Recurrent Neural Networks

Instructor: Mark Kramer

Parts motivated from Schmidt, Recurrent Neural Networks (RNNs): A gentle Introduction and Overview, 2019

Goal

Describe and simulate a recurrent neural network (RNN)

Models so far

Perceptron

$$x = \sum_i input_i \, w_i + \theta$$

Feedforward networks

binary activation function

$$output = 0 \text{ for x} \leq 0$$

$$output = 1 \text{ for x} > 0$$

Hopfield

$$\begin{matrix} V_i \rightarrow 1 \\ V_i \rightarrow 0 \end{matrix} \quad \text{if} \quad \sum_{j \neq i} T_{ij} V_j \quad \begin{matrix} > U_i \\ < U_i \end{matrix}$$

No self-connections

Recurrent Neural Network intuition

… pass information back to itself.

Perceptron

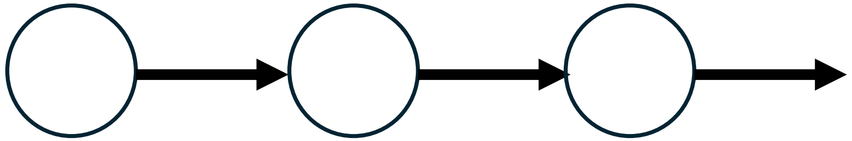

Hopfield

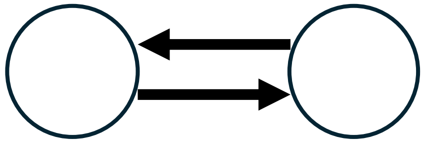

Recurrent Neural Network (RNN)

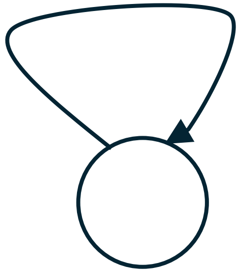

Incorporate <u>previous activity</u> of the neuron

When is this useful?

Sequences of data

"*We will, we will, rock* ____"

What is this word?

Useful to look back at the previous word(s)

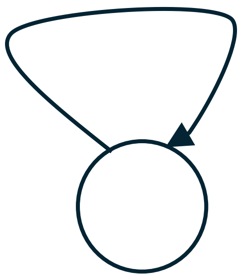

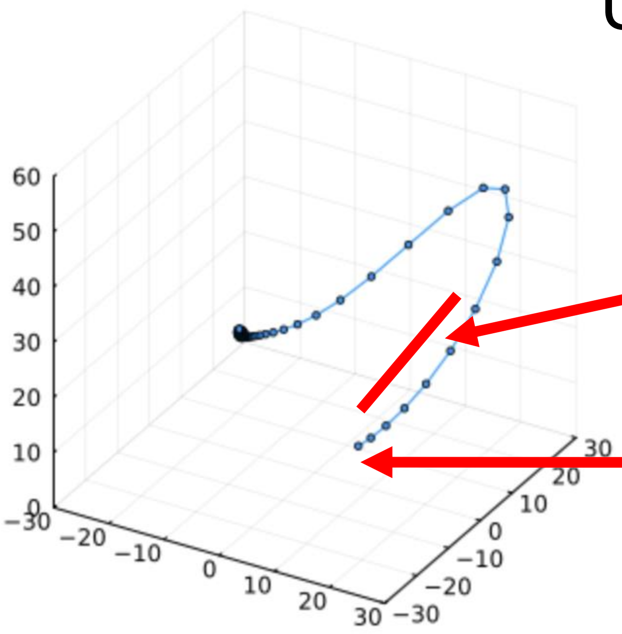

Useful to look back at the previous time(s)

What happens next?

What is the RNN model?

Let's start simple: consider the dynamics of a single neuron

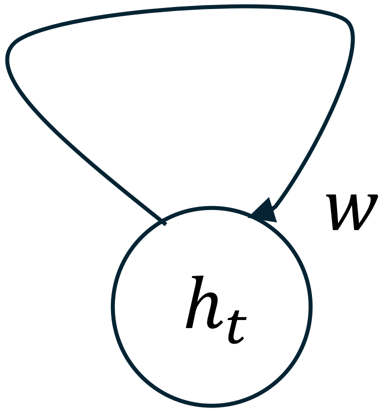

$$h_t = w\, h_{t-1}$$

Q: What does this system do?

Python

RNN dynamics (simplified)

$$h_t = w\ h_{t-1}$$

Q: What does this system do?

$$h_1 = \ w\ h_0$$

$$h_2 = \ w\ h_1 = w(wh_0) = w^2 h_0$$

$$h_3 = w\ h_2 \ = w(w^2 h_0) = w^3 h_0$$

$$\boxed{h_t = w^t h_0}$$

so, given h_0 and w, we can compute h_t

RNN dynamics (simplified)

Ex: Given $h_0 = 0.1$ and $w = 2$, what is h_{10}? $h_{10} = 2^{10}\ 0.1$

Q. In the distant future what happens to h_t? $h_t = w^t h_0$

Case 1. $1 < \mathrm{w}$ Then $w^\mathrm{t} \rightarrow \infty$ as t$\rightarrow\infty$ So, h_t explodes

Note, activation function protects against this.

Case 2. $\mathrm{w} = 1$ Then $w^\mathrm{t} \rightarrow 1$ as t$\rightarrow\infty$ So, h_t does not change

Case 3. $0 < \mathrm{w} < 1$ Then $w^\mathrm{t} \rightarrow 0$ as t$\rightarrow\infty$ So, h_t disappears

Case 4. $\mathrm{w} < 0$ Then $w^\mathrm{t} \rightarrow \pm$ as t is even or odd So, h_t “flickers”

RNN dynamics (simplified)

Q. In the distant future what happens to h_t?

$$h_t = w^t h_0$$

Case 1. $1 < w$ — So, h_t explodes

Case 2. $w = 1$ — So, h_t at equilibrium

Case 3. $0 < w < 1$ — So, h_t disappears

Case 4. $w < 0$ — So, h_t “flickers”

Python

What is the RNN model?

Initially, dynamics of a <u>single node</u>

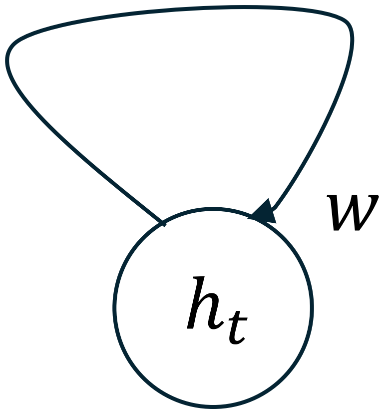

<u>Q.</u> What about a <u>network</u> of these nodes, with input & output?

