Local Homeostatic Regulation of the Spectral Radius of Echo-State Networks

Fabian Schubert* and Claudius Gros

Institute for Theoretical Physics, Goethe University Frankfurt am Main, Frankfurt am Main, Germany

Recurrent cortical networks provide reservoirs of states that are thought to play a crucial role for sequential information processing in the brain. However, classical reservoir computing requires manual adjustments of global network parameters, particularly of the spectral radius of the recurrent synaptic weight matrix. It is hence not clear if the spectral radius is accessible to biological neural networks. Using random matrix theory, we show that the spectral radius is related to local properties of the neuronal dynamics whenever the overall dynamical state is only weakly correlated. This result allows us to introduce two local homeostatic synaptic scaling mechanisms, termed flow control and variance control, that implicitly drive the spectral radius toward the desired value. For both mechanisms the spectral radius is autonomously adapted while the network receives and processes inputs under working conditions. We demonstrate the effectiveness of the two adaptation mechanisms under different external input protocols. Moreover, we evaluated the network performance after adaptation by training the network to perform a time-delayed XOR operation on binary sequences. As our main result, we found that flow control reliably regulates the spectral radius for different types of input statistics. Precise tuning is however negatively affected when interneural correlations are substantial. Furthermore, we found a consistent task performance over a wide range of input strengths/variances. Variance control did however not yield the desired spectral radii with the same precision, being less consistent across different input strengths. Given the effectiveness and remarkably simple mathematical form of flow control, we conclude that self-consistent local control of the spectral radius via an implicit adaptation scheme is an interesting and biological plausible alternative to conventional methods using set point homeostatic feedback controls of neural firing.

Keywords: recurrent networks, homeostasis, synaptic scaling, echo-state networks, reservoir computing, spectral radius

1. INTRODUCTION

Cortical networks are highly recurrent, a property that is considered to be crucial for processing and storing temporal information. For recurrent networks to remain stable and functioning, the neuronal firing activity has to be kept within a certain range by autonomously active homeostatic mechanisms. It is hence important to study homeostatic mechanisms on the level of single neurons, as well as the more theoretic question of characterizing the dynamic state that is to be attained on a global network level. It is common to roughly divide adaptation mechanisms into intrinsic homeostasis, synaptic homeostasis, and metaplasticity.

Synaptic scaling was identified as a mechanism that can postsynaptically regulate neural firing by adjusting synaptic efficacies in a proportional, multiplicative way. This finding has led to numerous studies investigating the role of synaptic scaling in controlling neural network activity (Turrigiano et al., 1998; Turrigiano and Nelson, 2000; Turrigiano, 2008) and in stabilizing other plasticity mechanisms (van Rossum et al., 2000; Stellwagen and Malenka, 2006; Tetzlaff, 2011; Toyoizumi et al., 2014). Indeed, synaptic scaling has proven successful in stabilizing activity in recurrent neural networks (Lazar et al., 2009; Remme and Wadman, 2012; Zenke et al., 2013; Effenberger and Jost, 2015; Miner and Triesch, 2016). However, these studies either used synaptic scaling as the sole homeostatic mechanism (Remme and Wadman, 2012; Zenke et al., 2013) or resorted to a variant of synaptic scaling where the scaling is not dynamically determined through a control loop using a particular target activity, but rather by a fixed multiplicative normalization rule (Lazar et al., 2009; Effenberger and Jost, 2015; Miner and Triesch, 2016). Therefore, these homeostatic models cannot account for higher moments of temporal activity patterns, i.e., their variance, as this would require at least the tuning of two parameters (Cannon and Miller, 2017).

Within more abstract models of rate encoding neurons, intrinsic homeostasis and synaptic scaling essentially correspond to adjusting a bias and gain factor on the input entering a nonlinear transfer function. Within this framework, multiple dual-homeostatic adaptation rules have been investigated concerning their effect on network performance. In this framework, the adaptation of the bias acts as an intrinsic plasticity mechanism for the control of the internal excitability of a neuron (Franklin et al., 1992; Abbott and LeMasson, 1993; Borde et al., 1995), while the gain factors functionally correspond to a synaptic scaling of the recurrent weights. Learning rules for these types of models were usually derived by defining a target output distribution that each neuron attempts to reproduce by changing neural gains and biases (Steil, 2007; Triesch, 2007; Schrauwen et al., 2008; Boedecker et al., 2009), or were directly derived from an information-theoretic measure (Bell and Sejnowski, 1995).

While these studies did indeed show performance improvements by optimizing local information transmission measures, apparently, optimal performance can effectively be traced back to a global parameter, the spectral radius of the recurrent weight matrix (Schrauwen et al., 2008). Interestingly, to our knowledge, theoretical studies on spiking neural networks did not explicitly consider the spectral radius as a parameter affecting network dynamics. Still, the theory of balanced states in spiking recurrent networks established the idea that synaptic strengths should scale with $1/\sqrt{k}$, where k is the average number of afferent connections (Van Vreeswijk and Sompolinsky, 1998). According to the circular law of random matrix theory, this scaling rule simply implies that the spectral radius of the recurrent weight matrix remains finite as the number of neurons N increases. More recent experiments on cortical cultures confirm this scaling (Barral and D'Reyes, 2016).

In the present study, we investigated whether the spectral radius of the weight matrix in a random recurrent network can be regulated by a combination of intrinsic homeostasis and synaptic scaling. Following the standard echo-state framework, we used rate encoding tanh-neurons as the model of choice. However, aside from their applications as efficient machine learning algorithms, echo state networks are potentially relevant as models of information processing in the brain (Nikolić et al., 2009; Hinaut et al., 2015; Enel et al., 2016). Note in this context that extensions to layered ESN architectures have been presented by Gallicchio and Micheli (2017), which bears a somewhat greater resemblance to the hierarchical structure of cortical networks than the usual shallow ESN architecture. This line of research illustrates the importance of examining whether local and biological plausible principles exist that would allow to tune the properties of the neural reservoir to the "edge of chaos" (Livi et al., 2018), particularly when a continuous stream of inputs is present. The rule has to be independent of both the network topology, which is not locally accessible information, and the distribution of synaptic weights.

We propose and compare two unsupervised homeostatic mechanisms, which we denote by flow control and variance control. Both regulate the mean and variance of neuronal firing such that the network works in an optimal regime concerning sequence learning tasks. The mechanisms act on two sets of node-specific parameters, the biases b_i, and the neural gain factors a_i. We restricted ourselves to biologically plausible adaptation mechanisms, viz adaptation rules for which the dynamics of all variables are local, i.e., bound to a specific neuron. Additional variables enter only when locally accessible. In a strict sense, this implies that local dynamics are determined exclusively by the neuron's dynamical variables and by information about the activity of afferent neurons. Being less restrictive, one could claim that it should also be possible to access aggregate or mean-field quantities that average a variable of interest over the population. For example, nitric oxide is a diffusive neurotransmitter that can act as a measure for the population average of neural firing rates (Sweeney et al., 2015).

Following a general description of the network model, we introduce both adaptation rules and evaluate their effectiveness in tuning the spectral radius in sections 2.4 and 2.5. We assess the performance of networks that were subject to adaptation in section 2.7, using a nonlinear sequential memory task. Finally, we discuss the influence of node-to-node cross-correlations within the population in section 2.8.

2. RESULTS

2.1. Model

A full description of the network model and parameters can be found in the methods section. We briefly introduce the network dynamics as

$$x_i(t) = x_{r,i}(t) + I_i(t) \tag{1}$$

$$x_{r,i}(t) := a_i \sum_{j=1}^{N} W_{ij} y_j(t-1) \tag{2}$$

$$y_i(t) = \tanh\left(x_i(t) - b_i\right). \tag{3}$$

Each neuron's membrane potential x_i consists of a recurrent contribution $x_{r,i}(t)$ and an external input $I_i(t)$. The biases b_i are subject to the following homeostatic adaptation:

$$b_i(t) = b_i(t - 1) + \epsilon_b \left[y_i(t) - \mu_t \right]. \tag{4}$$

Here, μ_t defines a target for the average activity and ϵ_b is the adaptation rate.

The local parameters a_i act as scaling factors on the recurrent weights. We considered two different forms of update rules. Loosely speaking, both drive the network toward a certain dynamical state which corresponds to the desired spectral radius. The difference between them lies in the variables characterizing this state: While flow control defines a relation between the variance of neural activity and the variance of the total recurrent synaptic current, variance control does so by a more complex relation between the variance of neural activity and the variance of the synaptic current from the external input.

2.1.1. Flow Control
The first adaptation rule, flow control, is given by

$$a_i(t) = a_i(t-1)\left[1 + \epsilon_a \Delta R_i(t)\right], \qquad \Delta R_i(t) = R_t^2 y_i^2(t-1) - x_{r,i}^2(t). \tag{5}$$

The parameter R_t is the desired target spectral radius and ϵ_a the adaptation rate of the scaling factor. The dynamical variables y_i^2 and $x_{r,i}^2$ have been defined before in Equations (1) and (2). We also considered an alternative global update rule where $\Delta R_i(t)$ is given by

$$\Delta R_i(t) = \frac{1}{N}\left[R_t^2 \, ||\mathbf{y}(t-1)||^2 - ||\mathbf{x}_r(t)||^2 \right], \tag{6}$$

where $|| \cdot ||$ denotes the Euclidean vector norm. However, since this is a non-local rule, it only served as a comparative model to Equation (5) when we investigated the effectiveness of the adaptation mechanism. Three key assumptions enter flow control, Equation (5):

- Represented by $x_{r,i}(t)$, we assume that there is a physical separation between the recurrent input that a neuron receives and its external inputs. This is necessary because $x_{r,i}(t)$ is explicitly used in the update rule of the synaptic scaling factors.
- Synaptic scaling only affects the weights of recurrent connections. However, this assumption is not crucial for the effectiveness of our plasticity rule, as we were mostly concerned with achieving a preset spectral radius for the recurrent weight matrix. If instead the scaling factors acted on both the recurrent and external inputs, this would lead to an "effective" external input $I_i'(t) := a_i I_i(t)$. However, a_i only affecting the recurrent input facilitated the parameterization of the external input by means of its variance (see section 2.7), a choice of convenience.
- For (5) to function, neurons need to able to represent and store squared neural activities.

Whether these three preconditions are satisfied by biological neurons needs to be addressed in future studies.

2.1.2. Variance Control
The second adaptation rule, variance control, has the form

$$a_i(t) = a_i(t-1) + \epsilon_a \left[\sigma_{t,i}^2(t) - \left(y_i(t) - \mu_i^y(t)\right)^2 \right] \tag{7}$$

$$\sigma_{t,i}^2(t) = 1 - \frac{1}{\sqrt{1 + 2R_t^2 y_i(t)^2 + 2\sigma_{\text{ext},i}^2(t)}}. \tag{8}$$

Equation (7) drives the average variance of each neuron toward a desired target variance $\sigma_{t,i}^2(t)$ at an adaptation rate ϵ_a by calculating the momentary squared difference between the local activity $y_i(t)$ and its trailing average $\mu_i^y(t)$. Equation (8) calculates the target variance as a function of the target spectral radius R_t, the current local square activity $y_i^2(t)$ and a trailing average $\sigma_{\text{ext},i}^2(t)$ of the local variance of the external input signal. When all $a_i(t)$ reach a steady state, the average neural variance equals the target given by (8). According to a mean-field approach that is described in section 5.6, reaching this state then results in a spectral radius R_a that is equal to the target R_t entering (8). Intuitively, it is to be expected that $\sigma_{t,i}^2$ is a function of both the spectral radius and the external driving variance: The amount of fluctuations in the network activity is determined by the dynamic interplay between the strength of the external input as well as the recurrent coupling.

A full description of the auxiliary equations and variables used to calculate $\mu_i^y(t)$ and $\sigma_{\text{ext},i}^2(t)$ can be found in section 5.1.

Similar to flow control, we also considered a non-local version for comparative reasons, where (8) is replaced with

$$\sigma_{t,i}^2(t) = 1 - \frac{1}{\sqrt{1 + 2R_t^2 ||\mathbf{y}(t)||^2/N + 2\sigma_{\text{ext},i}^2(t)}}. \tag{9}$$

Again, $||\cdot||$ denotes the Euclidean norm. Before proceeding to the results, we discuss the mathematical background of the proposed adaptation rules in some detail.

2.2. Autonomous Spectral Radius Regulation
There are some interesting aspects to the theoretical framework at the basis of the here proposed regulatory scaling mechanisms.

The circular law of random matrix theory states that the eigenvalues λ_j are distributed uniformly on the complex unit disc if the elements of a real $N \times N$ matrix are drawn from distributions having zero mean and standard deviation $1/\sqrt{N}$ (Tao and Vu, 2008). Given that the internal weight matrix $\widehat{W}$ ($\widehat{\cdot}$ denoting matrices) with entries W_{ij} has $p_r N$ non-zero elements per row (p_r is the connection probability), the circular law implies that the spectral radius of $a_i W_{ij}$, the maximum of $|\lambda_j|$, is unity when the synaptic scaling factors a_i are set uniformly to $1/\sigma_w$, where σ_w is the standard deviation of W_{ij}. Our goal is to investigate adaptation rules for the synaptic scaling factors that are based on dynamic quantities, which includes the membrane potential x_i, the neural activity y_i and the input I_i.

