris* **97**, 659–681 (2003).
14. Del Giudice, P. & Mattia, M. Long and short-term synaptic plasticity and the formation of working memory: A case study. *Neurocomp.* **38**, 1175–1180 (2001).
15. Lichtsteiner, P., Posch, C. & Delbruck, T. A 128x128, 120 dB 15 micro s latency asynchronous temporal contrast vision sensor. *IEEE Jour. of Solid St. Circ.* **43**, 566–576 (2008).
16. Giulioni, M. et al. A VLSI network of spiking neurons with plastic fully configurable “stop-learning” synapses. *Inter. Conf. on Elec., Circ., and Syst., ICECS*. 678–681; doi: 10.1109/ICECS.2008.4674944 (2008).
17. Giulioni, M., Pannunzi, M., Badoni, D., Dante, V. & Del Giudice, P. Classification of correlated patterns with a configurable analog VLSI neural network of spiking neurons and self-regulating plastic synapses. *Neural Comput.* **21**, 3106–3129 (2009).
18. Fusi, S., Annunziato, M., Badoni, D., Salamon, A. & Amit, D. J. Spike-driven synaptic plasticity: theory, simulation, VLSI implementation. *Neural Comput.* **12**, 2227–2258 (2000).
19. Mahowald, M. *An analog VLSI system for stereoscopic vision* (Springer, 1994).
20. Dante, V., Del Giudice, P. & Whatley, A. M. PCI-AER - hardware and software for interfacing to address-event based neuromorphic systems. *The Neuromorphic Engineer* **2**, 5–6 (2005).
21. Mascaro, M. & Amit, D. J. Effective neural response function for collective population states. *Network: Computation in Neural Systems*. **10**, 351–373 (1999).
22. Fusi, S. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cyber.* **87**, 459–470 (2002).
23. Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Nat. Acad. Sci., (PNAS) USA* **79**, 2554–2558 (1982).
24. Curti, E., Mongillo, G., La Camera, G. & Amit, D. J. Mean field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Comput.* **16**, 2597–2637 (2004).
25. Fusi, S. & Abbott, L. Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience*. **10**, 485–493 (2007).
26. Neftci, E. & Indiveri, G. A device mismatch compensation method for VLSI neural networks. *IEEE Biom. Circ. and Sys. Conf. (BioCAS)*. 262–265; doi: 10.1109/BIOCAS.2010.5709621 (2010).
27. Brader, J. M., Senn, W. & Fusi, S. Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* **19**, 2881–2912 (2007).
28. Mitra, S., Fusi, S. & Indiveri, G. Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. *IEEE Trans. on Biom. Circ. and Syst.* **3**, 32–42 (2009).
29. Widrow, B. & Hoff, M. E. *Adaptive switching circuits. Neurocomputing: foundations of research.* (The MIT press, Cambridge, USA, 1988).
30. Alspector, J., Gupta, B. & Allen, R. B. Performance of a stochastic learning microchip. *Advances in neural information processing systems.* Touretzky, D. S. (ed.) 748–760 (Morgan-Kaufmann, 1989).
31. Cauwenberghs, G. *Neuromorphic Learning VLSI Systems: a Survey* (eds Sejnowski, T. J.) 381–408 (Springer, 1998).
32. Hafliger, P. Adaptive WTA with an analog VLSI neuromorphic learning chip. *IEEE Trans. on Neural Networks*. **18**, 551–572 (2007).
33. Serrano-Gotarredona, R. et al. CAVIAR: A 45k neuron, 5M synapse, 12G connects/s AER hardware sensory processing learning actuating system for high-speed visual object recognition and tracking. *IEEE Trans. on Neural Networks*. **20**, 1417–1438 (2009).
34. Hafliger, P. & Kolle Riis, H. A multi-level static memory cell. *IEEE Int. Symp. on Circ. and Sys. (ISCAS)*. **1**, 25–28 (2003).
35. Gordon, C. & Hasler, P. Biological learning modeled in an adaptive floating-gate system. *IEEE Int. Symp. on Circ. and Sys.* **5**, 609–612 (2002).
36. Liu, S. C. & Mockel, R. Temporally learning floating-gate VLSI synapses. *IEEE Int. Symp. on Circ. and Sys. (ISCAS)*. 2154–2157; doi: 10.1109/ISCAS.2008.4541877 (2008).
37. Ramakrishnan, S., Hasler, P. E. & Gordon, C. Floating gate synapses with spike-time-dependent plasticity. *IEEE Trans. on Biom. Circ. and Sys. (BIOCAS)*. **5**, 244–252 (2011).
