


Implementation of a Minimal Recurrent Spiking Neural Network in a Solid-State Device

P. Stoliar^{1,†}, O. Schneegans^{2,†} and M. J. Rozenberg^{3,*}

¹*National Institute of Advanced Industrial Science and Technology (AIST), 305-8565 Tsukuba, Japan*

²*CentraleSupélec, CNRS, Université Paris-Saclay, Sorbonne Université, Laboratoire de Génie Electrique et Electronique de Paris, 91192 Gif-sur-Yvette, France*

³*Université Paris-Saclay, CNRS Laboratoire de Physique des Solides, 91405 Orsay, France*

 (Received 15 December 2020; revised 30 March 2021; accepted 30 July 2021; published 16 September 2021)

We study the minimal recurrent spiking neural network of a single neuron with an autaptic synapse. We implement the neural system in the solid state with a recently introduced ultracompact neuron (UCN) model, which is based on the memristive properties of a thyristor. The UCN is supplemented by a self-synaptic, *autaptic*, connection, where we control the feedback. Both excitatory and inhibitory cases are considered. We explore the systematic behavior as a function of autaptic intensity and feedback time delay. We realize a tunable dynamic memory, showing graded persistent activity, where short excitatory and inhibitory pulses allow the firing rate to be controlled. We finally reproduce recent experimentally observed behavior of a biological autapse measured *in vivo*, finding excellent qualitative agreement. Our work opens the way into the field of solid-state neuroscience, with the UCN as an accessible platform to implement and experimentally study the dynamic behavior of spiking neural networks.

DOI: [10.1103/PhysRevApplied.16.034030](https://doi.org/10.1103/PhysRevApplied.16.034030)

I. INTRODUCTION

One of the most challenging frontiers in artificial intelligence is an understanding of recurrent neural networks. This type of network allows for unprecedented success in the processing of time-dependent information, such as speech. In recurrent networks, information flows across the networks from input to output, but some of that flow reverberates back. This feedback makes the theoretical analysis of their dynamic behavior a particularly difficult challenge [1–4]. It becomes an even greater one for the case of *spiking neural networks*, which are relevant to the ongoing quest to understand the origin and functions of the dynamic behavior observed in brains. Examples of relevant issues include brain waves, epilepsy, dynamic working memory, avalanches, criticality, and chaotic behavior [5–8].

Since the possibility of systematic *in vivo* or *in vitro* exploration of neural networks remains a difficult challenge [9], and numerical simulations may rely on the specific details of the model and approximations [10,11], other methodologies are needed. In this context, our goal here is to illustrate that a different approach is readily

available to systematically study *physical* realizations of artificial spiking neural networks. This methodology, which we term solid-state neuroscience, consists of realizing neuronal model circuits, the electric spiking behavior of which can be systematically modified and studied in detail. This approach relies on a key recent breakthrough, namely, the implementation of an extremely simple solid-state neuron model that is based on a conventional electronic component, the thyristor. The spiking neuron circuit, or *ultracompact* neuron (UCN) can be likened to Lego blocks for the construction of networks [12]. It can also be easily extended to realize a multitude of biologically relevant spiking behaviors [13]. Here, we shall adopt the recently introduced circuit of a leaky-integrate-and-fire UCN and supplement it with a self-synaptic connection that produces a self-feedback effect [Fig. 1(a)]. This minimal neural system is called the *autapse* and may be considered as the smallest instance of a *recursive spiking neural network*. Here, we implement and study the behavior of the circuit of a UCN autapse (UCNA). Our main aims are (i) we demonstrate that, despite the circuit's simplicity, its dynamic behavior is rich, likely including chaotic states; (ii) we implement the working memory model proposed by Seung *et al.* [15], finding graded persistent activity, with no need for fine-tuning; and (iii) we use recent experimental data obtained in neocortical pyramidal cells, in both rodents and humans [16], to demonstrate that our physical

*marcelo.rozenberg@universite-paris-saclay.fr

†P. Stoliar, O. Schneegans, and M. J. Rozenberg contributed equally to this work.




FIG. 1. (a) Schematic diagram of a neuron with an autapse (yellow) and input excitation from an upstream neuron synapse, I_{in0} (brown). Autapse input, I_A , provides feedback from the output axon, I_O , with a time delay Δt . Synaptic connection may be either excitatory ($\beta > 0$) or inhibitory ($\beta < 0$). Total input to the cell body (soma) is $I_{in} = I_{in0} + I_A$. (b) Electronic circuit implementation of the UCN unit (gray) with autaptic connection (yellow) implemented as a current mirror and a time delay. Action potential voltage, V_{AP} . List of electronic components is provided in the Supplemental Material [14]. (c) Implementation of the delay line as a chain of UCN mimicking a long axon.

model can qualitatively reproduce actual biological neural behavior.

Besides its conceptual relevance in theoretical neuroscience, as the simplest instance of a recurrent spiking neural network, the autapse is also relevant in biological neuroscience. Neurons with self-synaptic connections were first reported more than 100 years ago and were termed autapses in the 1970s by van der Loos and Glaser [17]. These authors were surprised to find them in pyramidal cells of the neocortex, which is the brain region in mammals that is involved in higher-order functions, such as sensory perception, cognition, generation of motor commands, spatial reasoning, and language. They proposed that autapses might be responsible for a significant gating mechanism, by which the output of a neuron might regulate its own inputs [17]. However, autapses remained mostly a curiosity and were dismissed as aberrations or wiring errors [18]. They were eventually observed in a multitude of cerebral regions and interest in these structures was revived in the 2000s [15,18,19]. Moreover, they were associated with specific functions, for instance, they were found to be responsible for the enhancement of spike-timing precision of neocortical inhibitory interneurons [20]. Despite the evident challenge of first finding and then measuring a neuron with self-interactions, recent advances in neurobiology are making their imaging and behavioral study more accessible, as in the recent work of Yi *et al.* [21]. These authors reported the observation of several

functions, such as enhanced burst firing, neuronal responsiveness, and coincidence detection in glutamatergic (i.e., excitatory) pyramidal cell autapses. Interestingly, they also reported finding 28% of pyramidal cells in layer 5 of the neocortex with autapses in epileptic adult human brains [21].

