PREPRINT

A Robust Model of Gated Working Memory

Anthony Strock[1,2,3], Xavier Hinaut[1,2,3,†], and Nicolas P. Rougier[1,2,3,†]

[1]INRIA Bordeaux Sud-Ouest, Bordeaux, France – [2]LaBRI, Université de Bordeaux, Institut Polytechnique de Bordeaux, Centre National de la Recherche Scientifique, UMR 5800, Talence, France – [3]Institut des Maladies Neurodégénératives, Université de Bordeaux, Centre National de la Recherche Scientifique, UMR 5293, Bordeaux, France – †These authors contributed equally

Published
March 27, 2019

License
CC BY 4.0

Version
1.0

Abstract Gated working memory is defined as the capacity of holding arbitrary information at any time in order to be used at a later time. Based on electrophysiological recordings, several computational models have tackled the problem using dedicated and explicit mechanisms. We propose instead to consider an implicit mechanism based on a random recurrent neural network. We introduce a robust yet simple reservoir model of gated working memory with instantaneous updates. The model is able to store an arbitrary real value at random time over an extended period of time. The dynamics of the model is a line attractor that learns to exploit reentry and a non-linearity during the training phase using only a few representative values. A deeper study of the model shows that there is actually a large range of hyper parameters for which the results hold (number of neurons, sparsity, global weight scaling, etc.) such that any large enough population, mixing excitatory and inhibitory neurons can quickly learn to realize such gated working memory. This suggests this property could be an implicit property of any random population, that can be acquired through learning. Furthermore, considering working memory to be a physically open but functionally closed system, we give account on some counter-intuitive electrophysiological recordings.

Keywords Working Memory, Line Attractor, Reservoir Computing, Computational Neuroscience, Prefrontal Cortex, Model

Contents

1 Introduction 2

2 Methods 4
- 2.1 The n-value p-gate scalar working memory tasks . . . 4
- 2.2 The digit 1-value 1-gate working memory task . . . 5
- 2.3 The reduced model . . . 6
- 2.4 The reservoir model . . . 7

3 Results 7
- 3.1 The reduced model . . . 7
- 3.2 The reservoir model . . . 8
- 3.3 Robustness . . . 11
- 3.4 Dynamics . . . 12

4 Discussion 15

5 References 16

6 Supplementary 19

Copyright © 2019 A. Strock, X. Hinaut and N.P. Rougier, released under a Creative Commons Attribution 4.0 International license.
Correspondence should be addressed to Nicolas P. Rougier (Nicolas.Rougier@inria.fr) The authors have declared that no competing interests exists. Code is available at https://github.com/rougier/ESN-WM.

1 Introduction

The prefrontal cortex (PFC), noteworthy for its highly recurrent connections[1], is involved in many high level capabilities, such as decision making[2], working memory[1], goal directed behavior[3], temporal organisation and reasoning[4]. In this article, we are more specifically interested in gated working memory[5] that is defined as the capacity of holding arbitrary information at a given random time t_0 such as to be accessible at a later random time t_1 (see figure 1). Between times t_0 and t_1, we make no assumption on the inner mechanisms of the working memory. The only measures we are interested in are the precision of the output (compared to the initial information) and the maximal delay during which this information can be accessed within a given precision range. One obvious and immediate solution to the task is to make an explicit copy (inside the memory) of the information at time t_0 and to hold it unchanged until it is read at time t_1, much like a computer program variable that is first assigned a value in order to be read later. Such solution can be easily characterized by a fixed pattern of sustained activities inside the memory. This is precisely what lead researchers to search for such sustained activity inside the frontal cortex[6,7], where an important part of our working memory capacities are believed to be located. Romo et al.[8] have shown that PFC neurons of non-human primate can maintain information about a stimulus for several seconds. Their firing rate was correlated with the coding of a specific dimension (frequency) of the stimulus maintained in memory. However, when Machens, Romo, and Brody[9] later re-analyzed the data of this experiment, they showed that the stimulus was actually encoded over a sub-population using a distributed representation. Similarly, when Rigotti et al.[10] analyzed single neuron activity recorded in the lateral PFC of monkeys performing complex cognitive tasks, they found several neurons displaying task-related activity. Once they discarded all the neurons that were displaying task-related activity, they were still able to decode task information with a linear decoder. They proposed that the PFC hosts high-dimensional linear and non-linear mixed-selectivity activity. The question is thus, if working memory is not encoded in the sustained activity, what can be the alternatives?

Before answering that question, let us first characterize the type and the properties of information we consider before defining what it means to access the information. The type of information that can be stored inside working memory has been characterized using different cognitive tasks such as for example the delayed matching-to-sample (DMTS), the N-back task or the Wisconsin Card Sorting Task (WCST). From these different tasks, we can assume that virtually any kind of information, be it auditory or visual stimulus, textual or verbal instruction, implicit or explicit cue, is susceptible to be memorized and processed inside the working memory. From a computational point of view, this can be abstracted into a set of categorical, discrete or continuous values. In this work, we are only interested in the most general case, the continuous one, which can be reduced into a single scalar value. The question we want to address is how a neural population can gate and maintain an arbitrary (e.g. random) value, in spite of noise and

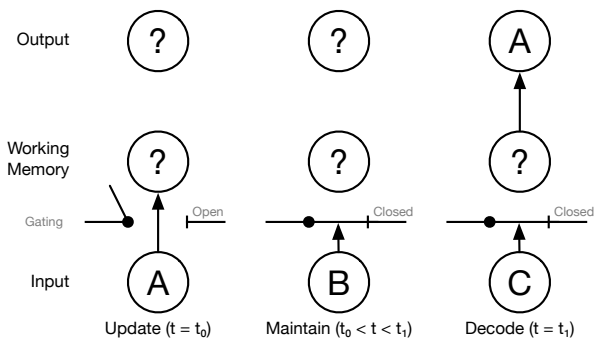

Figure 1. Gated working memory is defined as the capacity of holding arbitrary information at a given random time t_0 such as to be accessible at a later random time t_1. Between times t_0 and t_1, we make no assumption on the activity inside the working memory. Note that a closed gate does not mean external activities cannot enter the memory.

distractors, such that it can be decoded non-ambiguously at a later time.

To answer this question we can search the extensive literature on computational models of working memory that have been extensively reviewed by Durstewitz, Seamans, and Sejnowski[11] in 2000, by Compte[12] in 2006 and more recently by Barak and Tsodyks[13] in 2014. More specifically, Compte[12] explains that the retention of information is often associated with attractors. In the simplest case, a continuous scalar information can be identified to a position on a line attractor. However, this kind of memory exhibits stability issues: unlike a point attractor (i.e. a stable fixed point), a line attractor is marginally stable such that it cannot be robust against all small perturbations. Such line attractor can be stable against orthogonal perturbations, but not against colinear perturbations (i.e. perturbations along the line). Furthermore, the design (and numerical implementation) of a line attractor is tricky because even small imperfections (e.g. numerical errors) can lead to instability. Nevertheless, there exist several models that can overcome these limitations.

This is the case for the theoretical model by Amari[14] who proved that a local excitation could persist in the absence of stimuli, in the form of a localized bump of activity in an homogeneous and isotropic neural field model, using long range inhibition and short range excitation. This model represents *de facto* a spatial attractor formed by a collection of localized bumps of activity over the plane. A few decades later, Compte[15] showed that the same lasting bump property can also be achieved using leaky integrate and fire neurons arranged on a ring, with short range excitation and constant inhibition (ring bump attractor). This model has since then been extended[16,17] with the characterization of the conditions allowing to have simultaneous bumps of activity. This would explain multi-item memorization where each bump represents a different information that is maintained simultaneously with the other bumps. Similarly, Bouchacourt and Buschman[18] proposed to handle multi-items memorization by duplicating bump attractor model. They explicitly limited the number of items to be maintained in memory through the interaction between the different bump attractor models (using a random layer of neurons). If all these models can cope with the memorization of a graded information, this information is precisely localized in the bumps of activity and corresponds to a sustained activity. Such patterns of activity have been identified in several cerebral structures (e.g. head direction cells[19] in mammals, superior colliculus[20] in primates) but it is not yet clear to what extent this can give account of a general working memory mechanism. Such sustained activity is also present in the model of Koulakov et al.[21] who consider a population of bistable units that encodes a (quasi) graded information using distributed encoding (percentage of units in high state). This solves both the robustness and stability issue of the line attractor by virtue of discretization. Finally, some authors[22,23] consider the encoding of the value to be correlated with the firing rate of a neuron or of a group of neurons. This is the case for the model proposed by Lim and Goldman[23] who obtain stability of the firing rate by adding a negative derivative self feedback (hence artificially increasing the time constant of neurons). They show how such mechanism can be implemented by the interaction of two populations of excitatory and inhibitory neurons evolving at different time scales. However, and independently of the encoding of the graded value, most of model authors are interested in characterizing the mechanism responsible for the maintenance property. They tend to consider the memory as an isolated system, not prone to external perturbations, with the noticeable exception of the model by Zipser et al.[22] which is constantly fed by an input.

In this work, we consider working memory to be an open system under the constant influence of external activities (i.e. even when the gate is *closed*). Thus, we can not rely on a dynamical system that hosts a line attractor in the absence of inputs. We have to design an input-dependent dynamical system that is robust against all kinds of perturbations (input, internal, output feedback). First, we will formalize a set of tasks that

will be used to study features and performances of the different models we have considered. Then, we will introduce a minimal model that will help us in explaining the mechanism needed for a more general model. For this general one, we will consider a particular instance of reservoir: namely an Echo State Network (ESN)[24]. The analysis of this model will allow us to show that reservoir activity is characterized by a combination of both sustained and transient activities. Moreover, we will show that none of these activities are critical for the decoding of the correct output. Finally, we will show that in the absence of input the dynamics of the model implement a segment attractor (i.e. a line attractor with bounding values).