The circular law, i.e., a $N \times N$ matrix with i.i.d. entries with zero mean and $1/N$ variance approaching a spectral radius of one as $N \to \infty$, can be generalized. Rajan and Abbott

(2006) investigated the case where the statistics of the columns of the matrix differ in their means and variances: given row-wise E-I balance for the recurrent weights, the square of the spectral radius of a random $N \times N$ matrix whose columns have variances σ_i^2 is $N \langle \sigma_i^2 \rangle_i$ for $N \to \infty$. Since the eigenvalues are invariant under transposition, this result also holds for row-wise variations of variances and column-wise E-I balance. While the latter is not explicitly enforced in our case, deviations from this balance are expected to tend to zero for large N given the statistical assumptions that we made about the matrix elements W_{ij}. Therefore, the result can be applied to our model, where node-specific gain factors a_i are applied to each row of the recurrent weight matrix. Thus, the spectral radius R_a of the *effective random matrix* $\widehat{W}_a$ with entries $a_i W_{ij}$ [as entering (2)] is

$$R_a^2 \cong \frac{1}{N} \sum_i R_{a,i}^2, \qquad R_{a,i}^2 := a_i^2 \sum_j (W_{ij})^2, \qquad (10)$$

for large N, when assuming that the distribution underlying the *bare weight matrix* $\widehat{W}$ with entries W_{ij} has zero mean. Note that R_a^2 can be expressed alternatively in terms of the Frobenius norm $\| \widehat{W}_a \|_F$, via

$$R_a^2 \cong \| \widehat{W}_a \|_F^2 / N. \qquad (11)$$

We numerically tested Equation (10) for $N = 500$ and heterogeneous random sets of a_i drawn from a uniform $[0, 1]$-distribution and found a very close match to the actual spectral radii (1–2% relative error). Given that the $R_{a,i}$ can be interpreted as per-site estimates for the spectral radius, one can use the generalized circular law (10) to regulate R_a on the basis of local adaptation rules, one for every a_i.

For the case of flow control, the adaptation rule is derived using a comparison between the variance of neural activity that is present in the network with the recurrent contribution to the membrane potential. A detailed explanation is presented in sections 2.6, 5.5. In short, we propose that

$$\left\langle ||\mathbf{x}_r(t)||^2 \right\rangle_t \cong R_a^2 \left\langle ||\mathbf{y}(t-1)||^2 \right\rangle_t, \qquad (12)$$

where $x_{r,i}$ is the recurrent contribution to the membrane potential x_i. This stationarity condition leads to the adaptation rule given in Equation (5). An analysis of the dynamics of this adaptation mechanisms can be found in section 5.5.

Instead of directly imposing Equation (12) via an appropriate adaptation mechanism, we also considered the possibility of transferring this condition into a set point for the variance of neural activities as a function the external driving. To do so, we used a mean-field approach to describe the effect of recurrent input onto the resulting neural activity variance. An in-depth discussion is given in section 5.6. This led to the update rule given by Equations (7) and (8) for variance control.

2.3. Testing Protocols
We used several types of input protocols for testing the here proposed adaptation mechanisms, as well as for assessing the task performance discussed in section 2.7. The first two variants concern distinct biological scenarios:

- *Binary*. Binary input sequences correspond to a situation when a neural ensemble receives input dominantly from a singular source, which itself has only two states, being either active or inactive. Using binary input sequences during learning is furthermore consistent with the non-linear performance test considered here for the echo-state network as a whole, the delayed XOR-task. See section 2.7. For symmetric binary inputs, as used, the source signal $u(t)$ is drawn from ± 1.
- *Gaussian*. Alternatively one can consider the situation that a large number of essentially uncorrelated input streams are integrated. This implies random Gaussian inputs signals. Neurons receive in this case zero-mean independent Gaussian noise.

Another categorical dimension concerns the distribution of the afferent synaptic weights. Do all neurons receive inputs with the same strength, or not? As a quantifier for the individual external input strengths, the variances $\sigma_{\mathrm{ext},i}^2$ of the local external input currents where taken into account. We distinguished two cases

- *Heterogeneous*. In the first case, the $\sigma_{\mathrm{ext},i}^2$ are quenched random variables. This means that each neuron is assigned a random value $\sigma_{\mathrm{ext},i}^2$ before the start of the simulation, as drawn from a half-normal distribution parameterized by $\sigma = \sigma_{\mathrm{ext}}$. This ensures that the expected average variance $\langle \sigma_{\mathrm{ext},i}^2 \rangle$ is given by σ_{ext}^2.
- *Homogeneous*. In the second case, all $\sigma_{\mathrm{ext},i}^2$ are assigned the identical global value σ_{ext}^2.

Overall, pairing "binary" vs. "Gaussian" and "heterogeneous" vs. "homogeneous," leads to a total of four different input protocols, i.e., "heterogeneous binary," "homogeneous binary," "heterogeneous Gaussian," and "homogeneous Gaussian." If not otherwise stated, numerical simulations were done using networks with $N = 500$ sites and a connection probability $p_r = 0.1$.
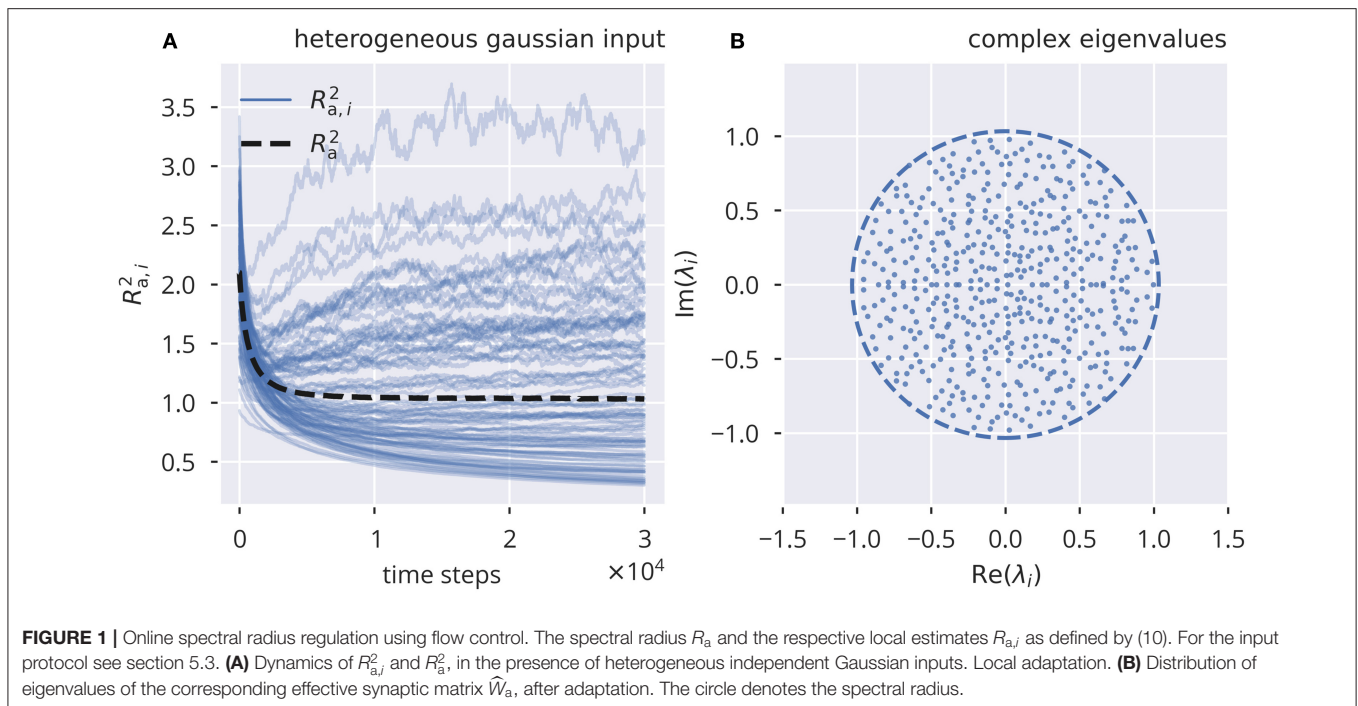
2.4. Performance Testing of Flow Control
In **Figure 1**, we present a simulation using flow control for heterogeneous Gaussian input with an adaptation rate $\epsilon_a = 10^{-3}$. The standard deviation of the external driving was set to $\sigma_{\mathrm{ext}} = 0.5$. The spectral radius of R_a of $\widehat{W}_a$ was tuned to the target $R_t = 1$ with high precision, even though the local, row-wise estimates $R_{a,i}$ showed substantial deviations from the target.

We further tested the adaptation with other input protocols (see section 2.3 and **Supplementary Figure 1**). We found that flow control robustly led to the desired spectral radius R_t under all Gaussian input protocols, while binary input caused R_a to converge to higher values than R_t. However, when using global adaptation, as given by Equation (6), all input protocols resulted in the correctly tuned spectral radius (see **Supplementary Figure 2**).

Numerically, we found that the time needed to converge to the stationary states depended substantially on R_t, slowing down when the spectral radius becomes small. It is then advantageous, as we have done, to scale the adaptation rate ϵ_a inversely with the trailing average $\bar{x}_r^2$ of $||x_r||^2$, viz as $\epsilon_a \to \epsilon_a / \bar{x}_r^2$. An exemplary

FIGURE 1 | Online spectral radius regulation using flow control. The spectral radius R_a and the respective local estimates $R_{a,i}$ as defined by (10). For the input protocol see section 5.3. **(A)** Dynamics of $R_{a,i}^2$ and R_a^2, in the presence of heterogeneous independent Gaussian inputs. Local adaptation. **(B)** Distribution of eigenvalues of the corresponding effective synaptic matrix $\widehat{W}_a$, after adaptation. The circle denotes the spectral radius.

plot showing the effect of this scaling is shown in **Figure 7** (see section 5.2 for further details).

To evaluate the amount of deviation from the target spectral radius under different input strengths and protocols, we plotted the difference between the resulting spectral radius and the target spectral radius for a range of external input strength, quantified by their standard deviation σ_{ext}. Results for different input protocols are shown in **Supplementary Figure 5**. For correlated binary input, increasing the input strength resulted in stronger deviations from the target spectral radius. On the other hand, uncorrelated Gaussian inputs resulted in perfect alignment for the entire range of input strengths that we tested.

2.5. Performance Testing of Variance Control

In comparison, variance control, shown in **Figure 2**, **Supplementary Figure 3**, resulted in notable deviations from R_t, for both uncorrelated Gaussian and correlated binary input. As for flow control, we also calculated the deviations from R_t as a function of σ_{ext} (see **Supplementary Figure 6**). For heterogeneous binary input, deviations from the target spectral radius did not increase monotonically as a function of the input strength (**Supplementary Figure 6A**), reaching a peak at $\sigma_{\text{ext}} \approx 0.4$ for target spectral radii larger than 1. For homogeneous binary input, we observed a substantial negative mismatch of the spectral radius for strong external inputs (see **Supplementary Figure 6C**).

