






ARTICLE

DOI: 10.1038/s41467-018-04933-y

OPEN

Neuromorphic computing with multi-memristive synapses

Irem Boybat ^{1,2}, Manuel Le Gallo ¹, S.R. Nandakumar ^{1,3}, Timoleon Moraitis¹, Thomas Parnell¹, Tomas Tuma¹, Bipin Rajendran ³, Yusuf Leblebici², Abu Sebastian ¹ & Evangelos Eleftheriou¹

Neuromorphic computing has emerged as a promising avenue towards building the next generation of intelligent computing systems. It has been proposed that memristive devices, which exhibit history-dependent conductivity modulation, could efficiently represent the synaptic weights in artificial neural networks. However, precise modulation of the device conductance over a wide dynamic range, necessary to maintain high network accuracy, is proving to be challenging. To address this, we present a multi-memristive synaptic architecture with an efficient global counter-based arbitration scheme. We focus on phase change memory devices, develop a comprehensive model and demonstrate via simulations the effectiveness of the concept for both spiking and non-spiking neural networks. Moreover, we present experimental results involving over a million phase change memory devices for unsupervised learning of temporal correlations using a spiking neural network. The work presents a significant step towards the realization of large-scale and energy-efficient neuromorphic computing systems.

¹IBM Research - Zurich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland. ²Microelectronic Systems Laboratory, EPFL, Bldg ELD, Station 11, CH-1015 Lausanne, Switzerland. ³Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ 07102, USA. Correspondence and requests for materials should be addressed to I.B. (email: ibo@zurich.ibm.com) or to A.S. (email: ase@zurich.ibm.com)

The human brain with less than 20 W of power consumption offers a processing capability that exceeds the petaflops mark, and thus outperforms state-of-the-art supercomputers by several orders of magnitude in terms of energy efficiency and volume. Building ultra-low-power cognitive computing systems inspired by the operating principles of the brain is a promising avenue towards achieving such efficiency. Recently, deep learning has revolutionized the field of machine learning by providing human-like performance in areas, such as computer vision, speech recognition, and complex strategic games¹. However, current hardware implementations of deep neural networks are still far from competing with biological neural systems in terms of real-time information-processing capabilities with comparable energy consumption.

One of the reasons for this inefficiency is that most neural networks are implemented on computing systems based on the conventional von Neumann architecture with separate memory and processing units. There are a few attempts to build custom neuromorphic hardware that is optimized to implement neural algorithms^{2–5}. However, as these custom systems are typically based on conventional silicon complementary metal oxide semiconductor (CMOS) circuitry, the area efficiency of such hardware implementations will remain relatively low, especially if in situ learning and non-volatile synaptic behavior have to be incorporated. Recently, a new class of nanoscale devices has shown promise for realizing the synaptic dynamics in a compact and power-efficient manner. These memristive devices store information in their resistance/conductance states and exhibit conductivity modulation based on the programming history^{6–9}. The central idea in building cognitive hardware based on memristive devices is to store the synaptic weights as their conductance states and to perform the associated computational tasks in place.

The two essential synaptic attributes that need to be emulated by memristive devices are the synaptic efficacy and plasticity. Synaptic efficacy refers to the generation of a synaptic output based on the incoming neuronal activation. In conventional non-spiking artificial neural networks (ANN), the synaptic output is obtained by multiplying the real-valued neuronal activation with the synaptic weight. In a spiking neural network (SNN), the synaptic output is generated when the presynaptic neuron fires and typically is a signal that is proportional to the synaptic conductance. Using memristive devices, synaptic efficacy can be realized using Ohm's law by measuring the current that flows through the device when an appropriate read voltage signal is applied. Synaptic plasticity, in contrast, is the ability of the synapse to change its weight, typically during the execution of a learning algorithm. An increase in the synaptic weight is referred to as potentiation and a decrease as depression. In an ANN, the weights are usually changed based on the backpropagation algorithm¹⁰, whereas in an SNN, local learning rules, such as spike-timing-dependent plasticity (STDP)¹¹ or supervised learning algorithms, such as NormAD¹² could be used. The implementation of synaptic plasticity in memristive devices is achieved by applying appropriate electrical pulses that change the conductance of these devices through various physical mechanisms^{13–15}, such as ionic drift^{16–20}, ionic diffusion²¹, ferroelectric effects²², spintronic effects^{23,24}, and phase transitions^{25,26}.

Demonstrations that combine memristive synapses with digital or analog CMOS neuronal circuitry are indicative of the potential to realize highly efficient neuromorphic systems^{27–33}. However, to incorporate such devices into large-scale neuromorphic systems without compromising the network performance, significant improvements in the characteristics of the memristive devices are needed³⁴. Some of the device characteristics that limit the system performance include the limited conductance range, asymmetric conductance response (differences in the manner in which the

conductance changes between potentiation and depression), nonlinear conductance response (nonlinear conductance evolution with respect to the number of pulses), stochasticity associated with conductance changes, and variability between devices.

Clearly, advances in materials science and device technology could play a key role in addressing some of these challenges^{35,36}, but equally important are innovations in the synaptic architectures. One example is the differential synaptic architecture³⁷, in which two memristive devices are used in a differential configuration such that one device is used for potentiation and the other for depression. This was proposed for synapses implemented using phase change memory (PCM) devices, which exhibit strong asymmetry in their conductance response. However, the device mismatch within the differential pair of devices, as well as the need to refresh the device conductance frequently to avoid conductance saturation could potentially limit the applicability of this approach³⁴. In another approach proposed recently³⁸, several binary memristive devices are programmed and read in parallel to implement a synaptic element, exploiting the probabilistic switching exhibited by certain types of memristive devices. However, it may be challenging to achieve fine-tuned probabilistic switching reliably across a large number of devices. Alternatively, pseudo-random number generators could be used to implement this probabilistic update scheme with deterministic memristive devices³⁹, albeit with the associated costs of increased circuit complexity.

In this article, we propose a multi-memristive synaptic architecture that addresses the main drawbacks of the above-mentioned schemes, and experimentally demonstrate an implementation using nanoscale PCM devices. First, we present the concept of multi-memristive synapses with a counter-based arbitration scheme. Next, we illustrate the challenges posed by memristive devices for neuromorphic computing by studying the operating characteristics of PCM fabricated in the 90 nm technology node and show how multi-memristive synapses can address some of these challenges. Using comprehensive and accurate PCM models, we demonstrate the potential of the multi-memristive synaptic concept in training ANNs and SNNs for the exemplary benchmark task of handwritten digit classification. Finally, we present a large-scale experimental implementation of training an SNN with multi-memristive synapses using more than one million PCM devices to detect temporal correlations in event-based data streams.

Results

The multi-memristive synapse. The concept of the multi-memristive synapse is illustrated schematically in Fig. 1. In such a synapse, the synaptic weight is represented by the combined conductance of N devices. By using multiple devices to represent a synaptic weight, the overall dynamic range and resolution of the synapse are increased. For the realization of synaptic efficacy, an input voltage corresponding to the neuronal activation is applied to all constituent devices. The sum of the individual device currents forms the net synaptic output. For the implementation of synaptic plasticity, only one out of N devices is selected and programmed at a time. This selection is done with a counter-based arbitration scheme where one of the devices is chosen according to the value of a counter (see Supplementary Note 1). This selection counter takes values between 1 and N , and each value corresponds to one device of the synapse. After the weight update, the counter is incremented by a fixed increment rate. Having an increment rate co-prime with the clock length N guarantees that all devices in each synapse will eventually get selected and will receive a comparable number of updates provided there is a sufficiently large number of updates. Moreover, if

a single selection clock is used for all synapses of a neural network, N can be chosen to be co-prime with the total number of synapses in the network to avoid updating the same device in a synapse repeatedly.

In addition to the global selection counter, additional independent counters, such as a potentiation counter or a depression counter, could be incorporated to control the frequency of potentiation/depression events (see Fig. 1). The value of the potentiation (depression) counter acts as an enable signal to the potentiation (depression) event; a potentiation (depression) event is enabled if the potentiation (depression) counter value is one, and is disabled otherwise (see Supplementary Note 2). The frequency of the potentiation (depression) events is controlled by the maximum value or length of the potentiation (depression) counter. The counters are incremented after the weight update. By controlling how often devices are programmed for a conductance increase or decrease, asymmetries in the device conductance response can be reduced.

The constituent devices of the multi-memristive synapse can be arranged in either a differential or a non-differential architecture. In the latter each synapse consists of N devices, and one device is selected and potentiated/depressed to achieve synaptic plasticity. In the differential architecture, two sets of devices are present, and the synaptic conductance is calculated as $G_{syn} = G_+ - G_-$, where G_+ is the total conductance of the set representing the potentiation of the synapse and G_- is the total conductance of the set representing the depression of the synapse. Each set consists of $N/2$ devices. When the synapse has to be potentiated, one device from the group representing G_+ is selected and potentiated, and when the synapse has to be depressed, one device from the group representing G_- is selected and potentiated.

An important feature of the proposed concept is its crossbar compatibility. In the non-differential architecture, by placing the devices that constitute a single synapse along the bit lines of a crossbar, it is possible to sum up the currents using Kirchoff's law and obtain the total synaptic current without the need for any additional circuitry (see Supplementary Note 3). The differential architecture can be implemented with a similar approach, where

one bit line contains devices of the group G_+ and another those of the group G_- . The total synaptic current can then be found by subtracting the current of these two bit lines. To alter the synaptic weight, one of the word lines is activated according to the value of the selection counter to program the selected device. The scheme can also be adapted to alter the weights of multiple synapses in parallel within the constraints of the maximum current that could flow through the bit line (see Supplementary Note 3).