Back to matrices & vectors …

Remember Feedforward

To build intuition, let's return to a previous model

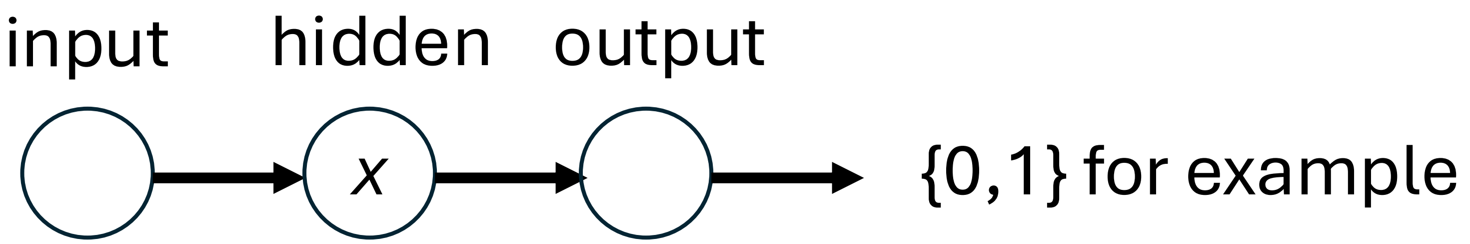

Activity in the hidden layer $$x = \sum_i input_i \, w_i + \theta$$

Remember Feedforward

Now, a more complicated model

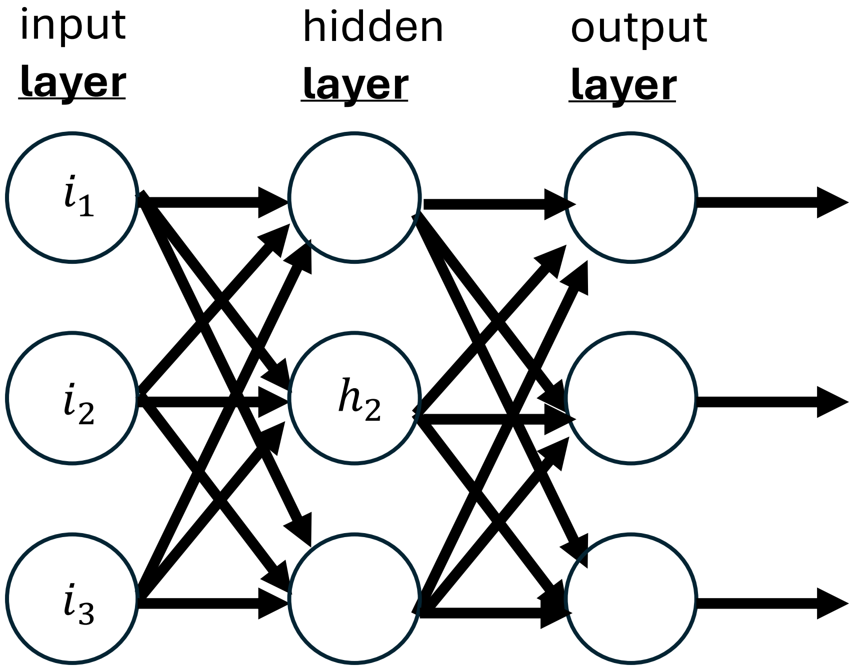

$$h_2 = ?$$

Remember Feedforward

More complicated model

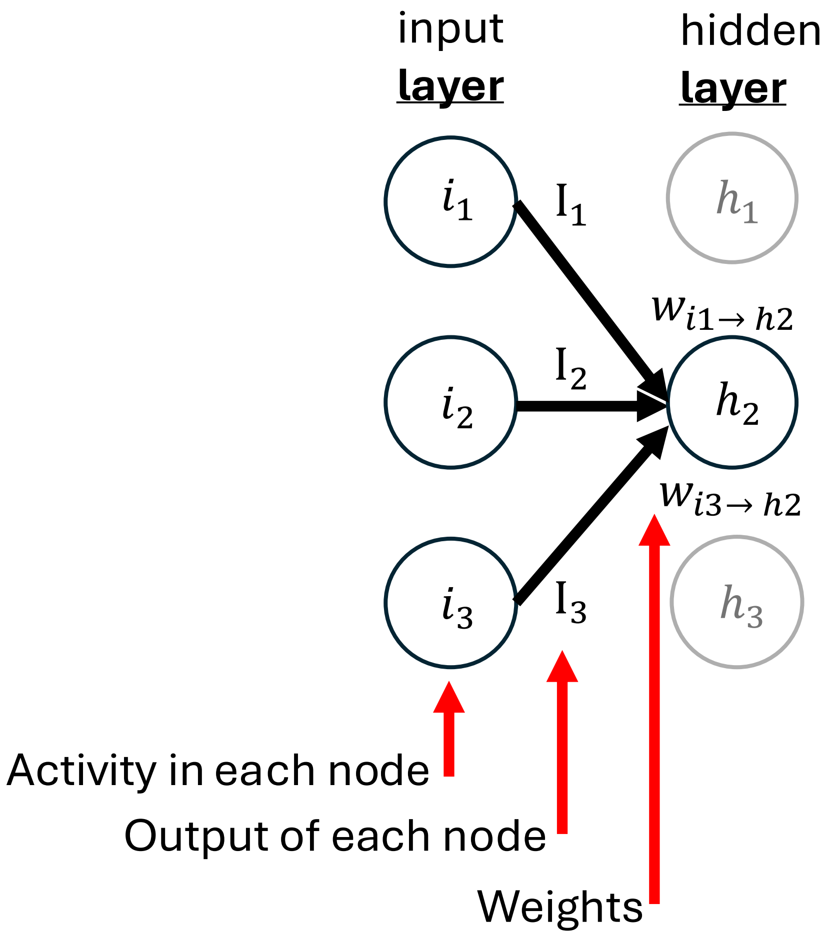

$$h_1 =$$

$$h_2 =$$

$$h_3 =$$

Program this … Python

Remember Feedforward

A more complicated model

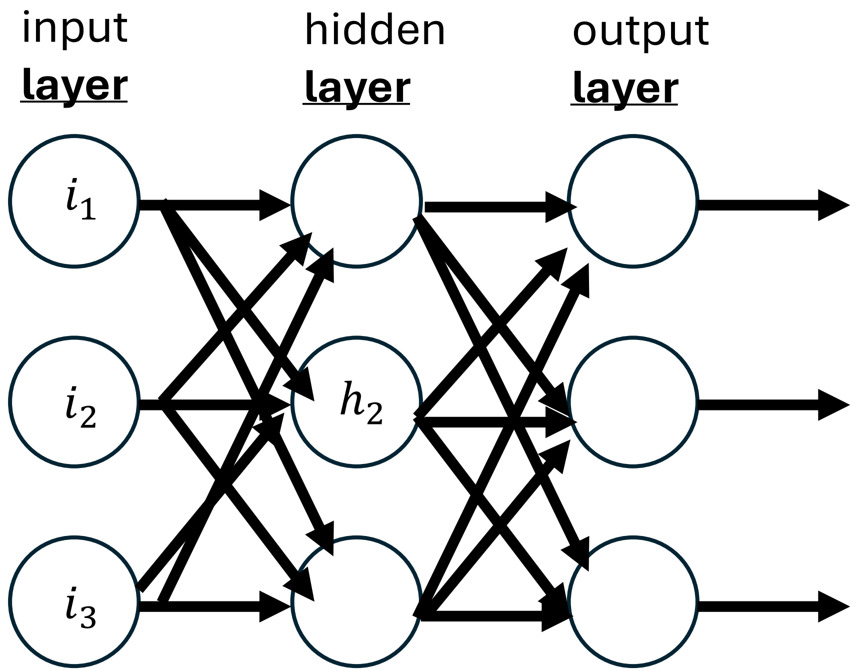

Q: What if you had N=100 neurons per layer?

Remember Feedforward

More complicated model

$$h_1 = w_{i1\to h1}\,\mathrm{I}_1 + w_{i2\to h1}\,\mathrm{I}_2 + w_{i3\to h1}\,\mathrm{I}_3$$

$$h_2 = w_{i1\to h2}\,\mathrm{I}_1 + w_{i2\to h2}\,\mathrm{I}_2 + w_{i3\to h2}\,\mathrm{I}_3$$

$$h_3 = w_{i1\to h3}\,\mathrm{I}_1 + w_{i2\to h3}\,\mathrm{I}_2 + w_{i3\to h3}\,\mathrm{I}_3$$

Q: What does this look like?

