Efficiently encoding noisy inputs in a recurrent neural network with adaptive representational dimensionality

Lyndon Duong[1], Teddy Yerxa[2], Eero Simoncelli[1,2], David Lipshutz[1,3,4]

[1]Center for Computational Neuroscience, Flatiron Institute, New York, NY
[2]Center for Neural Science, New York University, New York, NY
[3]Department of Neuroscience, Baylor College of Medicine, Houston, TX
[4]Neuroengineering Initiative, Rice University, Houston, TX

Efficient coding & representational dimensionality

Efficient coding hypothesis: early sensory systems are optimized to transmit as much information as possible given limited available resources[1].

Trade-off:
- High-dimensional, decorrelated codes are optimal for high SNR inputs
- Low-dimensional, redundant codes are optimal for low SNR inputs[2,3].

Question: what circuit mechanisms could *flexibly* modulate dimensionality?

Modulating dimensionality via adaptive recurrence

Contribution: we propose an RNN model in which recurrence sign and strength, parameterized by α, flexibly adjusts the dimensionality of the neural population code.

Our starting point is a novel unsupervised learning objective that balances (a) input reconstruction, and (b) a regularizer based on the participation ratio[4].

Optimization of this objective can be achieved in an RNN with adaptive sign and strength of recurrent weights.