Multi-memristive synapses based on PCM devices. In this section, we will demonstrate the concept of multi-memristive synapses using nanoscale PCM devices. A PCM device consists of a layer of phase change material sandwiched between two metal electrodes (Fig. 2(a))⁴⁰, which can be in a high-conductance crystalline phase or in a low-conductance amorphous phase. In an as-fabricated device, the material is typically in the crystalline phase. When a current pulse of sufficiently high amplitude (referred to as the depression pulse) is applied, a significant portion of the phase change material melts owing to Joule heating. If the pulse is interrupted abruptly, the molten material quenches into the amorphous phase as a result of the glass transition. To increase the conductance of the device, a current pulse (referred to as the potentiation pulse) is applied such that the temperature reached via Joule heating is above the crystallization temperature but below the melting point, resulting in the recrystallization of part of the amorphous region⁴¹. The extent of crystallization depends on the amplitude and duration of the potentiation pulse, as well as on the number of such pulses. By progressively crystallizing the amorphous region by applying potentiation pulses, a continuum of conductance levels can be realized.

First, we present an experimental characterization of single-device PCM-based synapses based on doped $Ge_2Sb_2Te_5$ (GST) and integrated into a prototype chip in 90 nm CMOS technology⁴² (see Methods). Figure 2(b) shows the evolution of the mean device conductance as a function of the number of potentiation pulses applied. A total of 9700 devices were used for the characterization, and the programming pulse amplitude I_{prog}

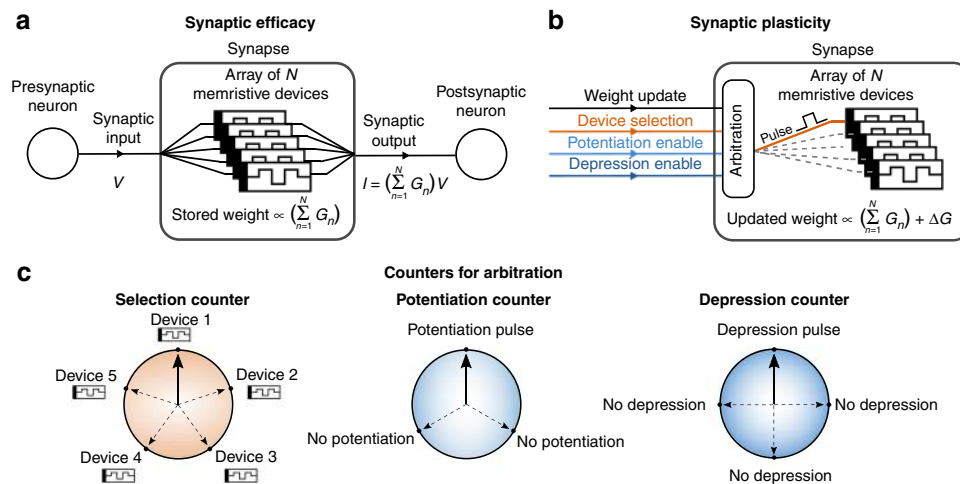


Fig. 1 The multi-memristive synapse concept. **a** The net synaptic weight of a multi-memristive synapse is represented by the combined conductance ($\sum G_n$) of multiple memristive devices. To realize synaptic efficacy, a read voltage signal, V , is applied to all devices. The resulting current flowing through each device is summed up to generate the synaptic output. **b** To capture synaptic plasticity, only one of the devices is selected at any instance of synaptic update. The synaptic update is induced by altering the conductance of the selected device as dictated by a learning algorithm. This is achieved by applying a suitable programming pulse to the selected device. **c** A counter-based arbitration scheme is used to select the devices that get programmed to achieve synaptic plasticity. A global selection counter whose maximum value is equal to the number of devices representing a synapse is used. At any instance of synaptic update, the device pointed to by the selection counter is programmed. Subsequently, the selection counter is incremented by a fixed amount. In addition to the selection counter, independent potentiation and depression counters can serve to control the frequency of the potentiation or depression events

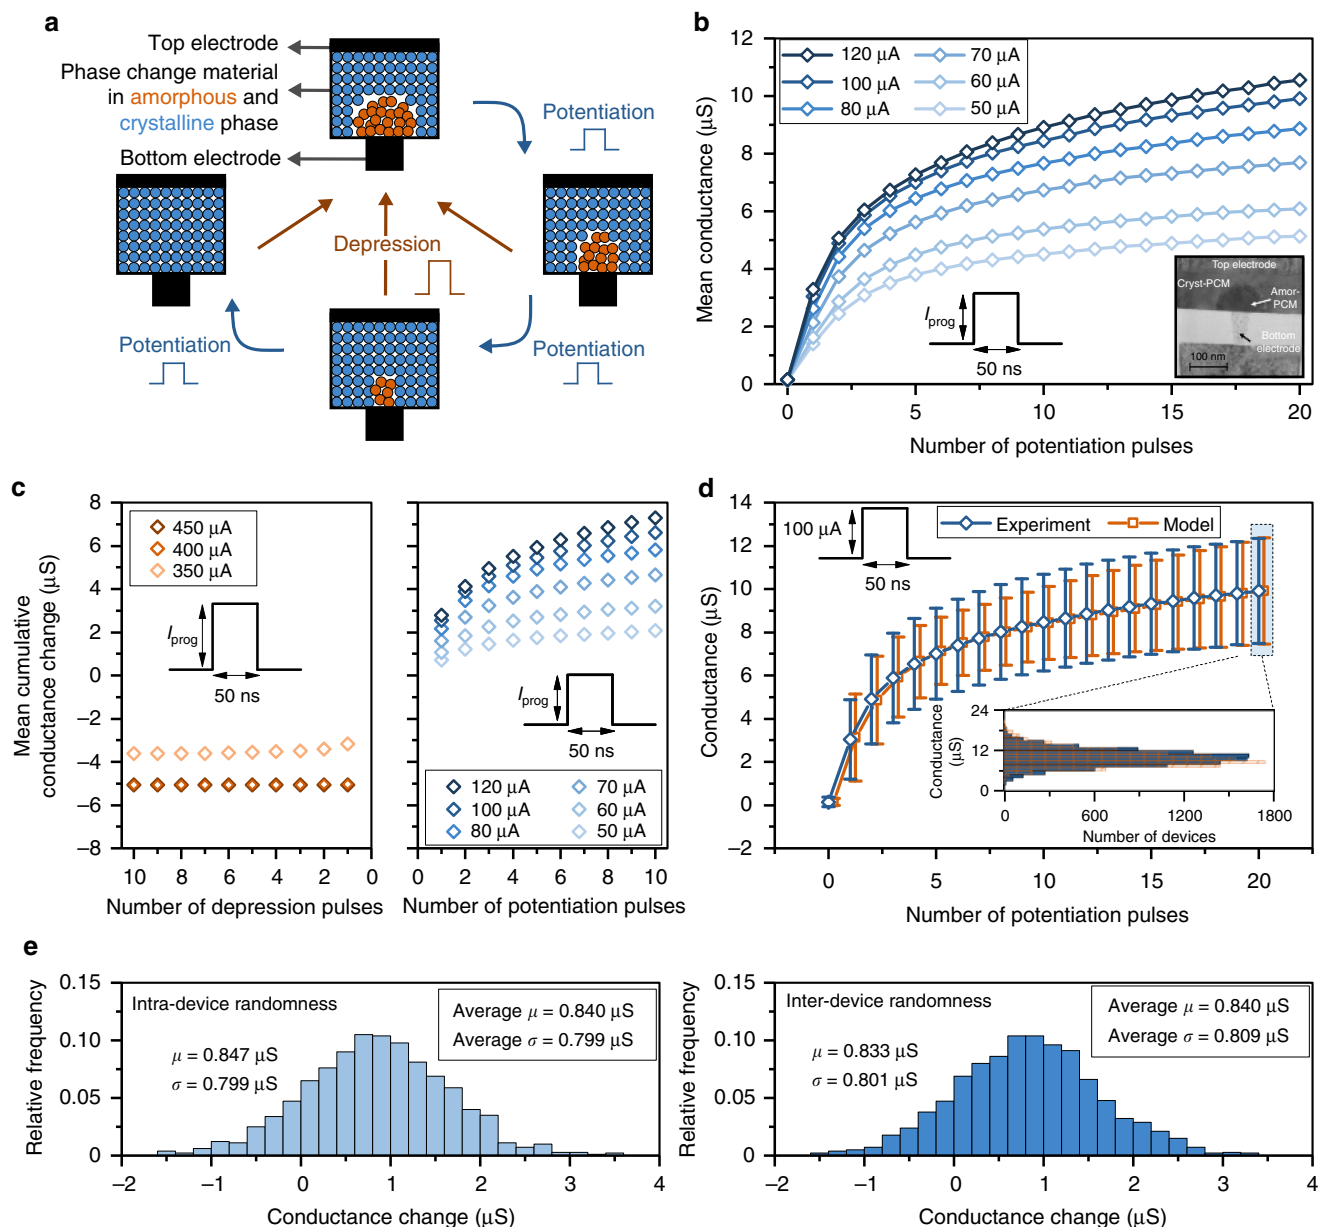


Fig. 2 Synapses based on phase change memory. **a** A PCM device consists of a phase-change material layer sandwiched between top and bottom electrodes. The crystalline region can gradually be increased by the application of potentiation pulses. A depression pulse creates an amorphous region that results in an abrupt drop in conductance, irrespective of the original state of the device. **b** Evolution of mean conductance as a function of the number of pulses for different programming current amplitudes (I_{prog}). Each curve is obtained by averaging the conductance measurements from 9700 devices. The inset shows a transmission electron micrograph of a characteristic PCM device used in this study. **c** Mean cumulative conductance change observed upon the application of repeated potentiation and depression pulses. The initial conductance of the devices is $\sim 5 \mu\text{S}$. **d** The mean and the standard deviation (1σ) of the conductance values as a function of number of pulses for $I_{\text{prog}} = 100 \mu\text{A}$ measured for 9700 devices and the corresponding model response for the same number of devices. The distribution of conductance after the 20th potentiation pulse and the corresponding distribution obtained with the model are shown in the inset. **e** The left panel shows a representative distribution of the conductance change induced by a single pulse applied at the same PCM device 1000 times. The pulse is applied as the 4th potentiation pulse to the device. The same measurement was repeated on 1000 different PCM devices, and the mean (μ) and standard deviation (σ) averaged over the 1000 devices are shown in the inset. The right panel shows a representative distribution of one conductance change induced by a single pulse on 1000 devices. The pulse is applied as the 4th potentiation pulse to the devices. The same measurement was repeated for 1000 conductance changes, and the mean and standard deviation averaged over the 1000 conductance changes are shown in the inset. It can be seen that the inter-device and the intra-device variability are comparable. The negative conductance changes are attributed to drift variability (see Supplementary Note 4)