The study of autapses has also received attention from other scientific communities, such as theoretical [22] and computational neuroscientists [23]. However, we are not able to find any reference to autaptic studies by neuromorphic engineers, despite significant progress in the development of neuroinspired hardware [24]. This is perhaps surprising, since an understanding of autapses should be conceptually relevant for the hot topic of recurrent spiking neural networks for artificial-intelligence applications [25,26].

The present work is organized as follows. In Sec. II, we describe the basic UCN circuit and its extension to the case of the autapse. In Sec. III, we present the results of the present study. In Sec. III A, we discuss the general behavior of the autapse for both excitatory and inhibitory feedback. We find evidence of chaotic behavior and very-slow-convergence rates, which are features common to nonlinear dynamic systems. In Sec. III B, we implement and control a dynamic state that realizes graded persistent activity, which is a relevant neuroscience concept for working memory. In Sec. III C, we establish an exciting connection to biological neuroscience, as we demonstrate

that our physical solid-state model can reproduce the qualitative behavior measured in recent experiments on biological autapses. Finally, Sec. IV contains our conclusions.

II. THE ULTRACOMPACT NEURON AUTAPSE

The UCN circuit can display electric spiking behavior analogous to a biological neuron and constitutes the basic building block of our methodology. In Ref. [12], we demonstrated a key feature, namely, that the spikes of a UCN could drive the spiking of a second downstream UCN. Here, we go beyond this scheme and show that a single UCN can self-excite (or self-inhibit). This is relevant, as it creates a reverberation of the output towards the input, dramatically affecting the dynamic behavior.

We next briefly describe the basic features of the UCN circuit; further details can be found in Refs. [12,13]. In Fig. 1(a), we show a schematic view of a neuron with an autapse, while in Fig. 1(b) we show the corresponding electronic circuit. The basic neuron is indicated in gray and the feedback autaptic connection in yellow. The UCN is a minimal physical implementation of a leaky-integrate-and-fire (LIF) neuron model that generates action potentials. We call this neuron circuit *ultracompact* [27] since the “soma,” which is the key part that generates the action potentials requires only three basic electronic components [see gray region in Fig. 1(b)]: a “membrane” capacitor (C_1) for the *integrate* function; a resistor ($R_1 + R_2$) for the *leaky* function; and, in contrast to conventional CMOS implementations [27], the *fire* function is realized by a thyristor or silicon control rectifier (SCR).

The SCR is a p - n - p - n device and one of the oldest solid-state electronic components. The fire function is most easily understood by noting that the thyristor can be considered as a diode with a voltage threshold V_{th} . This is a voltage that needs to be overcome to switch-*on* its conductive (direct) state. Thus, the SCR is normally *off* with a large resistance, much bigger than the leak resistor pair ($R_1 + R_2$). When a constant current, I_{in0} , inputs the neuron, the capacitor gets leaky charged with a time constant of about $\tau = (R_1 + R_2)C_1$. The firing event takes place when the SCR suddenly commutes to its *on* low-resistance state. This occurs when the capacitor accumulates enough charge, so that V_C drives the voltage at the gate of the SCR above its threshold value, V_{th} , which is a parameter of the SCR. Thus, the firing event can be tuned with the resistive voltage divider (R_1, R_2), such that the condition $V_{th} = V_C R_2 / (R_1 + R_2)$ is met. This sets the critical value, V_C , for the generation of a spike. At that point, the resistance of the SCR collapses, as its p - n diodelike junctions become forward polarized. The SCR resistance becomes much smaller than the leak pair ($R_1 + R_2$), and the membrane capacitor rapidly discharges through it. This current produces a spike of voltage, or action potential, on the small resistor, R_3 . The timescale for the duration of the spike is

the discharge time, which is about $(R_3 + R_{SCRLO})C_1$, where R_{SCRLO} is the low-resistance value of the SCR. The action potential terminates when the discharge current decreases beneath a low “holding” value, I_{hold} , which is another parameter of the SCR. At that point, the resistance of the SCR returns to high, and the cycle restarts. Therefore, upon a constant input, I_{in0} , that is strong enough to charge the capacitor beyond the threshold, the basic UCN [gray circuit in Fig. 1(b)] generates a succession of action-potential spikes [12]. We should emphasize here that it is the hysteresis of the SCR resistance, or its *memristive* property, that is the key feature behind the striking simplicity of the UCN circuit.

The output of the UCN circuit implements an “axon,” which strengthens the action-potential spike on R_3 . This is done so that the UCN spikes may drive other neurons that can be directly connected downstream, allowing networks to be built like construction with Legolike blocks [12,28]. The axon can be simply implemented by means of a transistor Q_1 (and a +5 V source). This completes the description of the basic spiking UCN circuit [gray circuit in Fig. 1(b)].

To implement an autaptic connection, i.e., a self-synapse, we need to feed the output back into the UCN input with a certain time delay, Δt . This can be achieved by a standard “current-mirror” configuration (Q_2, Q_3) and a delay line [yellow block in Fig. 1(b)]. Thus, current mirror plus delay line has two effects: (i) it multiplies the $I_O(t)$ signal by a factor β and (ii) it delays the signal by Δt . Therefore, the autaptic input current is $I_A(t) = \beta I_O(t - \Delta t)$, where β can be both positive or negative, for an excitatory or inhibitory autapse, respectively. Thus, the total input current is $I_{in}(t) = I_{in0} + \beta I_O(t - \Delta t)$. The delay line can be implemented with a “long-axon neuron” [Fig. 1(c)] that consists of a chain of identical UCNs, as described in Ref. [28]. The long-axon configuration is a transmission line, much in the spirit of the Hodgkin-Huxley model for the propagation of the action potential along the squid giant axon [29]. It permits a delay to be implemented in the form of $n\delta t$, where δt is the delay of the signal introduced by each block of the long axon, and n is the position of the block in the chain (δt is essentially the discharge time discussed above, also see Ref. [28]). Thus, to implement a delay of $\Delta t = n\delta t$, it is sufficient to take the spike signal from the n th Ranvier node of the axon chain [for instance, $\Delta t = 4\delta t$ in Fig. 1(c)] [28]. Here, however, we do not implement the long axon but, for convenience, we implement the delay with a standard electronic-data-acquisition instrument, as we describe in the Supplemental Material [14,30]. This is merely a practical way to explore the systematic behavior of the autapse as a function of the parameter Δt .