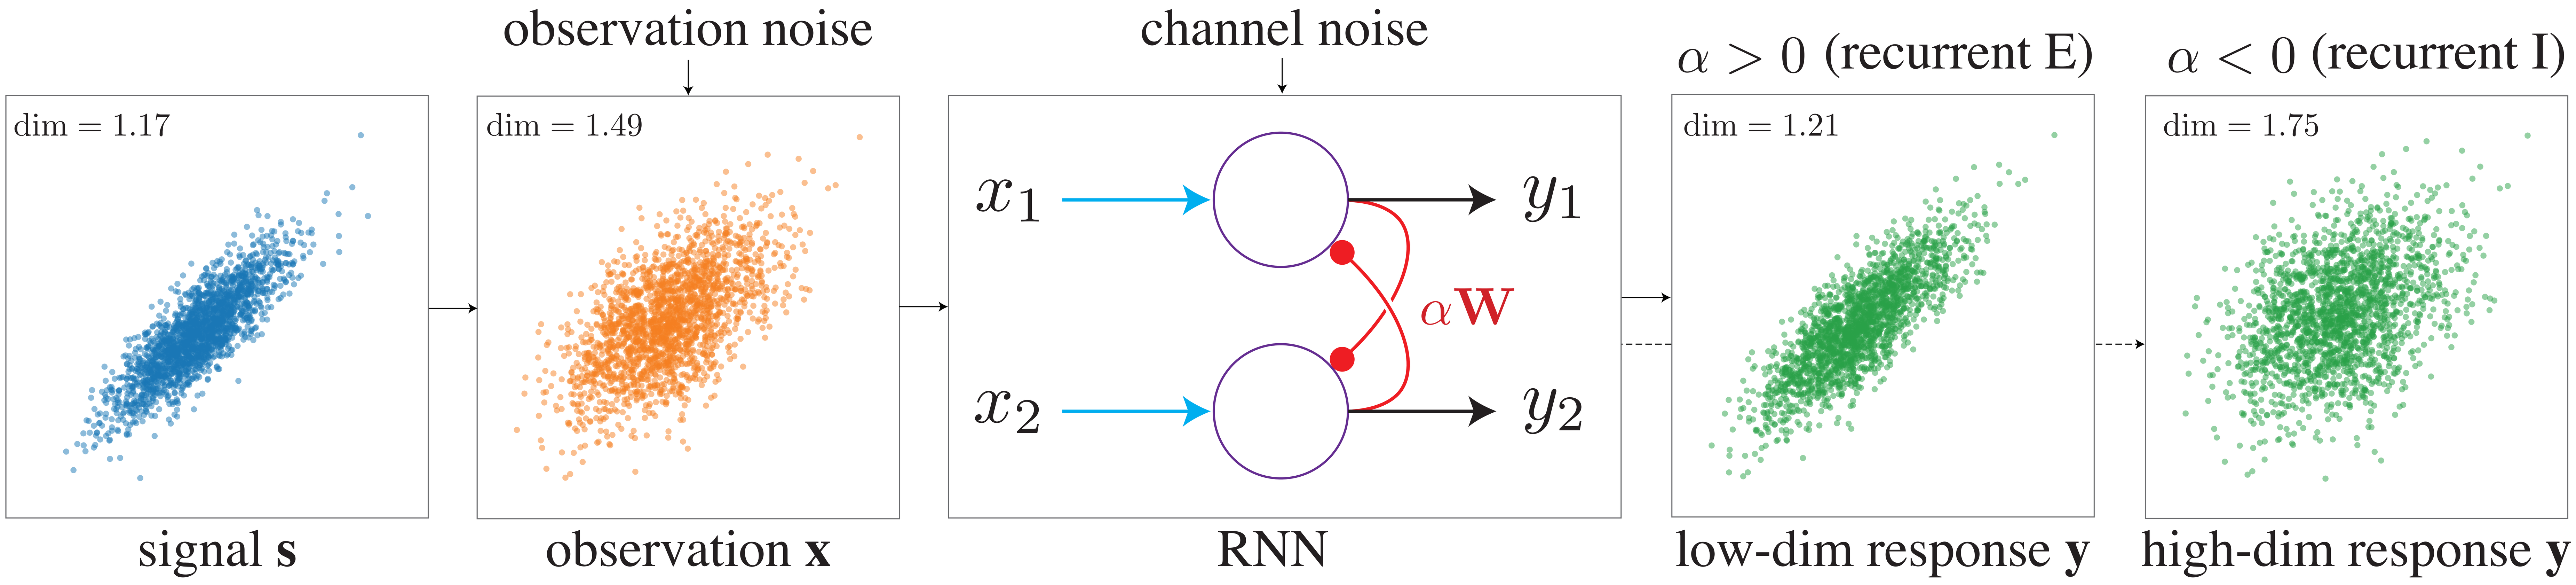

1. RNN with learned weights

Objective: balances reconstruction with response dimimension (participation ratio[4]):

$$\min_{\mathbf{y}} \left\{ \mathbb{E}\left[\|\mathbf{x}-\mathbf{y}\|^2\right] - \frac{\alpha}{2}\frac{N^2}{\dim(\mathbf{y})} \right\} \quad \text{s.t.} \quad \mathrm{Var}(y_i)=1$$

reconstruction — regularizer — capacity constraint

Circuit algorithm (see below for derivation)

1. Neural dynamics:

$$\dot{\mathbf{y}} = \mathbf{x} + \alpha\mathbf{W}\mathbf{y} - \boldsymbol{\lambda}\circ\mathbf{y}$$

input drive — recurrent inputs — neural leak

2. Hebbian plasticity:

$$\Delta W_{ij} = \eta_W(y_i y_j - W_{ij}) \qquad \Delta\lambda_i = \eta_\lambda(y_i^2 - 1)$$

Simulations. Gaussian **s** with $1/f^2$ spectrum corrupted by noise with variance σ_x^2.