2 Methods

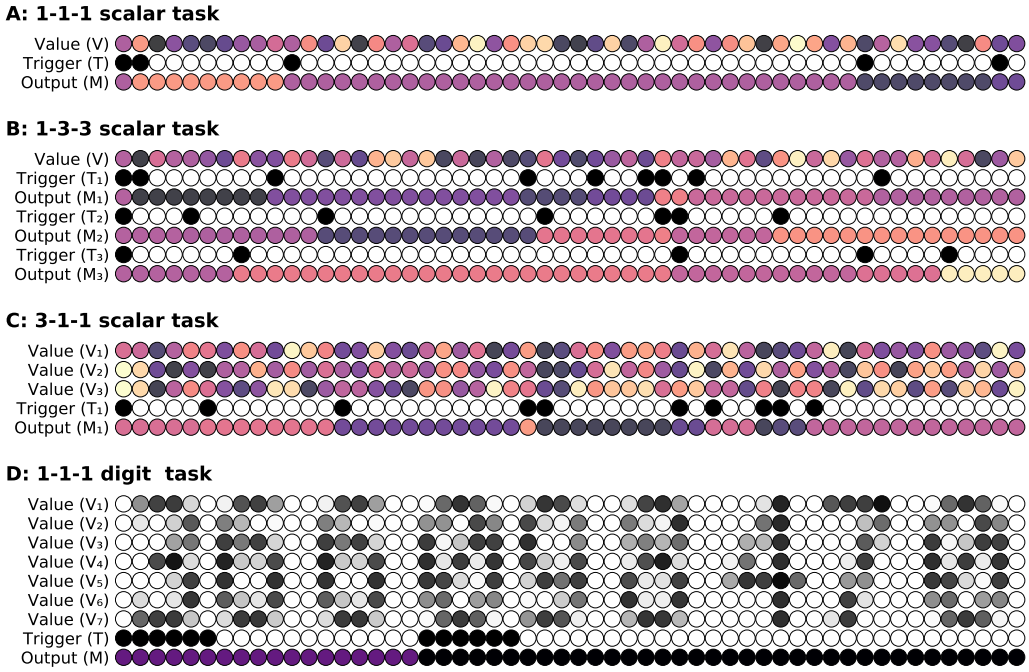

Figure 2. **Working memory tasks** Each column represents a time step (time increases from left to right), colored discs represent activity in the input (V or V_i and T or T_i) and the output (M or M_i). **A.** Gated working memory task with one gate. **B.** Gated working memory task with three gates. **C.** Gated working memory task with one gate but three inputs. To solve the task, it is necessary to ignore the two irrelevant inputs. **D.** Gated working memory task with one gate where the scalar input V has been replaced by a streamed input of bitmap digits. To solve the task, it is thus necessary to recognize the digits and to transform them into a normalized value.

In this section, we formalize and extend the gated working memory task that has been described in the introduction (see Figure 1). We consider four different tasks that will illustrate the generic capacity of the reservoir model to store continuous values or a discrete set of values. The first three are variations of the WM task for continuous values with various number of input values (n-value) and gating WM units (n-gate). The last task includes a non-linear computation (i.e. digit recognition from pixels) in addition to the gating task for discrete values.

2.1 The n-value p-gate scalar working memory tasks

In the 1-value 1-gate scalar version of the task, we consider at time $t \in \mathbb{N}$ an input signal $V(t)$ in [-1,+1], an input trigger $T(t)$ in {0, 1} and an output $M(t)$ that is defined as the

value of $V(t^*)$ where t^* is the most recent time such that $T(t^*) = 1$ (see figure 2A). This can be written as:

$$M(t) = V(t^*) \text{ with } t^* = \max_{0 \leq t' \leq t} (t' \mid T(t') = 1) \tag{1}$$

Said differently, the output is the value of the input when the gate was open for the last time. Note that this can be also rewritten as a simple select operator between the input V and the output M:

$$M(t) = T(t)V(t) + (1 - T(t))M(t-1) \tag{2}$$

In the 1-value 3-gate version of the task, we consider at time $t \in \mathbb{N}$ an input signal $V(t)$ in [-1,+1], three input triggers $T_{\{1,2,3\}}(t)$ in {0, 1} and three outputs $M_{\{1,2,3\}}(t)$ that are respectively defined as the value of $V(t^*)$ where t^* is the most recent time such that $T_{\{1,2,3\}}(t^*) = 1$ (see figure 2B). This can be written as:

$$M_i(t) = M_i(t^*) \text{ with } t^* = \max_{0 \leq t' \leq t} (t' \mid T_i(t') = 1), i \in \{1, 2, 3\} \tag{3}$$

In the 3-value 1-gate scalar version of the task, we consider at time $t \in \mathbb{N}$ three input signals $V_{\{1,2,3\}}(t)$ in [-1,+1], an input trigger $T(t)$ in {0, 1} and an output $M(t)$ that is defined as the value of $V_1(t^*)$ where t^* is the most recent time such that $T(t^*) = 1$ (see figure 2C). This can be written as:

$$M(t) = V_1(t^*) \text{ with } t^* = \max_{0 \leq t' \leq t} (t' \mid T(t') = 1) \tag{4}$$

This can be easily generalized to a *n-value p-gate scalar task* with n input signals, p input triggers, and p outputs. Only the first input signal and the t input triggers determines the p outputs. The other n-1 input signals are additional inputs irrelevant to the task.

2.2 The digit 1-value 1-gate working memory task

In the digit version of the 1-value 1-gate working memory task, we define a slightly more complex version of the V input by considering a bitmap representation of it (see figure 2D). Using a monotype font (Inconsolata Regular[1]) at size 11 points, we draw a grayscaled bitmap representation of a sequence of random digits (0 to 9), each digit being of size 6×7 pixels (after having cropped top and bottom empty lines) and the trigger signal being expanded to the width of a glyph. The output is defined as a discrete and normalized value. It is to be noted that there is no possible liner interpolation between the different inputs, as it was the case for the scalar version. Formally, we can define the output as:

$$M(t) = f(V_{i \in [1,7]}(t^*), V_{i \in [1,7]}(t^* - 1), ..., V_{i \in [1,7]}(t^* - 5)) \text{ with } t^* = \max_{0 \leq t' \leq t} (t' \mid T(t') = 1) \tag{5}$$

with f a function that maps the representation of a digit to a normalized value. Since there are 10 values, the digit n is associated to $\frac{n}{10}$. This function has to be learned by the model in order to solve the task. It would have been possible to use the MNIST database[25] instead of a regular font but this would have also complexified the task, and make the training period much longer, because a digit is processed only when a trigger is present. If we consider for example a sequence of 25,000 digits and a trigger probability of 0.01, this represents (in average) 250 triggers for the whole sequence and consequently only 25 presentations per digit. In comparison, MNIST train data set has 60,000 digits, which would have required as much triggers, and a hundred times more digits.

[1]The Inconsolata font is available from https://www.levien.com/type/myfonts/inconsolata.html

2.3 The reduced model

It is possible to define a minimal degenerated reservoir model (if we consider X_1, X_2 and X_3 to be the reservoir) that takes explicitly advantage of the non-linearity to perform the gating mechanism, as shown on figure 3. There is no learning in this model and it

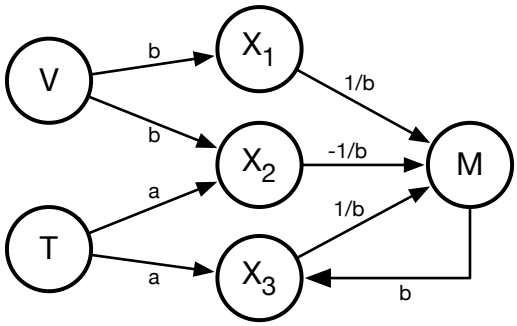

Figure 3. A minimal gated working memory model is able to solve the 1-value 1-gate working memory task by taking advantage of the non-linearity and asymptotic behavior of tanh units. Performance is controlled with parameters a and b. This architecture can be generalized to the *n-value p-gate task* by adding more than one reservoir unit (corresponding to X_3) for each supplementary trigger/output couple.

is fully defined by the following set of equations:

$$\begin{aligned} X_1[n+1] &= \tanh(b.V[n]) \\ X_2[n+1] &= \tanh(b.V[n] + a.T[n]) \\ X_3[n+1] &= \tanh(b.M[n] + a.T[n]) \\ M[n+1] &= (X_1[n+1] - X_2[n+1] + X_3[n+1])/b \end{aligned} \tag{6}$$

In order to understand how this minimal model works, we can write the output $M[n]$ relatively to the value of the trigger $T[n]$ which lies in $\{0, 1\}$:

$$\text{If } T[n] = 0 \begin{cases} X_1[n+1] & = \tanh(b.V[n]) \\ X_2[n+1] & = \tanh(b.V[n]) \\ X_3[n+1] & = \tanh(b.M[n]) \\ M[n+1] & = \tanh(b.M[n])/b \end{cases} \tag{7}$$

When $T[n] = 0$, if we assign b a small value (e.g. $b = 10^{-3}$) and considering that $\tanh(x) \underset{x \to 0}{\sim} x$, from equation (7), we have $M[n+1] \approx M[n]$.

$$\text{If } T[n] = 1 \begin{cases} X_1[n+1] & = \tanh(b.V[n]) \\ X_2[n+1] & = \tanh(b.V[n] + a) \\ X_3[n+1] & = \tanh(b.M[n] + a) \\ M[n+1] & = \tanh(b.V[n])/b \\ & \quad -(\tanh(b.V[n] + a) + \tanh(b.M[n] + a))/b \end{cases} \tag{8}$$

When $T[n] = 0$, if we assign a to a large value (e.g. $a = 1000$) and considering that b is small and that $\lim_{x \to \infty} \tanh(x) = 1$, we have $tanh(b.V[n] + a) \approx \tanh(b.M[n] + a) \approx 1$. From equation (8), we then have $M[n+1] \approx \tanh(b.V[n])/b \approx V[n]$.

Consequently, the trigger $T[n]$ fully dictates the output. When $T[n] = 0$, the output $M[n+1]$ is unchanged and corresponds to the current memory ($M[n]$). When $T[n] = 1$, the output $M[n+1]$ is assigned the value $V[n]$ of the input. We think this model represents the minimal model that is able to solve the gating working memory using such simple neurons (with tanh activation function). By taking advantage of the linear regime around 0 and the asymptotic behavior around infinity, this model gracefully solves the task using only two parameters (a and b). However, we have no formal proof that a

model having only two neurons in the reservoir part cannot solve the task. Note that this minimal model is similar to the Gated Recurrent Unit[26] (GRU) without the forget switch, but using only simple tanh neurons, in comparison to hand-crafted LSTM-like cells.