Remember Feedforward

New notation

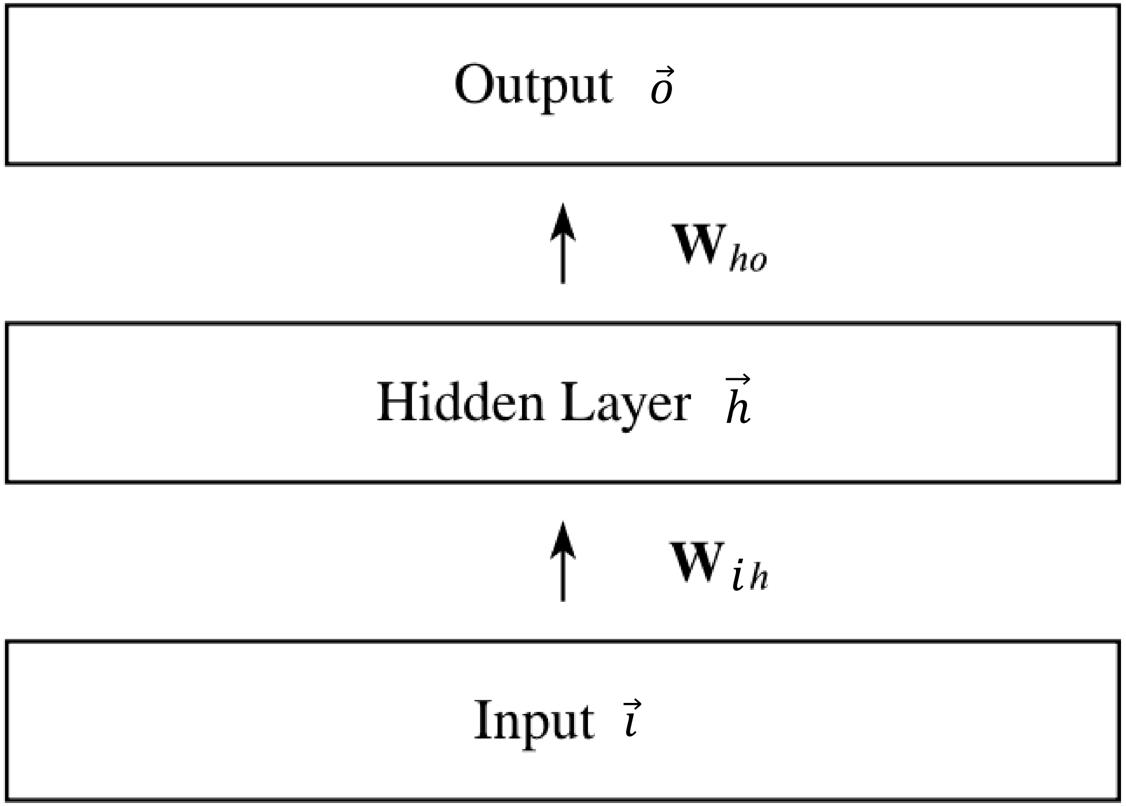

Feedforward Neural Network

Remember Feedforward

START HERE

More complicated model

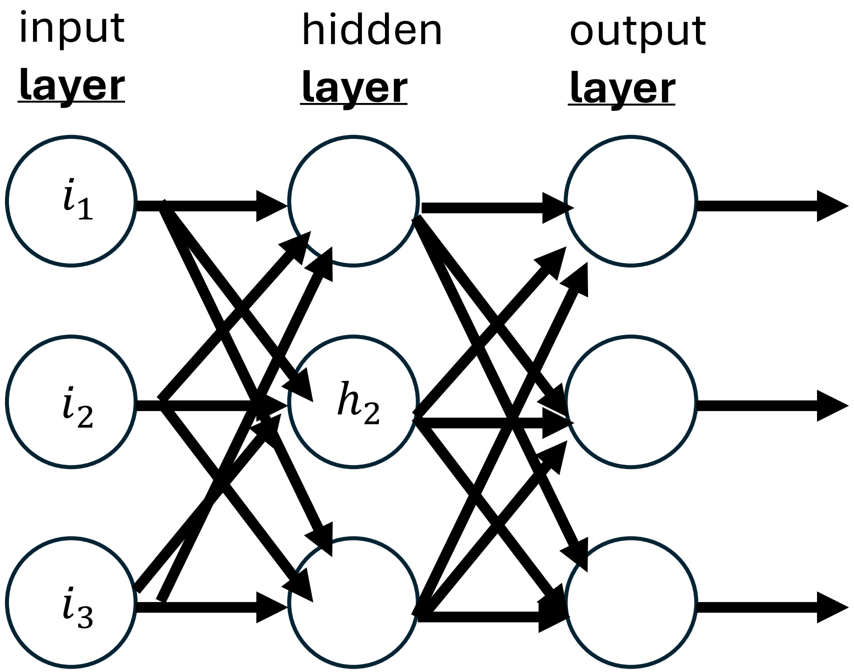

Set $I = \begin{bmatrix} 1 \\ 3 \\ -3 \end{bmatrix}$

$W_{ih} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$

$W_{ho} = \begin{bmatrix} 1 & 1 & 1 \\ -2 & -2 & -2 \\ 3 & 3 & 3 \end{bmatrix}$

Python (redo)

Remember Feedforward

What if the input layer has two neurons and the hidden layer has 10 neurons?

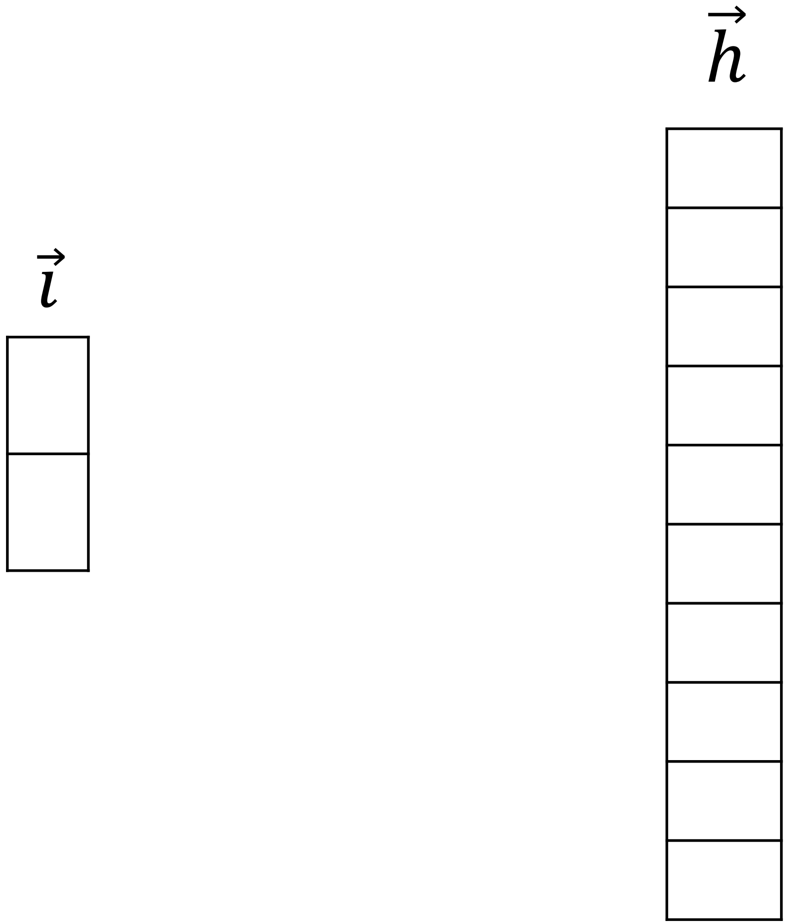

What is the shape of $\mathbf{W}_{ih}$?