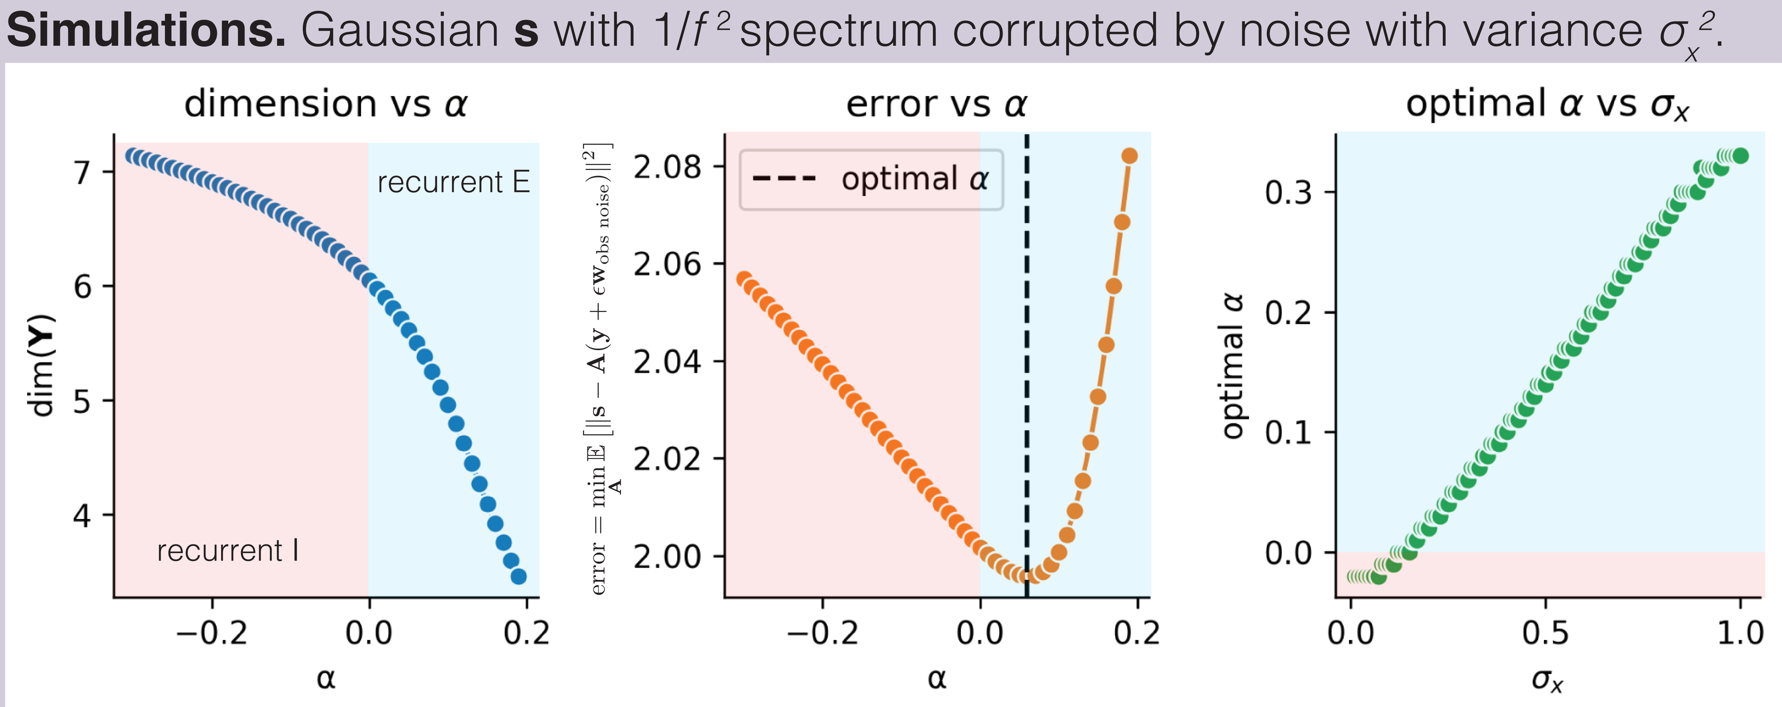

Take away. The "optimal" α depends on the input SNR.

Hypothesis: input SNR is estimated downstream of the RNN (see, e.g., ref. 6) and α represents a top-down input that is a function of SNR.