Overall, we found that variance control did not exhibit the same level of consistency in tuning the system toward a desired spectral radius, even though it did perform better in some particular cases (compare **Supplementary Figure 5A**
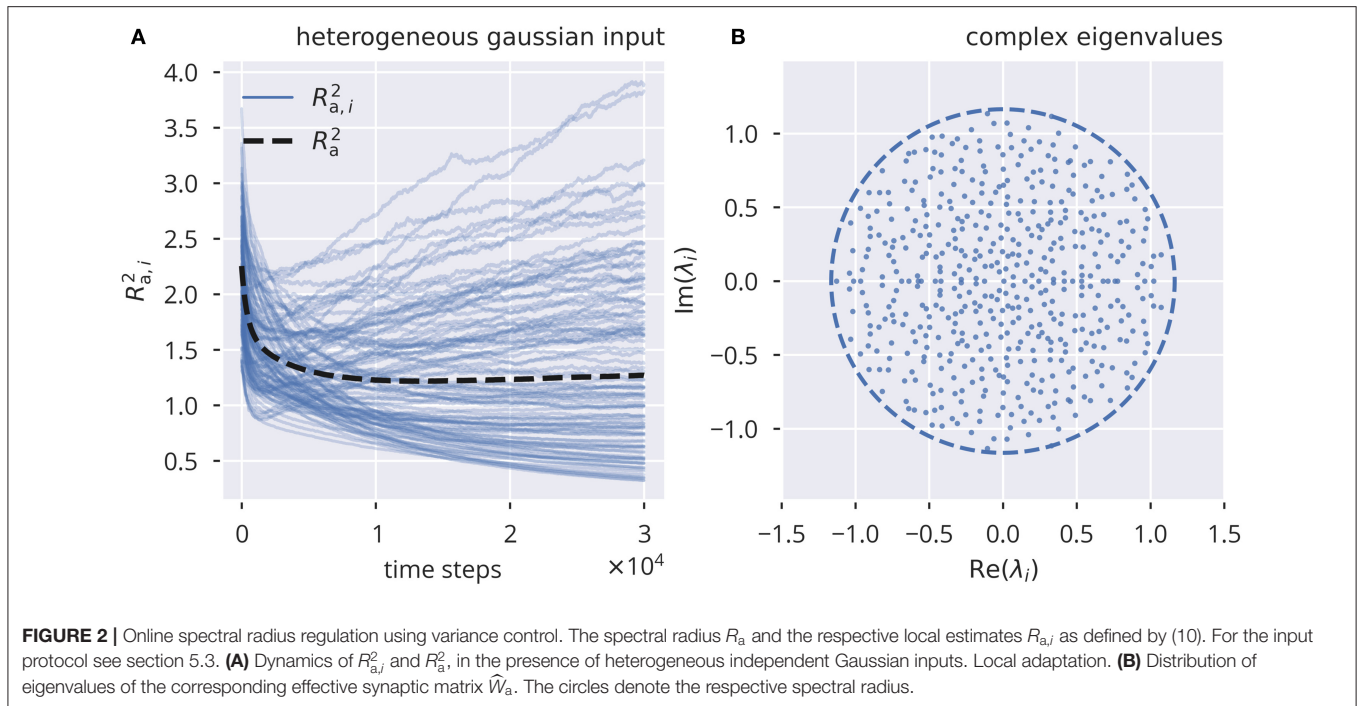
for large σ_{ext} with **Supplementary Figure 6**). Moreover, variance control exhibited deviations from the target (shown in **Supplementary Figure 4**) even when a global adaptation rule was used, as defined in (9). This is in contrast to the global variant of flow control, which, as stated in the previous section, robustly tuned the spectral radius to the desired value even in the presence of strongly correlated inputs.

2.6. Spectral Radius, Singular Values, and Global Lyapunov Exponents

Apart from the spectral radius R_a of the matrix $\widehat{W}_a$, one may consider the relation between the adaptation dynamics and the respective singular values σ_i of $\widehat{W}_a$. We recall that the spectrum of $\hat{U}_a = \widehat{W}_a^\dagger \widehat{W}_a$ is given by the squared singular values, σ_i^2, and that the relation $||\mathbf{x}_r||^2 = \mathbf{y}^\dagger \widehat{W}_a^\dagger \widehat{W}_a \mathbf{y}$ holds. Now, assume that the time-averaged projection of neural activity $\mathbf{y} = \mathbf{y}(t)$ onto all eigenvectors of $\hat{U}_a$ is approximately the same, that is, there is no preferred direction of neural activity in phase space. From this idealized case, it follows that the time average of the recurrent contribution to the membrane potential can be expressed with

$$\left\langle ||\mathbf{x}_r||^2 \right\rangle_t \approx \frac{\left\langle ||\mathbf{y}||^2 \right\rangle_t}{N} \sum_i \sigma_i^2 = \frac{\left\langle ||\mathbf{y}||^2 \right\rangle_t}{N} \sum_{i,j} \left(a_i W_{ij} \right)^2 \quad (13)$$

as the rescaled average of the σ_i^2. For the second equation, we used the fact that the $\sum_i \sigma_i^2$ equals the sum of all matrix elements squared (Sengupta and Mitra, 1999; Shen, 2001). With (10), one finds that (13) is equivalent to $\left\langle ||\mathbf{x}_r||^2 \right\rangle_t = R_a^2 \left\langle ||\mathbf{y}||^2 \right\rangle_t$ and hence to the flow condition (12). This result can be generalized, as done in section 5.5, to the case that the neural activities have inhomogeneous variances, while still being uncorrelated with

FIGURE 2 | Online spectral radius regulation using variance control. The spectral radius R_a and the respective local estimates $R_{a,i}$ as defined by (10). For the input protocol see section 5.3. **(A)** Dynamics of $R_{a,i}^2$ and R_a^2, in the presence of heterogeneous independent Gaussian inputs. Local adaptation. **(B)** Distribution of eigenvalues of the corresponding effective synaptic matrix $\widehat{W}_a$. The circles denote the respective spectral radius.

zero mean. We have thus shown that the stationarity condition leads to a spectral radius of (approximately) unity.

It is worthwhile to note that the singular values of $\hat{U}_a = \widehat{W}_a^\dagger \widehat{W}_a$ do exceed unity when $R_a = 1$. More precisely, for a random matrix with i.i.d. entries, one finds in the limit of large N that the largest singular value is given by $\sigma_{\max} = 2R_a$, in accordance with the Marchenko-Pastur law for random matrices (Marčenko and Pastur, 1967). Consequently, directions in phase space exist in which the norm of the phase space vector is elongated by factors greater than one. Still, this does not contradict the fact that a unit spectral radius coincides with the transition to chaos for the non-driven case. The reason is that the global Lyapunov exponents are given by

$$\lim_{n\to\infty} \frac{1}{2n} \ln\left(\left(\widehat{W}_a^n \right)^\dagger \widehat{W}_a^n \right) \qquad (14)$$

which eventually converge to $\ln\|\lambda_i\|$, see **Supplementary Figure 7** and Wernecke et al. (2019), where λ_i is the ith eigenvalue of $\widehat{W}_a$. The largest singular value of the nth power of a random matrix with a spectral radius R_a scales like R_a^n in the limit of large powers n. The global Lyapunov exponent goes to zero as a consequence when $R_a \to 1$.

2.7. XOR-Memory Recall

To this point, we presented results regarding the effectiveness of the introduced adaptation rules. However, we did not account for their effects onto a given learning task. Therefore, we tested the performance of locally adapted networks under the delayed XOR task, which evaluates the memory capacity of the echo state network in combination with a non-linear operation. For the task, the XOR operation is to be taken with respect to a

delayed pair of two consecutive binary inputs signals, $u(t-\tau)$ and $u(t-\tau-1)$, where τ is a fixed time delay. The readout layer is given by a single unit, which has the task of reproducing

$$f_\tau(t) = \mathrm{XOR}\left[u(t-\tau), u(t-\tau-1) \right], \qquad t, \tau = 1, 2, \ldots, \qquad (15)$$

where $\mathrm{XOR}[u, u']$ is 0/1 if u and u' are identical/not identical.

The readout vector $\mathbf{w}_{\mathrm{out}}$ is trained with respect to the mean squared output error,

$$\|\widehat{Y}\mathbf{w}_{\mathrm{out}} - \mathbf{f}_\tau\|^2 + \alpha\|\mathbf{w}_{\mathrm{out}}\|^2, \qquad (16)$$

using ridge regression on a sample batch of $T_{\mathrm{batch}} = 10N$ time steps, here for $N = 500$, and a regularization factor of $\alpha = 0.01$. The batch matrix $\widehat{Y}$, of size $T_{\mathrm{batch}} \times (N + 1)$, holds the neural activities as well as one node with constant activity serving as a potential bias. Similarly, the readout (column) vector $\mathbf{w}_{\mathrm{out}}$ is of size $(N + 1)$. The T_{batch} entries of $\mathbf{f}_\tau$ are the $f_\tau(t)$, viz the target values of the XOR problem. Minimizing (16) leads to

$$\mathbf{w}_{\mathrm{out}} = \left(\widehat{Y}^\dagger \widehat{Y} + \alpha^2 \hat{\mathbb{1}} \right)^{-1} \widehat{Y}^\dagger \mathbf{f}_\tau. \qquad (17)$$

The learning procedure was repeated independently for each time delay τ. We quantified the performance by the total memory capacity, $\mathrm{MC}_{\mathrm{XOR}}$, as

$$\mathrm{MC}_{\mathrm{XOR}} = \sum_{k=1}^{\infty} \mathrm{MC}_{\mathrm{XOR},k} \qquad (18)$$

$$\mathrm{MC}_{\mathrm{XOR},k} = \frac{\mathrm{Cov}^2\left[f_k(t), y_{\mathrm{out}}(t) \right]_t}{\mathrm{Var}\left[f_k(t) \right]_t \mathrm{Var}\left[y_{\mathrm{out}}(t) \right]_t}. \qquad (19)$$

This is a simple extension of the usual definition of short term memory in the echo state literature (Jaeger, 2002). The activity $y_{\text{out}} = \sum_{i=1}^{N+1} w_{\text{out},i} \, y_i$ of the readout unit is compared in (19) with the XOR prediction task, with the additional neuron, $y_{N+1} = 1$, corresponding to the bias of the readout unit. Depending on the mean level of the target signal, this offset might actually be unnecessary. However, since it is a standard practice to use an intercept variable in linear regression models, we decided to include it into the readout variable y_{out}. The variance and covariance are calculated with respect to the batch size T_{batch}.

The results for flow control presented in **Figure 3** correspond to two input protocols, heterogeneous Gaussian and binary inputs. Shown are sweeps over a range of σ_{ext} and R_t. The update rule (5) was applied to the network for each pair of parameters until the a_i values converged to a stable configuration. We then measured the task performance as described above. Note that in the case of Gaussian input, this protocol was only used during the adaptation phases. Due to the nature of the XOR task, binary inputs with the corresponding variances are to be used during performance testing. See **Supplementary Figure 8** for a performance sweep using the homogeneous binary and Gaussian input protocol. Optimal performance was generally attained around the $R_a \approx 1$ line. A spectral radius R_a slightly smaller than unity was optimal when using Gaussian input, but not for binary input signals. In this case the measured spectral radius R_a deviated linearly from the target R_t, with increasing strength of the input, as parameterized by the standard deviation σ_{ext}. Still, the locus of optimal performance was essentially independent of the input strength, with maximal performance attained roughly at $R_t \approx 0.55$. Note that the line $R_a = 1$ joins $R_t = 1$ in the limit $\sigma_{\text{ext}} \to 0$.