was varied from 50 to 120 μA . It can be seen that the mean conductance value increases as a function of the number of potentiation pulses. The dynamic range of conductance response is limited as the change in the mean conductance decreases and eventually saturates with increasing number of potentiation

pulses. Figure 2(c) shows the mean cumulative change in conductance as a function of the number of pulses for different values of I_{prog} . A well-defined nonlinear monotonic relationship exists between the mean cumulative conductance change and the number of potentiation pulses. In addition, there is a granularity

that is determined by how small a conductance change can be induced by applying a single potentiation pulse. Large conductance change granularities, as well as nonlinear conductance responses, both observed in the PCM characterization performed here, have been shown to degrade the performance of neural networks trained with memristive synapses^{34,43}. Moreover, when a conductance decrease is desired, a single high-amplitude depression pulse applied to a PCM device has an all-or-nothing effect that fully depresses the device conductance to (almost) 0 μS . Such a strongly asymmetric conductance response is undesirable in memristive-device-based implementations of neural networks⁴⁴, and this is a significant challenge for PCM-based synapses. Depression pulses with smaller amplitude could be applied to achieve higher conductance values. However, unlike the potentiation pulses, it is not possible to achieve a progressive depression by applying successive depression pulses.

There are also significant intra-device and inter-device variabilities associated with the conductance response in PCM devices as evidenced by the distribution of conductance values upon application of successive potentiation pulses (see Fig. 2(d)). Note that the variability observed in these devices fabricated in the 90 nm technology node is also found to be higher than that of those fabricated in the 180 nm node as reported elsewhere³⁴. Both the mean and variance associated with the conductance change depend on the mean conductance value of the devices. We capture this behavior in a PCM conductance response model that relies on piece-wise linear approximations to the functions that link the mean and variance of the conductance change to the mean conductance value⁴⁵. As shown in Fig. 2(d), this model approximates the experimental behavior fairly well.

The intra-device variability in PCM is attributed to the differences in atomic configurations associated with the amorphous phase change material created during the melt-quench process⁴⁶. Inter-device variability, on the other hand, arises predominantly from the variability associated with the fabrication process across the array and results in significant differences in the maximum conductance and conductance response across devices (see Supplementary Fig. 1). To investigate the intra-device variability, we measured the conductance change on the same PCM device induced by a single potentiation pulse of amplitude $I_{\text{prog}} = 100 \mu\text{A}$ over 1000 trials (Fig. 2(e), left panel). To quantify the inter-device variability, we monitored the conductance change induced by a single potentiation pulse across the 1000 devices (Fig. 2(e), right panel). These experiments show that the standard deviation of the conductance change due to intra-device variability is almost as large as that due to the inter-device variability. The finding that the randomness in the conductance change is to a large extent intrinsic to the physical characteristic of the device implies that improvements in the array-level variability will not necessarily be effective in reducing the randomness.

The characterization work presented so far highlights the challenges associated with synaptic realizations using PCM devices and these can be generalized to other memristive technologies. The limited dynamic range, the asymmetric and nonlinear conductance response, the granularity and the randomness associated with conductance changes all pose challenges for realizing neural networks using memristive synapses. We now show how our concept of multi-memristive synapses can help in addressing some of those challenges. Experimental characterizations of multi-memristive synapses comprising 1, 3, and 7 PCM devices per synapse arranged in a non-differential architecture are shown in Fig. 3(a). The conductance change is averaged over 1000 synapses. One selection counter with an increment rate of one arbitrates the device selection. As the total conductance is the sum of the individual conductance values, the dynamic range

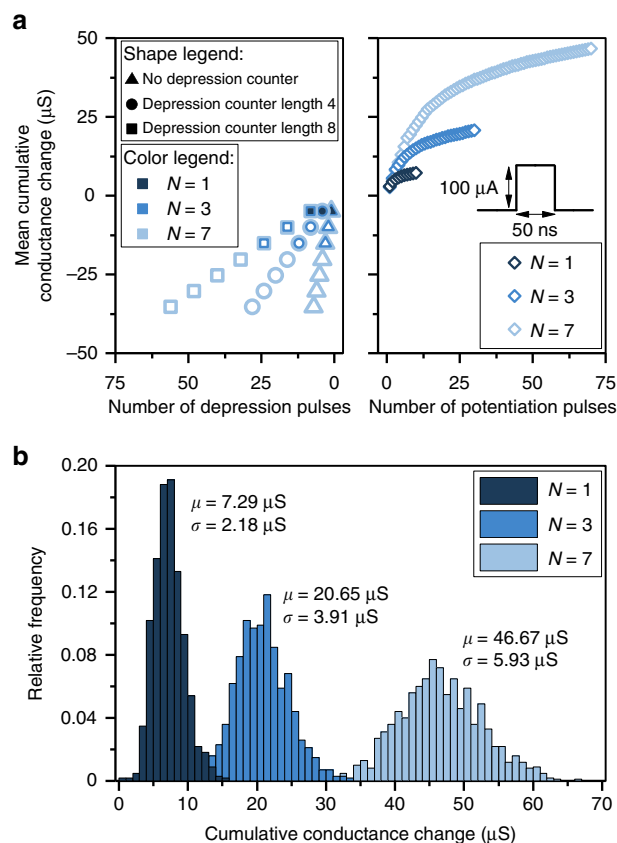


Fig. 3 Multi-memristive synapses based on phase change memory. **a** The mean cumulative conductance change is experimentally obtained for synapses comprising 1, 3, and 7 PCM devices. The measurements are based on 1000 synapses, whereby each individual device is initialized to a conductance of $\sim 5 \mu\text{S}$. For potentiation, a programming pulse of $I_{\text{prog}} = 100 \mu\text{A}$ was used, whereas for depression, a programming pulse of $I_{\text{prog}} = 450 \mu\text{A}$ was used. For depression, the conductance response can be made more symmetric by adjusting the length of the depression counter. **b** Distribution of the cumulative conductance change after the application of 10, 30, and 70 potentiation pulses to 1, 3, and 7 PCM synapses, respectively. The mean (μ) and the variance (σ^2) scale almost linearly with the number of devices per synapse, leading to an improved weight update resolution

scales linearly with the number of devices per synapse. Alternatively, for a learning algorithm requiring a fixed dynamic range, multi-memristive synapses can improve the effective conductance change granularity. In addition, in contrast to a single device, the mean cumulative conductance change here is linear over an extended range of potentiation pulses. With multiple devices, we can also partially mitigate the challenge of an asymmetric conductance response. At any instance, only one device is depressed, which implies that the effective synaptic conductance decreases gradually in several steps instead of the abrupt decrease observed in a single device. Moreover, using the depression counter, the cumulative conductance changes for potentiation and depression can be made approximately symmetric by adjusting the frequency of depression events. Finally, Fig. 3(b) shows that both the mean and the variance of the conductance change scale linearly with the number of devices per synapse. Hence, the smallest achievable mean weight change decreases by a factor of N , whereas the standard deviation of the weight change decreases by \sqrt{N} , leading to an overall increase in weight update resolution by \sqrt{N} (see Supplementary Fig. 2).

Simulation results on handwritten digit classification. In this section, we study the impact of PCM-based multi-memristive synapses in the context of training ANNs and SNNs. For synaptic potentiation, the PCM conductance response model presented above was used (see Fig. 2(d)). The depression pulses are assumed to cause an abrupt conductance drop to zero in a deterministic manner, modeling the PCM asymmetry. One selection counter is used for all synapses of the network, and the weight updates are done sequentially through all synapses in the same order at every pass. Potentiation and depression counters are used to balance the frequency of potentiation and depression events for $N > 1$.

First, we present simulation results that show the performance of an ANN trained with multi-memristive synapses based on the nonlinear conductance response model of the PCM devices. The feedforward fully-connected network with three neuron layers is trained with the backpropagation algorithm to perform a classification task on the MNIST data set of handwritten digits⁴⁷ (see Fig. 4(a) and Methods). The ideal classification performance of this network, assuming double-precision floating-point accuracy for the weights, is 97.8%. The synaptic weights are represented using the conductance values of a multi-memristive synapse model. In the non-differential architecture, a depression counter is used to improve the asymmetric conductance response and a potentiation counter to reduce the frequency of the

potentiation events. As shown in Fig. 4(a), the classification accuracy improves with the number of devices per synapse. With the conventional differential architecture with two devices, the classification accuracy is below 15%. With multi-memristive synapses in the differential architecture, we can achieve test accuracies exceeding 88.9%, a performance better than the state-of-the-art in situ learning experiments on PCM despite a significantly more nonlinear and stochastic conductance response due to technology scaling³⁴. Remarkably, accuracies exceeding 90% are possible even with the non-differential architecture, which clearly illustrates the efficacy of the proposed scheme.