Now, back to RNN

RRN vs Feedforward

<u>Q.</u> What is different?

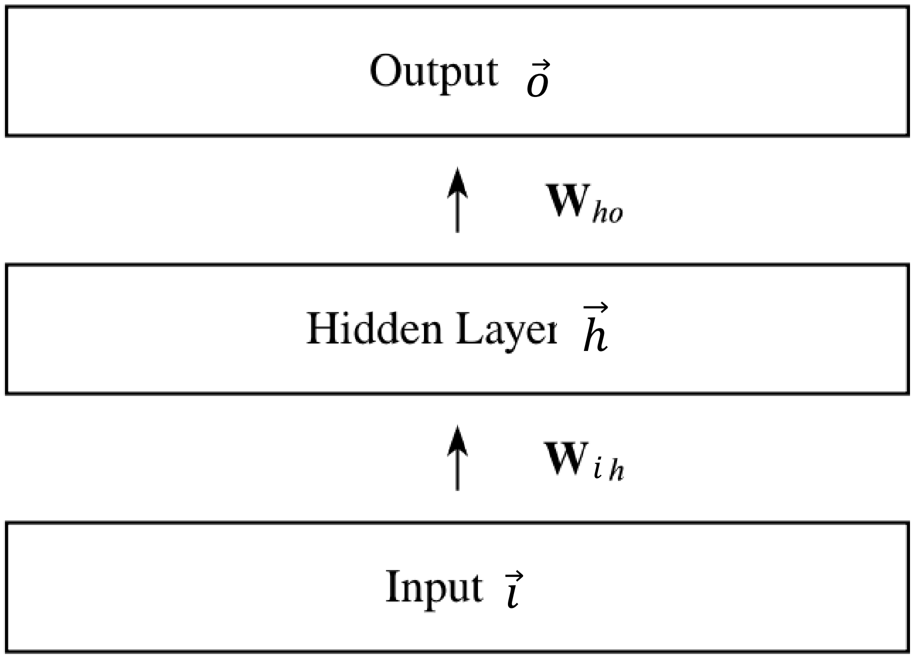

Feedforward Neural Network

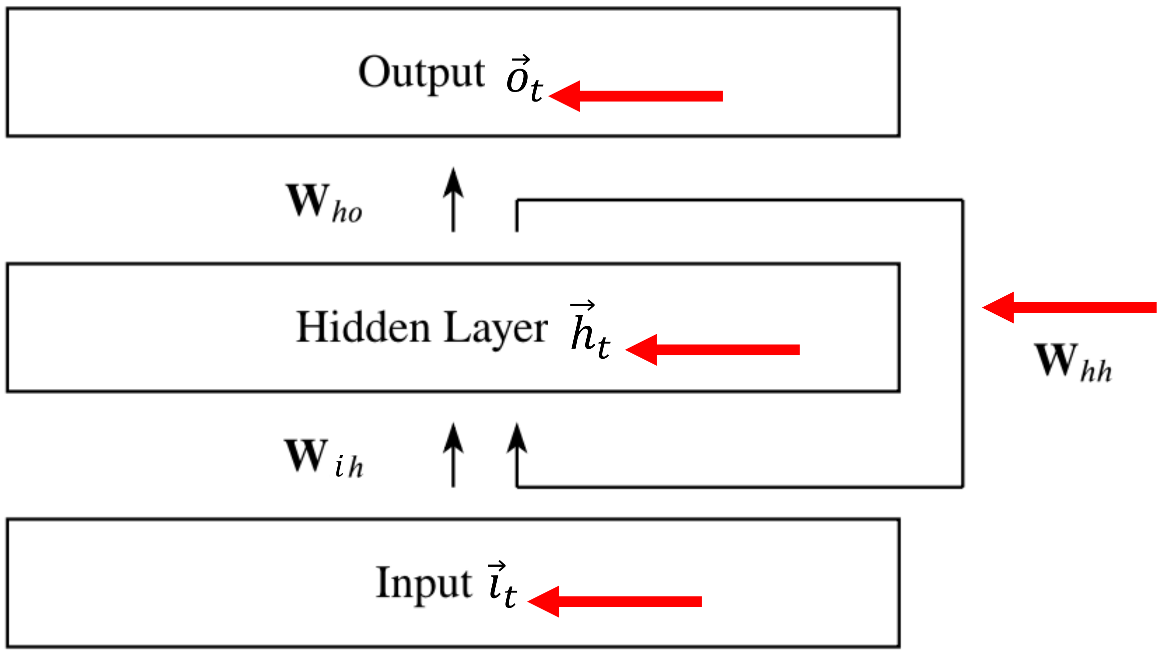

Recurrent Neural Network

Consider activity in the hidden layer $\vec{h}_t$

RRN hidden layer activity

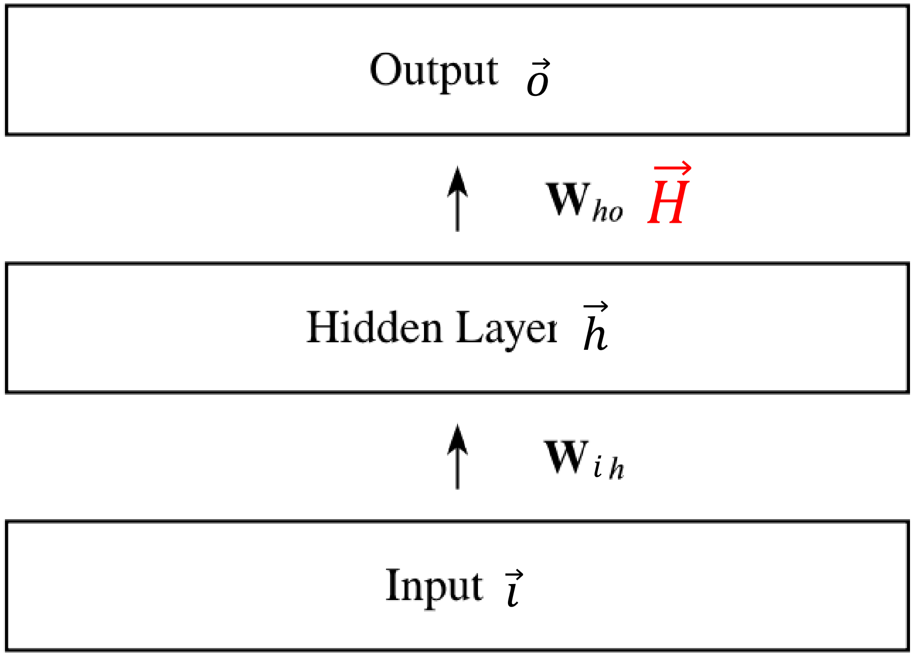

Feedforward Neural Network

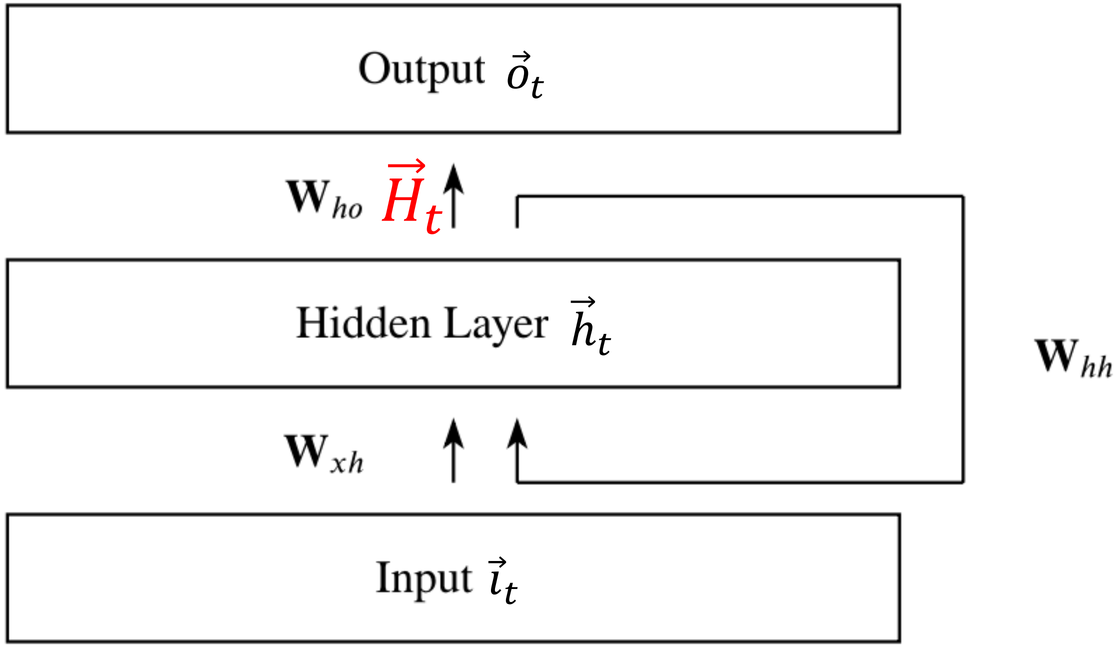

Recurrent Neural Network

$$\vec{h} = \boldsymbol{W}_{ih}\vec{\mathrm{I}} + \vec{b}_h \qquad \vec{H} = \phi(\vec{h})$$

activity — input — bias — activation function

$$\vec{h}_t = \boldsymbol{W}_{ih}\vec{I}_t + \boldsymbol{W}_{hh}\vec{h}_{t-1} + \vec{b}_h \qquad \vec{H}_t = \phi(\vec{h}_t)$$

activity at time t — input at time t — activity at time t-1 — bias — activation function