Importantly, and for the sake of promoting the use of UCN blocks as a physical platform to build arbitrary spiking neural networks for solid-state neuroscience, we

would like to emphasize that all components are readily available off the shelf. This is in sharp contrast to other compact neurons based on memristive properties, such as Mott or phase-change materials [21,31–35], which may require complex material deposition and low-temperature techniques. These features make our methodology readily accessible for adoption by a potentially large community of physicists, electronic engineers, neurobiologists, and computer scientists interested in neuroscience and neurocomputing. We explicitly illustrate this point in the following section.

III. RESULTS

We now turn to the presentation of the results of our investigation into the basic spiking behavior of the minimal recurrent neural network, namely, the UCNA, depicted in Fig. 1.

A. Systematic behavior as a function of feedback parameters β and Δt

We begin to explore the basic systematic behavior of the UCNA as a function of its two parameters, gain β and time delay Δt . A useful observable, which is often adopted in neurobiological studies, is the interspike interval (ISI), which is simply defined as the time between two subsequent output spikes. It is sometimes given in terms of the instantaneous frequency, which is defined as $1/\text{ISI}$. In the absence of feedback (i.e., a spiking neuron with no autapse), $\beta = 0$ and the ISI is denoted as ISI_0 . In a simple LIF neuron (in gray in Fig. 1), a constant input current, $I_{\text{in}0}$, controls the spiking rate, and thus, $\text{ISI}_0(I_{\text{in}0})$.

When β is nonzero, we have feedback from the output to the input and the spiking rate has to adapt until the neuron finds a self-consistent spiking rate. In practice, the spiking rate eventually converges; however, the system may sometimes be close to a chaotic situation, as we shall see later in the present section.

Intuitively, we may expect that, with positive (i.e., excitatory) feedback, the rate would increase, since the input current is enhanced. In contrast, for negative (i.e., inhibitory) feedback, the input current is decreased, and then one may expect the initial rate to decrease. This intuitive behavior is, in fact, born out from experiments, but several nontrivial features also emerge.

In Fig. 2, we show the self-consistent ISI for the cases of an excitatory and an inhibitory autaptic synapse. In both panels, we indicate with a blue line the reference value of ISI_0 , which is set by the constant input, $I_{\text{in}0}$. We observe that, as expected, positive β (excitatory) produces a reduction of the ISI, while negative β (inhibitory) gives an increase in this quantity. However, the behavior of the ISI data as a function of Δt at fixed β and $I_{\text{in}0}$ is more complex, as we describe next.

A common feature of both cases is an approximate periodicity of the ISI as a function of the delay Δt with a period of about 20 ms [36]. This period roughly corresponds to the ISI_0 value (blue line) that is set by the input current, $I_{\text{in}0}$. While the system is very nonlinear, which renders intuition very limited, we can still present a qualitative argument to understand the periodicity. The feedback, $I_A(t)$, is obtained from the output, which, at least initially, is also a periodic function with period ISI_0 . The first feedback spike will arrive at the soma after a delay of Δt . During Δt , the soma itself is periodically undergoing charging and discharging of the membrane capacitor due to the action of the injected constant current, $I_{\text{in}0}$. Thus, if the first I_A spike arrives at a certain time t_0 after, say, the n th discharge, its effect will *not depend on the number, n , of previous cycles* of the membrane capacitor. The effect of the arriving spike will mainly depend on the value of t_0 , and it will then feedback continuously without further delay. Hence, the periodicity with Δt , which one may expect in both excitatory and inhibitory cases. It is interesting to mention that the approximate periodicity observed in our experimental results are consistent with the numerical

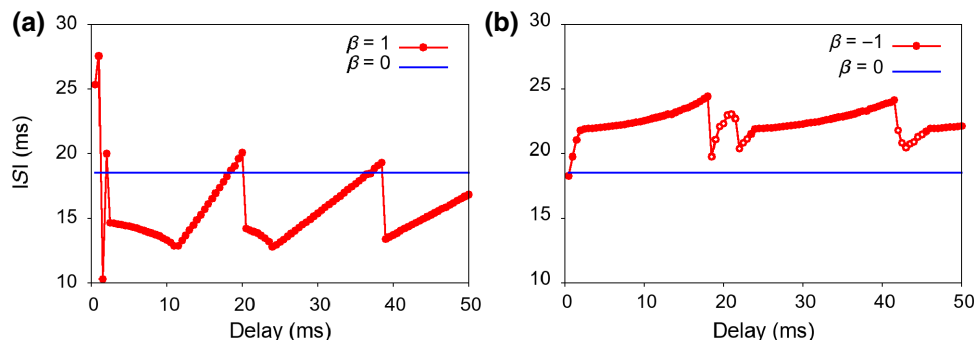


FIG. 2. (a) Self-consistent ISI as a function of delay Δt for the case of excitatory feedback (red dots). Blue line indicates ISI_0 , i.e., ISI value at zero feedback. (b) Idem for inhibitory feedback. Notice that the open dots indicate the *average* ISI value, as they correspond to Δt regions of complex dynamic behavior without a single value of self-consistent ISI (see Fig. 3). Supplemental Material [14] contains waveforms of I_{in} , V_C , and I_{out} for selected points of these two plots.

simulations of Wang *et al.* in a Hodgkin-Huxley autapse model [36].

Both the excitatory and inhibitory cases share the common feature of approximate periodicity in Δt ; however, they also have marked qualitative differences, as we discuss next.