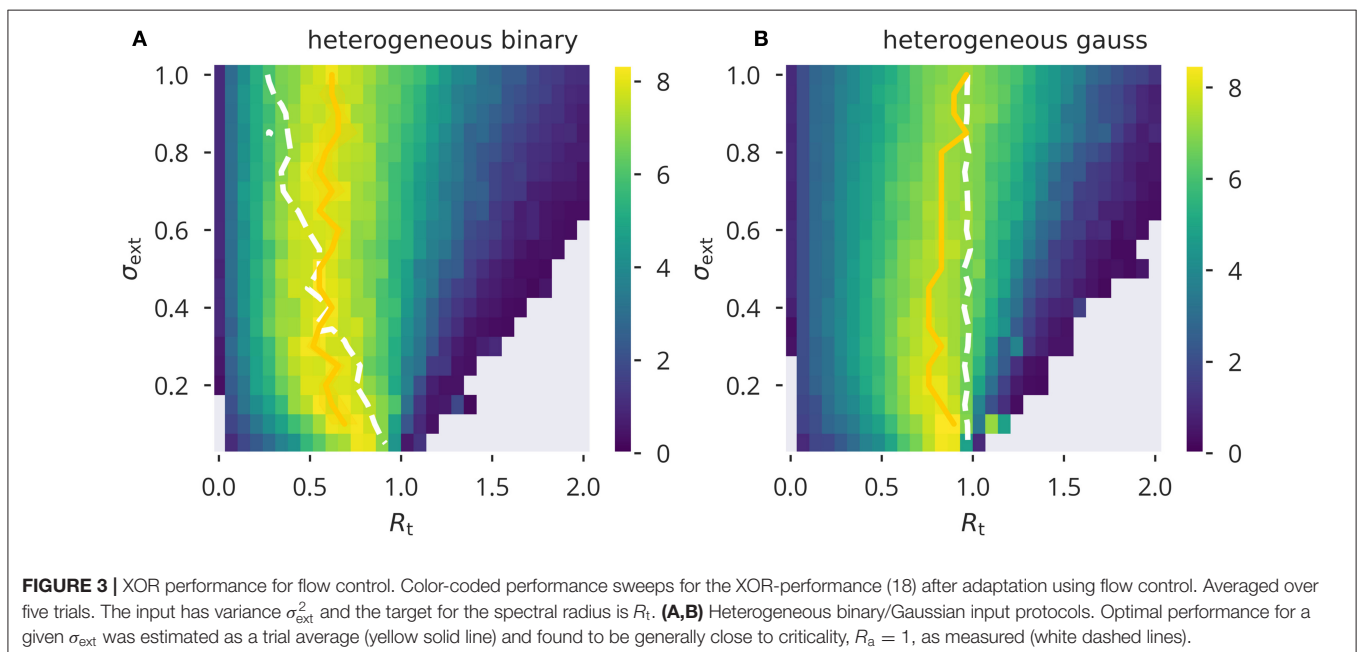Comparing these results to variance control, as shown in **Figure 4**, we found that variance control led to an overall lower performance. To our surprise, for external input with a large variance, Gaussian input caused stronger deviations from the desired spectral radius as compared to binary input. Therefore, in a sense, it appeared to behave opposite to what we found for flow control. However, similar to flow control, the value of R_t giving optimal performance under a given σ_{ext} remained relatively stable over the range of external input strength measured. On the other hand, using homogeneous input (see **Supplementary Figure 9**), did cause substantial deviations from the target spectral radius when using binary input.
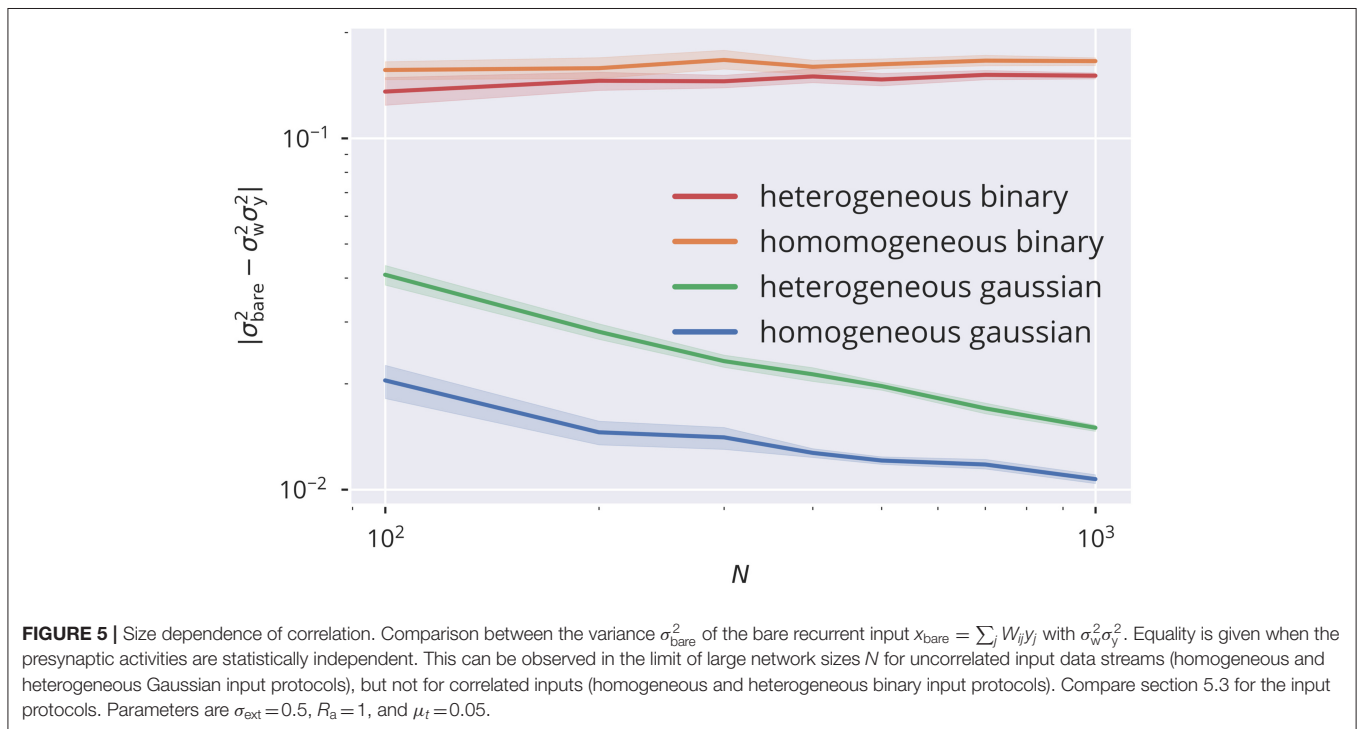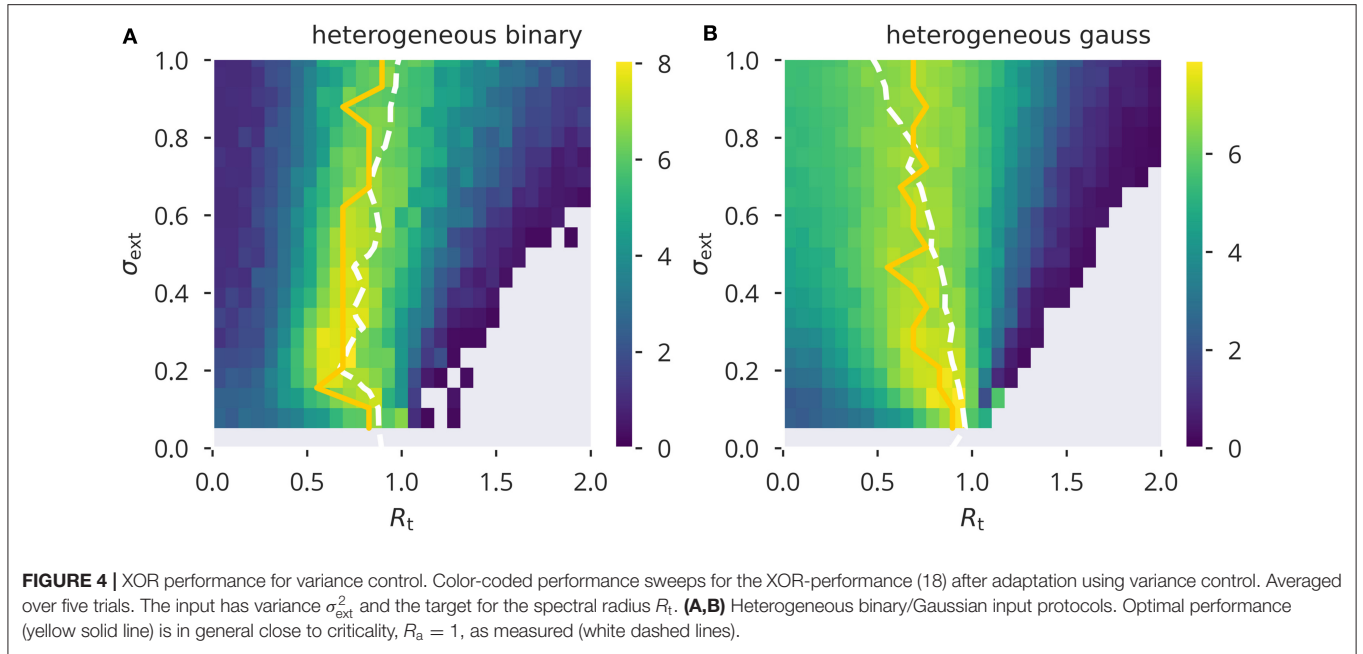
2.8. Input Induced Correlations

A crucial assumption leading to the proposed adaptation rules is the statistical independence of neural activity for describing the statistical properties of the bare recurrent contribution to the membrane potential, $x_{\text{bare}} = \sum_j W_{ij} y_j$. In particular, the variance σ_{bare}^2 of x_{bare} enters the mean-field approach described in section 5.6. Assuming statistical independence across the population for $y_i(t)$, it is simply given by $\sigma_{\text{bare}}^2 = \sigma_{\text{w}}^2 \sigma_{\text{y}}^2$, where

$$\sigma_{\text{w}}^2 \equiv \text{Var}\left[\sum_{j=1}^{N} W_{ij} \right] \tag{20}$$

is the variance of the sum of the bare afferent synaptic weights (see also section 5.1). Being a crucial element of the proposed rules, deviations from the prediction of σ_{bare}^2 would also negatively affect the precision of tuning the spectral radius. In **Figure 5**, a comparison of the deviations $|\sigma_{\text{bare}}^2 - \sigma_{\text{w}}^2 \sigma_{\text{y}}^2|$ is presented for the four input protocols introduced in section 5.3. For the Gaussian protocols, for which neurons receive statistically uncorrelated external signals, one observes

FIGURE 3 | XOR performance for flow control. Color-coded performance sweeps for the XOR-performance (18) after adaptation using flow control. Averaged over five trials. The input has variance σ_{ext}^2 and the target for the spectral radius is R_t. **(A,B)** Heterogeneous binary/Gaussian input protocols. Optimal performance for a given σ_{ext} was estimated as a trial average (yellow solid line) and found to be generally close to criticality, $R_a = 1$, as measured (white dashed lines).

FIGURE 4 | XOR performance for variance control. Color-coded performance sweeps for the XOR-performance (18) after adaptation using variance control. Averaged over five trials. The input has variance σ_{ext}^2 and the target for the spectral radius R_t. **(A,B)** Heterogeneous binary/Gaussian input protocols. Optimal performance (yellow solid line) is in general close to criticality, $R_a = 1$, as measured (white dashed lines).

FIGURE 5 | Size dependence of correlation. Comparison between the variance σ_{bare}^2 of the bare recurrent input $x_{bare} = \sum_j W_{ij} y_j$ with $\sigma_w^2 \sigma_y^2$. Equality is given when the presynaptic activities are statistically independent. This can be observed in the limit of large network sizes N for uncorrelated input data streams (homogeneous and heterogeneous Gaussian input protocols), but not for correlated inputs (homogeneous and heterogeneous binary input protocols). Compare section 5.3 for the input protocols. Parameters are $\sigma_{ext} = 0.5$, $R_a = 1$, and $\mu_t = 0.05$.

that $\sigma_{bare}^2 \rightarrow \sigma_w^2 \sigma_y^2$ in the thermodynamic limit $N \rightarrow \infty$ via a power law, which is to be expected when the presynaptic neural activities are decorrelated. On the other side, binary 0/1 inputs act synchronous on all sites, either with site-dependent or site-independent strengths (heterogeneous/homogeneous). Corresponding activity correlations are induced and a finite and only weakly size-dependent difference between σ_{bare}^2 and $\sigma_w^2 \sigma_y^2$ shows up. Substantial corrections to the analytic theory are to

be expected in this case. To this extend we measured the cross-correlation $C(y_i, y_j)$, defined as

$$\bar{C} = \frac{1}{N(N-1)} \sum_{i \neq j} |C(y_i, y_j)|,$$

$$C(y_i, y_j) = \frac{\text{Cov}(y_i, y_j)}{\sqrt{\text{Cov}(y_i, y_i)\text{Cov}(y_j, y_j)}}, \quad (21)$$

with the covariance given by $Cov(y_i, y_j) = \langle(y_i - \langle y_i\rangle_t)(y_j - \langle y_j\rangle_t)\rangle_t$. For a system of $N = 500$ neurons the results for the averaged absolute correlation $\bar{C}$ are presented in **Figure 6** (see **Supplementary Figure 10** for homogeneous input protocols). Autonomous echo-state layers are in chaotic states when supporting a finite activity level, which implies that correlations vanish in the thermodynamic limit $N \to \infty$. The case $\sigma_{ext} = 0$, as included in **Figure 6**, serves consequently as a yardstick for the magnitude of correlations that are due to the finite number of neurons.
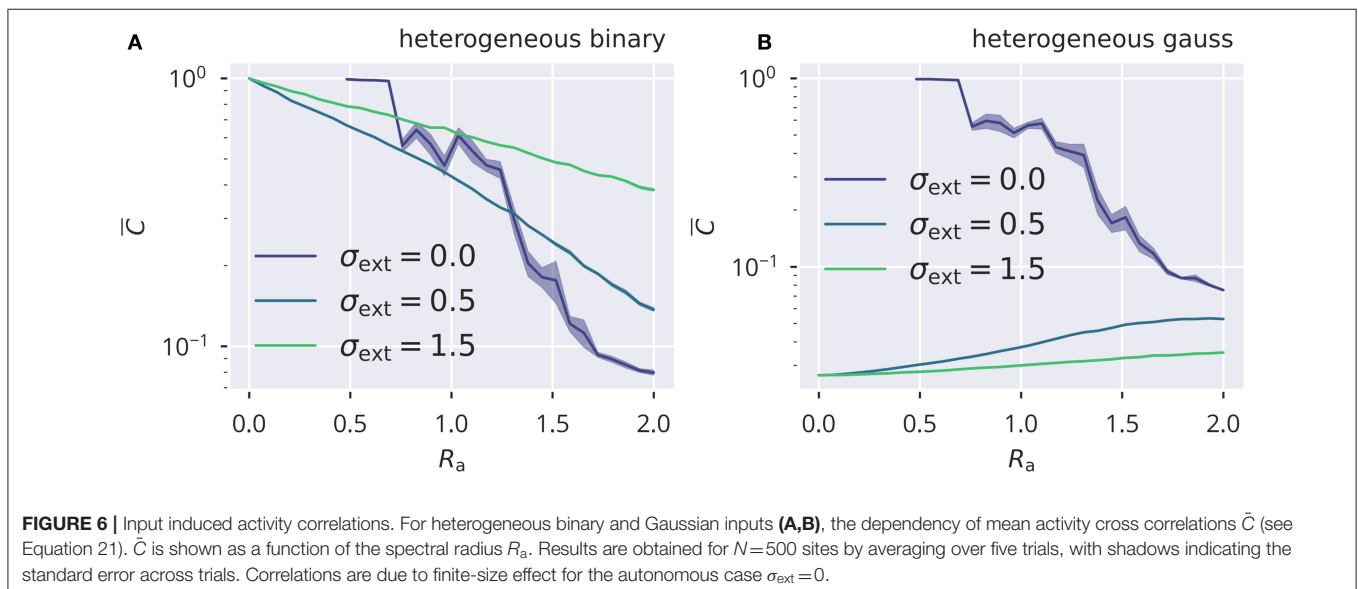
Input correlations were substantially above the autonomous case for correlated binary inputs, with the magnitude of $\bar{C}$ decreasing when the relative contribution of the recurrent activity increased. This was the case for increasing R_a. The effect was opposite for the Gaussian protocol, for which the input did not induce correlations, but contributed to decorrelating neural activity. In this case, the mean absolute correlation $\bar{C}$ was suppressed when the internal activity became small in the limit $R_a \to 0$. For larger R_a, the recurrent input gained more impact on neural activity relative to the external drive and thus drove $\bar{C}$ toward an amount of correlation that would be expected in the autonomous case.