RRN output layer activity

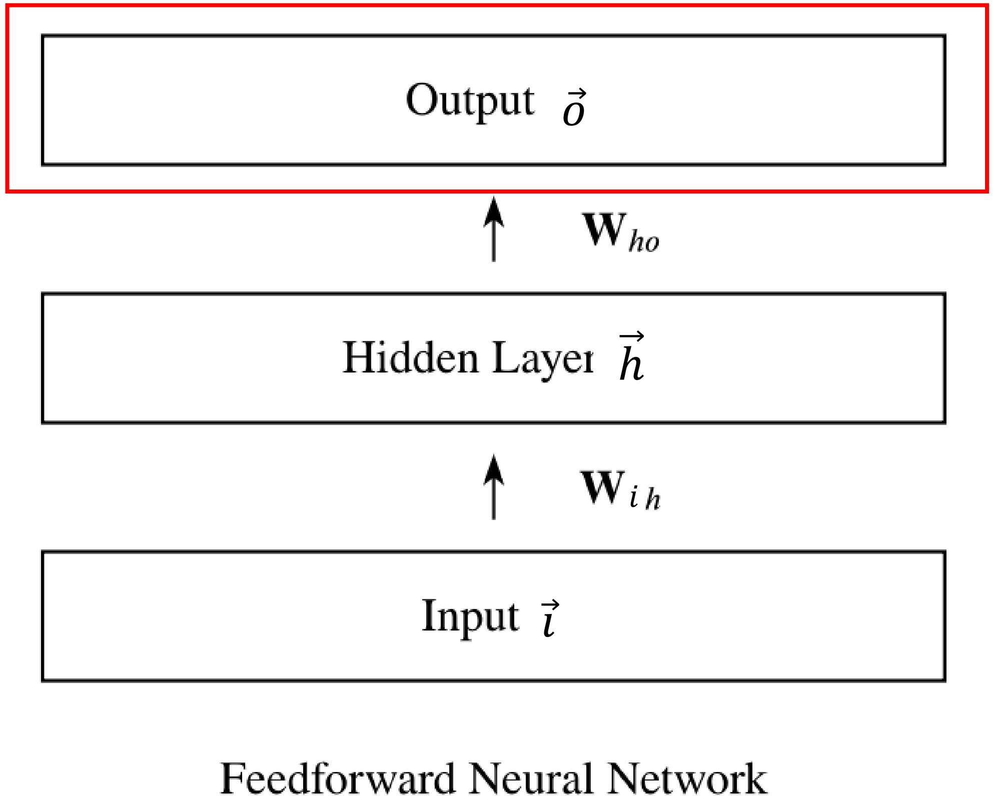

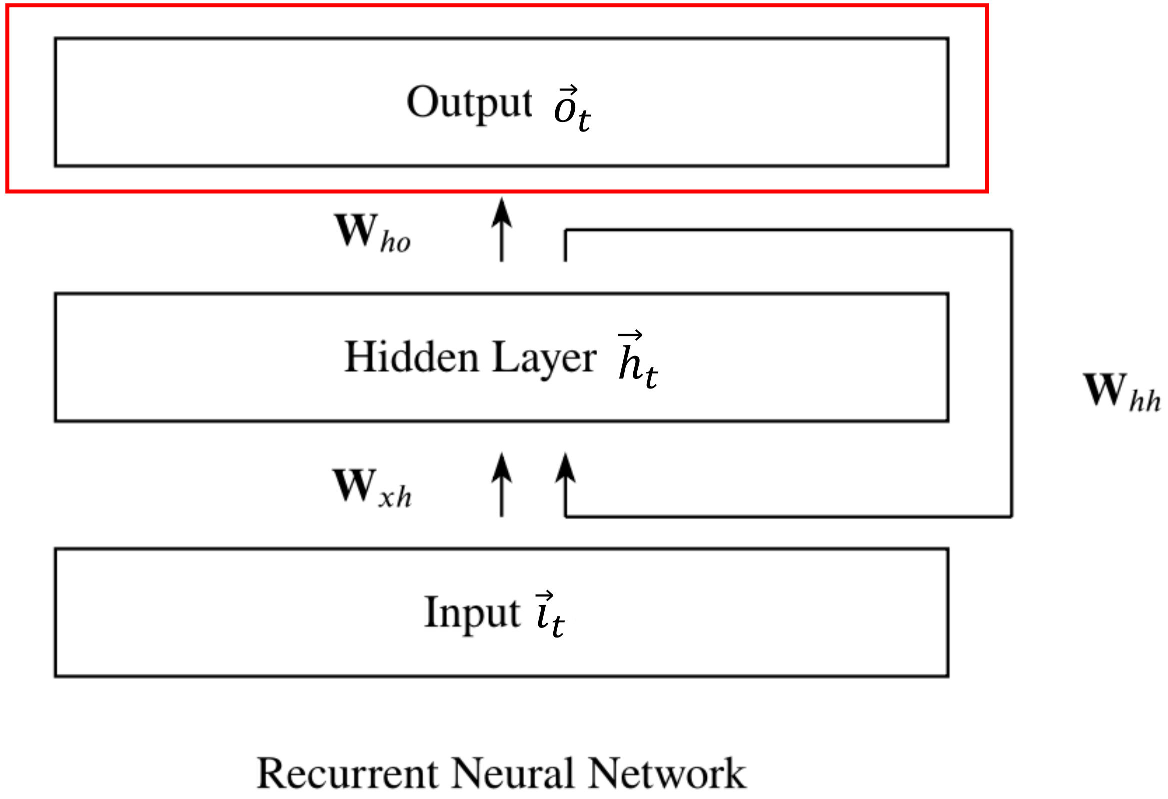

$\vec{o} = ?$

$\vec{o}_t = ?$

RRN model

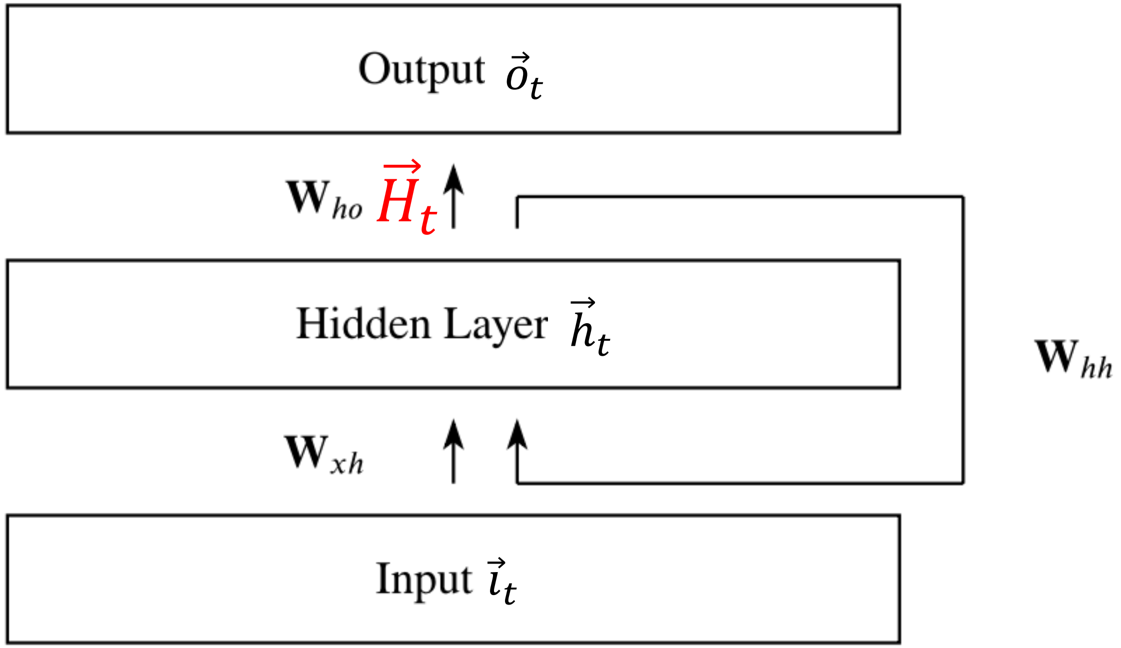

Recurrent Neural Network

$$\vec{o}_t =$$

$$\vec{h}_t = \boldsymbol{W}_{ih}\vec{I}_t + \boldsymbol{W}_{hh}\vec{h}_{t-1} + \vec{b}_h$$
$$\vec{H}_t = \phi(\vec{h}_t)$$

Python (Simulate the RNN)

RNN additional notes

h_t recursively includes h_{t-1} & this process occurs for every time step
→ the RNN includes influence of all preceding hidden states h_{t-1} to h_0.

To train a RNN → backpropagation through time (BPTT)

Long short-term memory (LSTM)
→ a type of RNN with more components