In the case of excitatory feedback, the ISI is globally shifted to shorter spiking intervals. In a first approximation, this is because the feedback current, I_A , adds to the input, I_{in0} , and therefore, the capacitor charges faster and the thyristor reaches the threshold voltage faster. Despite the strong nonlinearities that are introduced by delayed feedback, we observe that the system always stabilizes at a single self-consistent ISI value. In a first approximation, we observe that the behavior of the ISI increases in a rather linear fashion as a function of Δt , after each of the jumps, which occur at approximate intervals of delay of about ISI_0 . This linear increase can be qualitatively understood by the following argument: as discussed before, the membrane capacitor discharges periodically as it emits spikes. If a feedback spike, I_A , arrives immediately after discharge, the full spike intensity will contribute to recharging of the empty capacitor. Thus, the spike will fully add its sudden contribution to constant I_{in0} , and hence, will produce the biggest reduction of the ISI. In contrast, if spike I_A arrives at a longer time, t_0 , after the discharge event, it will find that the membrane capacitor is already slowly and partially filled by constant I_{in0} . The arriving contribution will then result in the neuron firing, but with a small effect on the reduction of the ISI. In other words, just a portion of spike I_A will be enough to allow the neuron to fire, while the remaining portion of the spike will be “wasted” to ground through the low resistance of the thyristor during the firing discharge. Clearly, a longer t_0 will lead to a lower current of the feedback spike being required, so the effect of feedback in the reduction of the ISI will be smaller. This is approximately linear in t_0 , since the membrane capacitor charges approximately linearly in time. It is interesting to notice also that there is an archlike behavior of the $ISI(\Delta t)$ for small delays, Δt , below 12 ms and just after the jump at 20 ms. This can be traced to spike I_A arriving during linear charging but suddenly bringing V_C to the nonlinear charging regime, where the thyristor resistance decreases before switching *on*. We provide traces in the Supplemental Material [14] that illustrate this effect.

On closer examination of $ISI(\Delta t)$, we also note that, near delays Δt corresponding to multiples of the ISI_0 , the ISI value slightly increases above the blue line of ISI_0 . This unexpected feature can be traced to the situation where reinjected spike $I_A(t)$ arrives during discharging of the capacitor, i.e., during the firing event. Thus, since the thyristor is *on*, this feedback current is again wasted to ground, as discussed before. However, there is an additional subtle effect, namely, the I_A spike also keeps the thyristor *on* (i.e., above I_{hold}) for a slightly longer time,

during the whole duration of the I_A spike. While the thyristor is *on*, the capacitor cannot be recharged by I_{in0} , which leads to the observed slight increase of the ISI. We may consider these subtle features as emerging from realistic “finite-size” effects. In fact, they originate in the finite duration of the action-potential spike. Therefore, they would be missed in a mathematical model simulation of an autapse spiking neuron, if the action potential were represented by a schematic δ function. Detailed traces of I_{in} , V_C , and I_{out} are shown in the Supplemental Material [14].

In the inhibitory case, the feedback current, $I_A(t)$, is always subtracted from the input, so the self-consistent ISI is, in general, increased, since the membrane capacitor needs more time to charge. Interestingly, and in stark contrast to the previous case, however, the system does not always converge to a single self-consistent ISI state. In fact, for delay Δt values that are approximately commensurate with ISI_0 , we observe complex dynamic behavior [this is indicated by open circles in Fig. 2(b), which denote the average ISI, so the nonmonotonic behavior of ISI is only apparent]. Similarly to that discussed before for the excitatory case, this feature may also be partially due to the effect of the finite duration of the action-potential spike and its eventual overlap in time with the capacitor discharging and the conductive state of the thyristor. However, in this case, we are not able to provide even a qualitative description, as the emergent behavior is likely chaotic. This phenomenon is explored in greater detail in Fig. 3.

The results of our measurements show interesting indications of chaotic behavior. For instance, decreasing Δt from 46 ms [Fig. 3(b)], we observe that the single self-consistent ISI splits in two, and then the two ISIs simultaneously split in two again. This is behavior reminiscent of the period-doubling route to chaos [37], which is reported in a mathematical model of an autapse with inhibitory self-interactions [38]. More generally, mathematical models of spiking neurons often show chaotic behavior [11]. This is the case, for instance, of the Adaptive Exponential (AdEx) model [39], which also has negative feedback [10].

We also note that the system may take a long time to converge. In the Supplemental Material [14], we show some instances of slow convergence. This feature is relevant for potential issues in numerical simulation methods, which may require significant computational effort to sort out. In contrast, the present circuit model deals with this feature in a natural physical manner, as it operates in continuous real time.

We now turn to explore the case where we vary the feedback gain parameter, β , keeping the time delay, Δt , fixed. Motivated by previous results, we consider the two cases in which Δt is different and similar to ISI_0 , as shown in Figs. 4(a) and 4(b), respectively. In the first case, we observe that, away from the commensurate delays, the ISI converges to a single self-consistent value. Consistent with




FIG. 3. Self-consistent ISI as a function of delay Δt for the case of inhibitory feedback, as in Fig. 2(b). Grayscale intensity in (a),(b) represents the relative frequency of the distribution of ISI values. Notice that, for delays near the ISI_0 -commensurate values, there are periodic sequences of multiple ISIs. Blue box in (a) indicates the region of complex behavior that is expanded for greater detail in (b). (c) Examples of periodic sequences at specific Δt values, as indicated by dashed lines in (b). The z scale in (a),(b) represents the number of events during 300-ms and 1.5-s time windows, respectively.

our previous discussion, we observe that positive and negative β increase and decrease the spiking rate, respectively. This may be qualitatively understood by the fact that the total input current is, respectively, increased or decreased by feedback, leading to either faster or slower charging of the capacitor. However, we also observe that, at high feedback intensities, for both positive and negative gains, the spike rate eventually saturates. This can be understood by the fact that the UCN reaches its maximal ($\beta > 0$) or minimal ($\beta < 0$) firing frequencies, which are set by the values of the circuit components and the *on* and *off* values of the thyristor resistance [12].

In contrast to this simple behavior, we observe in Fig. 4(b) that, for a Δt value of 20 ms, which is similar to ISI_0 , complex dynamic behavior emerges for the inhibitory

case. This is consistent with the previous discussion; however, we may also observe that the qualitative nature of the complexity, such as duplication of the period, seems rather insensitive to the intensity of the inhibitory-feedback parameter, β . This indicates that it is the timing, Δt , which has a bigger influence on the behavior.