3. DISCUSSION

The mechanisms for tuning the spectral radius via a local homeostatic adaptation rule introduced in the present study require neurons to have the ability to distinguish and locally measure both external and recurrent input contributions. For flow control, neurons need to be able to compare the recurrent membrane potential with their own activity, as assumed in section 2.2. On the other hand, variance control directly measures the variance of the external input and derives the activity target variance accordingly. The limiting factor to a successful spectral radius control is the amount

of cross-correlation induced by external driving statistics. As such, the functionality and validity of the proposed mechanisms depended on the ratio between external input, i.e., feed-forward or feedback connections, with respect to recurrent, or lateral connections. In general, it is not straightforward to directly connect experimental evidence regarding the ratio between recurrent and feed-forward contributions to the effects observed in the model. It is, however, worthwhile to note that the fraction of synapses associated with interlaminar loops and intralaminar lateral connections are estimated to make up roughly 50% (Binzegger et al., 2004). Relating this to our model, it implies that the significant interneural correlations that we observed when external input strengths were of the same order of magnitude as the recurrent inputs, can not generally be considered an artifact of biologically implausible parameter choices. Synchronization (Echeveste and Gros, 2016) is in fact a widely observed phenomenon in the brain (Usrey and Reid, 1999), with possible relevance for information processing (Salinas and Sejnowski, 2001).

On the other hand, correlations due to shared input reduces the amount of information that can be stored in the neural ensemble (Bell and Sejnowski, 1995). Maximal information is achieved if neural activities or spikes trains form an orthogonal ensemble (Földiak, 1990; Bell and Sejnowski, 1995; Tetzlaff et al., 2012). Furthermore, neural firing in cortical microcircuits was found to be decorrelated across neurons, even if common external input was present (Ecker et al., 2010), that is, under a common orientation tuning. Therefore, the correlation we observed in our network due to shared input might be significantly reduced by possible modifications/extensions of our model: First, a strict separation between inhibitory and excitatory nodes according to Dale's law might help actively decorrelating neural activity (Tetzlaff et al., 2012; Bernacchia and Wang, 2013). Second, if higher dimensional input was used, a combination of plasticity mechanisms in the recurrent and feed-forward connections could lead to the formation of

FIGURE 6 | Input induced activity correlations. For heterogeneous binary and Gaussian inputs **(A,B)**, the dependency of mean activity cross correlations $\bar{C}$ (see Equation 21). $\bar{C}$ is shown as a function of the spectral radius R_a. Results are obtained for $N = 500$ sites by averaging over five trials, with shadows indicating the standard error across trials. Correlations are due to finite-size effect for the autonomous case $\sigma_{ext} = 0$.

an orthogonal representation of the input (Földiak, 1990; Bell and Sejnowski, 1995; Wick et al., 2010), leading to richer, "higher dimensional" activity patterns, i.e., a less dominant largest principal component. Ultimately, if these measures helped in reducing neural cross-correlations in the model, we thus would expect them to also increase the accuracy of the presented adaptation mechanisms. We leave these modifications to possible future research.

Overall, we found flow control to be generally more robust than variance control in the sense that, while still being affected by the amount of correlations within the neural reservoir, the task performance was less so prone to changes in the external input strength. Comparatively stable network performance could be observed, in spite of certain deviations from the desired spectral radius (see **Figure 3**). A possible explanation may be that flow control uses a distribution of samples from only a restricted part of phase space, that is, from the phase space regions that are actually visited or "used" for a given input. Therefore, while a spectral radius of unity ensures—statistically speaking—the desired scaling properties in all phase-space directions, it seem to be enough to control the correct scaling for the subspace of activities that is actually used for a given set of input patters. Variance control, on the other hand, relies more strictly on the assumptions that neural activities are statistical independent. In consequence, the desired results could only be achieved under a rather narrow set of input statistics (independent Gaussian input with small variance). In addition, the approximate expression derived for the nonlinear transformation appearing in the mean field approximation adds another layer of potential source of systematic error to the control mechanism. This aspect also speaks in favor of flow control, since its rules are mathematically more simple. In contrast to variance control, the stationarity condition stated in Equation (12) is independent of the actual nonlinear activation function used and could easily be adopted in a modified neuron model. It should be noted, however, that the actual target R_t giving optimal performance might then also be affected.

Interestingly, flow control distinguishes itself from a conventional local activity-target perspective of synaptic homeostasis: There is no predefined set point in Equation (5). This allows heterogeneities of variances of neural activity to develop across the network, while retaining the average neural activity at a fixed predefined level.

We would like to point out that, for all the results presented here, only stationary processes were used for generating the input sequences. Therefore, it might be worth considering the potential effects of non-stationary, yet bounded, inputs on the results in future work. It should be noted, however, that the temporal domain enters both adaptation mechanisms only in the form of trailing averages of first and second moments. As a consequence, we expect the issue of non-stationarity of external inputs to present itself simply as a trade-off between slower adaptation, i.e., longer averaging time scales, and the mitigation of the effects of non-stationarities. Slow adaptation is, however, completely in line with experimental results on the dynamics of synaptic scaling, which is taking place on the time scale of hours to days (Turrigiano et al., 1998; Turrigiano, 2008).

4. CONCLUSION

Apart from being relevant from a theoretical perspective, we propose that the separability of recurrent and external contributions to the membrane potential is an aspect that is potentially relevant for the understanding of local homeostasis in biological networks. While homeostasis in neural compartments has been a subject of experimental research (Chen et al., 2008), to our knowledge, it has not yet been further investigated on a theoretical basis, although it has been hypothesized that the functional segregation within the dendritic structure might also affect (among other intraneural dynamical processes) homeostasis (Narayanan and Johnston, 2012). The neural network model used in this study lacks certain features characterizing biological neural networks, like strict positivity of the neural firing rate or Dale's law, viz E-I balance (Trapp et al., 2018). Future research should therefore investigate whether the here presented framework of local flow control can be implemented within more realistic biological neural network models. A particular concern regarding our findings is that biological neurons are spiking. The concept of an underlying instantaneous firing rate is, strictly speaking, a theoretical construct, let alone the definition of higher moments, such as the "variance of neural activity." It is however acknowledged that the variability of the neural activity is central for statistical inference (Echeveste et al., 2020). It is also important to note that real-world biological control mechanisms, e.g., of the activity, rely on physical quantities that serve as measurable correlates. A well-known example is the intracellular calcium concentration, which is essentially a linearly filtered version of the neural spike train (Turrigiano, 2008). On a theoretical level, Cannon and Miller showed that dual homeostasis can successfully control the mean and variance of this type of spike-averaging physical quantities (Cannon and Miller, 2017). An extension of the flow control to filtered spike trains of spiking neurons could be an interesting subject of further investigations. However, using spiking neuron models would have shifted the focus of our research toward the theory of liquid state machines (Maass et al., 2002; Maass and Markram, 2004), exceeding the scope of this publication. We therefore leave the extension to more realistic network/neuron models to future work.

5. MATERIALS AND METHODS

5.1. Model

We implemented an echo state network with N neurons, receiving D_{in} inputs. The neural activity is $y_i \in [-1, 1]$, x_i the membrane potential, u_i the input activities, W_{ij} the internal synaptic weights and I_i the external input received. The output layer will be specified later. The dynamics

$$x_i(t) = a_i \sum_{j=1}^{N} W_{ij} y_j(t-1) + I_i(t), \qquad y_i(t) = \tanh\left(x_i(t) - b_i\right)$$

(22)

is discrete in time, where the input I_i is treated instantaneously. A tanh-sigmoidal has been used as a nonlinear activation function.

The synaptic renormalization factor a_i in (22) can be thought of as a synaptic scaling parameter that neurons use to regulate the overall strength of the recurrent inputs. The strength of the inputs I_i is unaffected, which is biologically plausible if external and recurrent signals arrive at separate branches of the dendritic tree (Spruston, 2008).

The W_{ij} are the bare synaptic weights, with $a_i W_{ij}$ being the components of the effective weight matrix $\widehat{W}_a$. Key to our approach is that the propagation of activity is determined by $\widehat{W}_a$, which implies that the spectral radius of the effective, and not of the bare weight matrix needs to be regulated.

The bare synaptic matrix W_{ij} is sparse, with a connection probability $p_r = 0.1$. The non-zero elements are drawn from a Gaussian with standard deviation

$$\sigma = \frac{\sigma_w}{\sqrt{N p_r}}, \qquad (23)$$

and vanishing mean μ. Here $N p_r$ corresponds to the mean number of afferent internal synapses, with the scaling $\sim 1/\sqrt{N p_r}$ enforcing size-consistent synaptic-weight variances. As discussed in the results section, we applied the following adaptation mechanisms:

$$b_i(t) = b_i(t-1) + \epsilon_b \left[y_i(t) - \mu_t \right] \qquad (24)$$

for the thresholds b_i.

- adaptation of gains, using flow control:

$$a_i(t) = a_i(t-1) \left[1 + \epsilon_a \Delta R_i(t) \right],$$
$$\Delta R_i(t) = R_t^2 |y_i(t-1)|^2 - |x_{r,i}(t)|^2. \qquad (25)$$

- adaptation of gains, with variance control:

$$a_i(t) = a_i(t-1) + \epsilon_a \left[\sigma_{t,i}^2(t) - \left(y_i(t) - \mu_i^y(t) \right)^2 \right] \qquad (26)$$

$$\sigma_{t,i}^2(t) = 1 - \frac{1}{\sqrt{1 + 2R_t^2 y_i(t)^2 + 2\sigma_{ext,i}^2(t)}} \qquad (27)$$

$$\mu_i^y(t) = \mu_i^y(t-1) + \epsilon_\mu \left[y_i(t) - \mu_i^y(t-1) \right] \qquad (28)$$

$$\sigma_{ext,i}^2(t) = \sigma_{ext,i}^2(t-1) + \epsilon_\sigma \left[\left(I_i(t) - \mu_{ext,i}(t) \right)^2 - \sigma_{ext,i}^2(t-1) \right] \qquad (29)$$

$$\mu_{ext,i}(t) = \mu_{ext,i}(t-1) + \epsilon_\mu \left[I_i(t) - \mu_{ext,i}(t-1) \right]. \qquad (30)$$

Note that Equations (28)–(30) have the same mathematical form

$$\langle trail \rangle(t) = \langle trail \rangle(t-1) + \epsilon \left[\langle var \rangle(t) - \langle trail \rangle(t-1) \right]$$

since they only serve as trailing averages that are used in the two main Equations (26) and (27).

For a summary of all model parameters (see **Table 1**).

5.2. Convergence Acceleration for Flow Control

For small values of R_t and weak external input, the average square activities and membrane potentials $y_i^2(t)$ and $x_{t,i}^2(t)$ can

TABLE 1 | Standard values for model parameters.

N	p_r	σ_w	μ_t	ϵ_b	ϵ_a	ϵ_μ	ϵ_σ
500	0.1	1	0.05	10^{-3}	10^{-3}	10^{-4}	10^{-3}

become very small. As a consequence, their difference entering $\Delta R_i(t)$ in (25) also becomes small in absolute value, slowing down the convergence process. To eliminate this effect, we decided to rescale the learning rate by a trailing average of the squared recurrent membrane potential, i.e., $\epsilon_a \to \epsilon_a / \bar{x}_r^2$. The effect of this renormalization is shown in **Figure 7**. Rescaling the learning rate effectively removes the significant rise of convergence times for small σ_{ext} and small R_t.

5.3. Input Protocols
Overall, we examined four distinct input protocols.

- *Homogeneous Gaussian.* Nodes receive inputs $I_i(t)$ that are drawn individually from a Gaussian with vanishing mean and standard deviation σ_{ext}.
- *Heterogeneous Gaussian.* Nodes receive stochastically independent inputs $I_i(t)$ that are drawn from Gaussian distributions with vanishing mean and node specific standard deviations $\sigma_{i,ext}$. The individual $\sigma_{i,ext}$ are normal distributed, as drawn from the positive part of a Gaussian with mean zero and variance σ_{ext}^2.
- *Homogeneous binary.* Sites receive identical inputs $I_i(t) = \sigma_{ext} u(t)$, where $u(t) = \pm 1$ is a binary input sequence.
- *Heterogeneous binary.* We define with

$$I_i = W_i^u u(t), \qquad u_j(t) = \pm 1 \qquad (31)$$

the afferent synaptic weight vector W_i^u, which connects the binary input sequence $u(t)$ to the network. All W_i^u are drawn independently from a Gaussian with mean zero and standard deviation σ_{ext}.