In a second investigation, we studied an SNN with multi-memristive synapses to perform the same task of digit recognition, but with unsupervised learning⁴⁸ (see Fig. 4(b) and Methods). The weight updates are performed using an STDP rule: the synapse is potentiated whenever a presynaptic neuronal spike appears prior to a postsynaptic neuronal spike, and depressed otherwise. The amount of weight increase (decrease) within the potentiation (depression) window is constant and independent of the timing difference between the spikes. This necessitates a certain weight update granularity, which can be achieved by the proposed approach. The classification performance of the network trained with this rule using double-precision floating-point accuracy for the network parameters is

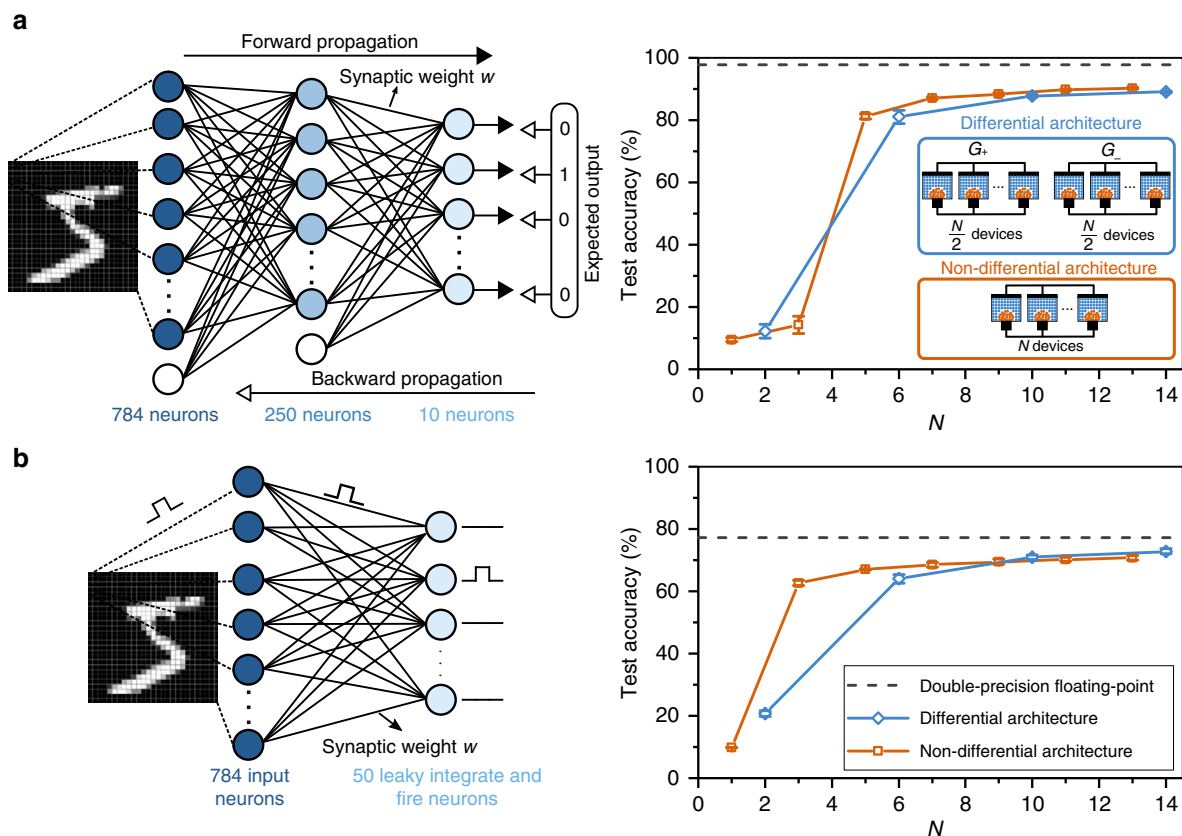


Fig. 4 Applications of multi-memristive synapses in neural networks. **a** An artificial neural network is trained using backpropagation to perform handwritten digit classification. Bias neurons are used for the input and hidden neuron layers (white). A multi-memristive synapse model based on the nonlinear conductance response of PCM devices is used to represent the synaptic weights in these simulations. Increasing the number of devices in multi-memristive synapses (both in the differential and the non-differential architecture) improves the test accuracy. Simulations are repeated for five different weight initializations. The error bars represent the standard deviation (1σ). The dotted line shows the test accuracy obtained from a double-precision floating-point software implementation. **b** A spiking neural network is trained using an STDP-based learning rule for handwritten digit classification. Here again, a multi-memristive synapse model is used to represent the synaptic weights in simulations where the devices are arranged in the differential or the non-differential architecture. The classification accuracy of the network increases with the number of devices per synapse. Simulations are repeated for five different weight initializations. The error bars represent the standard deviation (1σ). The dotted line shows the test accuracy obtained from a double-precision floating-point implementation

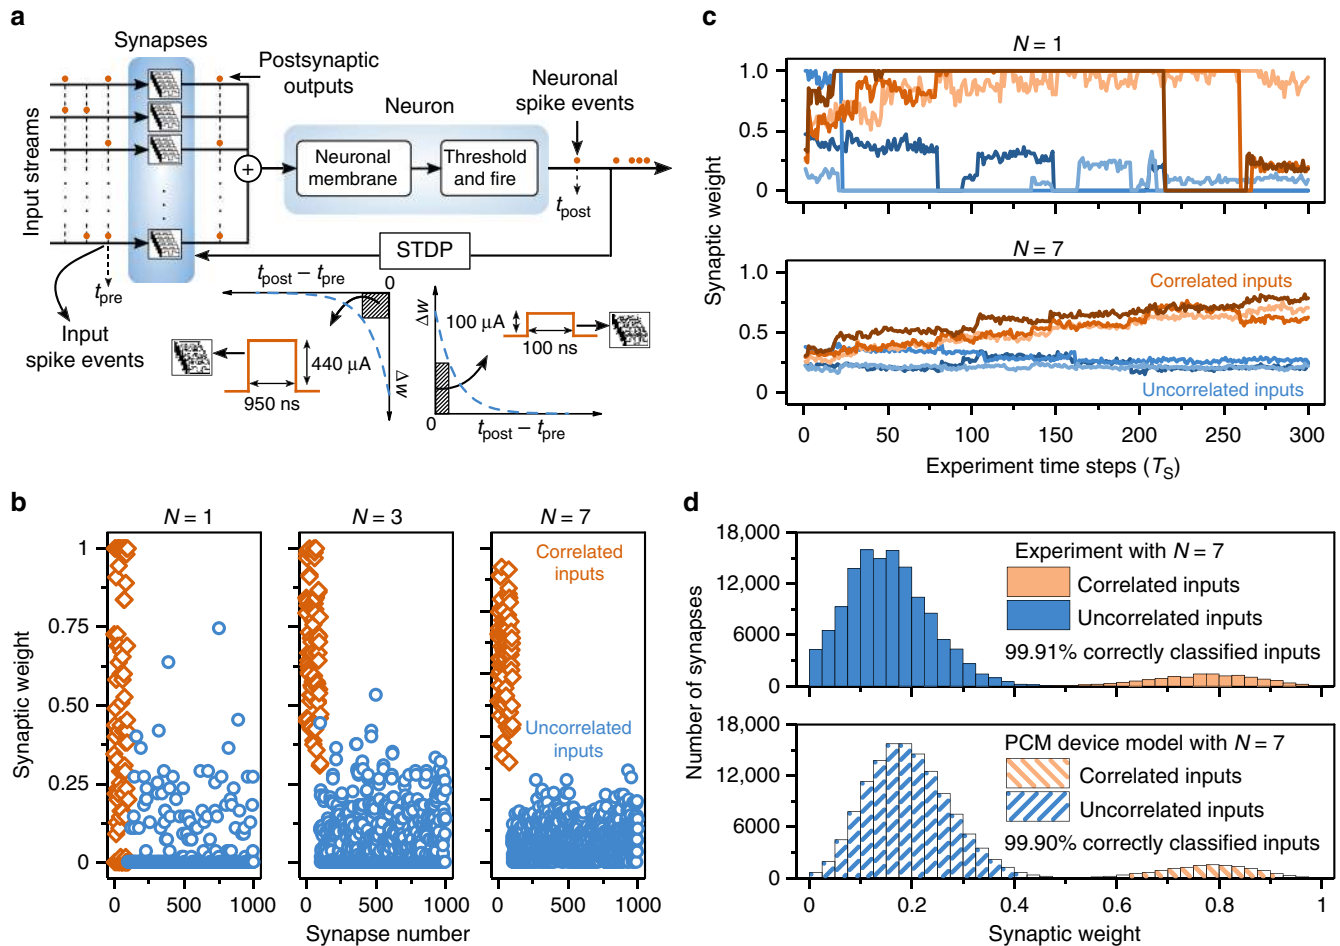


Fig. 5 Experimental demonstration of multi-memristive synapses used in a spiking neural network. **a** A spiking neural network is trained to perform the task of temporal correlation detection through unsupervised learning. Our network consists of 1000 multi-PCM synapses (in hardware) connected to one integrate-and-fire (I&F) software neuron. The synapses receive event-based data streams generated with Poisson distributions as presynaptic input spikes. 100 of the synapses receive correlated data streams with a correlation coefficient of 0.75, whereas the rest of the synapses receive uncorrelated data streams. The correlated and the uncorrelated data streams both have the same rate. The resulting postsynaptic outputs are accumulated at the neuronal membrane. The neuron fires, i.e., sends an output spike, if the membrane potential exceeds a threshold. The weight update amount is calculated using an exponential STDP rule based on the timing of the input spikes and the neuronal spikes. A potentiation (depression) pulse with fixed amplitude is applied if the desired weight change is higher (lower) than a threshold. **b** The synaptic weights are shown for synapses comprising $N = 1, 3,$ and 7 PCM devices at the end of the experiment (5000 time steps). It can be seen that the weights of the synapses receiving correlated inputs tend to be larger than the weights of those receiving uncorrelated inputs. The weight distribution shows a clearer separation with increasing N . **c** Weight evolution of six synapses in the first 300 time steps of the experiment. The weight evolves more gradually with the number of devices per synapse. **d** Synaptic weight distribution of an SNN comprising 144,000 multi-PCM synapses with $N = 7$ PCM devices at the end of an experiment (3000 time steps) (upper panel). 14,400 synapses receive correlated input data streams with a correlation coefficient of 0.75. A total of 1,008,000 PCM devices are used for this large-scale experiment. The lower panel shows the synaptic weight distribution predicted by the PCM device model

77.2%. A potentiation counter is used to reduce the frequency of the potentiation events in both the differential and non-differential architectures, and a depression counter is used in the non-differential architecture to improve the asymmetric conductance response. The network could classify more than 70% of the digits correctly for $N > 9$ with both the differential and the non-differential architecture, whereas the network with the conventional differential architecture with two devices has a classification accuracy below 21%.