2.4 The reservoir model

We consider an echo state network[24] (see figure 4) with leaky neurons, and feedback from readout units to the reservoir, which is described by the following update equations:

$$\begin{aligned} x[n] &= (1-\alpha)x[n-1] \\ &\quad + \alpha \tanh\left(W_{in}u[n] + W(x[n-1]+\xi) + W_{fb}(y[n-1])\right) \\ y[n] &= W_{out}x[n] \end{aligned} \tag{9}$$

where $u[n]$, $x[n]$ and $y[n]$ are respectively the input, the reservoir and the output at time n. W, W_{in}, W_{fb} and W_{out} are respectively the recurrent, the input, the feedback and the output weight matrix and α is the leaking rate. ξ is a uniform white noise term added to the reservoir: ξ is independent for each neuron. During initialization, matrices W_{in}

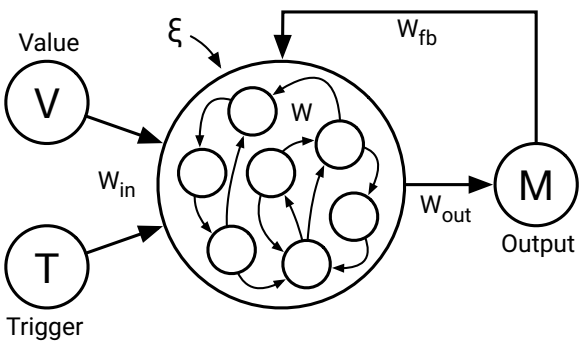

Figure 4. The reservoir model receives a random signal V in $[-1,+1]$ and a trigger signal T in $\{0,1\}$. The reservoir is made of non-linear units (tanh) and the output M is fed back to the reservoir at each iteration. For *n-gate p-value task*, we use n triggers, p values and n outputs. For a generic approach, we use notation u for the input, x for the reservoir and y for the output, independently of the task.

and W_{fb} are sampled randomly and uniformly between -1 and $+1$ and multiplied by the *input scaling factor* and the *feedback scaling factor* respectively. Matrix W is sampled randomly between -1 and $+1$ and we ensure the matrix enforces the defined *sparsity* (ratio of non null values). The resulting matrix is then divided by its largest absolute eigenvalue and multiplied by the specified *spectral radius*. Matrix W_{out} is initially undefined and is learned through teacher forcing[27] and a linear regression:

$$W_{out} = YX^T(XX^T)^{-1} \tag{10}$$

where Y and X corresponds to the respective concatenation of all the desired outputs and reservoir states during a run, each row corresponding to a time step.

3 Results

3.1 The reduced model

The reduced model displays a quasi-perfect performance (RMSE = 2e-6 with $a = 10$ and $b = 10^{-3}$) as shown on figure 5 and the three neurons $X1$,$X2$ and $X3$ behave as expected. $X1$ is strongly correlated with V (figure 5B), $X2$ is strongly correlated with V and saturates in the presence of a tick in T (figure 5C) and $X3$ is strongly correlated with M and

saturates in the presence of a tick in T (figure 5D). This reduced model is actually a very good approximation of a line attractor (i.e. a line of points with very slow dynamics) even though we can prove that, due to the tanh non-linearity, in the absence of inputs, the model will converge to a null state (possibly after a very long time), independently of parameters a and b and the initial state. Nonetheless, figure 5 clearly shows that information can be maintained provided b is small enough. There is a drift, but this drift is so small that it can be considered negligible relatively to the system time constants: these slow points can be considered as a *line* or *segment attractor*[28,29]. As explained by Seung[30], *the reader should be cautioned that the term "continuous attractor" is an idealization and should not be taken too literally. In real networks, a continuous attractor is only approximated by a manifold in state space along which drift is very slow.*

Nevertheless, it is worth mentioning that in order to have a true line attractor, one can replace the tanh activity function with a linear function saturated to 1 and -1[28].

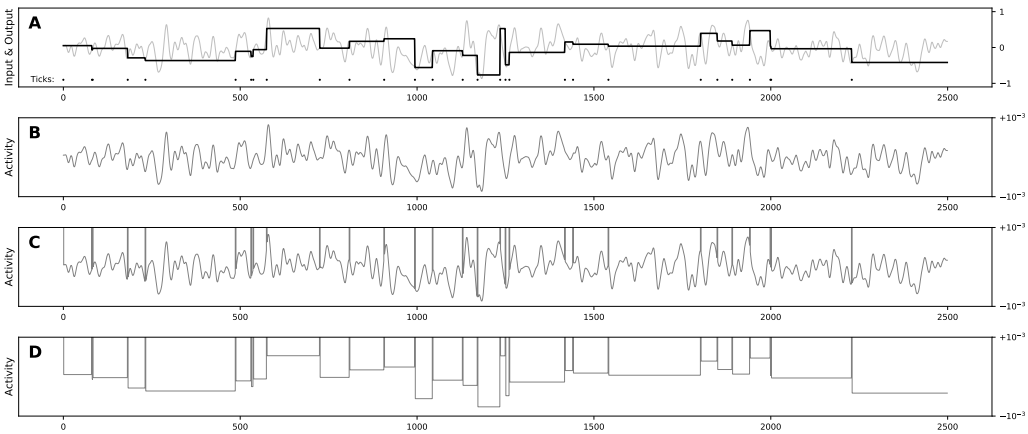

Figure 5. Performance of the reduced model on the 1-gate task. A The light gray line is the input signal and the thick black one is the output of the model. Black dots at the bottom represents the trigger signal (when to memorize a new value). **B, C, D** Respective activity of X_1, X_2 and X_3 units.

3.2 The reservoir model

Unless specified otherwise, all the reservoir models were parameterized using values given in table 1. These values were chosen to be simple and do not really impact the performance of the model as it will be explained in the Robustness section. All simulations and figures were produced using the Python scientific stack, namely, SciPy[31], Matplotlib[32] and NumPy[33]. Sources are available at github.com/rougier/ESN-WM.

Parameter	Value
Spectral radius	0.1
Sparsity	0.5
Leak	1.0 (no leak)
Input scaling	1.0
Feedback scaling	1.0
Number of units	1000
Noise	0.0001
Training timesteps	25,000
Testing timesteps	2,500
Trigger probability	0.01

Table 1. Default parameters Unless specified otherwise, these are the parameters used in all the simulations.

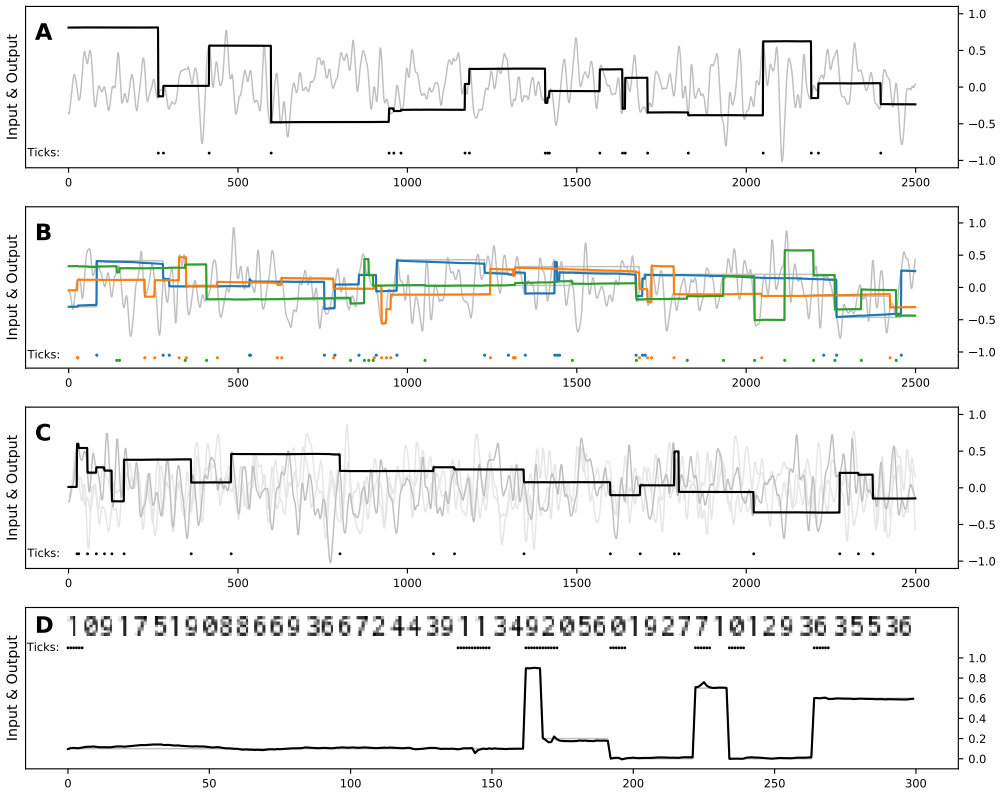

Figure 6. Performance of the reservoir model on working memory tasks. The light gray line is the input signal and the thick black (or colored) one is the output of the model. Dots at the bottom represents the trigger signal (when to memorize a new value). For the digit task, the input containing the value to maintain is shown as a picture instead. **A** *1-value 1-gate scalar task* **B** *1-value 3-gate scalar task* **C** *3-value 1-gate scalar task* **D** *1-value 1-gate digit task*

Results for the reservoir model show a very good generalization performance with a precision of the order of 10^-3 for the level of noise considered (10^-4). Better precision can be obtained for lower noise levels, as show in figure 7. Surprisingly, this generalization property stands with as few as only four random training values where we can achieve a 10^-3 level of precision.

1-value 1-gate scalar task — The model has been trained using parameters given in table 1. The V signal is made of 25,000 random values sampled from a pseudo-random uniform distribution between -1 and +1. The T signal is built from 25,000 random binary values with probability 0.01 of having $T = 1$, and $T = 0$ otherwise. During training, each of the input is presented to the model and the output is forced with the last triggered input. All the input (u) and internal (x) states are collected and the matrix W_{out} is computed according to equation (10). The model has been then tested using a V signal made of 2500 random values sampled from a pseudo-random uniform distribution between -1 and +1. For readability of the figure (see figure 6A), this signal has been smoothed using a fixed-size Hanning window filter. The corresponding T signal has been generated following the same procedure as during the training stage. Figure 6A displays an illustrative test run of the model (for a more thorough analysis of the performance, see the Robustness section) where the error in the output is always kept below 10^{-2} and the RMSE is about $3 * 10^{-3}$.

1-value 3-gate scalar task — We trained the model on the *1-value 3-gate task* using the same protocol as for the 1-value 1-gate task, using a single value input, three input triggers and three corresponding outputs. Since there are now three feedbacks, we divided respective feedback scaling by 3. Figure 6B shows that maintaining information simultaneously impacts the performance of the model (illustrative test run). There is no catastrophic effect but performances are clearly degraded when compared to the *1-value 1-gate task*. The RMSE on this test run increased by one order of magnitude and is about $2 * 10^{-2}$. Nevertheless, in the majority of the cases we tested, the error does stay below 10^{-2}. However in a few cases, one memory (and not necessary all) degrades faster and gains one order of magnitude.

3-value 1-gate scalar task — We used the same protocol as for the *1-value 1-gate scalar task* but there are now two additional inputs not related to the task and that can be considered as noise. Adding such irrelevant inputs had no effect on the solving of the task as illustrated on figure 6C that shows an illustrative test run of the model. The error in the output is also always kept below 10^{-2} and the RMSE is about the same ($3 * 10^{-3}$). This is an interesting result, because it means the network is not only able to deal with "background noise" at 10^{-5}, but it is also able to deal with noise that have the same amplitude as the input. This is an important property to be considered for the modelling of the prefrontal cortex: being an integrative area, the PFC is typically dealing with multimodal information, many of which being not relevant for the working memory task at hand.