We conclude this section by mentioning that the UCN considered here realizes a leaky-integrate-and-fire spiking neuron model, which is by far the most common one. However, it is an interesting question to explore the effects of recursion in a minimal configuration of a spiking neuron that realizes other dynamic behavior. Different biologically relevant models were implemented as variations of the UCN circuit in Ref. [13]. In the Supplemental Material [14], we consider the effect of excitatory recursion




FIG. 4. Self-consistent ISI as a function of delay β at fixed values of delay Δt . (a) Behavior away from delay commensuration; (b) $\Delta t = 20$ ms \sim ISI_0 case, where a doubling period is observed for the inhibitory case. Grayscale indicates the relative frequency of ISI values (number of events during a 200-ms window). Supplemental Material [14] contains waveforms of I_{in} , V_C , and I_{out} for selected points of these plots.




FIG. 5. Left (a), schematic circuit of the dynamic working memory with an UCNA. Additional synapse (hashed brown) arriving at the soma acts with short excitatory or inhibitory pulses. Externally applied input current to the neuron is $I_{in} = I_{in0} + I_{ex-in}$. Right, experimental demonstration of graded persistent activity in the UCNA. (b) Stepwise increase of the activity and stability of the instantaneous spiking frequency. Note that the first excitatory pulse initiates neuron activity, and the second inhibitory pulse stops it. (c) Input current to the UCNA. Excitatory (inhibitory) input pulses are added (subtracted) to the constant current input, $I_{in0} = 90 \mu A$. (d) The (soma) membrane voltage $V_C(t)$. (e) Output current of the UCNA with graded spiking activity. Experimental parameters are as follows I_{in} has a baseline of $90 \mu A$ (I_{in0}) and 4 pulses (2 excitatory pulses followed by 2 inhibitory pulses) of $\pm 500 \mu A$ (with respect to the baseline) 5 ms wide, which are applied at times 150 ms, 346, 605, and 834 ms. Feedback has a delay of 30 ms and gain of $\beta = 2$.

in two such models to illustrate this point. A more thorough exploration is beyond the present scope and is left for future work.

B. Dynamic working memory

Working memory is an emergent feature of brain activity, but its cellular mechanism remains elusive. Working, or short-term, memory is the persistent dynamic state that allows one to remember the words that one reads before reaching the period punctuation mark at the end of this sentence. The mechanism is also related to the fixation mechanism in eye movement [40]. Key features of this function are stability and tunability. This is called graded persistent neural activity. While it has been historically assumed that this function emerges from the reverberating behavior of recurrent neural networks [41,42], impressive experimental work has demonstrated that graded persistent

activity can be observed at the single-neuron level [43,44]. Mathematical models of graded persistent activity in single neurons struggle with the issue of stability [15,36,40].

Here, we consider the model introduced by Seung *et al.* [15] as the starting point for our investigation into graded persistent activity. We schematically depict the model in Fig. 5(a). It consists of a neuron autapse that is subject to a constant input current, with persistent dynamic states that are tuned by short excitatory or inhibitory bursts. The former should induce and increase the spiking rate, while the latter should decrease or terminate it.

We physically implement the model by means of an UCNA and excite it by applying an input current with excitatory-inhibitory synapse pulses [Figs. 5(a) and 5(c)]. As shown in Figs. 5(d) and 5(e), we succeed at inducing and suppressing the spiking activity and, crucially, achieve stable and tunable dynamic memory states. Our main result is shown in Fig. 5(b), where

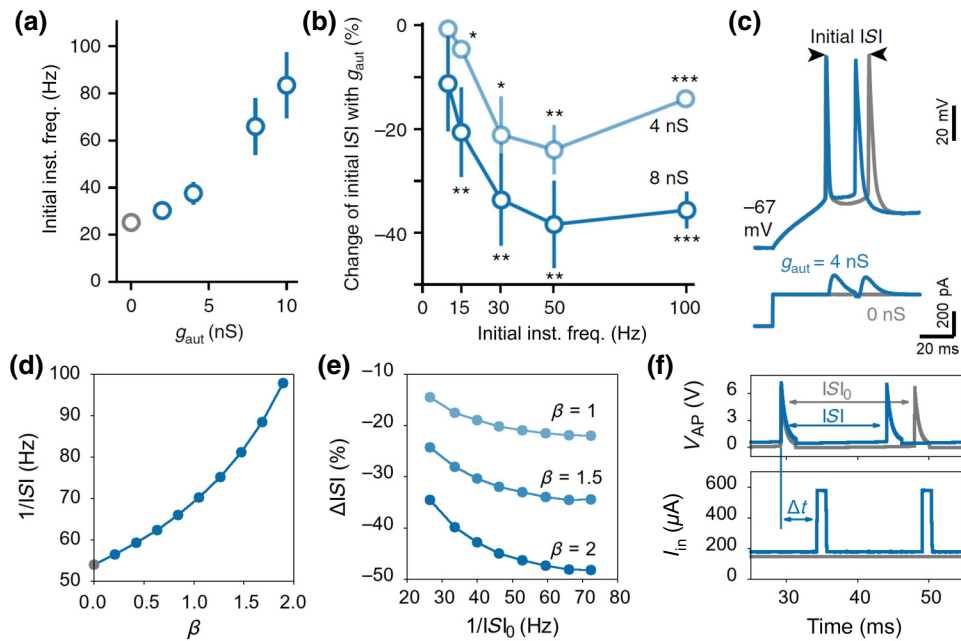


FIG. 6. Top panels, experimental data measured in a pyramidal cell autapse by Yin *et al.* [16]. (a) Initial instantaneous frequency, $1/ISI$, as a function of applied conductance (i.e., current intensity/67 mV). (b) Relative change in ISI as a function of initial instantaneous frequency, $1/ISI_0$, measured at two different constant input conductances of 4 and 8 nS. (c) Line shape of measured action potentials without and with applied autapse conductance (gray and blue, respectively). Autapse input (hump) is delayed by 1.4 ms. Lower panels, experimental data measured in the UCNA circuit. (d) Initial instantaneous frequency, $1/ISI$, as a function of intensity of applied feedback, β , at a fixed delay of $\Delta t = 5$ ms. (e) Relative change in ISI as a function of initial instantaneous frequency, $1/ISI_0$, measured at three different feedback intensities and at a fixed delay of $\Delta t = 5$ ms. (f) Top, line shape of action-potential spikes without and with applied autapse current, as shown in the bottom panel (gray and blue, respectively); $I_{in0} = 150$ mA, $I_A = 400$ mA, width of 2 ms. Lines are minimally vertically shifted for the sake of clarity. Similar to (c), top panel shows output spikes and lower panel shows input current.

we observe the graded persistent activity in the step-like $1/ISI$ rates [15]. In fact, we are able to qualitatively reproduce the basic features observed in biological neuroscience experiments [43,44]. Namely, short excitatory pulses increase the activity in a discrete manner, and inhibitory pulses decrease it likewise, and eventually can stop firing activity altogether. Moreover, our data also show that the behavior is stable on the scale of seconds. For clarity, we limit the results in Fig. 5 to 1 s, so the individual spikes can be resolved. A much longer interval is shown in the Supplemental Material [14].