In both cases, we see that the multi-memristive synapse significantly outperforms the conventional differential architecture with two devices, clearly illustrating the effectiveness of the proposed architecture. Moreover, the fact that the non-differential architecture achieves a comparable performance to that of the differential architecture is promising for synaptic realizations using highly asymmetric devices. A non-differential

architecture would have a lower implementation complexity than its differential counterpart because the refresh operation^{34,37}, which requires reading and reprogramming G_+ and G_- , can be completely avoided.

Experimental results on temporal correlation detection. Next, we present an experimental demonstration of the multi-memristive synapse architecture using our prototype PCM chip (see Methods) to train an SNN that detects temporal correlations in event-based data streams in an unsupervised way. Unsupervised learning is widely perceived as a key computational task in neuromorphic processing of big data. It becomes increasingly important given today's variety of big data sources, for which often neither labeled samples nor reliable training sets are available. The key task of unsupervised learning is to reveal the

statistical features of big data, and thereby shed light on its internal correlations. In this respect, detecting temporal and spatial correlations in the data is essential.

The SNN comprises a neuron interfaced to plastic synapses, with each one receiving an event-based data stream as presynaptic input spikes^{49,50} (see Fig. 5(a) and Methods). A subset of the data streams are mutually temporally correlated, whereas the rest are uncorrelated (see Supplementary Note 5). When the input streams are applied, postsynaptic outputs are generated at the synapses that received a spike. The resulting postsynaptic outputs are accumulated at the neuron. When the neuronal membrane potential exceeds a threshold, the output neuron fires, generating a spike. The synaptic weights are updated using an STDP rule; synapses that receive an input spike within a time window before (after) the neuronal spike get potentiated (depressed). As it is more likely that the temporally correlated inputs will eventually govern the neuronal firing events, the conductance of synapses receiving correlated inputs is expected to increase, whereas that of synapses whose input are uncorrelated is expected to decrease. Hence, the final steady-state distribution of the weights should display a separation between synapses receiving correlated and uncorrelated inputs.

First, we perform small-scale experiments in which multi-memristive synapses with PCM devices are used to store the synaptic weights. The network comprises 1000 synapses, of which only 100 receive temporally correlated inputs with a correlation coefficient c of 0.75. The difficulty in detecting whether an input is correlated or not increases both with decreasing c and decreasing number of correlated inputs. Hence, detecting only 10% correlated inputs with $c < 1$ is a fairly difficult task and requires precise synaptic weight changes for the network to be trained effectively⁵¹. Each synapse comprises N PCM devices organized in a non-differential architecture. During the weight update of a synapse, a single potentiation pulse or a single depression pulse is applied to one of the devices the selection counter points to. A depression counter with a maximum value of 2 is incorporated for $N > 1$ to balance the PCM asymmetry. Figure 5(b) depicts the synaptic weights at the end of the experiment for different values of N . To quantify the separation of the weights receiving correlated and uncorrelated inputs, we set a threshold weight that leads to the lowest number of misclassifications. The number of misclassified inputs were 49, 8, and 0 for $N = 1, 3,$ and $7,$ respectively. This demonstrates that the network's ability to detect temporal correlations increases with the number of devices. This holds true even for lower values of the correlation coefficient as shown in Supplementary Note 6. With $N = 1,$ there are strong abrupt fluctuations in the evolution of the conductance values because of the abrupt depression events as shown in Fig. 5(c). With $N = 7,$ a more gradual potentiation and depression behavior is observed. For $N = 7,$ the synapses receiving correlated and uncorrelated inputs can be perfectly separated at the end of the experiments. In contrast, the weights of correlated inputs display a wider weight distribution and there are numerous misclassified weights for $N = 1.$

The multi-memristive synapse architecture is also scalable to larger network sizes. To demonstrate this, we repeated the above correlation experiment with 144,000 input streams, and with seven PCM devices per synapse, resulting in more than one million PCM devices in the network. As shown in Fig. 5(d), well-separated synaptic distributions have been achieved in the network at the end of the experiment. Moreover, a simulation was performed with the nonlinear PCM device model (see Methods). The simulation captures the separation of weights receiving correlated and uncorrelated inputs. In both experiment and simulation, $\sim 0.1\%$ of the inputs were misclassified after training.

Discussion

The proposed synaptic architecture bears similarities to several aspects of neural connectivity in biology, as biological neural connections also comprise multiple sub-units. For instance, in the central nervous system, a presynaptic neuron may form multiple synaptic terminals (so-called boutons) to connect to a single postsynaptic neuron⁵². Moreover, each biological synapse contains a plurality of presynaptic release sites⁵³ and postsynaptic ion channels⁵⁴. Furthermore, our implementation of plasticity through changes in the individual memristors is analogous to individual plasticity of the synaptic connections between a pair of biological neurons⁵⁵, which is also true for the individual ion channels of a synaptic connection^{55,56}. The involvement of progressively larger numbers of memristive devices during potentiation is analogous to the development of new ion channels in a potentiated synapse^{53,54}.

A significant advantage of the proposed multi-memristive synapse is its crossbar compatibility. In memristive crossbar arrays, matrix–vector multiplications associated with the synaptic efficacy can be implemented with a read operation achieving $O(1)$ complexity. Note that memristive devices can be read with low energy (10–100 fJ for our devices), which leads to a substantially lower energy consumption than in conventional von Neumann systems^{57–59}. Furthermore, the synaptic plasticity is realized in place without having to read back the synaptic weights. Even though, the power dissipation associated with programming the memristive devices is at least 10 times higher than that required for the read operation, as only one device of the multi-memristive synapse is programmed at each instance of synaptic update, our scheme does not introduce a significant energy overhead. Memristive crossbars can also be fabricated with very small areal footprint^{27,29,60}. The neuron circuitry of the crossbar array, which typically consumes a larger area than the crossbar array, only increases marginally owing to the additional circuitry needed for arbitration. Finally, because even a single global counter can be used for arbitrating a whole array, the additional area/power overhead is expected to be minimal.

The proposed architecture also offers several advantages in terms of reliability. The other constituent devices of a synapse could compensate for the occasional device failure. In addition, each device in a synapse gets programmed less frequently than if a single device were used, which effectively increases the overall lifetime of a multi-memristive synapse. The potentiation and depression counters reduce the effective number of programming operations of a synapse, further improving endurance-related issues.

Device selection in the multi-memristive synapse is performed based on the arbitration module alone, without knowledge of the conductance values of the individual devices, thus there is a non-zero probability that a potentiation (depression) pulse will not result in an actual potentiation (depression) of the synapse. This would effectively translate into a weight-dependent plasticity whereby the probability to potentiate reduces with increasing synaptic weight and the probability to depress reduces with decreasing synaptic weight (see Supplementary Notes 7, 8). This attribute could affect the overall performance of a neural network. For example, weight-dependent plasticity has been shown to impact the classification accuracy negatively in an ANN⁶¹. In contrast, a study suggests that it can stabilize an SNN intended to detect temporal correlations⁴⁹.

The ANN and SNN simulations in section “Simulation results on handwritten digit classification” with the PCM model perform worse, even in the presence of multi-memristive synapses with $N > 10,$ than the simulations with double-precision floating-point weights. We show that this behavior does not arise from the weight-dependent plasticity of the multi-memristive synapse

scheme, but from the nonlinear PCM conductance response (see Supplementary Note 9). Using a uni-directional linear device model where the conductance change is modeled as a Gaussian random number with mean (granularity) and standard deviation (stochasticity) of 0.5 μS , accuracies exceeding 96.7% are possible in ANN with only 1% performance loss compared with double-precision floating-point weights. Similarly, the network can classify more than 77% of the digits correctly in the SNN using the linear device model, reaching the accuracy of the double-precision floating-point weights.

Note also that the drift in conductance states, which is unique to PCM technology, does not appear to have a significant impact on our studies. As described recently⁶², as long as the drift exponent is small enough (<0.1 ; in our devices it is on average 0.05, see Supplementary Note 4), it is not very detrimental for neural network applications. Our own experimental results on SNNs presented in section “Experimental results on temporal correlation detection” point in this direction, as the network seems to maintain the classification accuracy despite drift. Although conductance drift is not intended to be countered using the multi-memristive concept, there are attempts to overcome it via advanced device-level ideas³⁵, which could be used in conjunction with a multi-memristive synapse.