1-value 1-gate digit task — The model has been trained using 25,000 random integer values between 0 and 9 sampled from an uniform distribution. Each of these values is drawn one after the other onto an image, each character covering 6 columns of the image. The input V consists then of the raw of this image. The T signal is sampled similarly than in the *1-value 1-gate task* and then expanded 6 times to fit the transformation of the value to the picture of the values. Which means that trigger lasts 6 time steps. Interestingly, as we show on figure 6D, even if the value to maintain is not explicit anymore, it can

still be extracted and maintained. On the test run we show the RMSE is about $4 * 10^{-2}$. It is to be noted that the recognition of a digit is not straightforward and may require a few timesteps before the digit is actually identified. However when the good value is maintained it seems to last, the absolute error stays below 0.05, which is the threshold from which we can distinguish between two values. The reservoir parameters that we found are robust enough to enable not only a pure memory task (i.e. gating), but also a discrimination task (i.e. digit recognition).

Dambre et al.[34] demonstrated the existence of a universal trade-off between the non-linearity of the computation and the short-term memory in the information processing capacity of any dynamical systems, including echo state networks. In other words, the hyperparameters used to generate an optimal reservoir for solving a given memory task would not be optimal for a non-linear computation task. Here we see that even if the reservoir is made to memorised a particular value, it is still able to do a non-linear task such as discriminating stream of digits. Pascanu and Jaeger[35] initiated the concept of such working memory (WM) units for reservoirs. They processed streams of characters using six binary WM-units to record the deepness of curly brackets that appeared. We made such reservoir-WM-units coupling more general: from binary to continuous values. Instead of relying on a collection of N binary WM-units to encode N values, or on the maximum encoding of 2^N values, we have shown that a reservoir can use only one WM-unit to encode a continuous value with a good precision.

3.3 Robustness

We analyzed the robustness of the model first by measuring its sensitivity for each of the hyper-parameters (Figure 7A-F), namely: input scaling, feedback scaling, spectral radius, sparsity of the reservoir weight matrix, leak term (α) in equation (9) and number of units in the reservoir. We also we measured its sensitivity to task hyper-parameters (Figure 7G-L): the noise level (ξ), number of discrete values used during training (when there is a trigger, V is sampled uniformly in uniformly sampled between -1 and 1 discrete values), the temporal delay between successive gating signals in training (T is built sampling its interval between triggers uniformly between 0 and a bound), the bound used to sample the input value in training (V is uniformly sampled between -b and b instead, where b is the bound), the total number of input gates (with a corresponding number of outputs), the number of input values. For most hyper-parameter, we set a minimum and maximum value and pick 20 values logarithmically spread inside this range. For each task and model hyperparater we ran 20 simulation instances for 25,000 timesteps and record the mean performance using 2,500 values. Results are shown on figure 7.

First, we can see a non-sensitivity to the sparsity (i.e. minor differences in performances when these parameters vary). Similarly we can see a non-sensitivity to leak term, input and feedback scaling, as long as it there are not too small. It is to be noted that input and feedback scaling should also not be too big. As expected, the performance increase with the number of neurons. Surprisingly, we can note that the performance decrease with the spectral radius. In fact, in supplementary figure 14 we analysed the behavior of the reservoir model with various spectral radius. We show that even with a bigger spectral radius, the reservoir keeps maintaining something relevant but less precise (the segment attractor is slowly degenerating). Globally, the reservoir model is very robust against model hyper-parameters changes as long as it is trained in this condition.
Concerning the task hyper-parameters, the trigger range do not have an impact. Performance naturally decrease with the noise level, the number of input gates and the number of input values. We can note that the number of discrete values used during training impacts the performance in a very particular way (see Figure 7): as soon as we have at least 4 to 7 training values, the performance does not improve as much. Interest-

ingly, the performance reached a plateau before to use the bound 1 to sample the input value.

3.4 Dynamics

A segment attractor On figure 8 we show how the model evolves after having been trained for the *1-value 1-gate task* using different starting positions and receiving no input. This results in the formation of a segment attractor even tough the model was only trained to gate and memorize continuous value. If we compute the principal component analysis (PCA) on the reservoir states and look at the principal components (PCs) in the absence of inputs, we can observe (supplementary figure 13) that all the reservoir states are organized on along a straight line on the first component (the one which explains the most the variance) and each point of this line can be associated with a corresponding memory value. Interestingly enough, there are points on this line that corresponds to value outside the $[-1; 1]$ range, i.e. values for which the model has not been trained for. However, these points are not stable and an any dynamics starting from these points converge towards the points associated to the values 1 or -1 (see figure 8).

V-like and M-like neurons Similarly as in the reduced model, when receiving input, the inner dynamic of the reservoir model is a combination of both sustained and highly variable activities. More precisely, in the *1-value 1-gate task* we notice two types of neurons that are similar to neurons $X1$ and $X3$ in the reduced model: (1) neurons which solely follows the input V (i.e. V-like neurons), and (2) neurons that mostly follow the output M (i.e. M-like neurons), respectively. We also notice that M-like neurons average activity is linearly linked with the M value and fluctuate around this mean activity according to the input V. On figure 9 we show M-like neurons for the different tasks. These neurons were found by taking the most correlated neurons with the output M[2]. From figure 9**A** to **D** we see that the M-like neurons link with M output goes weaker and weaker: the average sustained activity goes from nearly flat to highly perturbed.
This "degradation" of maintained activity is explained by the change in the distribution of correlations of the whole reservoir population with M output: in figure 9 **(right)** we see that the correlation with M output is quickly shrinking from **A** to **D**. For the *1-value 1-gate task* (9A) mostly all neurons are maintained. However, when more values have to be maintained (figure 9B) or more inputs are received (figure 9C), most of the activities are not maintained anymore. In fact, in the *1-value 1-gate digit task* the neurons do not display a maintained activity anymore (figure 9D). Interestingly, similar behavior (no maintained activity) can be obtained by lowering the feedback scaling (supplementary figure 15) or by enforcing the input weights to be big enough (supplementary figure 17).

Linear decoder To go further in the understanding of the role played by sustained activities, we wanted to know how much of these sustained activities were necessary to decode the output memory M. For the *1-value 1-gate task*, we trained a separate decoder based on a subsample of the reservoir population. We increasingly subsampled neurons based on three conditions: either by choosing the most correlated one first, the least correlated one first or just randomly selecting them. On figure 10 we can see two interesting facts. First, there is no need of the whole reservoir population to decode well enough the memory M: taking a hundred neurons among one thousand is sufficient. Second, if we take enough neurons, there is no advantage in taking the most correlated one first, random ones are enough. Surprisingly, it seems better to rely on randomly selected units than most correlated ones. This suggests that distributed information contained more information than only using most correlated units.

[2] Because a trigger input has a substantial influence on the reservoir states, we made the categorization by ignoring the time steps when there is a trigger.

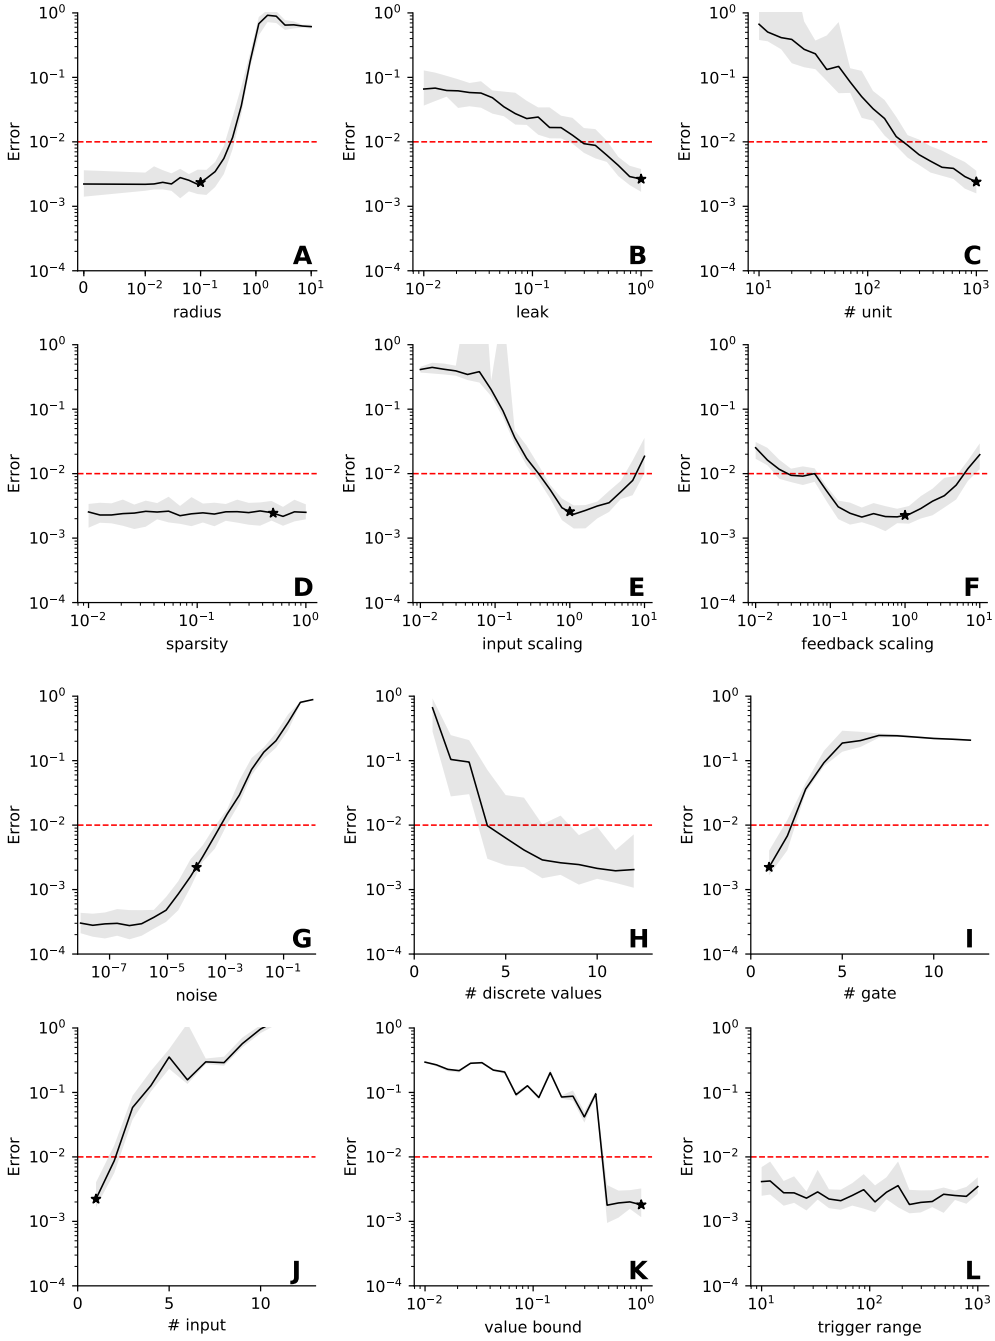

Figure 7. Robustness of the model to hyper-parameters The performance of the model has been measured when a single hyper-parameter varies while all others are kept constant and equal to the value given in table 1. For each plot, we ran 20 iterations of the model using different random seeds for initialization. Most of the time (**A-GK-L**), the varying parameter has been split into 20 log-uniformly spread values between the minimum and maximum values. **A-F.** Model hyper-parameters. **G-L.** Task hyper-parameters. **A.** Spectral radius ranging from 0.01 to 10. with an additional 0. **B.** Leak term ranging from 0.01 to 1 (no leak). **C.** Number of units in the reservoir (from 1 to 1000). **D.** Sparsity level ranging from 0.01 to 1.0 (fully connected). **E.** Input scaling ranging from 0.01 to 10. **F.** Feedback scaling ranging from 0.01 to 10. **G.** Noise level inside the reservoir, ranging from 10^{-8} to 1. **H.** Number of discrete values used to train the model, ranging from 1 to 12. **I.** Number of input gates and output channels, ranging from 1 to 12. (1-value n-gate scalar task, RMSE has been divided by the square root of the number of gates) **J.** Number of input values, ranging from 1 to 12. (n-value 1-gate scalar task) **K.** Bound for the input value during training, ranging from 0.01 to 1. **L.** Maximal interval (number of steps) between consecutive ticks during training, ranging from 10 to 1000. ⋆. Performance for default parameters (Table 1).