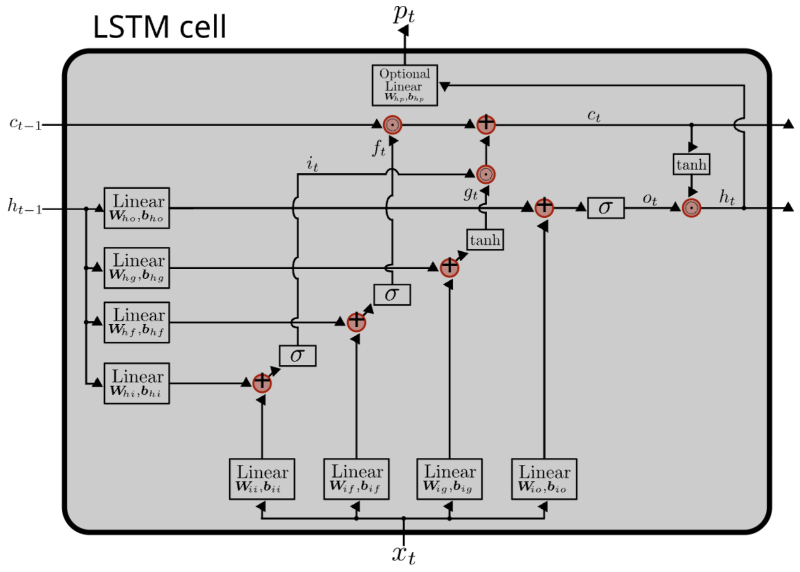

RNN Summary

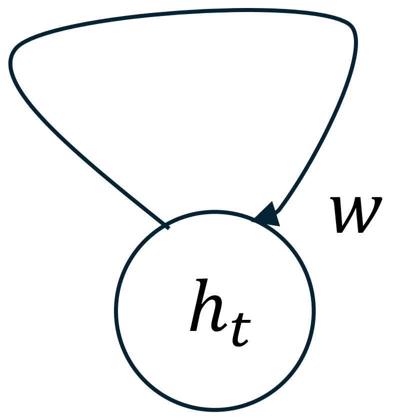

An RNN passes information back to itself.

The activity in an RNN depends on its past activity.

RNN dynamics are relatively "simple"

$$h_t = w\, h_{t-1}$$

… but collections of RNNs can be much more complicated (LSTM).

What about rhythms?

Q. Isn't this course about rhythms?

Biological brains generate rhythms

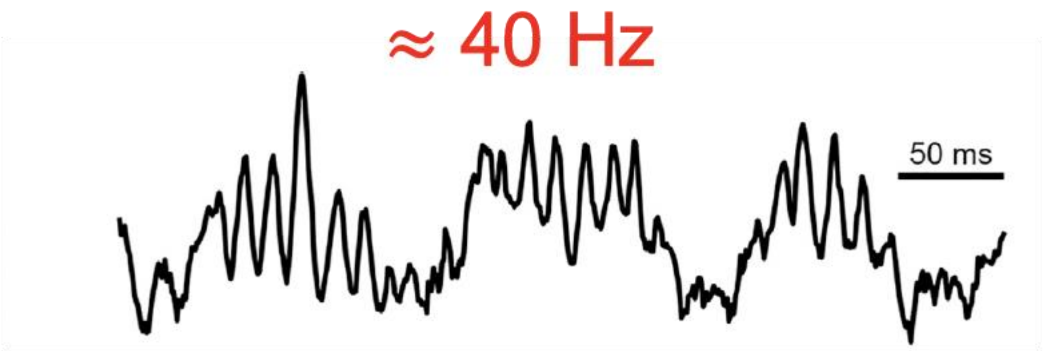

[Fernandez-Ruiz et al., Neuron, 2023]

None of the models considered so far can generate a rhythm (on its own).

These models are the fundamental elements of all modern LLMs.

But a small modification can support rhythms ...