In the presence of significant nonlinear conductance response and drift, one could envisage an alternate multi-memristive synaptic architecture in which multiple devices are used to store the weights, but with varying significance. For instance, if N -bit synaptic resolution is required, N memory devices could be used, with each device programmed to the maximum (fully potentiated) or minimum (fully depressed) conductance states to represent a number in binary format. In such a binary system, for synaptic efficacy, each device needs to be read independently, which could be accomplished by reading each of the N bits one by one, or alternatively, N amplifiers could be used to read the N bits in parallel. For synaptic plasticity, the desired weight update should be done with prior knowledge of the stored weight. Otherwise, a blind update could have a large detrimental effect, especially if the error is associated with devices representing the most significant bits. However, a direct comparison between these alternate architectures and our proposed scheme requires a detailed system-level investigation, which is beyond the scope of this paper.

In summary, we propose a novel synaptic architecture comprising multiple memristive devices with non-ideal characteristics to efficiently implement learning in neural networks. This architecture is shown to overcome several significant challenges that are characteristic to nanoscale memristive devices proposed for synaptic implementation, such as the asymmetric conductance response, limitations in resolution and dynamic range, as well as device-level variability. The architecture is applicable to a wide range of neural networks and memristive technologies and is crossbar-compatible. The high potential of the concept is demonstrated experimentally in a large-scale SNN performing unsupervised learning. The proposed architecture and its experimental demonstration are a significant step towards the realization of highly efficient, large-scale neural networks based on memristive devices with typical, experimentally observed non-ideal characteristics.

Methods

Experimental platform. The experimental hardware platform is built around a prototype PCM chip with 3 million devices with a four-bank inter-leaved architecture. The mushroom-type PCM devices are based on doped $\text{Ge}_2\text{Sb}_2\text{Te}_5$ (GST) and were integrated into the prototype chip in 90 nm CMOS technology, based on an existing fabrication recipe⁴². The radius of the bottom electrode is ~ 20 nm, and the thickness of the phase change material is ~ 100 nm. A thin oxide n-type field-effect transistor (FET) enables access to each PCM device. The chip also integrates

the circuitry for addressing, an eight-bit on-chip analog-to-digital converter (ADC) for readout, and voltage-mode or current-mode programming. An analog-front-end (AFE) board is connected to the chip and accommodates digital-to-analog converters (DACs) and ADCs, discrete electronics, such as power supplies, voltage and current reference sources. An FPGA board with embedded processor and Ethernet connection implements the overall system control and data management.

PCM characterization. For the experiment of Fig. 2(b), measurements were done on 10,000 devices. All devices were initialized to ~ 0.1 μS with an iterative procedure. In the experiment, 20 potentiation pulses with a duration of 50 ns and varying amplitudes were applied. After each potentiation pulse, the devices were read 50 times in ~ 5 s intervals. The reported device conductance for a potentiation pulse is the average conductance obtained by the 50 consecutive read operations. This method is used to minimize the impact of drift⁶³ and read noise⁴². At the end of the experiment, ~ 300 devices were excluded because they had an initial conductance of less than 0.1 μS or a final conductance after 20 potentiation pulses of more than 30 μS .

In the measurements for Fig. 2(c), 10,000 devices were used. The data was obtained after initializing the device conductances to ~ 5 μS by an iterative procedure. Next, potentiation (depression) pulses of varying amplitude and 50 ns duration were applied. Every potentiation (depression) pulse was followed by 50 read operations done ~ 5 s apart. The device conductance was averaged for the 50 read operations.

In the experiments of Fig. 2(e), 1000 devices were used. All devices were initialized to ~ 0.1 μS with an iterative procedure. This was followed by four potentiation pulses of amplitude $I_{\text{prog}} = 100$ μA and width 50 ns. After the last two potentiation pulses, devices were read 20 times with the reads ~ 1.5 s apart. The device conductances for 20 read operations were averaged. The difference between the averaged conductances for the third and fourth potentiation pulses is defined as the conductance change. This experimental sequence was repeated on the same devices for 1000 times so that 1000 conductance changes were measured for each device.

For the experiments of Fig. 3, measurements were done on 1000, 3000, and 7000 devices for $N = 1, 3$, and 7, respectively. Device conductances were initialized to ~ 5 μS by an iterative procedure. Next, for potentiation, programming pulses of amplitude 100 μA and width 50 ns were applied. For depression, programming pulses of 450 μA amplitude and 50 ns width were applied. After each potentiation (depression) pulse, device conductances were read 50 times and averaged. The delay between each read event was ~ 5 s.

In all measurements, device conductances were obtained by applying a fixed voltage of 0.3 V amplitude and measuring the corresponding current.

Simulation of neural networks. The ANN contains 784 input neurons, 250 hidden layer neurons, and 10 output neurons. In addition, there is one bias neuron at the input layer and one bias neuron at the hidden layer. For training, all 60,000 images from the MNIST training set are used in the order they appear in the database over 10 epochs. Subsequently, all 10,000 images from the MNIST test set are shown for testing. The test set is applied at every 1000th example for the last 20,000 images of the 10th epoch of training, and the results are averaged. The input images from the MNIST set are grayscale pixels with values ranging from 0 to 255 and have a size of 28 times 28. Each of the input layer neurons receives input from one image pixel, and the input is the pixel intensity scaled by 255 in double-precision floating point. The neurons of the hidden and the output layers are sigmoid neurons. Synapses are multi-memristive, and each synapse comprises N devices. The devices in a synapse are arranged using either a non-differential or a differential architecture. In the non-differential architecture, we scale the device conductance of 0 μS to weight $-\frac{1}{N}$ and that of 10 μS to weight $\frac{1}{N}$. The weight is not incremented further if it exceeds $\frac{1}{N}$ to model the PCM saturation behavior. The minimum weight is $-\frac{1}{N}$ because the minimum device conductance is 0 μS . The weight of each device w_n is initialized randomly with a uniform distribution in the interval $[\frac{-1}{2N}, \frac{1}{2N}]$. The total synaptic weight is calculated as $\sum_{n=1}^N w_n$. In the differential architecture, N devices are arranged in two sets, where $\frac{N}{2}$ devices represent G_+ and $\frac{N}{2}$ devices represent G_- . We scale the device conductance of 0 μS to weight 0 and that of 10 μS to weight $\frac{2}{N}$. The weight is not incremented if it exceeds $\frac{2}{N}$ and the minimum weight is 0. The weight of each device $w_{n+}, n-$ for $n = 1, 2, \dots, \frac{N}{2}$ is initialized randomly with a uniform distribution in the interval $[\frac{1}{N}, \frac{2}{N}]$. The total synaptic weight w is $(\sum_n w_{n+}) - (\sum_n w_{n-})$. For double-precision floating-point simulations, the synaptic weights are initialized with a uniform distribution in the interval $[-0.5, 0.5]$. The weight updates Δw are done sequentially to synapses, and the selection counter is incremented by one after each weight update. If $\Delta w > 0$, the synapse will undergo potentiation. In both architectures, each potentiation pulse on average would induce a weight change of size $\epsilon = \frac{0.1}{N}$ if a linear model was used; the number of potentiation pulses to be applied are calculated by rounding $\frac{\Delta w}{\epsilon}$. Then, for each potentiation pulse, an independent Gaussian random number with mean and standard deviation according to the model of Fig. 2(d) is added. This weight change is applied to the device to which the selection counter points. If $\Delta w < 0$, the synapse will undergo a depression. In the differential architecture, a potentiation pulse is applied to a device from the set representing G_+ using the above-mentioned methodology. In the non-differential architecture, a depression pulse is applied to one of the devices pointed at by the selection counter if $\Delta w < -0.5\epsilon$. The

weight of the device drops to 0. For $N > 1$, we used a depression counter of length 5 and a potentiation counter of length 2. No depression or potentiation counter is used for $N = 1$. In the differential architecture, after the weight change has been applied for potentiation and depression, synapses are checked for the refresh operation. If there is a synapse which has $w_+ > 0.9$ or $w_- > 0.9$, a refresh is done on that synapse; w is recorded, and all devices in the synapse are set to 0. The programming will be done to devices of the set w_+ if $w > 0$ or to devices of the set w_- if $w < 0$. The number of potentiation pulses is calculated by rounding $\frac{\Delta w}{\epsilon}$. The pulses are applied to all devices, starting with the first device of the set. One independent Gaussian random number with mean and standard deviation according to the model of Fig. 2(d) is calculated for each of the potentiation pulses. The learning rate is 0.4 for all simulations.