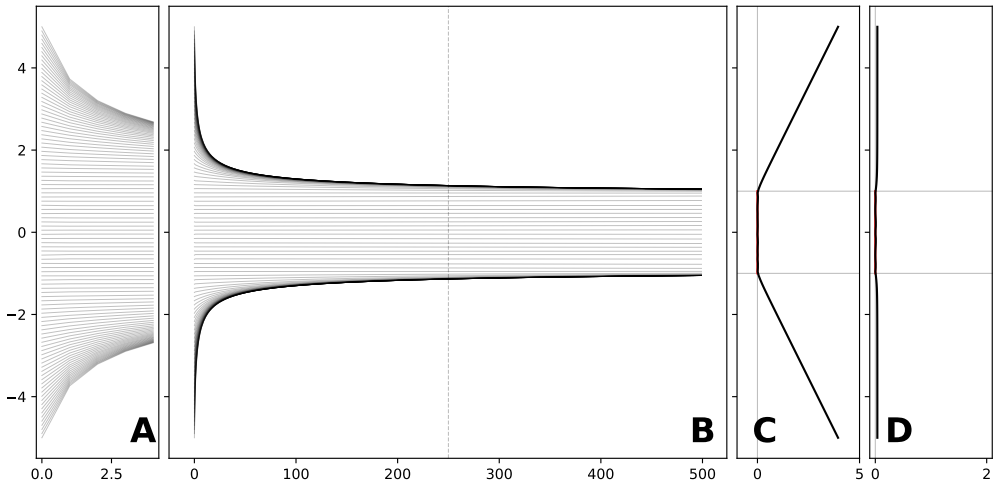

Figure 8. An approximate segment attractor The same model trained on the 1-value 1-gate was tested for 500 iterations without inputs, starting with an initial trigger along with values linearly distributed between -5 and +5. **A** and **B** Output trajectories. Each line corresponds to one trajectory (output value) of the model. **A** is a zoom of **B** on the first few time steps. **C** Absolute difference in the output between initial time and and final time. **D** Maximal absolute difference in the neurons between intermediate (dashed line) and final time.

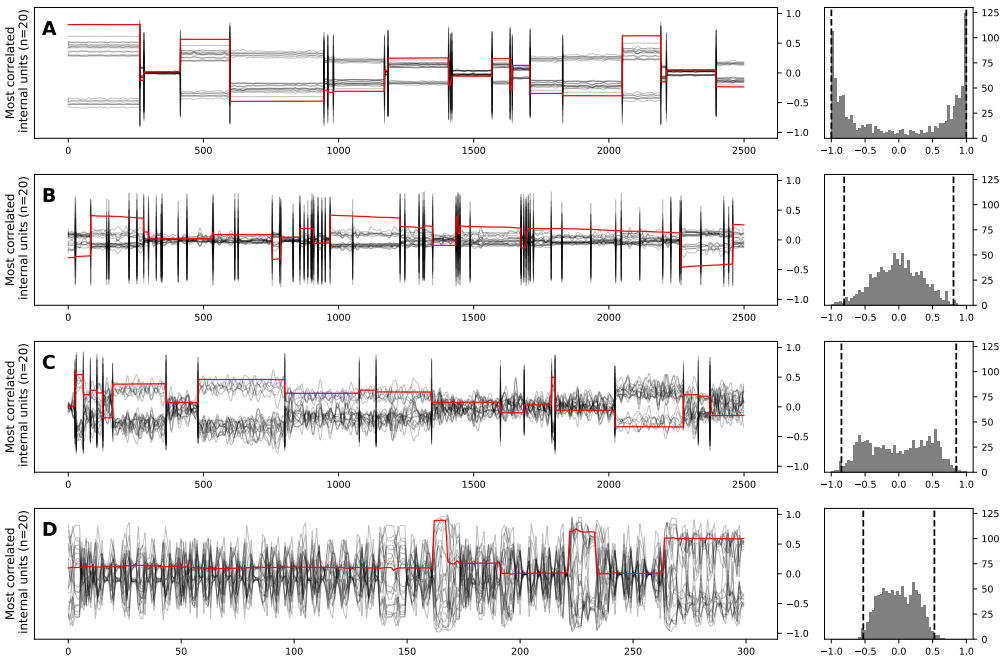

Figure 9. Most correlated reservoir units displaying various degrees of maintained activity. Left: in black the activities of the 20 neurons the most correlated with the output, in red the output of the model. Right: histogram of the correlation between the neurons and the output. The correlation of the 20th and its opposite is shown in dashed. **A** *1-value 1-gate scalar task* **B** *1-value 3-gate scalar task* **C** *3-value 1-gate scalar task* **D** *1-value 1-gate digit task*. Most correlated reservoir units do not necessarily show clear sustained activity: we can see a degradation of sustained activities from **A** to **D** according to the different tasks. Thus, maintained activity is not compulsory to perform working memory tasks.

This assumes a complementary perspective when comparing it to Rigotti et al.[10] decoding of neural population recorded in monkeys: when task-related neurons are not kept for decoding, information (but fewer) can still be decoded.

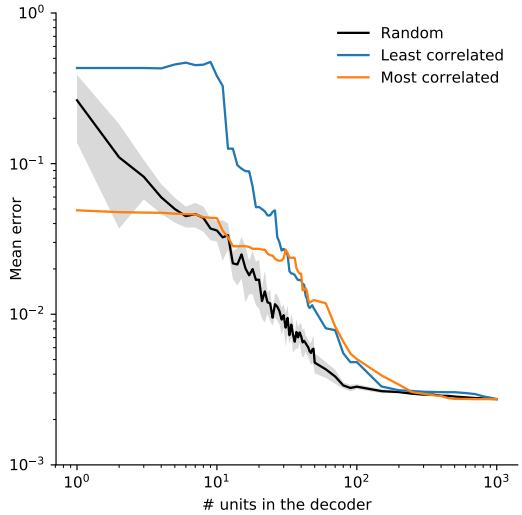

Figure 10. Performance of variable size linear decoders. After training, we ran the model on the training set and recorded all the internal states. We selected a subset (from 1 to 1000 units) of these internal states and find the associated best linear decoder on the training data set. Finally, we measure the performance of these decoders on the training sets. Units composing a subset have been sample either using the least or the most correlated units (with the output) or simply randomly. Performance for subsets of size 1000 is equal to the performance of the original model.

4 Discussion

In computational neuroscience, the reservoir computing paradigm (RC)[24,36,37], originally proposed independently by Dominey[38], and Buonomano and Merzenich[39], is often used as a model of canonical microcircuits[36,40,41]. It is composed of a random recurrent neural network (i.e. a reservoir) from which readout units (i.e. outputs) are trained to linearly extract information from the high-dimensional non-linear dynamics of the reservoir. Several authors have taken advantage of this paradigm to model cortical areas such as PFC[42,43,44,45] because most of the connections are not trained, especially the recurrrent ones. Another reason to use reservoir for PFC modelling is because PFC also hosts high-dimensional non-linear dynamics [10]. RCs offers a neuroanatomically plausible view of how cortex-to-basal-ganglia (i.e. cortico-basal) connections could be trained with dopamine: the reservoir plays the role of a cortical area (e.g. trained with unsupervised learning), and the read-out units play the role of basal ganglia input (i.e. striatum).

However, in many dynamical systems, reservoir included, there exists a trade-off between memory capacity and non-linear computation[34][3]. This is why some studies have focused on reservoirs with dedicated readout units acting as working memory (WM) units[40,35,47]. These WM units have feedback connections projecting to the reservoir and are trained to store binary values that are input-dependent. This somehow simplifies the task and enables the reservoir to access and use such long-time dependency information to perform a more complex task, freeing the system from constraining reservoir short-term dynamics. Pascanu and Jaeger[35] used up to six binary WM units to store information in order to solve a nested bracketing levels task. Using a principal component analysis, they showed that these binary WM units constrain the reservoir in lower dimensional "attractors". Additionally, Hoerzer, Legenstein, and Maass[40] showed that analog WM units (encoding binary information) also drive the reservoir into a lower dimensional space (i.e. 99% of the variability of the reservoir activities are explained by fewer principal components). More recently, Strock, Rougier, and Hinaut[48], Beer and Barak[49] tried to use such WM units in order to store analog values (as opposed to binary

[3]For reservoirs this trade-off depends on the hyper-parameters (HP) chosen: some HP sets will give more memory, others more computations capacity[46].

ones) in order to build a *line attractor*[28,29]. In particular, Beer and Barak[49] explored how a line attractor can be built online, by comparing FORCE[50] and LMS algorithms, using a WM unit to maintain a continuous value in *the absence* of input perturbations.

In that context, the minimal model of three neurons we have proposed helps to understand the mechanisms that allows a reservoir model to gate and maintain scalar values, in *the presence* of input perturbations, instead of just binary values. As explained previously, this minimal model exploits the non-linearity and the asymptotic behavior of the three tanh units and mimics a select operator between the input signal and the output. In the case of the reservoir model, there is no precise architecture or crafted weights, but this is compensated by the size of the population inside the reservoir along with the training of the output weights. More precisely, we have shown that virtually any population of randomly connected units is able to maintain an analog value at any time and for an arbitrary delay. Taking advantage of the non-linearity of the neuron transfer function, we have shown how such population can learn a set of weights during the training phase, using only a few representative values. Given the random nature of the architecture and the large set of hyper-parameters, for which the precision of the output remains acceptable, this suggests that this property could be a structural property of any population that could be acquired through learning. To achieve such property we mainly used offline training in our analyses (for efficiency reasons), but we have shown that it also works with online FORCE learning (see supplementary materials and figure 12).