It is interesting to note that the key feature to realize this graded behavior was to “polarize” the UCNA with a constant input current, which placed it right beneath its spiking threshold. Notice that the firing threshold of the UCN is independent of both, the gain parameter β and the delay Δt , since at subthreshold there is no spiking and obviously there is no feedback. The detailed exploration of the graded persistent activity is both, important and interesting, but goes beyond the scope of the present work and will be discussed elsewhere.

C. Modeling biological autapse experiments with an excitatory UCNA

We now turn to another exciting topic of autapse behavior that is enabled by our implementation of a physical spiking neuron. Namely, we shall try to replicate the observations reported in the recent experimental study by Yin *et al.* on autaptic neocortical pyramidal cells [16], which we reproduce in the top panels of Fig. 6. In that remarkable work, an autaptic neuron was initially identified and its activity recorded. Subsequently, the measured behavior was exploited to study, in a systematic manner, the effect of the autaptic feedback intensity by means of dynamic clamp methods [16]. In biological neuroscience, it is customary to specify the intensity of feedback in units of conductance (nS), which actually refers to the magnitude input currents. This is done by a conversion of units using the membrane potential (67 mV), which is fixed by the ionic equilibrium of the cell [16]. As shown in Fig. 6(c), typical input currents are thus in the order of hundreds of pA, which is about 67 mV nS. The intensity of the feedback conductance plays a role analogous to that of parameter β in our model. On the other hand, the delay in feedback is also experimentally

controlled and fixed at 1.4 ms in the experiments of Yin *et al.* [16]. This delay plays an analogous role to our model parameter Δt .

Our goal here is to capture some of the behavior of biological autaptic neurons with our UCNA physical model. Thus, we focus our attention on two particular experiments, where the initial instantaneous frequency, $1/\text{ISI}$, is measured and its systematic changes are observed. In the first experiment [Fig. 6(a)], the rate of $1/\text{ISI}$ is observed to increase supralinearly as a function of the excitatory feedback intensity, from the baseline value of $1/\text{ISI}_0$ at zero feedback (0 nS). In a second experiment [Fig. 6(b)], the percentage reduction of the ISI is measured as a function of the unperturbed rate of $1/\text{ISI}_0$, which itself is set by the input current injected to the neuron. Experimental data show a significant decrease that can reach up to about 40% and is observed to level off at the highest instantaneous rates.

In the lower panels of Fig. 6, we show the corresponding measured data from our UCNA physical circuit model. From the discussion in Sec. III A, we have a clear intuition of the qualitative behavior, which shows that excitatory feedback increases the firing rates by increasing the total input current. However, we observe that the systematic behavior of our artificial autapse is in remarkably good qualitative agreement with that of the biological autapse. Moreover, the spiking rates and the delay are at least in order-of-magnitude, or semiquantitative, agreement. While the comparison of the experimental magnitude of the autapse intensity, g_{aut} , and our model parameter β may *a priori* do not seem to be direct, we show next that an interesting quantitative comparison can be established.

The respective lower panels of Figs. 6(c) and 6(f) show the total input current, $I_{\text{in}} = I_{\text{in}0} + I_A$. In both cases, the input consists of a constant component [0 nS in Fig. 6(c) and $I_{\text{in}0}$ in gray in Fig. 6(f)] plus spikes that are mounted on it (g_{aut} and I_A in blue). From these current traces, we can be more quantitative and obtain the *injected charge* by considering the integral of the respective current components. In fact, the charge is independent of the specific shape of the current signal and is a direct measure of the contributions to membrane potential (since $Q = CV$). From biological experimental data [Fig. 6(c)], we can estimate, for the autapse, $100 \text{ pA} \times 20 \text{ ms}/2 = 1 \text{ pA s}$ (the factor $1/2$ approximates the triangular shape of the bump) and, for constant input, $200 \text{ pA} \times 20 \text{ msec} = 4 \text{ pA s}$, thus a ratio of about $1/4$ or about 25%. As we shall see below, this ratio is consistent with that of the artificial neuron model. Following a similar procedure, we estimate the amount of charge contributed by the constant input current during the two spikes to be $I_{\text{in}0} \times \text{ISI} = 150 \text{ mA} \times 20 \text{ ms} = 3 \text{ mA s}$. On the other hand, the amount of charge contributed by the autapse is obtained from the spike intensity times its duration: $I_A \times 2 \text{ ms} = 400 \text{ mA} \times 2 \text{ ms} = 0.8 \text{ mA s}$. Thus, we obtain, for the UCNA model, an autapse-to-constant-input-

contribution ratio of about $0.8/3$ or 26%, which is in remarkable agreement with the estimate from biological autapse data.