Rhythmic RNN

Consider this updated equation

$$h_t = w_1 \, h_{t-1} + w_2 \, h_{t-2}$$

Activity now

Activity at time $t-1$

Activity at time $t-2$

Q. What does this do?

Rhythmic RNN

Consider this updated equation

$$h_t = w_1 \, h_{t-1} + w_2 \, h_{t-2}$$

Equivalent to a damped harmonic oscillator

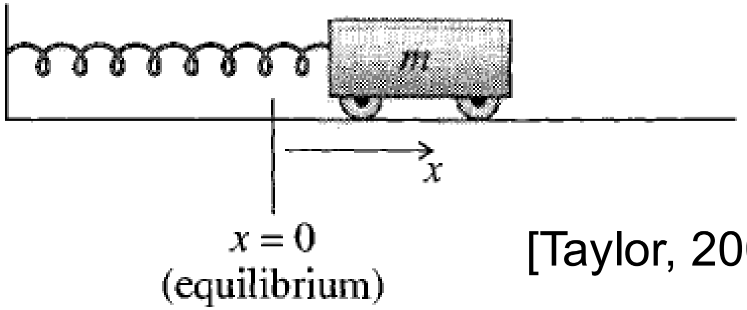

[Taylor, 2005]

Rhythmic RNN

Consider this updated equation $h_t = w_1\, h_{t-1} + w_2\, h_{t-2}$

Equivalent to a damped harmonic oscillator

$$\ddot{x} + 2\beta\dot{x} + \omega^2 x = F$$

damping or friction — spring — forcing

$$\omega = 2\pi f$$

f = natural frequency of the spring

Rhythmic RNN

$$\ddot{x} + 2\beta\boxed{\dot{x}} + \omega^2 x = F$$

Replace each derivative with a discrete approximation.

index: ***k*** index: ***k-1***

$$\dot{x} = \frac{dx}{dt} = \frac{\text{change in x}}{\text{change in time}} = \frac{x(\text{now}) - x(\text{past})}{\boxed{t(\text{now}) - t(\text{past})}} = \boxed{\frac{x_k - x_{k-1}}{\Delta}}$$

Δ: a small step forward in time

Rhythmic RNN

$$\boxed{\ddot{x}} + 2\beta\dot{x} + \omega^2 x = F$$

Replace each derivative with a discrete approximation.

index: ***k*** — index: ***k-1*** — index: ***k-2***

$$\ddot{x} = \frac{d^2x}{dt^2} \quad \cdots \quad = \boxed{\frac{x_k - 2x_{k-1} + x_{k-2}}{\Delta^2}}$$

Rhythmic RNN

$$\ddot{x} + 2\beta\dot{x} + \omega^2 x = F$$

Replace each derivative with a discrete approximation.

Assume no forcing

$$\frac{x_k - 2x_{k-1} + x_{k-2}}{\Delta^2} + 2\beta\frac{x_k - x_{k-1}}{\Delta} + \omega^2 x_k = 0$$

Collect terms ...

Rhythmic RNN

$$\ddot{x} + 2\beta\dot{x} + \omega^2 x = F$$

Replace each derivative with a discrete approximation.

$$(1 + 2\beta\Delta + \omega^2\Delta^2)x_k - 2(1 + \beta\Delta)x_{k-1} + x_{k-2} = 0$$

Activity now

Activity at time $k - 1$

Activity at time $k - 2$

$$x_k = \alpha_1 x_{k-1} + \alpha_2 x_{k-2}$$

Rhythmic RNN

$$\ddot{x} + 2\beta\dot{x} + \omega^2 x = F$$

Replace each derivative with a discrete approximation.

After some steps: $x_k = \alpha_1 \, x_{k-1} + \alpha_2 \, x_{k-2}$

Original rhythmic RNN proposal:

$$h_t = w_1 \, h_{t-1} + w_2 \, h_{t-2}$$

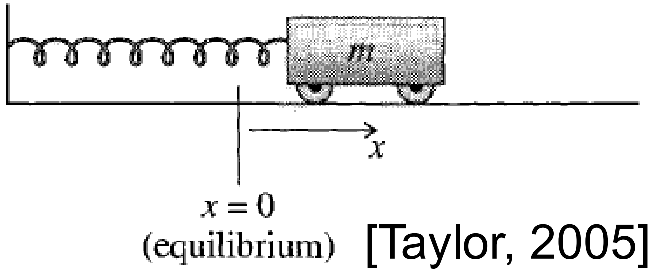

[Taylor, 2005]

Equivalent to a damped harmonic oscillator

Rhythmic RNN Summary

Original RNN

$$h_t = w_1 \, h_{t-1}$$

Rhythmic RNN

$$h_t = w_1 \, h_{t-1} + w_2 \, h_{t-2}$$

→ damped oscillator

→ more like the biological brain

<u>**Q (big):**</u> Is the rhythmic RNN better?

Artificial neural networks (Perceptron, Hopfield, RNNs) <u>do not intrinsically oscillate</u>.

Does adding rhythms improve performance?

Rhythmic RNN in practice

<u>Q (big):</u> Is the rhythmic RNN better? **<u>A:</u>** Maybe?

Modified National Institute of Standards and Technology (**MNIST**) database of handwritten digits

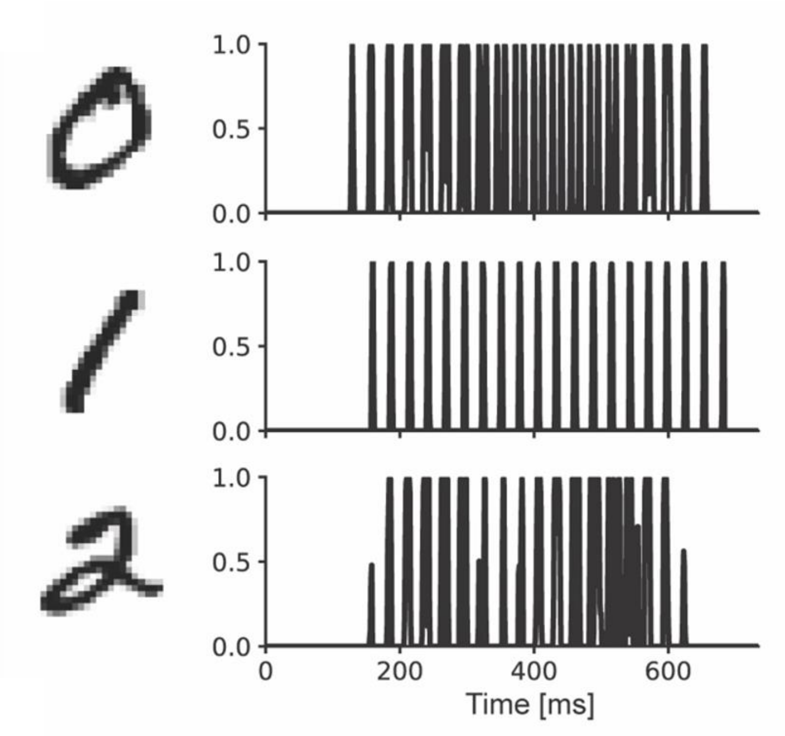

→ Rhythmic RRN →

	Accuracy Mean (SEM)
Rhythmic (n=11)	0.759 (0.007)
RNN (n=11)	0.617 (0.018)
LSTM (n=6)	0.715 (0.019)

→ Rhythmic RNN outperforms standard RNN architectures

Rhythmic RNN in practice

<u>Q (big):</u> Is the rhythmic RNN better?

<u>A:</u> Maybe?