One important feature of the model is the fact that it is actually an open system and as such, it is continuously under the direct influence of external activities. More precisely, the model is able to retain an information when the gate is closed, but this *closed gate* corresponds to a functionnal state rather than a physical state where input signals would have been blocked and information continues to enter the system. This is illustrated quite clearly when one looks at the internal activities inside the reservoir: a lot of neurons are directly (and only) correlated with the input signals. This draws consequences on the analysis of the dynamics of the population because this population is partly driven by the task (working memory) and partly driven by uncorrelated external activities. If we go back to biology, this makes perfect sense: a population of neurons is never isolated from the rest of the brain. When studying population electrophysiological recordings, we have to keep in mind that such activities can be fully uncorrelated with the task we are observing. This might be one of the reasons for the variability of hypotheses about working memory encoding mechanisms.

5 References

1. P. S. Goldman-Rakic. **Circuitry of Primate Prefrontal Cortex and Regulation of Behavior by Representational Memory**. In: Comprehensive Physiology (1987).
2. A. Bechara, H. Damasio, D. Tranel, and S. W. Anderson. **Dissociation Of Working Memory from Decision Making within the Human Prefrontal Cortex**. In: The Journal of Neuroscience 18.1 (Jan. 1998), pp. 428–437.
3. E. K. Miller and J. D. Cohen. **An Integrative Theory of Prefrontal Cortex Function**. In: Annual Review of Neuroscience 24.1 (Mar. 2001), pp. 167–202.
4. J. M. Fuster. **The Prefrontal Cortex—An Update**. In: Neuron 30.2 (May 2001), pp. 319–333.
5. R. C. O'Reilly and M. J. Frank. **Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and Basal Ganglia**. In: Neural Computation 18.2 (Feb. 2006), pp. 283–328.
6. S. Funahashi. **Working Memory in the Prefrontal Cortex**. In: Brain Sciences 7.12 (Apr. 2017), p. 49.
7. C. Constantinidis, S. Funahashi, D. Lee, J. D. Murray, X.-L. Qi, M. Wang, and A. F. Arnsten. **Persistent Spiking Activity Underlies Working Memory**. In: The Journal of Neuroscience 38.32 (Aug. 2018), pp. 7020–7028.
8. R. Romo, C. D. Brody, A. Hernández, and L. Lemus. In: Nature 399.6735 (June 1999), pp. 470–473.
9. C. K. Machens, R. Romo, and C. D. Brody. **Functional, But Not Anatomical, Separation of "What" and "When" in Prefrontal Cortex**. In: Journal of Neuroscience 30.1 (Jan. 2010), pp. 350–360.

10. M. Rigotti, O. Barak, M. R. Warden, X.-J. Wang, N. D. Daw, E. K. Miller, and S. Fusi. **The importance of mixed selectivity in complex cognitive tasks**. In: Nature 497.7451 (May 2013), pp. 585–590.
11. D. Durstewitz, J. K. Seamans, and T. J. Sejnowski. **Neurocomputational models of working memory**. In: Nature Neuroscience 3 (Nov. 2000), pp. 1184–1191.
12. A. Compte. **Computational and in vitro studies of persistent activity: Edging towards cellular and synaptic mechanisms of working memory**. In: Neuroscience 139.1 (Apr. 2006), pp. 135–151.
13. O. Barak and M. Tsodyks. **Working models of working memory**. In: Current Opinion in Neurobiology 25 (Apr. 2014), pp. 20–24.
14. S.-i. Amari. **Dynamics of pattern formation in lateral-inhibition type neural fields**. In: Biological Cybernetics 27.2 (1977), pp. 77–87.
15. A. Compte. **Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model**. In: Cerebral Cortex 10.9 (Sept. 2000), pp. 910–923.
16. F. Edin, T. Klingberg, P. Johansson, F. McNab, J. Tegner, and A. Compte. **Mechanism for top-down control of working memory capacity**. In: Proceedings of the National Academy of Sciences 106.16 (Apr. 2009), pp. 6802–6807.
17. Z. Wei, X.-J. Wang, and D.-H. Wang. **From Distributed Resources to Limited Slots in Multiple-Item Working Memory: A Spiking Network Model with Normalization**. In: Journal of Neuroscience 32.33 (Aug. 2012), pp. 11228–11240.
18. F. Bouchacourt and T. J. Buschman. **A Flexible Model of Working Memory**. In: (Sept. 2018).
19. K. Zhang. **Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory**. In: The Journal of Neuroscience 16.6 (Mar. 1996), pp. 2112–2126.
20. N. J. Gandhi and H. A. Katnani. **Motor Functions of the Superior Colliculus**. In: Annual Review of Neuroscience 34.1 (July 2011), pp. 205–231.
21. A. A. Koulakov, S. Raghavachari, A. Kepecs, and J. E. Lisman. **Model for a robust neural integrator**. In: Nature Neuroscience 5.8 (July 2002), pp. 775–782.
22. D. Zipser, B. Kehoe, G. Littlewort, and J. Fuster. **A spiking network model of short-term active memory**. In: The Journal of Neuroscience 13.8 (Aug. 1993), pp. 3406–3420.
23. S. Lim and M. S. Goldman. **Balanced cortical microcircuitry for maintaining information in working memory**. In: Nature Neuroscience 16.9 (Aug. 2013), pp. 1306–1314.
24. H. Jaeger. **The "echo state" approach to analysing and training recurrent neural networks**. Tech. rep. 148. Bonn, Germany: German National Research Center for Information Technology GMD, Jan. 2001.
25. N. Schaetti, M. Salomon, and R. Couturier. **Echo State Networks-Based Reservoir Computing for MNIST Handwritten Digits Recognition**. In: 2016 IEEE Intl Conference on Computational Science and Engineering (CSE) and IEEE Intl Conference on Embedded and Ubiquitous Computing (EUC) and 15th Intl Symposium on Distributed Computing and Applications for Business Engineering (DCABES). IEEE, Aug. 2016.
26. K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. **Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation**. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 1724–1734.
27. M. Lukoševičius. **A Practical Guide to Applying Echo State Networks**. In: Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, pp. 659–686.
28. H. S. Seung. **How the brain keeps the eyes still**. In: Proceedings of the National Academy of Sciences 93.23 (Nov. 1996), pp. 13339–13344.
29. D. Sussillo and O. Barak. **Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks**. In: Neural Computation 25.3 (Mar. 2013), pp. 626–649.
30. H. S. Seung. **Learning Continuous Attractors in Recurrent Networks**. In: Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10. NIPS '97. Denver, Colorado, USA: MIT Press, 1998, pp. 654–660.
31. E. Jones, T. Oliphant, and P. Peterson. SciPy: Open source scientific tools for Python. 2001.
32. J. D. Hunter. **Matplotlib: A 2D graphics environment**. In: Computing In Science & Engineering 9.3 (2007), pp. 90–95.
33. S. van der Walt, S. C. Colbert, and G. Varoquaux. **The NumPy Array: A Structure for Efficient Numerical Computation**. In: Computing in Science & Engineering 13.2 (Mar. 2011), pp. 22–30.
34. J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar. **Information Processing Capacity of Dynamical Systems**. In: Scientific Reports 2.1 (July 2012).
35. R. Pascanu and H. Jaeger. **A neurodynamical model for working memory**. In: Neural Networks 24.2 (Mar. 2011), pp. 199–207.

36. W. Maass, T. Natschläger, and H. Markram. **Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations**. In: Neural Computation 14.11 (Nov. 2002), pp. 2531–2560.
37. D. Verstraeten, B. Schrauwen, M. D'Haene, and D. Stroobandt. **An experimental unification of reservoir computing methods**. In: Neural Networks 20.3 (Apr. 2007), pp. 391–403.
38. P. F. Dominey. **Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning**. In: Biological Cybernetics 73.3 (Aug. 1995), pp. 265–274.
39. D. V. Buonomano and M. M. Merzenich. **Temporal information transformed into a spatial code by a neural network with realistic properties**. In: Science 267.5200 (1995), pp. 1028–1030.
40. G. M. Hoerzer, R. Legenstein, and W. Maass. **Emergence of complex computational structures from chaotic neural networks through reward-modulated Hebbian learning**. In: Cerebral cortex 24.3 (2012), pp. 677–690.
41. D. Sussillo. **Neural circuits as computational dynamical systems**. In: Current opinion in neurobiology 25 (2014), pp. 156–163.
42. X. Hinaut and P. F. Dominey. **Real-Time Parallel Processing of Grammatical Structure in the Fronto-Striatal System: A Recurrent Network Simulation Study Using Reservoir Computing**. In: PLoS ONE 8.2 (Feb. 2013). Ed. by T. Boraud, e52946.
43. F. Mannella and G. Baldassarre. **Selection of cortical dynamics for motor behaviour by the basal ganglia**. In: Biological Cybernetics 109.6 (Nov. 2015), pp. 575–595.
44. X. Hinaut, F. Lance, C. Droin, M. Petit, G. Pointeau, and P. F. Dominey. **Corticostriatal response selection in sentence production: Insights from neural network simulation with reservoir computing**. In: Brain and Language 150 (Nov. 2015), pp. 54–68.
45. P. Enel, E. Procyk, R. Quilodran, and P. F. Dominey. **Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex**. In: PLOS Computational Biology 12.6 (June 2016). Ed. by J. X. O'Reilly, e1004967.
46. R. Legenstein and W. Maass. **Edge of chaos and prediction of computational performance for neural circuit models**. In: Neural Networks 20.3 (Apr. 2007), pp. 323–334.
47. T. Nachstedt and C. Tetzlaff. **Working Memory Requires a Combination of Transient and Attractor-Dominated Dynamics to Process Unreliably Timed Inputs**. In: Scientific Reports 7.1 (May 2017).
48. A. Strock, N. P. Rougier, and X. Hinaut. **A Simple Reservoir Model of Working Memory with Real Values**. In: 2018 International Joint Conference on Neural Networks (IJCNN). Sao Paulo, Brazil: IEEE, July 2018.
49. C. Beer and O. Barak. **One step back, two steps forward: interference and learning in recurrent neural networks**. In: arXiv preprint arXiv:1805.09603v2 (2019).
50. D. Sussillo and L. Abbott. **Generating Coherent Patterns of Activity from Chaotic Neural Networks**. In: Neuron 63.4 (Aug. 2009), pp. 544–557.

6 Supplementary

6.1 Reduced model

The effect of parameters a and b are of distinct nature as illustrated on figure 11. b controls the time the memory takes to fade to zero whereas a controls how the model will mix the previous output and the new desired one. Interestingly, the time the information takes to fade out to zero is not directly linked to a time constant of neurons. Moreover, the mixing factor can actually be computed for b small enough and is equal to $\tanh(a)^2$. Consequently, by feeding T with a (instead of 1) and fixing the input weights coming from T to 1 (instead of a), one can obtain a model which mimic the update mechanism of a Gated Recurrent Unit (GRU) but with regular neurons.