The Gaussian input variant simulates external noise. We used it in particular to test predictions of the theory developed in section 5.6. In order to test the performance of the echo state network with respect to the delayed XOR task, the binary input protocols are employed. A generalization of the here defined protocols to the case of higher-dimensional input signals would be straightforward.
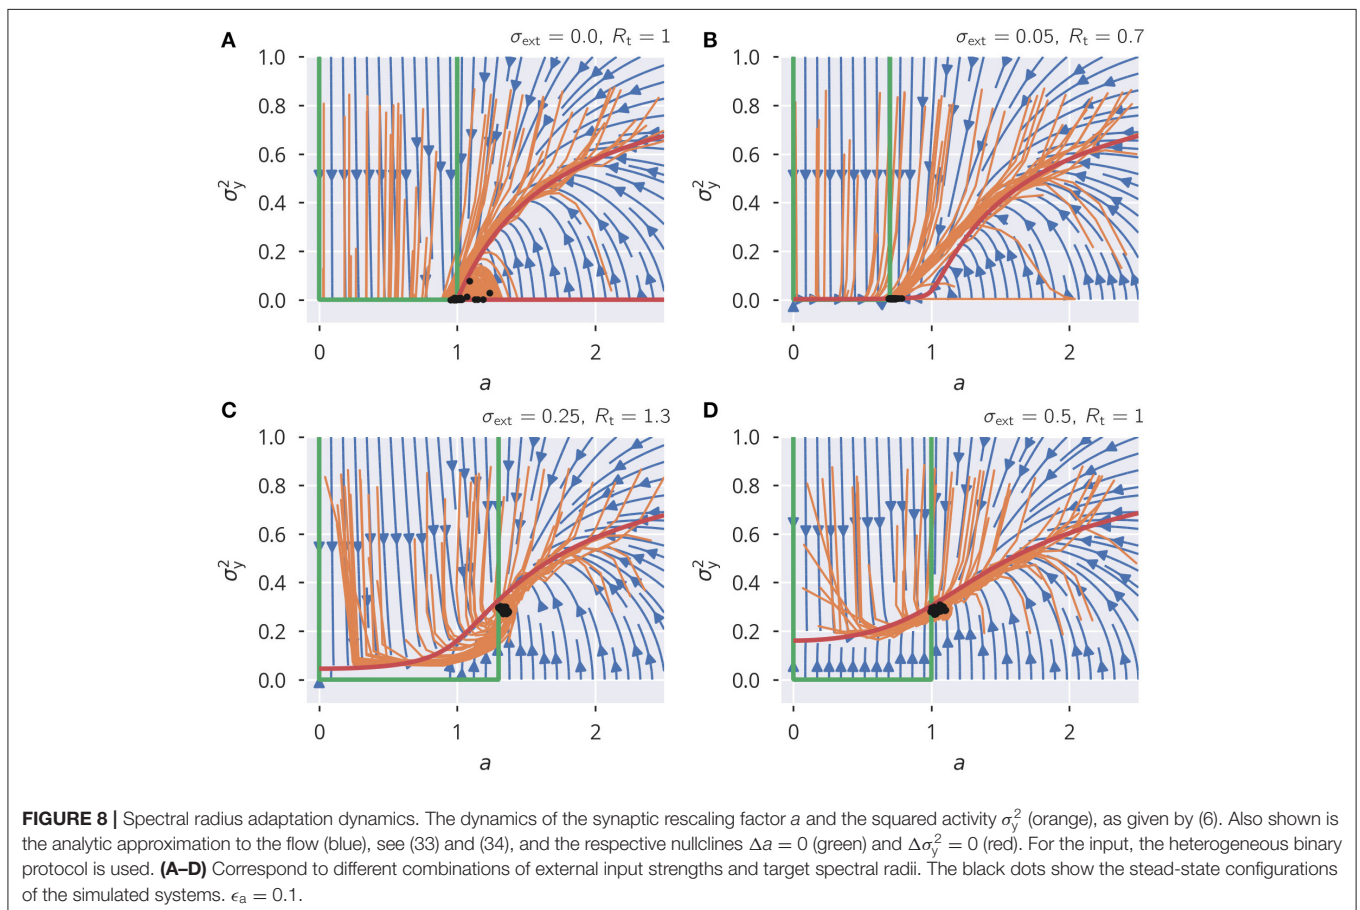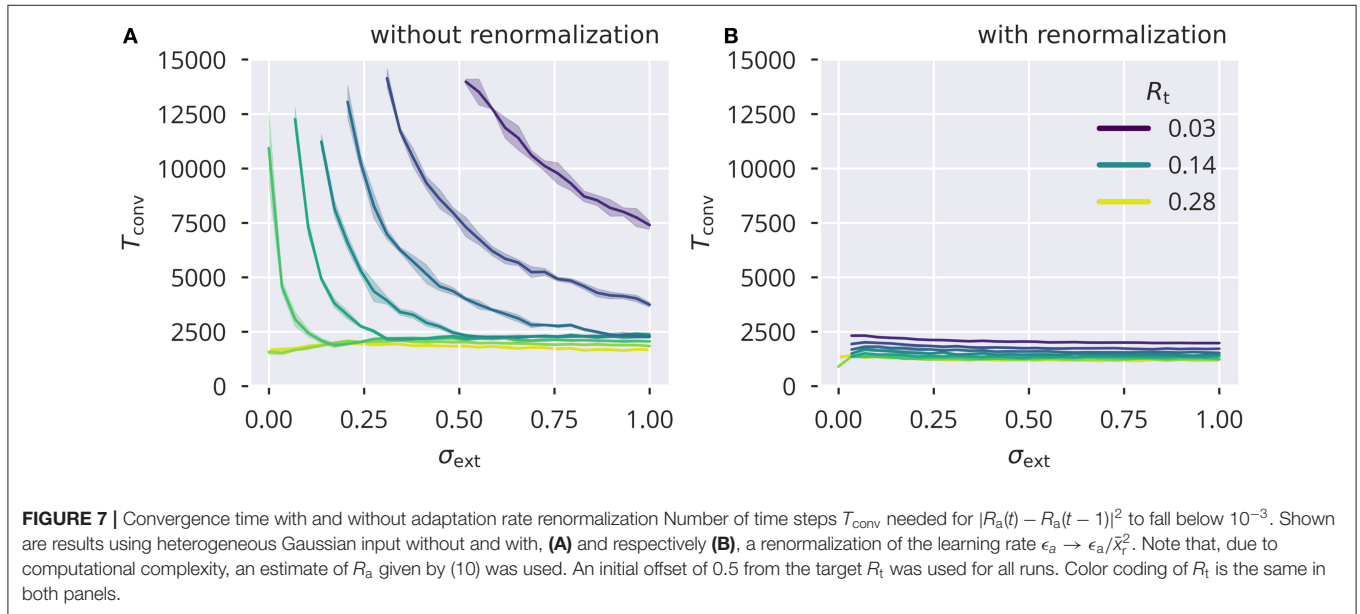
5.4. Spectral Radius Adaptation Dynamics
For an understanding of the spectral radius adaptation dynamics of flow control, it is of interest to examine the effect of using the global adaptation constraint

$$\Delta R_i(t) = \frac{1}{N} \left[R_t^2 ||y(t-1)||^2 - ||x_r(t)||^2 \right] \qquad (32)$$

in (5). The spectral radius condition (12) is then enforced directly, with the consequence that (32) is stable and precise

even in the presence of correlated neural activities (see **Supplementary Figures 5A–C**). This rule, while not biologically plausible, provides an opportunity to examine the dynamical

flow, besides the resulting state. There are two dynamic variables, $a = a_i \; \forall i$, where, for the sake of simplicity, we assumed that all a_i are homogeneous, and the activity

FIGURE 7 | Convergence time with and without adaptation rate renormalization Number of time steps T_{conv} needed for $|R_a(t) - R_a(t-1)|^2$ to fall below 10^{-3}. Shown are results using heterogeneous Gaussian input without and with, **(A)** and respectively **(B)**, a renormalization of the learning rate $\epsilon_a \rightarrow \epsilon_a / \bar{x}_r^2$. Note that, due to computational complexity, an estimate of R_a given by (10) was used. An initial offset of 0.5 from the target R_t was used for all runs. Color coding of R_t is the same in both panels.

FIGURE 8 | Spectral radius adaptation dynamics. The dynamics of the synaptic rescaling factor a and the squared activity σ_y^2 (orange), as given by (6). Also shown is the analytic approximation to the flow (blue), see (33) and (34), and the respective nullclines $\Delta a = 0$ (green) and $\Delta \sigma_y^2 = 0$ (red). For the input, the heterogeneous binary protocol is used. **(A–D)** Correspond to different combinations of external input strengths and target spectral radii. The black dots show the steady-state configurations of the simulated systems. $\epsilon_a = 0.1$.

variance $\sigma_y^2 = ||\mathbf{y}||^2/N$. The evolution of (a, σ_y^2) resulting from the global rule (6) is shown in **Figure 8**. For the flow, $\Delta a = a(t+1) - a(t)$ and $\Delta \sigma_y^2 = \sigma_y^2(t) - \sigma_y^2(t-1)$, the approximation

$$\Delta a = \epsilon_a a \left(R_t^2 - a^2 \sigma_w^2 \right) \sigma_y^2 \tag{33}$$

$$\Delta \sigma_y^2 = 1 - \sigma_y^2 - \frac{1}{\sqrt{1 + 2a^2 \sigma_w^2 \sigma_y^2 + 2\sigma_{ext}}} \tag{34}$$

is obtained. For the scaling factor a, this leads to a fixed point of R_t/σ_w. We used the mean-field approximation for neural variances that is derived in section 5.6. The analytic flow compares well with numerics, as shown in **Figure 8**. For a subcritical rescaling factor a and $\sigma_{ext} = 0$, the system flows toward a line of fixpoints defined by a vanishing σ_y^2 and a finite $a \in [0, 1]$ (see **Figure 8A**). When starting with $a > 0$, the fixpoint is instead $(a, \sigma_y^2) = (1, 0)$. The situation changes qualitatively for finite external inputs, viz when $\sigma_{ext} > 0$, as shown in **Figures 8B–D**. The nullcline $\Delta \sigma_y^2 = 0$ is now continuous and the system flows to the fixed point, as shown in **Figures 8B–D**, with the value of σ_y^2 being determined by the intersection of the two nullclines. In addition, we also varied the target spectral radius (see **Figures 8B,C**). This caused a slight mismatch between the flow of the simulated systems and the analytic flow. It should be noted, however, that this is to be expected anyhow since we used an approximation for the neural variances, again (see section 5.6).

This analysis shows that external input is necessary for a robust flow toward the desired spectral weight, the reason being that the dynamics dies out before the spectral weight can be adapted when the isolated systems starts in the subcritical regime.

5.5. Extended Theory of Flow Control for Independent Neural Activity

We would like to show that the stationarity condition in Equation (12) results in the correct spectral radius, under the special case of independently identically distributed neural activities with zero mean.

We start with Equation (12) as a stationarity condition for a given R_t:

$$\left\langle ||\mathbf{x}_r(t)||^2 \right\rangle_t \stackrel{!}{=} R_t^2 \left\langle ||\mathbf{y}(t-1)||^2 \right\rangle_t. \tag{35}$$

We can express the left side of the equation as

$$E \left[\mathbf{y}^\dagger(t) \widehat{W}_a^\dagger \widehat{W}_a \mathbf{y}(t) \right]_t. \tag{36}$$

We define $\widehat{U}_a == \widehat{W}_a^\dagger \widehat{W}_a$ with $\{\sigma_k^2\}$ being the set of eigenvalues, which are also the squared singular values of $\widehat{W}_a$, and $\{\mathbf{u}_k\}$ the respective set of orthonormal (column) eigenvectors. We insert

the identity $\sum_{k=1}^N \mathbf{u}_k \mathbf{u}_k^\dagger$ and find

$$E \left[\mathbf{y}^\dagger(t) \widehat{U}_a \sum_{k=1}^N \mathbf{u}_k \mathbf{u}_k^\dagger \mathbf{y}(t) \right]_t \tag{37}$$

$$= E \left[\sum_{k=1}^N \sigma_k^2 \mathbf{y}^\dagger(t) \mathbf{u}_k \mathbf{u}_k^\dagger \mathbf{y}(t) \right]_t \tag{38}$$

$$= \sum_{k=1}^N \sigma_k^2 \mathbf{u}_k^\dagger E \left[\mathbf{y}(t) \mathbf{y}^\dagger(t) \right]_t \mathbf{u}_k \tag{39}$$

$$= \sum_{k=1}^N \sigma_k^2 \mathbf{u}_k^\dagger \widehat{C}_{yy} \mathbf{u}_k \tag{40}$$

$$= \text{Tr} \left(\widehat{D}_{\sigma^2} \widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u \right). \tag{41}$$

Given zero mean neural activity, $\widehat{C}_{yy} = E[\mathbf{y}(t)\mathbf{y}^\dagger(t)]_t$ is the covariance matrix of neural activities. $\widehat{D}_{\sigma^2}$ is a diagonal matrix holding the $\{\sigma_k^2\}$ and $\widehat{S}_u$ is a unitary matrix whose columns are $\{\mathbf{u}_k\}$. $\widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u$ is expressing $\widehat{C}_{yy}$ in the diagonal basis of $\widehat{U}_a$.