IV. CONCLUSIONS

One of the main goals of the present work is to demonstrate that the UCN is a sound experimental platform for a model of spiking neurons. To this end, we focus on the behavior of the simplest recurrent spiking neural network, namely, a single neuron with an autaptic synapse. This minimal network is, nevertheless, of biological relevance, as autapses are observed in the cortex. We explore both excitatory and inhibitory feedback and observe that, consistent with previous numerical studies of mathematical models, the behavior is qualitatively different in these two cases. Moreover, to count now with a physical model implementation of a spiking neuron allows us to elucidate some matters of previous debate. For instance, our study sheds light on the issue of chaos, which is ubiquitously present in mathematical models of neurons with negative feedback, such as the AdEx and recurrent networks [11,36,38,39,45–47]. From data of our UCNA model, we do observe traces of chaoticlike instabilities, which only approximately develop, but ultimately find converged nontrivial states. This may be traced to the fact that mathematical models describe action potentials by δ functions [45], while, in contrast, a physical model fires action potentials of a finite width. Consistent with that observation, we find that chaotic features emerge for the longest delays, $\Delta t \approx 40 \text{ ms}$, which are much longer than the width of the spike pulses, about 2 ms. An even more attractive advantage of counting with a physical model of a spiking neuron is that one may explore, with the same amount of effort, the long-time-limit behavior of the networks. This limit is challenging to reach in a mathematical simulation. The time step, δt , in a numerical calculation is set by the duration of a fast timescale, such as the leaky time constant, τ , so that $\delta t \ll \tau$ is required. On the other hand, the duration of the action potential, which in mathematical models is represented by a δ function, is therefore about δt . Thus, one should adopt a time step, δt , of at least 2 orders of magnitude shorter than τ , which in neurons is in the order of ms. Then, to explore the dynamic behavior of a neuron at timescales in the order of just 1 s, one needs about 10^5 time steps. If one wants to simulate a network with multiple interconnected neurons, the computational time requirements grow with the number of neurons, and the multiplicity of spikes also makes keeping error propagations under control a taxing task. Moreover, a recurrent spiking neuron is a strongly nonlinear dynamic system and, even for the simple autapse, we observe multiple instances of slow convergence to reach asymptotic behavior in the order of several seconds. One may expect that this type of slow-convergence problem may become a formidable

challenge in the numerical simulation of large networks [48]. In contrast, the evolution of a UCN-based network is a physical circuit, and it takes place in real continuous physical time. Therefore, once the circuit is setup, the time to run it does not depend on the number of neurons in the network, their connectivity, or their level of activity. One does not need to worry about error propagation nor mathematical approximations.

While the previous aspects emphasize the methodological interest of our approach, we also explicitly show how we can address long-standing problems in neuroscience. Specifically, we demonstrate that a single neuron autapse model of working memory can realize stable graded persistent activity, which remains an issue from approximate mathematical treatments [43,44]. Moreover, we also demonstrate that we can make comparisons with actual neurobiological experiments. Indeed, we demonstrate that our artificial neuron autapse can semiquantitatively capture recent current-clamp experiments on a biological neuron autapse [16]. The systematic exploration of this behavior as a function of parameters may serve to establish a dialogue with experimental neuroscience and is an exciting possibility left open for future work.

From a wider perspective, our results are a strong motivation to move on to consider multineuron recurrent networks and begin to address some of the fundamental questions that are posed by the challenge of understanding the nature of the neural code.

ACKNOWLEDGMENTS

M.J.R. acknowledges support from the French ANR Grant No. ANR-19-CE30-0020 “MoMA.” This work is supported, in part, by the JSPS KAKENHI Grant No. JP18H05911. This work is part of the project 234270 CCR LPS-AIST “Bioinspired electronic systems.”

P.S. performed all experiments. O.S. performed all numerical simulations. M.R. conceived the idea. All authors contributed to writing of the manuscript.

-
- [1] A. Graves, A. Mohamed, and G. Hinton, in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, (2013), pp. 6645–6649.
 - [2] X.-J. Wang, Decision making in recurrent neuronal circuits, *Neuron* **60**, 215 (2008).
 - [3] K. Rajan, C. D. Harvey, and D. W. Tank, Recurrent network models of sequence generation and memory, *Neuron* **90**, 128 (2016).
 - [4] H. F. Song, G. R. Yang, and X.-J. Wang, Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework, *PLoS Comput. Biol.* **12**, 1 (2016).

- [5] D. Sussillo and L. F. Abbott, Generating coherent patterns of activity from chaotic neural networks, *Neuron* **63**, 544 (2009).
- [6] L. Abbott, B. DePasquale, and R. M. Memmesheimer, Building functional networks of spiking model neurons, *Nat. Neurosci.* **19**, 350 (2016).
- [7] J. H. Lee, T. Delbruck, and M. Pfeiffer, Training deep spiking neural networks using backpropagation, *Front. Neurosci.* **10**, 508 (2016).
- [8] D. Thalmeier, M. Uhlmann, H. J. Kappen, and R. M. Memmesheimer, Learning universal computations with spikes, *PLoS Comput. Biol.* **12**, 1 (2016).
- [9] A. D. Reyes, Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro, *Nat. Neurosci.* **6**, 593 (2003).
- [10] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition* (Reading: Cambridge University Press, Cambridge, UK, 2014).
- [11] M. I. Rabinovich, P. Varona, A. I. Selverston, and H. D. I. Abarbanel, Dynamical principles in neuroscience, *Rev. Mod. Phys.* **78**, 1213 (2006).
- [12] M. Rozenberg, O. Schneegans, and P. Stoliar, An ultra-compact leaky-integrate-and-fire model for building spiking neural networks, *Sci. Rep.* **9**, 11123 (2019).
- [13] P. Stoliar, O. Schneegans, and M. Rozenberg, Biologically relevant dynamical behaviors realized in an ultra-compact neuron model, *Front. Neurosci.* **14**, 421 (2020).
- [14] See the Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevApplied.16.034030> for a list of components, details on autapse working memory, details on the convergence time, additional waveforms, behavior in different neuron types, and experimental details.
- [15] H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank, The autapse: A simple illustration of short-term analog memory storage by tuned synaptic feedback, *J. Comput. Neurosci.* **9**, 171 (2000).
- [16] L. Yin, R. Zheng, W. Ke, Q. He, Y. Zhang, J. Li, B. Wang, Z. Mi, Y. S. Long, M. J. Rasch, T. Li, G. Luan, and Y. Shu, Autapses enhance bursting and coincidence detection in neocortical pyramidal cells, *Nat. Commun.* **9**, 4890 (2018).
- [17] H. van der Loos and E. M. Glaser, Autapses in neocortex cerebri: Synapses between a pyramidal cell’s axon and its own dendrites, *Brain Res.* **48**, 355 (1972).
- [18] J. M. Bekkers, Synaptic transmission: Functional autapses in the cortex, *Curr. Biol.* **13**, R433 (2003).
- [19] A. Bacci, J. R. Huguenard, and D. A. Prince, Functional autaptic neurotransmission in fast-spiking interneurons: A novel form of feedback inhibition in the neocortex, *J. Neurosci.* **23**, 859 (2003).
- [20] A. Bacci and J. R. Huguenard, Enhancement of spike-timing precision by autaptic transmission in neocortical inhibitory interneurons, *Neuron* **49**, 119 (2006).
- [21] W. Yi, K. K. Tsang, S. K. Lam, X. Bai, J. A. Crowell, and E. A. Flores, Biological plausibility and stochasticity in scalable VO₂ active memristor neurons, *Nat. Commun.* **9**, 4661 (2018).
- [22] P. Dayan and L. F. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (The MIT Press, Cambridge, MA, 2005).