The SNN comprises 784 input neurons and 50 output neurons. These synapses are multi-memristive, and each synapse consists of N memristive devices. The network is trained with all 60,000 images from the MNIST set over three epochs and tested with all 10,000 test images from the set. The test set is applied at every 1000th example for the last 20,000 images, and the results are averaged. The simulation time step is 5 ms. Each input neuron receives input from one pixel of the input image. Each input image is presented for 350 ms, and the information regarding the intensity of each pixel is in the form of spikes. We create the input spikes using a Poisson distribution, where independent Bernoulli trials are conducted to determine whether there is a spike at a time step. A spike rate is calculated as $\frac{\text{pixel intensity}}{255} \times 20$ Hz. A spike is generated if spike rate $\times 5$ ms $> x$, where x is a uniformly distributed random number between 0 and 1. The input spikes create a current with the shape of a delta function at the corresponding synapse. The magnitude of this current is equal to the weight of the synapse. The synaptic weights w are learned with an STDP rule⁴⁸. The synapses are arranged in a non-differential or a differential architecture. In both architectures, we scale the device conductance of 0 μS to weight 0 and that of 10 μS to weight $\frac{1}{N}$. The weight is not incremented further if it exceeds $\frac{1}{N}$. The minimum weight is 0 because the minimum device conductance is 0 μS . In the non-differential architecture, the weight of each device w_n is initialized randomly with a uniform distribution in the interval $[\frac{2}{5N}, \frac{3}{5N}]$. The total synaptic weight is calculated as $\sum_{n=1}^N w_n$. In the differential architecture, N devices are arranged in two sets. The weight of each device w_{n+} for $n = 1, 2, \dots, \frac{N}{2}$ is initialized randomly with a uniform distribution in the interval $[\frac{3}{5N}, \frac{4}{5N}]$. The total synaptic weight is $(\sum_n w_{n+}) - (\sum_n w_{n-}) + 0.5$. For double-precision floating-point simulations, the synaptic weights are initialized with a uniform distribution in the interval [0.25, 0.75]. At each simulation time step, the synaptic currents are summed at the output neurons and accumulated using a state variable X . The output neurons are of the leaky integrate-and-fire type and have a leak constant of $\tau = 200$ ms. Each output neuron has a spiking threshold. This spiking threshold is set initially to 0.125 (note that the sum of the currents is normalized by the number of input neurons) and is altered by homeostasis during training. An output neuron spikes when X exceeds the neuron threshold. Only one output neuron is allowed to spike at a single time step, and if the state variables of several neurons exceed their threshold, then the neuron whose state variable exceeds its threshold the most is the winner. The state variables of all other neurons are set to 0 if there is a spiking output neuron. If there is a postsynaptic neuronal spike, the synapses that received a presynaptic spike in the past 30 ms are potentiated. If there is a presynaptic spike the synapses that had a postsynaptic neuronal spike in the past 1.05 s are depressed. The weight change amount is constant for potentiation ($\Delta w = 0.01$) and depression ($\Delta w = -0.006$), following a rectangular STDP rule. The weight updates are done using the scheme described above with $\epsilon = \frac{0.05}{N}$. For the non-differential architecture, a depression pulse is applied when $\Delta w < 0$. The depression counter length is set to the floor of $\frac{1}{N \times 0.006}$ for $N > 1$. In the non-differential and the differential architecture, a potentiation counter of length 3 and 2 is used, respectively. After the 1000th input image, upon presentation of every two images, the spiking thresholds of the output neurons are adjusted through homeostasis. The threshold increase for every output neuron is calculated as $0.0005 \times (A - T)$, where A is the activity of the neuron and T is the target firing rate. A is calculated as $\frac{S}{350 \text{ ms} \times 100}$, where S is the sum of the neuron's firing event in the past 100 examples. We define the T as $\frac{5}{350 \text{ ms} \times 50}$, where 50 is the number of output neurons in the network. After training, the synaptic weights and the neuron thresholds are kept constant. To quantify how well the training is, we show all 60,000 images to the network, and the neuron that spikes the most often during the presentation of an image for 350 ms is recorded. The neuron is mapped to the class, i.e., to one of the 10 digits, for which it spiked the most. This mapping is then used to detect the classification accuracy when the test set is presented.

Correlation detection experiment. The network for correlation detection comprises 1000 plastic synapses connected to an output neuron. Each synapse is multi-memristive and consists of N devices. The synaptic weights $w \in [0, 1]$ are learned with an STDP rule⁶⁴. Because of the hardware latency, we will use normalized units to describe time in the experiment. The experiment time steps are of size $T_s = 0.1$. Each synapse receives a stream of spikes, and the spikes have the shape of the delta function. 100 of the input spike stream are correlated. The correlated and the uncorrelated spike streams have equal rates of $r_{\text{cor}} = r_{\text{uncor}} = 1$. The correlated inputs share the outcome of a Bernoulli trial. This Bernoulli trial is described as $B = x > 1 - r_{\text{cor}} \times T_s$, where x is a uniformly distributed random number. By using this event, the input spikes for the correlated streams are generated as $B \times (r_{\text{cor}} \times T_s + \sqrt{c} \times (1 - r_{\text{cor}} \times T_s) > x_1) + \sim B \times (r_{\text{cor}} \times T_s \times (1 - \sqrt{c}) > x_2)$,

where x_1 and x_2 are uniformly distributed random numbers, c is the correlation coefficient of value 0.75, and \sim denotes the negation operation^{49,51}. The uncorrelated processes are generated as $x_3 > 1 - r_{\text{uncor}} \times T_s$, where x_3 is a uniformly distributed random variable. Note that the probability of generating a spike is low because $r_{\text{cor,uncor}} \times T_s \ll 1$. These input spikes generate a current of the size of the synaptic weights. At every time step, the currents are summed and accumulated at the neuronal membrane variable X . The neuronal firing events in any given time step are determined only by the spiking events that occur in that time step. If X exceeds a threshold of 52, the output neuron fires. The weight update calculation follows an exponential STDP rule where the amount of potentiation is calculated as $A_+ e^{-|\Delta t|/\tau_+}$ and that of depression is calculated as $-A_- e^{-|\Delta t|/\tau_-}$. A_+ , A_- are the learning rates, τ_+ , τ_- are time constants, and Δt is the time difference between the input spikes and the neuronal spikes. We set $2 \times A_+ = A_- = 0.004$ and $\tau_+ = \tau_- = 3 \times T_s$. The higher-order pairs of spikes are also considered in our algorithm.

The weight storage and weight update operations are done on PCM devices. We access each PCM device sequentially for reading and programming. For device initialization, an iterative procedure is used to program the device conductances to 0.1 μS and this is followed by one potentiation pulse of amplitude $I_{\text{prog}} = 120$ μA and 50 ns width. Although the weight update is calculated using an exponential STDP rule, it is applied following a rectangular STDP rule. For potentiation, a single potentiation pulse of amplitude $I_{\text{prog}} = 100$ μA and 100 ns width is applied when $\Delta w \geq 0.001$. For depression, a single depression pulse of amplitude $I_{\text{prog}} = 440$ μA and 950 ns width is applied when $\Delta w \leq -0.001$. The potentiation and depression pulses are sent to one device from the multi-memristive synapse the selection counter points to. When applying depression pulses, a depression counter of length 2 is used for $N > 1$. After each programming operation, the device conductances are read by applying a fixed voltage of amplitude 0.2 V and measuring the corresponding current. The conductance value G of a device is converted to its synaptic weight as $w_n = \frac{G}{N \times 9.5 \mu\text{S}}$. The weights of the devices in a multi-memristive synapse are summed to calculate the total synaptic weight $\sum_{n=1}^N w_n$.

For the large-scale experiment, 144,000 synapses are trained, of which 14,400 receive correlated inputs. Each multi-memristive synapse comprises $N = 7$ devices, and a total of 1,008,000 PCM devices are used for this experiment. The same network parameters as in the small-scale experiment are used, except for the neuron threshold. The neuron threshold is scaled with the number of synapses and is set to 7488. The learning algorithm and conductance-to-weight conversion are identical to those in the small-scale experiment.

The nonlinear PCM model used for the accompanying simulation study is based on the conductance evolution of PCM devices with $I_{\text{prog}} = 100$ μA pulse amplitude and a pulse width of 50 ns. Two potentiation pulses are applied consecutively to capture the conductance change behavior of one potentiation pulse with pulse width 100 ns of the experiments. A depression pulse is assumed to set the device conductance to 0 μS , irrespective of the conductance value prior to the application of the pulse.

Data availability. The data that support the findings of this study are available from the corresponding authors upon request.

Received: 8 November 2017 Accepted: 4 June 2018

Published online: 28 June 2018

References

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
2. Schemmel, J. et al. A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Proc. IEEE International Symposium on Circuits and Systems (ISCAS)*, 1947–1950 (IEEE, Paris, France, 2010).
3. Painkras, E. et al. SpiNNaker: a 1-W 18-core system-on-chip for massively-parallel neural network simulation. *IEEE J. Solid-State Circuits* **48**, 1943–1953 (2013).
4. Merolla, P. A. et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* **345**, 668–673 (2014).
5. Qiao, N. et al. A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses. *Front. Neurosci.* **9**, 141 (2015).
6. Beck, A., Bednorz, J., Gerber, C., Rossel, C. & Widmer, D. Reproducible switching effect in thin oxide films for memory applications. *Appl. Phys. Lett.* **77**, 139–141 (2000).
7. Strukov, D. B., Snider, G. S., Stewart, D. R. & Williams, R. S. The missing memristor found. *Nature* **453**, 80–83 (2008).
8. Chua, L. Resistance switching memories are memristors. *Appl. Phys. A* **102**, 765–783 (2011).
9. Wong, H. S. P. & Salahuddin, S. Memory leads the way to better computing. *Nat. Nanotech.* **10**, 191–194 (2015).

10. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1: Foundations, (eds Rumelhart, D. E. & McClelland, J. L.) 318–362 (MIT Press, Cambridge, MA, 1986).
11. Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).
12. Anwani, N. & Rajendran, B. NormAD - normalized approximate descent based supervised learning rule for spiking neurons. In *Proc. International Joint Conference on Neural Networks (IJCNN)*, 1–8 (IEEE, Killarney, Ireland, 2015).
13. Saighi, S. et al. Plasticity in memristive devices for spiking neural networks. *Front. Neurosci.* **9**, 51 (2015).
14. Burr, G. W. et al. Integration of nanoscale memristor synapses in neuromorphic computing architectures. *Adv. Phys. X* **2**, 89–124 (2016).
15. Rajendran, B. & Alibart, F. Neuromorphic computing based on emerging memory technologies. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **6**, 198–211 (2016).
16. Ohno, T. et al. Short-term plasticity and long-term potentiation mimicked in single inorganic synapses. *Nat. Mater.* **10**, 591–595 (2011).
17. Yu, S., Wu, Y., Jeyasingh, R., Kuzum, D. & Wong, H. S. P. An electronic synapse device based on metal oxide resistive switching memory for neuromorphic computation. *IEEE Trans. Electron Dev.* **58**, 2729–2737 (2011).
18. Ambrogio, S. et al. Neuromorphic learning and recognition with one-transistor-one-resistor synapses and bistable metal oxide RRAM. *IEEE Trans. Electron Dev.* **63**, 1508–1515 (2016).
19. Covi, E. et al. Analog memristive synapse in spiking networks implementing unsupervised learning. *Front. Neurosci.* **10**, 482 (2016).
20. van de Burgt, Y. et al. A non-volatile organic electrochemical device as a low-voltage artificial synapse for neuromorphic computing. *Nat. Mater.* **16**, 414–418 (2017).
21. Wang, Z. et al. Memristors with diffusive dynamics as synaptic emulators for neuromorphic computing. *Nat. Mater.* **16**, 101–108 (2017).
22. Boyn, S. et al. Learning through ferroelectric domain dynamics in solid-state synapses. *Nat. Commun.* **8**, 14736 (2017).
23. Wu, W., Zhu, X., Kang, S., Yuen, K. & Gilmore, R. Probabilistically programmed STT-MRAM. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **2**, 42–51 (2012).
24. Vincent, A. F. et al. Spin-transfer torque magnetic memory as a stochastic memristive synapse for neuromorphic systems. *IEEE Trans. Biomed. Circ. Syst.* **9**, 166–174 (2015).
25. Kuzum, D., Jeyasingh, R. G. D., Lee, B. & Wong, H. S. P. Nanoelectronic programmable synapses based on phase change materials for brain-inspired computing. *Nano Lett.* **12**, 2179–2186 (2012).
26. Ambrogio, S. et al. Unsupervised learning by spike timing dependent plasticity in phase change memory (PCM) synapses. *Front. Neurosci.* **10**, 56 (2016).
27. Alibart, F., Zamanidoost, E. & Strukov, D. B. Pattern classification by memristive crossbar circuits using ex situ and in situ training. *Nat. Commun.* **4**, 2072 (2013).
28. Indiveri, G., Linares-Barranco, B., Legenstein, R., Deligeorgis, G. & Prodromakis, T. Integration of nanoscale memristor synapses in neuromorphic computing architectures. *Nanotechnology* **24**, 384010 (2013).
29. Prezioso, M. et al. Training and operation of an integrated neuromorphic network based on metal-oxide memristors. *Nature* **521**, 61–64 (2015).
30. Kim, S. et al. NVM neuromorphic core with 64k-cell (256-by-256) phase change memory synaptic array with on-chip neuron circuits for continuous in-situ learning. In *Proc. IEEE International Electron Devices Meeting (IEDM)*, 17-1 (IEEE, Washington, DC, USA, 2015).
31. Mostafa, H. et al. Implementation of a spike-based perception learning rule using TiO_{2-x} memristors. *Front. Neurosci.* **9**, 357 (2015).
32. Tuma, T., Le Gallo, M., Sebastian, A. & Eleftheriou, E. Detecting correlations using phase-change neurons and synapses. *IEEE Electron Dev. Lett.* **37**, 1238–1241 (2016).
33. Wozniak, S., Tuma, T., Pantazi, A. & Eleftheriou, E. Learning spatio-temporal patterns in the presence of input noise using phase-change memristors. In *Proc. IEEE International Symposium on Circuits and Systems (ISCAS)*, 365–368 (IEEE, Montreal, QC, Canada, 2016).
34. Burr, G. W. et al. Experimental demonstration and tolerancing of a large-scale neural network (165 000 synapses) using phase-change memory as the synaptic weight element. *IEEE Trans. Electron Dev.* **62**, 3498–3507 (2015).
35. Koelmans, W. W. et al. Projected phase-change memory devices. *Nat. Commun.* **6**, 8181 (2015).
36. Fuller, E. J. et al. Li-ion synaptic transistor for low power analog computing. *Adv. Mater.* **29**, 1604310 (2017).
37. Suri, M. et al. Phase change memory as synapse for ultra-dense neuromorphic systems: Application to complex visual pattern extraction. In *Proc. IEEE International Electron Devices Meeting (IEDM)*, 4.4.1–4.4.4 (IEEE, Washington, DC, USA, 2011).
38. Bill, J. & Legenstein, R. A compound memristive synapse model for statistical learning through STDP in spiking neural networks. *Front. Neurosci.* **8**, 412 (2014).
39. Garbin, D. et al. HfO₂-based OxRAM devices as synapses for convolutional neural networks. *IEEE Trans. Electron Dev.* **62**, 2494–2501 (2015).
40. Burr, G. W. et al. Recent progress in phase-change memory technology. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **6**, 146–162 (2016).
41. Sebastian, A., Le Gallo, M. & Krebs, D. Crystal growth within a phase change memory cell. *Nat. Commun.* **5**, 4314 (2014).
42. Close, G. F. et al. Device, circuit and system-level analysis of noise in multi-bit phase-change memory. In *Proc. IEEE Int. Electron Devices Meeting (IEDM)*, 29.5.1–29.5.4 (IEEE, San Francisco, CA, USA, 2010).
43. Boybat, I. et al. Stochastic weight updates in phase-change memory-based synapses and their influence on artificial neural networks. In *Proc. Ph.D. Research in Microelectronics and Electronics (PRIME)*, 13–16 (IEEE, Giardini Naxos, Italy, 2017).
44. Gokmen, T. & Vlasov, Y. Acceleration of deep neural network training with resistive cross-point devices: design considerations. *Front. Neurosci.* **10**, 333 (2016).
45. Nandakumar, S. R. et al. Supervised learning in spiking neural networks with MLC PCM synapses. In *Proc. Device Research Conference (DRC)*, 1–2 (IEEE, South Bend, IN, USA, 2017).
46. Le Gallo, M., Tuma, T., Zipoli, F., Sebastian, A. & Eleftheriou, E. Inherent stochasticity in phase-change memory devices. In *Proc. European Solid-State Device Research Conference (ESSDERC)*, 373–376 (IEEE, Lausanne, Switzerland, 2016).
47. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
48. Querlioz, D., Bichler, O. & Gamrat, C. Simulation of a memristor-based spiking neural network immune to device variations. In *Proc. International Joint Conference on Neural Networks (IJCNN)*, 1775–1781 (IEEE, San Jose, CA, USA, 2011).
49. Gütiğ, R., Aharonov, R., Rotter, S. & Sompolinsky, H. Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J. Neurosci.* **23**, 3697–3714 (2003).
50. Tuma, T., Pantazi, A., Le Gallo, M., Sebastian, A. & Eleftheriou, E. Stochastic phase-change neurons. *Nat. Nanotech.* **11**, 693–699 (2016).
51. Sebastian, A. et al. Temporal correlation detection using computational phase-change memory. *Nat. Commun.* **8**, 1115 (2017).
52. Edwards, F. LTP is a long term problem. *Nature* **350**, 271–272 (1991).
53. Bolshakov, V. Y., Golan, H., Kandel, E. R. & Siegelbaum, S. A. Recruitment of new sites of synaptic transmission during the cAMP-dependent late phase of LTP at CA3-CA1 synapses in the hippocampus. *Neuron* **19**, 635–651 (1997).
54. Malenka, R. C. & Nicoll, R. A. Long-term potentiation—a decade of progress? *Science* **285**, 1870–1874 (1999).
55. Malenka, R. C. & Bear, M. F. LTP and LTD: an embarrassment of riches. *Neuron* **44**, 5–21 (2004).
56. Benke, T. A., Luthi, A., Isaac, J. T. R. & Collingridge, G. L. Modulation of AMPA receptor unitary conductance by synaptic activity. *Nature* **393**, 793–797 (1998).
57. Le Gallo, M., Sebastian, A., Cherubini, G., Giefers, H. & Eleftheriou, E. Compressed sensing recovery using computational memory. In *Proc. IEEE International Electron Devices Meeting (IEDM)*, 28.3.1–28.3.4 (IEEE, San Francisco, CA, USA, 2017).
58. Li, C. et al. Analogue signal and image processing with large memristor crossbars. *Nat. Electron.* **1**, 52–59 (2018).
59. Le Gallo, M. et al. Mixed-precision in-memory computing. *Nat. Electron.* **1**, 246–253 (2018).
60. Eryilmaz, S. B. et al. Brain-like associative learning using a nanoscale non-volatile phase change synaptic device array. *Front. Neurosci.* **8**, 205 (2014).
61. Sidler, S. et al. Large-scale neural networks implemented with non-volatile memory as the synaptic weight element: Impact of conductance response. In *Proc. European Solid-State Device Research Conference (ESSDERC)*, 440–443 (IEEE, Lausanne, Switzerland, 2016).
62. Fumarola, A. et al. Accelerating machine learning with non-volatile memory: exploring device and circuit tradeoffs. In *Proc. IEEE International Conference on Rebooting Computing (ICRC)*, 1–8 (IEEE, San Diego, CA, USA, 2016).
63. Sebastian, A., Krebs, D., Le Gallo, M., Pozidis, H. & Eleftheriou, E. A collective relaxation model for resistance drift in phase change memory cells. In *Proc. IEEE International Reliability Physics Symposium (IRPS)*, MY-5 (IEEE, Monterey, CA, USA, 2015).
64. Song, S., Miller, K. D. & F., A. L. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* **3**, 919–926 (2000).

Acknowledgements

We would like to thank N. Papandreou, U. Egger, S. Wozniak, S. Sidler, A. Pantazi, M. BrightSky, and G. Burr for technical input. I. B. and T. M. would like to acknowledge financial support from the Swiss National Science Foundation. A. S. would like to acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant agreement no. 682675).

Author contributions

I.B., M.L.G., T.T., and A.S. designed the concept. I.B., N.S.R., and T.M. performed the simulations. I.B., N.S.R., M.L.G., and A.S. performed the experiments. T.M. and T.P. provided critical insights. I.B., M.L.G., and A.S. co-wrote the manuscript with input from the other authors. A.S., Y.L., B.R. and E.E. supervised the work.

Additional information

Supplementary Information accompanies this paper at <https://doi.org/10.1038/s41467-018-04933-y>.

Competing interests: The authors declare no competing interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018