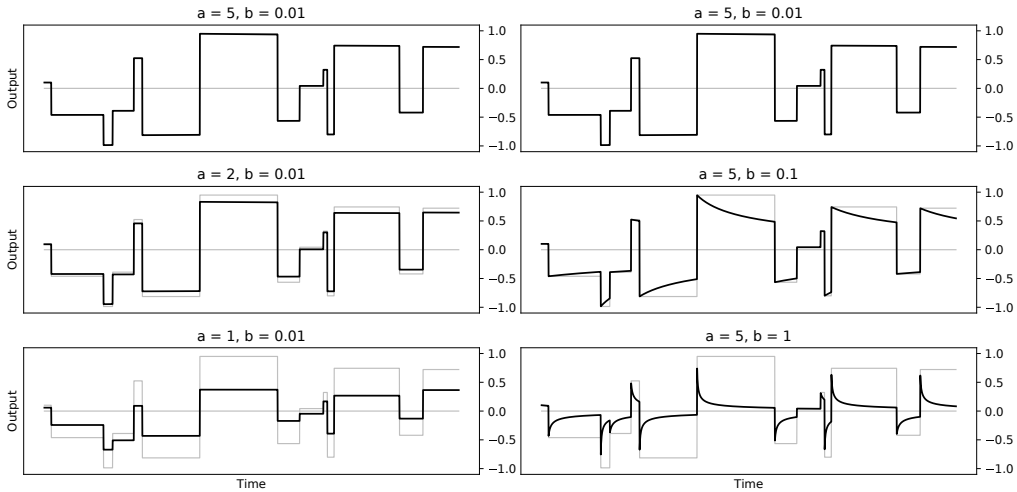

Figure 11. Influence of parameters a and b on the reduced model. The parameter b controls the amount of feedback fed to the model. If b is too high ($b = 0.1$, $b = 1.0$), the memory quickly fade to zero after a few timesteps. The parameter a controls the gate behavior and the ratio of the memorized and the new value that will constitute the new memorized value. If a is too low ($a = 1$, $a = 2$), there is a systematic undershoot.

6.2 Online learning

We tested online learning by using FORCE[50] described by the following update equations:

$$W_{out}[n] = W_{out}[n-1] - e_-[n]P[n]x[n] \tag{11}$$

$$P[n] = P[n-1] - \frac{P[n-1]x[n]x^T[n]P[n-1]}{1 + x^T[n]P[n]x[n]} \tag{12}$$

where $e_-[n]$ represents the error made at time n if there were no update of W_{out} and $P[n]$ is a squared matrix computing an online inverse of $XX^T - \alpha I$, α a regularization term and P initialized to $P[0] = \frac{I}{\alpha}$.

In figure 12, one can see that the behavior obtain with online learning is similar to offline learning.

6.3 More on the segment attractor

On Figure 13 we performed a principal component analysis on the reservoir state ob-

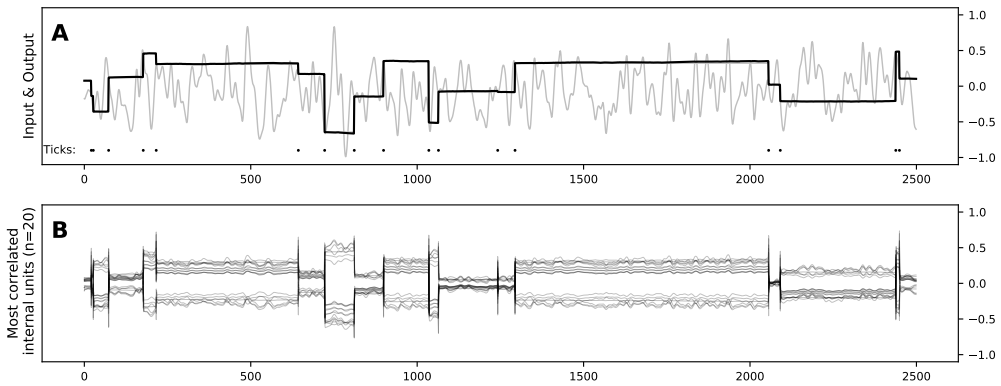

Figure 12. Performance of the model on the 1-gate task with online FORCE learning. Same training/testing protocol than for figure 12. The regularization parameter α has been fixed to 0.0001. **A** The light gray line is the input signal and the thick black one is the output of the model. Black dots at the bottom represents the trigger signal (when to memorize a new value). **B** Activity of 20 units that are the most correlated with the output.

tained in Figure 8. We can note two interesting facts: (1) The activity evolves in a very low dimensional space, the first component explains more than 99% of the variance by itself, and (2) this component is linked to the memory maintained, linearly between -1 and 1 and non-linearly outside. With both combined together we can say that the segment attractor we builded is actually a straight line.

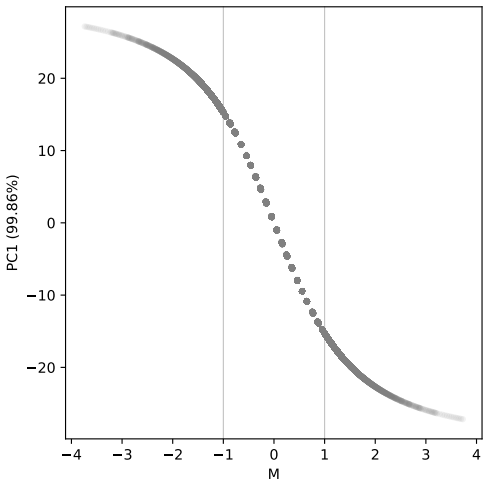

Figure 13. The line attractor. The reservoir state and output analyzed are the one of Figure 8. x-axis: output (memory). y-axis: first principal component of the reservoir state.

6.4 Influence of the spectral radius on the segment attractor

By watching how the segment attractor is evolving against the spectral radius, we can better explain its influence on the solving of the 1-value 1-gate task. The smaller the spectral radius is, the more precise is the approximation of the segment attractor. For a 0.1 spectral radius the approximate segment attractor seems continuous. The difference is quasi-null for starting values in [-1,+1] while for values outside this range, the different correspond to the difference with the nearest bound of the [-1,+1] segment.

For a 0.5 spectral radius we clearly can see that the segment attractor is discretized. In the very first step the outputs concentrates around discretized value between -1 and 1 and is kept constant. In the extreme case for a 1.0 spectral radius it is discretized into two points corresponding to -1 and 1.

6.5 Killing neurons having a sustained activity

On figure 15 we show how the neurons are correlated with the output (M), the value (V) and the trigger (T) for different feedback scaling. Because some irregularities come with triggers, we removed the trigger time steps while computing the correlation with the output and the value. We can note that the profile of neurons seems to change according to the feedback scaling. When it is high (1.0), there is mostly neurons following the output, when it is low (0.1) few neurons continue to follow the output while most of them vary according to the value they receive as input. To better identify the neurons which are following the output we looked at where were located the weights of these neurons on figure 16. We can note that the most correlated neurons with the output are mainly the one with relatively small input weights coming from the value. To go further with this idea we changed the sampling of the input weights to remove small values. In practice the sampling was done uniformly in $[-1,-0.5]\cup[0.5,1.0]$ instead. We show how behave this modified model on figure 17. We can note that the model is still able to perform the task relatively well, just a little worse than before, but there is no more neuron displaying a sustained activity. Interestingly, on figure 16, we also see that the output weights of the neurons the most correlated with the trigger seems to be correlated to their feedback weights.