- [23] H. Fan, Y. Wang, H. Wang, Y. C. Lai, and X. Wang, Autapses promote synchronization in neuronal networks, *Sci. Rep.* **8**, 580 (2018).
- [24] C. S. Thakur, J. L. Molin, G. Cauwenberghs, G. Indiveri, K. Kumar, N. Qiao, J. Schemmel, R. Wang, E. Chicca, J. Olson Hasler, and J. S. Seo, Large-Scale neuromorphic spiking array processors: A quest to mimic the brain, *Front. Neurosci.* **12**, 891 (2018).
- [25] G. Bellec, F. Scherr, A. Subramoney, E. Hajek, D. Salaj, R. Legenstein, and W. Maass, A solution to the learning dilemma for recurrent networks of spiking neurons, *Nat. Commun.* **11**, 3625 (2020).
- [26] R. Kim, Y. Li, and T. J. Sejnowski, Simple framework for constructing functional spiking recurrent neural networks, *PNAS* **116**, 22811 (2019).
- [27] G. Indiveri, B. Linares-Barranco, T. J. Hamilton, A. Van Schaik, R. Etienne-Cummings, T. Delbruck, S. C. Liu, P. Dudek, P. Häfliger, S. Renaud, and J. Schemmel, Neuromorphic silicon neuron circuits, *Front. Neurosci.* **5**, 73 (2011).
- [28] P. Stoliar, O. Schneegans, and M. Rozenberg, A functional spiking neural network of ultra compact neurons, *Front. Neurosci.* **15**, 635098 (2021).
- [29] A. L. Hodgkin and A. F. Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve, *J. Physiol. London* **117**, 500 (1952).
- [30] A. Strong, *The LT1167: Precision, Low Cost, Low Power Instrumentation Amplifier Requires a Single Gain-Set Resistor* (Linear Technology Magazine, Milpitas, CA, USA, 1998).
- [31] J. Del Valle, J. G. Ramirez, M. J. Rozenberg, and I. K. Schuller, Challenges in materials and devices for resistive-xxswitching-based neuromorphic computing, *J. Appl. Phys.* **124**, 211101 (2018).
- [32] P. Stoliar, L. Cario, E. Janod, B. Corraze, C. Guillot-Deudon, S. Salmon-Bourmand, V. Guiot, J. Tranchant, and M. Rozenberg, Universal electric-field-driven resistive transition in narrow-gap mott insulators, *Adv. Mater.* **25**, 3222 (2013).
- [33] E. Janod, J. Tranchant, B. Corraze, M. Querré, P. Stoliar, M. Rozenberg, T. Cren, D. Roditchev, V. T. Phuoc, M. P. Besland, and L. Cario, Resistive switching in Mott insulators and correlated systems, *Adv. Func. Mater.* **25**, 6287 (2015).
- [34] J. del Valle, P. Salev, F. Tesler, N. M. Vargas, Y. Kalcheim, P. Wang, J. Trastoy, M. H. Lee, G. Kassabian, J. G. Ramirez, and M. J. Rozenberg, Subthreshold firing in Mott nanodevices, *Nature* **569**, 388 (2019).
- [35] T. Tuma, A. Pantazi, M. Le Gallo, A. Sebastian, and E. Eleftheriou, Stochastic phase-change neurons. *Nat. Nanotech.* **11**, 693 (2016)
- [36] H. Wang, L. Wang, Y. Chen, and Y. Chen, Effect of autaptic activity on the response of a Hodgkin-Huxley neuron, *Chaos* **24**, 033122 (2014).
- [37] M. J. Feigenbaum, The universal metric properties of non-linear transformations, *J. Stat. Phys.* **21**, 669 (1979).
- [38] F. Pasemann, A simple chaotic neuron, *Physica D: Nonlinear Phenomena* **104**, 205 (1997).
- [39] W. Gerstner and R. Brette, Adaptive exponential integrate-and-fire model, *Scholarpedia* **4**, 8427 (2009).
- [40] H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank, Stability of the memory of eye position in a recurrent network of conductance-based model neurons, *Neuron* **26**, 259 (2000).
- [41] R. Lorente de Nó, Vestibulo-ocular reflex arc, *Arch. Neuropsych.* **30**, 245 (1933).
- [42] D. Zipser, B. Kehoe, G. Littlewort, and J. Fuster, A spiking network model of short-term active memory, *J Neurosci.* **13**, 3406 (1993).
- [43] A. V. Egorov, B. N. Hamam, E. Fransén, M. E. Hasselmo, and A. A. Alonso, Graded persistent activity in entorhinal cortex neurons, *Nature* **420**, 173 (2002).
- [44] E. Fransén, B. Tahvildari, A. V. Egorov, M. E. Hasselmo, and A. A. Alonso, Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons, *Neuron* **49**, 735 (2006).
- [45] S. Nobukawa, H. Nishimura, and T. Yamanishi, Routes to chaos induced by a discontinuous resetting process in a hybrid spiking neuron model, *Sci. Rep.* **8**, 379 (2018).
- [46] C. van Vreeswijk and H. Sompolinsky, Chaotic balanced state in a model of cortical circuits, *Neural Comput.* **10**, 1321 (1998).
- [47] E. M. Izhikevich, Which model to use for cortical spiking neurons?, *IEEE Trans. Neural Netw.* **15**, 1063 (2004).
- [48] E. M. Izhikevich and G. M. Edelman, Large-scale model of mammalian thalamocortical systems, *PNAS* **105**, 3593 (2008).