Algorithm derivation. First, rewrite the participation ratio:

$$\frac{N^2}{\dim(\mathbf{y})} \overset{\text{def of PR, } \mathbf{C}=\mathrm{cov}(\mathbf{y})}{=} \frac{N^2\,\mathrm{Tr}(\mathbf{C}^2)}{\mathrm{Tr}(\mathbf{C})^2} \overset{\text{capacity constraint}}{=} \mathrm{Tr}(\mathbf{C}^2) \overset{\text{Legendre transform}}{=} \max_{\mathbf{W}} \mathrm{Tr}(2\mathbf{C}\mathbf{W}-\mathbf{W}^2) \overset{\mathbf{C}=\langle\mathbf{y}\mathbf{y}^\top\rangle \text{ for centered } \mathbf{y}}{=} \max_{\mathbf{W}}\{2\mathbb{E}[\mathbf{y}^\top\mathbf{W}\mathbf{y}] - \|\mathbf{W}\|_F^2\}$$

Substitute into the objective, add Lagrange multipliers:

$$\operatorname*{opt}_{\mathbf{W}} \max_{\boldsymbol{\lambda}} \left\{ \mathbb{E}\left[\min_{\mathbf{y}} E_{\text{Hopfield}}(\alpha\mathbf{W},\boldsymbol{\lambda},\mathbf{x},\mathbf{y})\right] + \frac{\alpha}{2}\|\mathbf{W}\|_F^2 - \|\boldsymbol{\lambda}\|_1 \right\}$$

where E_{Hopfield} is the Hopfield energy function when activations are linear[5]:

$$E_{\text{Hopfield}}(\mathbf{W},\boldsymbol{\lambda},\mathbf{x},\mathbf{y}) = -2\mathbf{x}\cdot\mathbf{y} - \mathbf{y}^\top\mathbf{W}\mathbf{y} + \mathbf{y}^\top\mathrm{diag}(\boldsymbol{\lambda})\mathbf{y}$$

2. RNN with fixed weights

Can the circuit rapidly adapt to input SNR without learning weights **W**?

Objective: balance reconstruction with regularizer (matrix **W** is fixed)

$$\min_{\mathbf{y}} \mathbb{E}\left[\|\mathbf{x}-\mathbf{y}\|^2 - \alpha\mathbf{y}^\top\mathbf{W}\mathbf{y}\right] \quad \text{s.t.} \quad \mathrm{Var}(y_i)=1$$

reconstruction — regularizer — capacity constraint

Circuit algorithm. Neural dynamics and leak update are identical to the case of learned weights (panel 1), but recurrent weights **W** are fixed.

Simulations. Fix weights: $\mathbf{W}^2 = \langle\mathbf{s}\mathbf{s}^\top\rangle$ (whitens responses at high SNR and $\alpha << 0$)