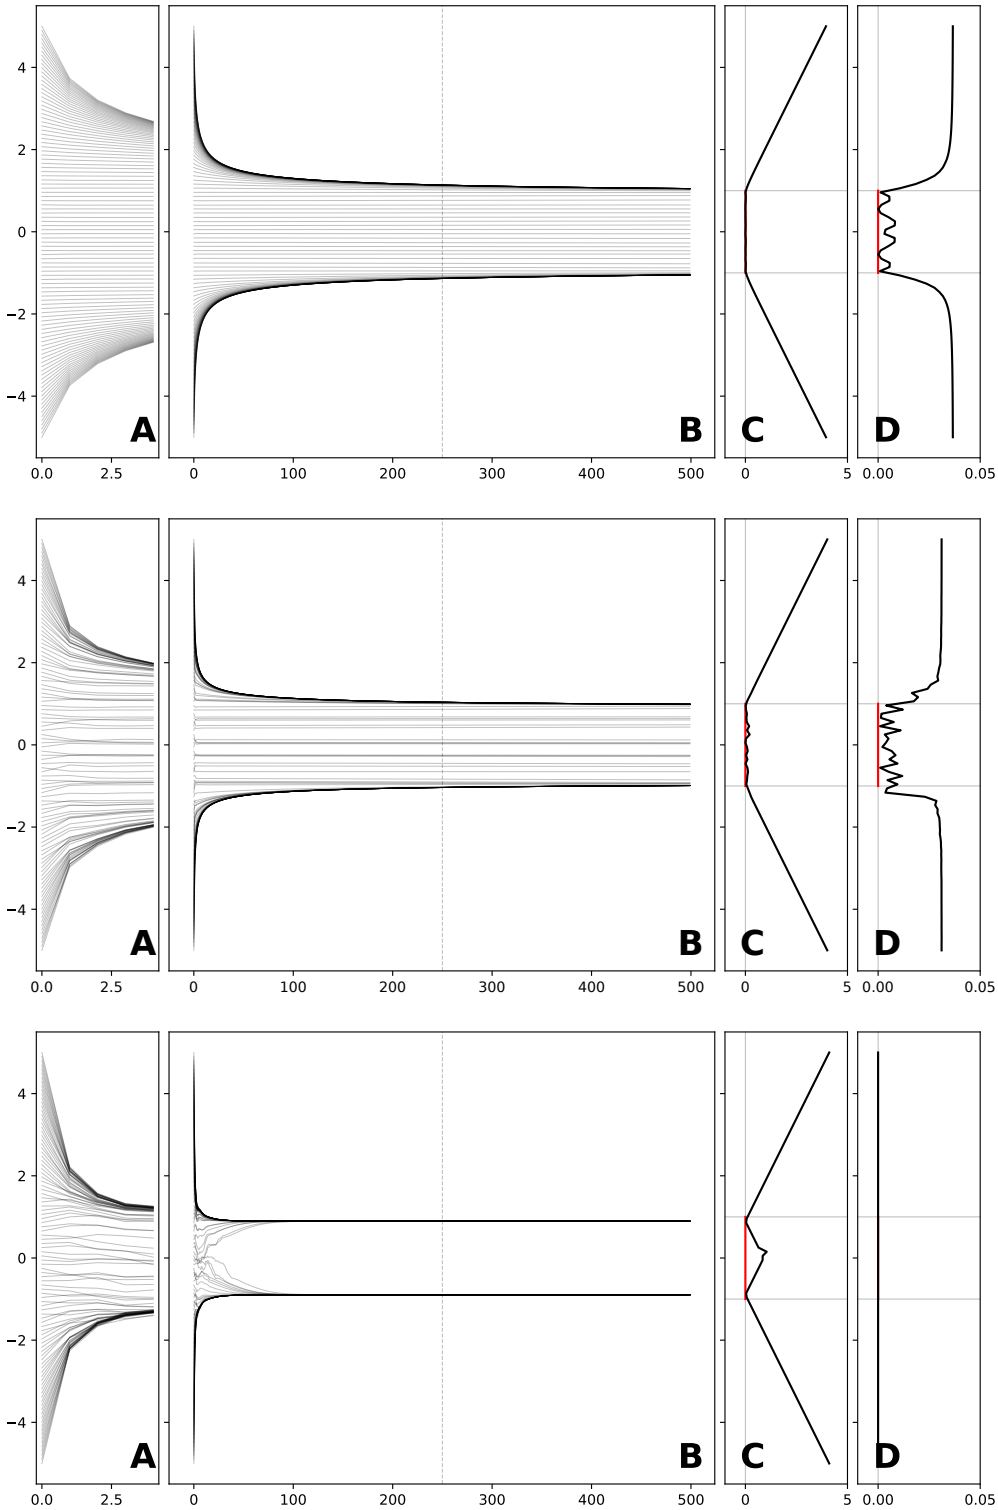

Figure 14. Influence of spectral radius on approximate line attractor The same trained models were tested for 500 iterations without inputs, only starting with an initial trigger along with values linearly distributed between -5 and +5. First line: Spectral radius 0.1. Second line: Spectral radius 0.5. Third line: Spectral radius 1.0. **A** and **B** Output trajectories. Each line corresponds to one trajectory (output value) of the model. **A** is a zoom of **B** on the first few time steps. **C** Asolute difference in the output between initial time and and final time. **D** Maximal absolute difference in the neurons between intermediate (dashed line) and final time.

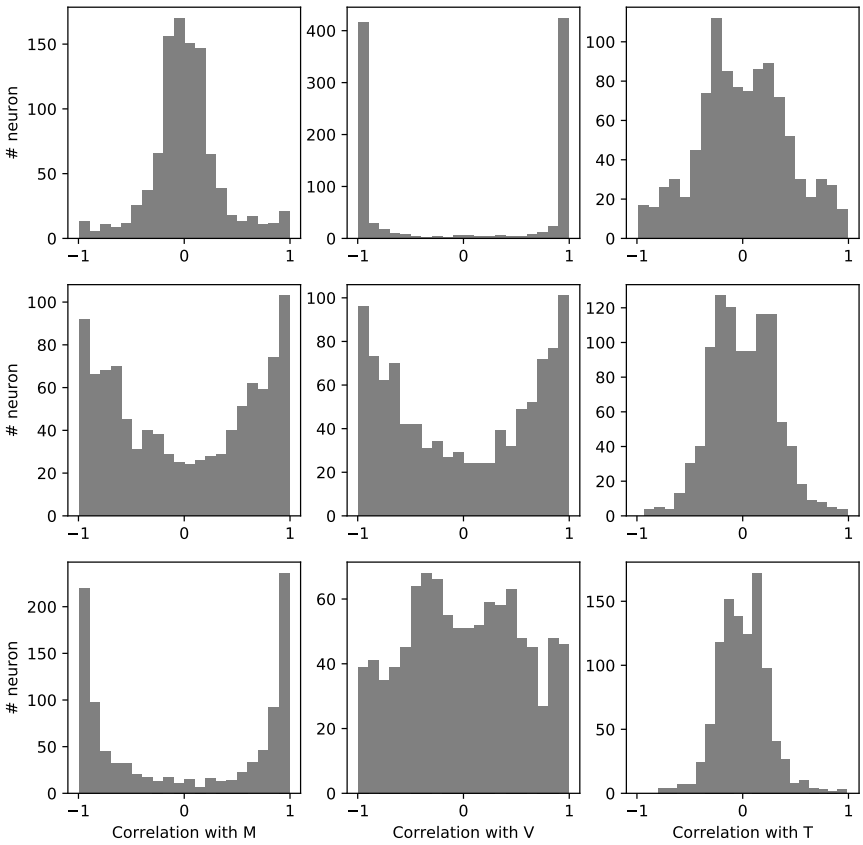

Figure 15. Influence of feedback scaling on maintenance of neurons First line: Feedback scaling 0.1. Second line: Feedback scaling 0.5. Third line: Feedback scaling 1.0. First column: Correlation with M. Second column: Correlation with V. Third column: Correlation with T.

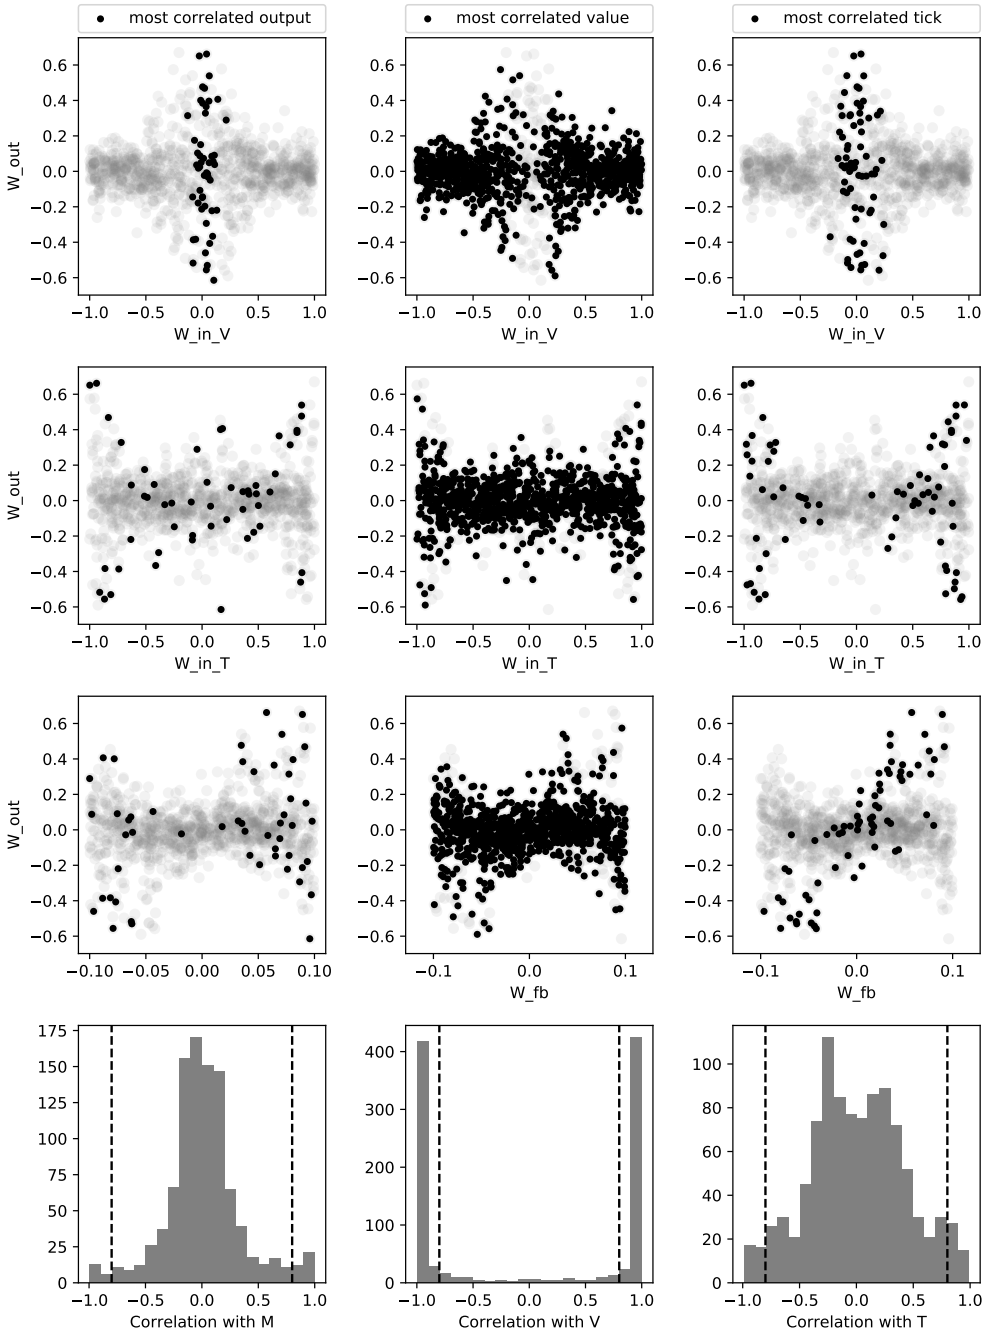

Figure 16. Link between weights and behavior of the neurons For each columns particular neurons are highlighted: the neurons most correlated with either M, V or T. First column: M. Second column: V. Third column: T. The three first lines show the link between the output weights and some other weights. First line: weights coming from V. Second line: weights coming from T. Third line: weights coming from M. The fourth line shows the correlation with the quantity used to highlight. The threshold used to decide which are the most correlated is shown in dotted black.

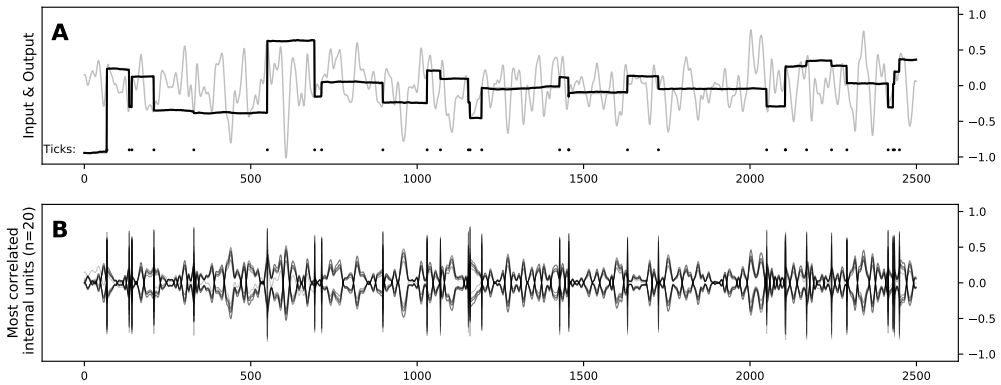

Figure 17. Performance of the model with high input value weights on the 1-gate task. A The light gray line is the input signal and the thick black one is the output of the model. Black dots at the bottom represents the trigger signal (when to memorize a new value). **B** Activity of 20 units that are the most correlated with the output.