38. Boffill-I-Petit, A. & Murray, A. F. Synchrony detection and amplification by silicon neurons with STDP synapses. *IEEE Trans. on Neural Networks*. **15**, 1296–1304 (2004).
39. Schemmel, J., Grubl, A., Meier, K. & Mueller, E. Implementing synaptic plasticity in a VLSI spiking neural network model. *IEEE Joint Conf. on Neural Networks (IJCNN)*. 1–6; doi: 10.1109/IJCNN.2006.246651 (2006).
40. Bamford, S. A., Murray, A. F. & Willshaw, D. J. Spike-timing-dependent plasticity with weight dependence evoked from physical constraints. *IEEE Trans. on Biom. Circ. and Sys.* **6**, 385–398 (2012).
41. Rachmuth, G., Shouval, H. Z., Bear, M. F. & Poon, C. S. A biophysically-based neuromorphic model of spike rate- and timing-dependent plasticity. *Proc. Nat. Acad. Sci., (PNAS)* **108**, E1266–E1274 (2011).
42. Indiveri, G., Chicca, E. & Douglas, R. A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. *IEEE Trans. on Neural Network*. **17**, 211–221 (2006).
43. Indiveri, G., Stefanini, F. & Chicca, E. Spike-based learning with a generalized integrate and fire silicon neuron. *IEEE Int. Symp. on Circ. and Sys., (ISCAS)*, 1951–1954; doi: 10.1109/ISCAS.2010.5536980 (2010).
44. Brink, S. et al. A learning-enabled neuron array is based upon transistor channel models of biological phenomena. *IEEE Trans. on Biomed. Circ. and Syst.* **7**, 71–81 (2013).
45. Suri, M. et al. Physical aspects of low power synapses based on phase change memory devices. *J. of Appl. Phys.* **112**, 054904 (2012).
46. Mayr, C. et al. Waveform driven plasticity in BiFeO₃ memristive devices: model and implementation. *Adv. in Neur. Infor. Proc. Syst., (NIPS)*. **25**, 2323–2326 (2012).
47. Serrano-Gotarredona, T., Masquelier, T., Prodromakis, T., Indiveri, G. & Linares-Barranco, B. STDP and STDP variations with memristors for spiking neuromorphic learning systems. *Front. in Neurosci.* **7**, doi: 10.3389/fnins.2013.00002 (2013).
48. Gelencser, A., Prodromakis, T., Toumazou, C. & Roska, T. Biomimetic model of the outer plexiform layer by incorporating memristive devices. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **85**, 041918 (2012).
49. Kuzum, D., Yu, S. & Wong, H. P. Synaptic electronics: materials, devices and applications. *Nanotechnology*. **24**, 382001 (2013).
50. Arthur, J. & Boahen, K. Learning in silicon: timing is everything. *Advances in Neural Information Processing Systems*. 75–82 (The MIT Press, USA, 2006).
51. Seo, J. et al. A 45nm CMOS neuromorphic chip with a scalable architecture for learning in networks of spiking neurons. *IEEE Cust. Integr. Circ. Conf., CICC* (Piscataway, NJ). 1–4; doi: 10.1109/CICC.2011.6055293 (2011).
52. Braun, J. & Mattia, M. Attractors end noise: twin drivers of decisions and multistability. *Neuroimage* **52**, 740–751 (2010).
53. Mattia, M. & Del Giudice, P. Finite-size dynamics of inhibitory and excitatory interacting spiking neurons. *Phys. Rev. E Stat. Nonlin. Soft. Matter Phys.* **70**, 052903 (2004).

Acknowledgements

This work was partially supported by the EU FET project Coronet. We are indebted to Tobi Delbruck for having made available to us early prototypes of the silicon retina, and for his assistance. We gratefully thank Maurizio Mattia for his precious contribution of ideas and suggestions along the way, and Stefano Fusi for a critical reading of the manuscript.

Author Contributions

M.G. and F.C. wrote the software, performed the experiments, analysed the results and wrote the paper. V.D. designed the test system and provided the firmware. P.d.G. supervised the entire work and contributed to data analysis and paper writing. All authors reviewed the manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Giulioni, M. *et al.* Real time unsupervised learning of visual stimuli in neuromorphic VLSI systems. *Sci. Rep.* **5**, 14730; doi: 10.1038/srep14730 (2015).

This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>