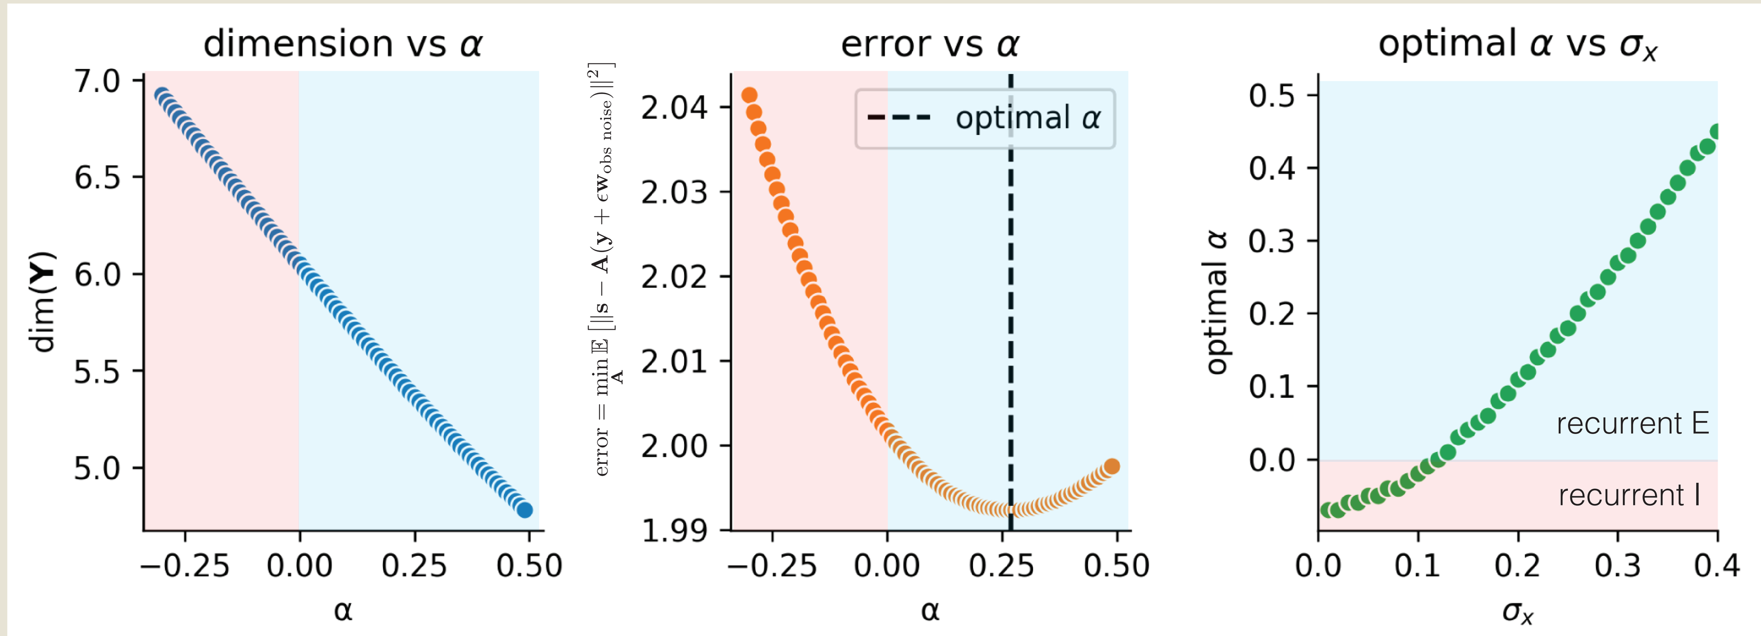

Take away. An RNN can adapt its response dimensionality by rapidly modulating the sign and strength of structured recurrence. But this violates Dale's law...

3. E-I RNN

The RNN can be restructured to satisfy Dale's law.

Objective: As in panel 1, but now

$$\alpha = \beta - \gamma$$

where $\beta > 0$ ($\gamma > 0$) controls recurrent E (I).

Circuit algorithm

1. Neural dynamics:

$$\dot{\mathbf{y}} = \mathbf{x} + \beta\mathbf{W}^{\text{EE}}\mathbf{y} - \gamma\mathbf{W}^{\text{IE}}\mathbf{z} - \boldsymbol{\lambda}\circ\mathbf{y}$$
$$\dot{\mathbf{z}} = \mathbf{W}^{\text{EI}}\mathbf{y} - \mathbf{W}^{\text{II}}\mathbf{z}$$

2. Hebbian plasticity:

$$\Delta W_{ij}^{\text{EE}} = \eta_W(y_i y_j - W_{ij}^{\text{EE}}) \quad \Delta W_{ij}^{\text{EI}} = \eta_W(z_i y_j - W_{ij}^{\text{EI}})$$
$$\Delta W_{ij}^{\text{II}} = \eta_W(z_i z_j - W_{ij}^{\text{II}}) \quad \Delta W_{ij}^{\text{IE}} = \eta_W(y_i z_j - W_{ij}^{\text{IE}})$$