Including the right hand side of (35), we get

$$\text{Tr} \left(\widehat{D}_{\sigma^2} \widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u \right) = R_t^2 \text{Tr} \left(\widehat{C}_{yy} \right). \tag{42}$$

However, since the trace is invariant under a change of basis, we find

$$\text{Tr} \left(\widehat{D}_{\sigma^2} \widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u \right) = R_t^2 \text{Tr} \left(\widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u \right). \tag{43}$$

Defining $\widehat{C}^u == \widehat{S}_u^\dagger \widehat{C}_{yy} \widehat{S}_u$, we get

$$\sum_{k=1}^N \sigma_k^2 C_{kk}^u = R_t^2 \sum_{k=1}^N C_{kk}^u. \tag{44}$$

If we assume that the node activities are independently identically distributed with zero mean, we get $(\widehat{C}_{yy})_{ij} = (\widehat{C}^u)_{ij} = \langle y^2 \rangle_t \delta_{ij}$. In this case, which was also laid out in section 2.6, the equation reduces to

$$\sum_{k=1}^N \sigma_k^2 = R_t^2 N. \tag{45}$$

The Frobenius norm of a square Matrix $\widehat{A}$ is given by $||\widehat{A}||_F^2 \equiv \sum_{i,j} \widehat{A}_{ij}^2$. Furthermore, the Frobenius norm is linked to the singular values via $||\widehat{A}||_F^2 = \sum_k \sigma_k^2(\widehat{A})$ (Sengupta and Mitra, 1999; Shen, 2001). This allows us to state

$$\sum_{i,j} \left(\widehat{W}_a \right)_{ij}^2 = R_t^2 N \tag{46}$$

which, by using (10), gives

$$R_a^2 = R_t^2. \tag{47}$$

A slightly less restrictive case is that of uncorrelated but inhomogeneous activity, that is $(\widehat{C}_{yy})_{ij} = \langle y_i^2 \rangle_t \delta_{ij}$. The statistical

properties of the diagonal elements C_{kk}^u then determine to which degree one can still claim that Equation (44) leads to Equation (45). **Supplementary Figure 11** shows an example of a randomly generated realization of $(\widehat{C}_{yy})_{ij} = \langle y_i^2 \rangle_t$ and the resulting diagonal elements of $\widehat{C}^u$, where the corresponding orthonormal basis $\widehat{S}_u$ was generated from the SVD of a random Gaussian matrix. As one can see, the basis transformation has a strong smoothing effect on the diagonal entries, while the mean over the diagonal elements is preserved. Note that this effect was not disturbed by introducing random row-wise multiplications to the random matrix from which the orthonormal basis was derived. The smoothing of the diagonal entries allows us to state that $C_{kk}^u \cong \langle y^2 \rangle$ is a very good approximation in the case considered, which therefore reduces (44) to the homogeneous case previously described. We can conclude that the adaptation mechanism also gives the desired spectral radius under uncorrelated inhomogeneous activity.

In the most general case, we can still state that if C_{kk}^u and σ_k^2 are uncorrelated, for large N, Equation (44) will tend toward

$$N \langle \sigma^2 \rangle \langle C^u \rangle = N R_t^2 \langle C^u \rangle \tag{48}$$

which would also lead to Equation (45). However, we can not generally guarantee statistical independence since the recurrent contribution on neural activities and the resulting entries of $\widehat{C}_{yy}$ and thus also C_{kk}^u are linked to $\widehat{S}$ and σ_k^2, being the SVD of the recurrent weight matrix.

5.6. Mean Field Theory for Echo State Layers

In the following, we deduce analytic expressions allowing to examine the state of echo-state layers subject to a continuous timeline of inputs. Our approach is similar to the one presented by Massar and Massar (2013).

The recurrent part of the input x_i received by a neuron is a superposition of $N p_r$ terms, which are assumed here to be uncorrelated. Given this assumption, the self-consistency equations

$$\sigma_{y,i}^2 = \int_{-\infty}^{\infty} dx \, \tanh^2(x) N_{\mu_i, \sigma_i}(x) - \mu_{y,i}^2 \tag{49}$$

$$\mu_{y,i} = \int_{-\infty}^{\infty} dx \, \tanh(x) N_{\mu_i, \sigma_i}(x) \tag{50}$$

$$\sigma_i^2 = a_i^2 \sigma_w^2 \langle \sigma_{y,j}^2 \rangle_j + \sigma_{ext,i}^2, \qquad \mu_i = \mu_{ext,i} - b_i \tag{51}$$

determine the properties of the stationary state. We recall that σ_w parameterizes the distribution of bare synaptic weights via (23). The general expressions (49) and (50) hold for all neurons, with the site-dependency entering exclusively via a_i, b_i, $\sigma_{ext,i}$ and $\mu_{ext,i}$, as in (51), with the latter characterizing the standard deviation and the mean of the input. Here, $a_i^2 \sigma_w^2 \sigma_y^2$ is the variance of the recurrent contribution to the membrane potential, x, and σ^2 the respective total variance. The membrane potential

is Gaussian distributed, as $N_{\mu,\sigma}(x)$, with mean μ and variance σ^2, which are both to be determined self-consistently. Variances are additive for stochastically independent processes, which has been assumed in (51) to be the case for recurrent activities and the external inputs. The average value for the mean neural activity is μ_i.

For a given set of a_i, $\sigma_{ext,i}$, and b_i, the means and variances of neural activities, $\sigma_{y,i}^2$ and $\mu_{y,i}$, follow implicitly.

We compared the numerically determined solutions of (49) and (50) against full network simulations using, as throughout this study, $N = 500$, $p_r = 0.1$, $\sigma_w = 1$, $\mu_t = 0.05$. In **Figure 9**, the spectral radius R_a is given for the four input protocols defined in section 5.3. The identical ensemble of input standard deviations $\sigma_{ext,i}$ enters both theory and simulations. Theory and simulations are in good accordance for vanishing input. Here, the reason is that finite activity levels are sustained in an autonomous random neural network when the ongoing dynamics is chaotic and hence decorrelated. For reduced activity levels, viz for small variances σ_y^2, the convergence of the network dynamics is comparatively slow, which leads to a certain discrepancy with the analytic prediction (see **Figure 9**).

5.6.1. Gaussian Approximation
The integral occurring in the self-consistency condition (49) can be evaluated explicitly when a tractable approximation to the squared transfer function $\tanh^2()$ is available. A polynomial approximation would capture the leading behavior close to the origin, however without accounting for the fact that $\tanh^2()$ converges to unity for large absolute values of the membrane potential. Alternatively, an approximation incorporating both conditions, the correct second-order scaling for small, and the correct convergence for large arguments, is given by the Gaussian approximation
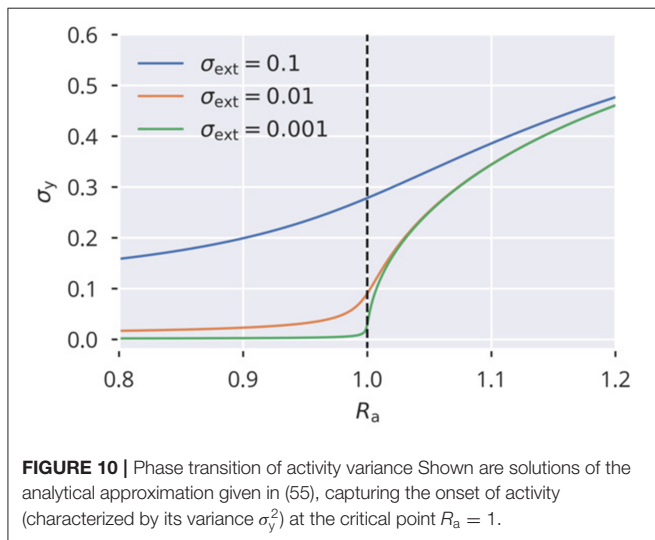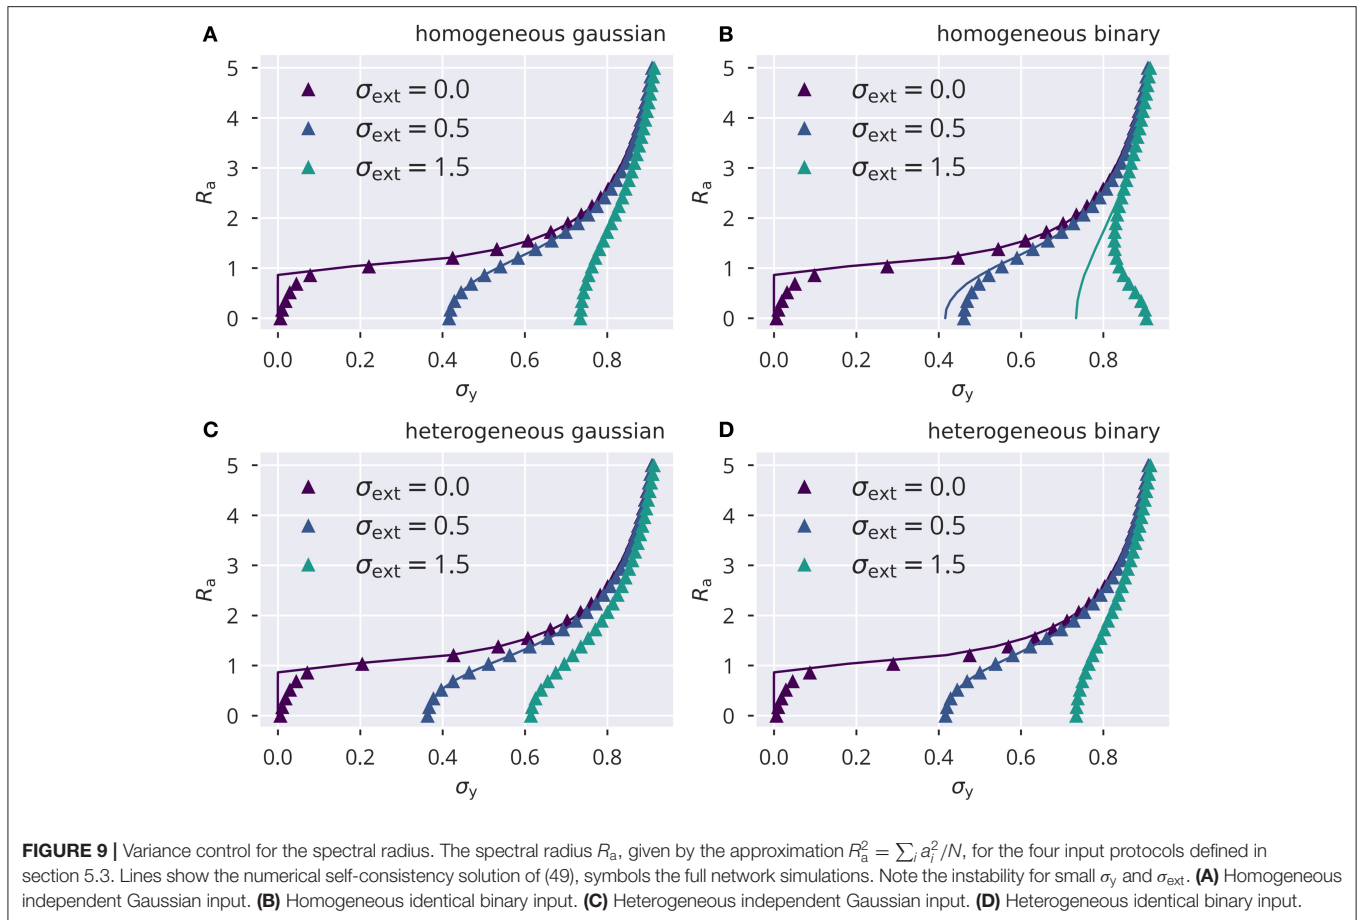
$$\tanh^2(x) \approx 1 - \exp\left(-x^2\right). \tag{52}$$

With this approximation the integral in (49) can be evaluated explicitly. The result is

$$\frac{1}{1 - \sigma_y^2 - \mu_y^2} = \sqrt{1 + 2\sigma^2} / \exp\left(-\mu^2 / \left(1 + 2\sigma^2\right)\right) \tag{53}$$

$$= \sqrt{1 + 2a^2 \sigma_w^2 \sigma_y^2 + 2\sigma_{ext}^2} /$$
$$\exp\left(-\mu^2 / \left(1 + 2a^2 \sigma_w^2 \sigma_y^2 + 2\sigma_{ext}^2\right)\right).$$

Assuming that $\mu \approx 0$ and $\mu_y \approx 0$, inverting the first equation yields a relatively simple analytic approximation for the variance self-consistency equation:

$$\sigma_y^2 = 1 - \frac{1}{\sqrt{1 + 2a^2 \sigma_w^2 \sigma_y^2 + 2\sigma_{ext}^2}}. \tag{54}$$

This equation was then used for the approximate update rule in (8) and (34).