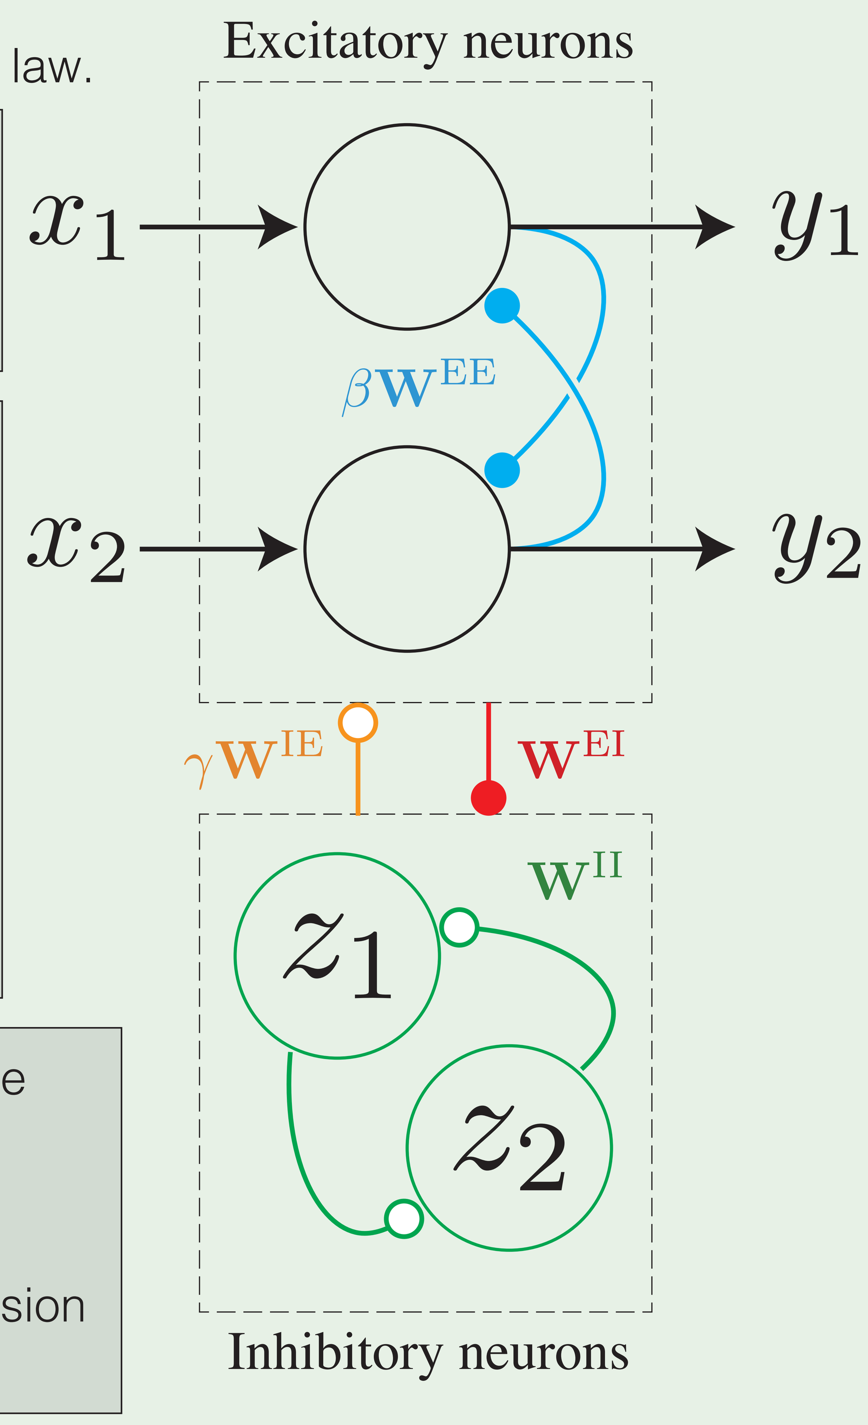

Take away. "E-I balance" $\alpha = \beta - \gamma$ controls the dimensionality of the neural code.

Prediction. Adaptive E-I balance is a neural substrate for maximizing information transmission across varying SNR regimes.

References

[1] Barlow 1961
[2] Doi & Lewicki 2007
[3] Tkačik *et al.* 2010
[4] Recanatesi *et al.* 2019
[5] Hopfield 1984
[6] Kiani and Shadlen 2009