FIGURE 9 | Variance control for the spectral radius. The spectral radius R_a, given by the approximation $R_a^2 = \sum_i a_i^2/N$, for the four input protocols defined in section 5.3. Lines show the numerical self-consistency solution of (49), symbols the full network simulations. Note the instability for small σ_y and σ_{ext}. **(A)** Homogeneous independent Gaussian input. **(B)** Homogeneous identical binary input. **(C)** Heterogeneous independent Gaussian input. **(D)** Heterogeneous identical binary input.

FIGURE 10 | Phase transition of activity variance Shown are solutions of the analytical approximation given in (55), capturing the onset of activity (characterized by its variance σ_y^2) at the critical point $R_a = 1$.

Alternatively, we can write (54) as a self-consistency equation between σ_y^2, σ_{ext}^2 $a^2\sigma_w^2 = R_a^2$, describing a phase transition at $R_a = 1$:

$$2R_a^2\sigma_y^2\left(1-\sigma_y^2\right)^2 = 1 - \left(1 + 2\sigma_{ext}^2\right)\left(1-\sigma_y^2\right)^2. \quad (55)$$

See **Figure 10** for solutions of (55) for different values of σ_{ext}^2. Note that for vanishing external driving and values of R_a above but close to the critical point, the standard deviation σ_y scales with $\sigma_y \propto (R_a - 1)^{1/2}$, which is the typical critical exponent for the order parameter in classical Landau theory of second-order phase transitions (Gros, 2008, p. 169). If combined with a slow homeostatic process, flow or variance control in our case, this constitutes a system with an absorbing phase transition (Gros, 2008, p. 182–183), settling at the critical point $R_a = 1$. This phase transition can also be observed in **Figure 9** for $\sigma_{ext} = 0$ as a sharp onset in σ_y.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: The datasets generated for this study can be found at: https://itp.uni-frankfurt.de/~fschubert/data_esn_frontiers. Simulation and plotting code is available at: https://github.com/FabianSchubert/ESN_Frontiers.

AUTHOR CONTRIBUTIONS

FS and CG contributed equally to the writing and review of the manuscript. FS provided the code, ran the simulations, and

prepared the figures. All authors contributed to the article and approved the submitted version.

This manuscript was published as a pre-print on biorxiv (Schubert and Gros, 2020).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fncom. 2021.587721/full#supplementary-material

REFERENCES

Abbott, L. F., and LeMasson, G. (1993). Analysis of neuron models with dynamically regulated conductances. *Neural Comput.* 5, 823–842. doi: 10.1162/neco.1993.5.6.823

Barral, J., and D'Reyes, A. (2016). Synaptic scaling rule preserves excitatory-inhibitory balance and salient neuronal network dynamics. *Nat. Neurosci.* 19, 1690–1696. doi: 10.1038/nn.4415

Bell, A. J., and Sejnowski, T. J. (1995). An Information-maximisation approach to blind separation and blind deconvolution. *Neural Comput.* 7, 1129–1159. doi: 10.1162/neco.1995.7.6.1129

Bernacchia, A., and Wang, X. J. (2013). Decorrelation by recurrent inhibition in heterogeneous neural circuits. *Neural Comput.* 25, 1732–1767. doi: 10.1162/NECO_a_00451

Binzegger, T., Douglas, R. J., and Martin, K. A. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* 24, 8441–8453. doi: 10.1523/JNEUROSCI.1400-04.2004

Boedecker, J., Obst, O., Mayer, N. M., and Asada, M. (2009). Initialization and self-organized optimization of recurrent neural network connectivity. *HFSP J.* 3, 340–349. doi: 10.2976/1.3240502

Borde, M., Cazalets, J. R., and Buno, W. (1995). Activity-dependent response depression in rat hippocampal CA1 pyramidal neurons *in vitro. J. Neurophysiol.* 74, 1714–1729. doi: 10.1152/jn.1995.74.4.1714

Cannon, J., and Miller, P. (2017). Stable control of firing rate mean and variance by dual homeostatic mechanisms. *J. Math. Neurosci.* 7:1. doi: 10.1186/s13408-017-0043-7

Chen, N., Chen, X., and Wang, J. H. (2008). Homeostasis established by coordination of subcellular compartment plasticity improves spike encoding. *J. Cell Sci.* 121, 2961–2971. doi: 10.1242/jcs.022368

Echeveste, R., Aitchison, L., Hennequin, G., and Lengyel, M. (2020). Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nat. Neurosci.* 1138–1149. doi: 10.1038/s41593-020-0671-1

Echeveste, R., and Gros, C. (2016). Drifting states and synchronization induced chaos in autonomous networks of excitable neurons. *Front. Comput. Neurosci.* 10:98. doi: 10.3389/fncom.2016.00098

Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., and Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* 327, 584–587. doi: 10.1126/science.1179867

Effenberger, F., and Jost, J. (2015). Self-organization in balanced state networks by STDP and homeostatic plasticity. *PLoS Comput. Biol.* 11:1004420. doi: 10.1371/journal.pcbi.1004420

Enel, P., Procyk, E., Quilodran, R., and Dominey, P. F. (2016). Reservoir computing properties of neural dynamics in prefrontal cortex. *PLOS Comput. Biol.* 12:e1004967. doi: 10.1371/journal.pcbi.1004967

Földiak, P. (1990). Forming sparse representations by local anti-Hebbian learning. *Biol. Cybern.* 64, 165–170. doi: 10.1007/BF02331346

Franklin, J. L., Fickbohm, D. J., and Willard, A. L. (1992). Long-term regulation of neuronal calcium currents by prolonged changes of membrane potential. *J. Neurosci.* 12, 1726–1735. doi: 10.1523/JNEUROSCI.12-05-01726.1992

Gallicchio, C., and Micheli, A. (2017). Echo state property of deep reservoir computing networks. *Cogn. Comput.* 9, 337–350. doi: 10.1007/s12559-017-9461-9

Gros, C. (2008). *Complex and Adaptive Dynamical Systems, a Primer.* Berlin; Heidelberg: Springer. doi: 10.1007/978-3-540-71874-1

Hinaut, X., Lance, F., Droin, C., Petit, M., Pointeau, G., and Dominey, P. F. (2015). Corticostriatal response selection in sentence production: insights from neural network simulation with reservoir computing. *Brain Lang.* 150, 54–68. doi: 10.1016/j.bandl.2015.08.002

Jaeger, H. (2002). *Short Term Memory in Echo State Networks.* Technical report, Fraunhofer Institute for Autonomous Intelligent Systems.

Lazar, A., Pipa, G., and Triesch, J. (2009). SORN: a self-organizing recurrent neural network. *Front. Comput. Neurosci.* 3:2009. doi: 10.3389/neuro.10.023.2009

Livi, L., Bianchi, F. M., and Alippi, C. (2018). Determination of the edge of criticality in echo state networks through Fisher information maximization. *IEEE Trans. Neural Netw. Learn. Syst.* 29, 706–717. doi: 10.1109/TNNLS.2016.2644268

Maass, W., and Markram, H. (2004). On the computational power of circuits of spiking neurons. *J. Comput. Syst. Sci.* 69, 593–616. doi: 10.1016/j.jcss.2004.04.001

Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 2531–2560. doi: 10.1162/089976602760407955

Marčenko, V. A., and Pastur, L. A. (1967). Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sbornik* 1, 457–483. doi: 10.1070/SM1967v001n04ABEH001994

Massar, M., and Massar, S. (2013). Mean-field theory of echo state networks. *Phys. Rev. E* 87:42809. doi: 10.1103/PhysRevE.87.042809

Miner, D., and Triesch, J. (2016). Plasticity-driven self-organization under topological constraints accounts for non-random features of cortical synaptic wiring. *PLoS Comput. Biol.* 12:e1004759. doi: 10.1371/journal.pcbi.1004759

Narayanan, R., and Johnston, D. (2012). Functional maps within a single neuron. *J. Neurophysiol.* 108, 2343–2351. doi: 10.1152/jn.00530.2012

Nikolić, D., Häusler, S., Singer, W., and Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol.* 7:e1000260. doi: 10.1371/journal.pbio.1000260

Rajan, K., and Abbott, L. F. (2006). Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.* 97:188104. doi: 10.1103/PhysRevLett.97.188104

Remme, M. W., and Wadman, W. J. (2012). Homeostatic scaling of excitability in recurrent neural networks. *PLoS Comput. Biol.* 8:1002494. doi: 10.1371/journal.pcbi.1002494

Salinas, E., and Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* 2, 539–550. doi: 10.1038/35086012

Schrauwen, B., Wardermann, M., Verstraeten, D., Steil, J. J., and Stroobandt, D. (2008). Improving reservoirs using intrinsic plasticity. *Neurocomputing* 71, 1159–1171. doi: 10.1016/j.neucom.2007.12.020

Schubert, F., and Gros, C. (2020). Local homeostatic regulation of the spectral radius of echo-state networks. *bioRxiv [Preprint].* doi: 10.1101/2020.07.21.213660

Sengupta, A. M., and Mitra, P. P. (1999). Distributions of singular values for some random matrices. *Phys. Rev. E* 60:3389. doi: 10.1103/PhysRevE.60.3389

Shen, J. (2001). On the singular values of Gaussian random matrices. *Linear Algeb. Appl.* 326, 1–14. doi: 10.1016/S0024-3795(00)00322-0

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nat. Rev. Neurosci.* 9, 206–221. doi: 10.1038/nrn2286

Steil, J. J. (2007). Online reservoir adaptation by intrinsic plasticity for backpropagation-decorrelation and echo state learning. *Neural Netw.* 20, 353–364. doi: 10.1016/j.neunet.2007.04.011

Stellwagen, D., and Malenka, R. C. (2006). Synaptic scaling mediated by glial TNF-α. *Nature* 440, 1054–1059. doi: 10.1038/nature04671

Sweeney, Y., Hellgren Kotaleski, J., and Hennig, M. H. (2015). A diffusive homeostatic signal maintains neural heterogeneity and responsiveness in cortical networks. *PLoS Comput. Biol.* 11:e1004389. doi: 10.1371/journal.pcbi.1004389

Tao, T., and Vu, V. (2008). Random matrices: the circular law. *Commun. Contemp. Math.* 10, 261–307. doi: 10.1142/S0219199708002788

Tetzlaff, C. (2011). Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Front. Comput. Neurosci.* 5:47. doi: 10.3389/fncom.2011.00047

Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLoS Comput. Biol.* 8:1002596. doi: 10.1371/journal.pcbi.1002596

Toyoizumi, T., Kaneko, M., Stryker, M. P., and Miller, K. D. (2014). Modeling the dynamic interaction of hebbian and homeostatic plasticity. *Neuron* 84, 497–510. doi: 10.1016/j.neuron.2014.09.036

Trapp, P., Echeveste, R., and Gros, C. (2018). Ei balance emerges naturally from continuous hebbian learning in autonomous neural networks. *Sci. Rep.* 8, 1–12. doi: 10.1038/s41598-018-27099-5

Triesch, J. (2007). Synergies between intrinsic and synaptic plasticity mechanisms. *Neural Comput.* 19, 885–909. doi: 10.1162/neco.2007.19.4.885

Turrigiano, G. G. (2008). The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* 135, 422–435. doi: 10.1016/j.cell.2008.10.008

Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896. doi: 10.1038/36103

Turrigiano, G. G., and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* 10, 358–364. doi: 10.1016/S0959-4388(00)00091-X

Usrey, W. M., and Reid, R. C. (1999). Synchronous activity in the visual system. *Annu. Rev. Physiol.* 61, 435–456. doi: 10.1146/annurev.physiol.61.1.435

van Rossum, M. C. W., Bi, G. Q., and Turrigiano, G. G. (2000). Stable hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821. doi: 10.1523/JNEUROSCI.20-23-08812.2000

Van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 10, 1321–1371. doi: 10.1162/089976698300017214

Wernecke, H., Sándor, B., and Gros, C. (2019). Chaos in time delay systems, an educational review. *arXiv preprint arXiv:1901.04826.* doi: 10.1016/j.physrep.2019.08.001

Wick, S. D., Wiechert, M. T., Friedrich, R. W., and Riecke, H. (2010). Pattern orthogonalization via channel decorrelation by adaptive networks. *J. Comput. Neurosci.* 28, 29–45. doi: 10.1007/s10827-009-0183-1

Zenke, F., Hennequin, G., and Gerstner, W. (2013). Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput. Biol.* 9:1003330. doi: 10.1371/journal.pcbi.1003330