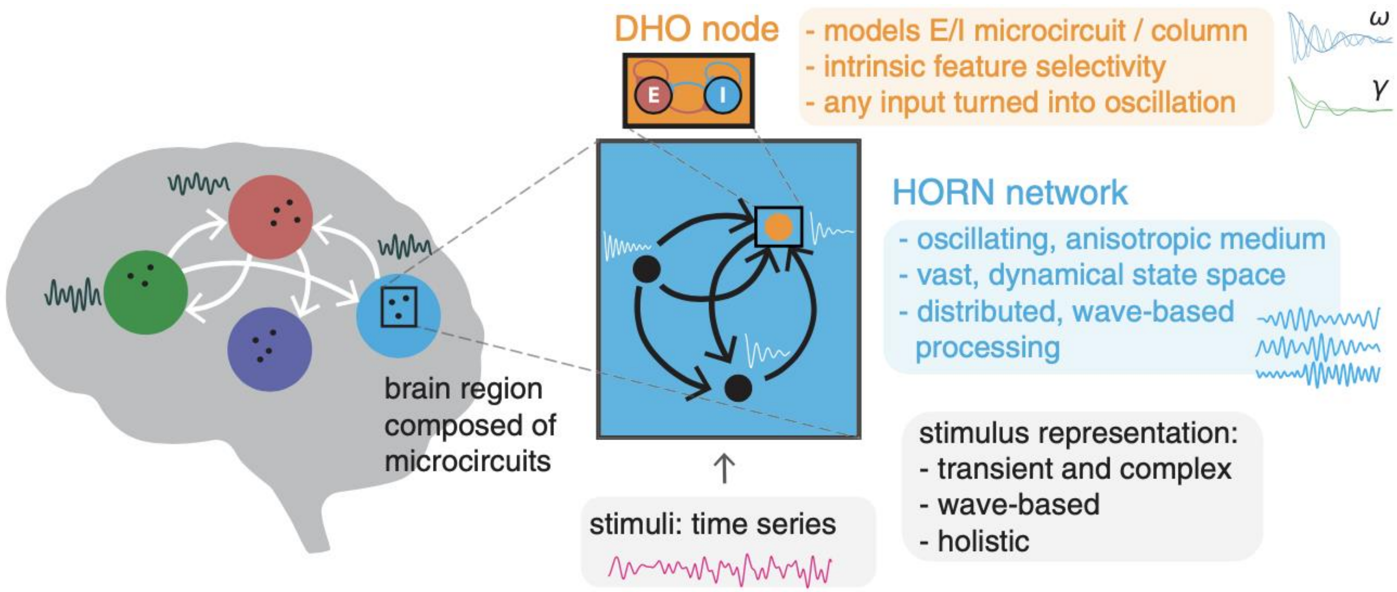

[Effenberger et al., PNAS, 2025]

Outperforms classic ANNs

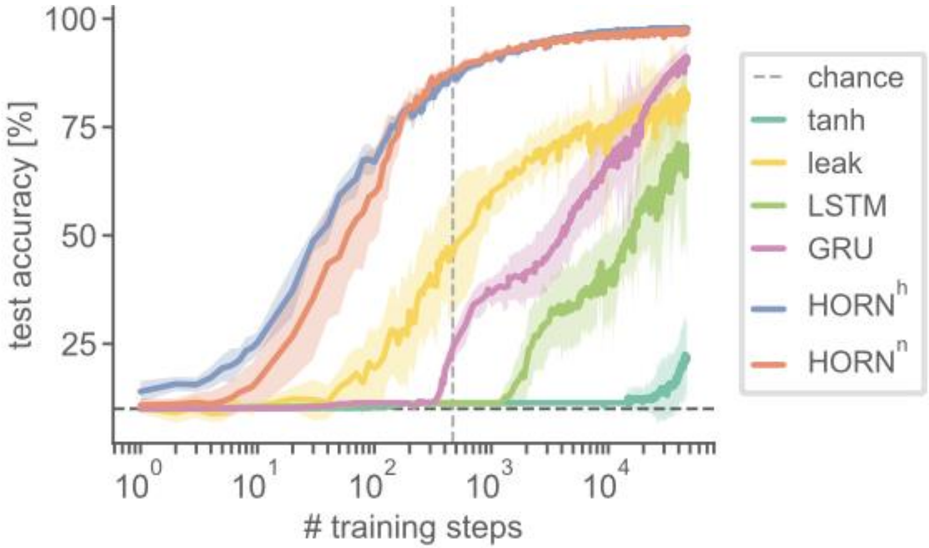

Rhythmic RNN in practice

Rhythmic RNNs outperform other artificial neural networks (maybe)

Q. What rhythms?

Brain rhythms

Remember ...

≈ 150 Hz

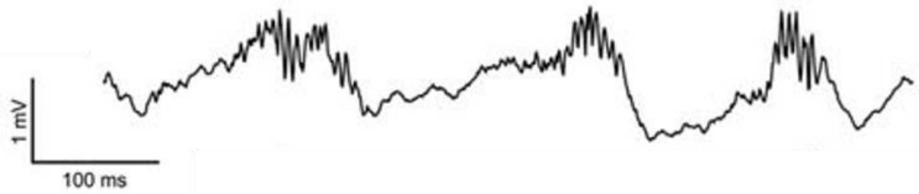

[Buzsaki, Hippocampus, 2015]

≈ 40 Hz

[Fernandez-Ruiz et al., Neuron, 2023]

≈ 12 Hz

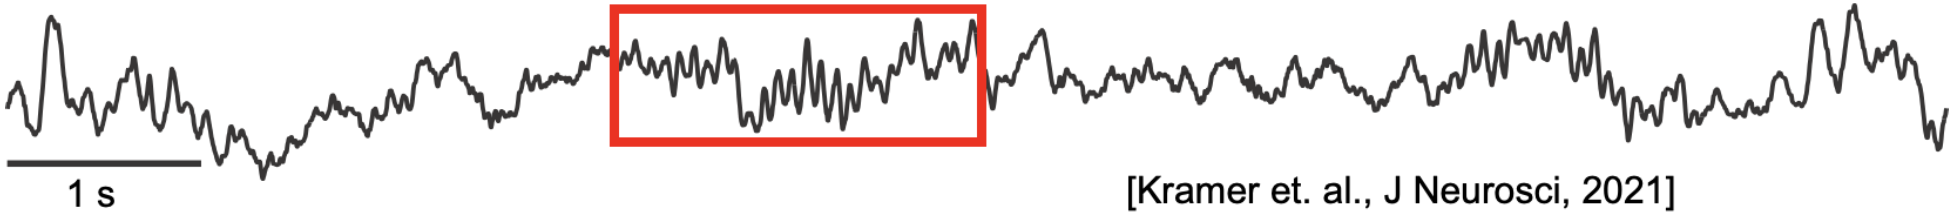

[Kramer et. al., J Neurosci, 2021]

≈ 2 Hz

[Carracedo et. al., PNAS, 2013] 1 s

< 1 Hz

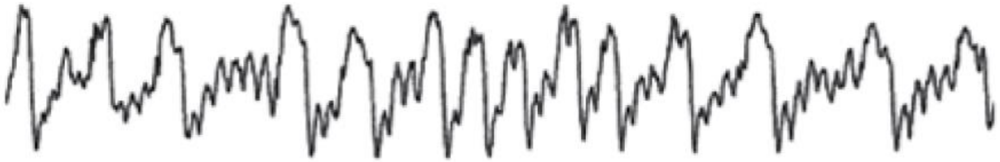

[Crunelli & Hughes, Nat Neuro, 2010] 5 s 250 μV

Brain rhhytms

Brain rhythms organize in separate bands

Different proposal for rhythm spacing:

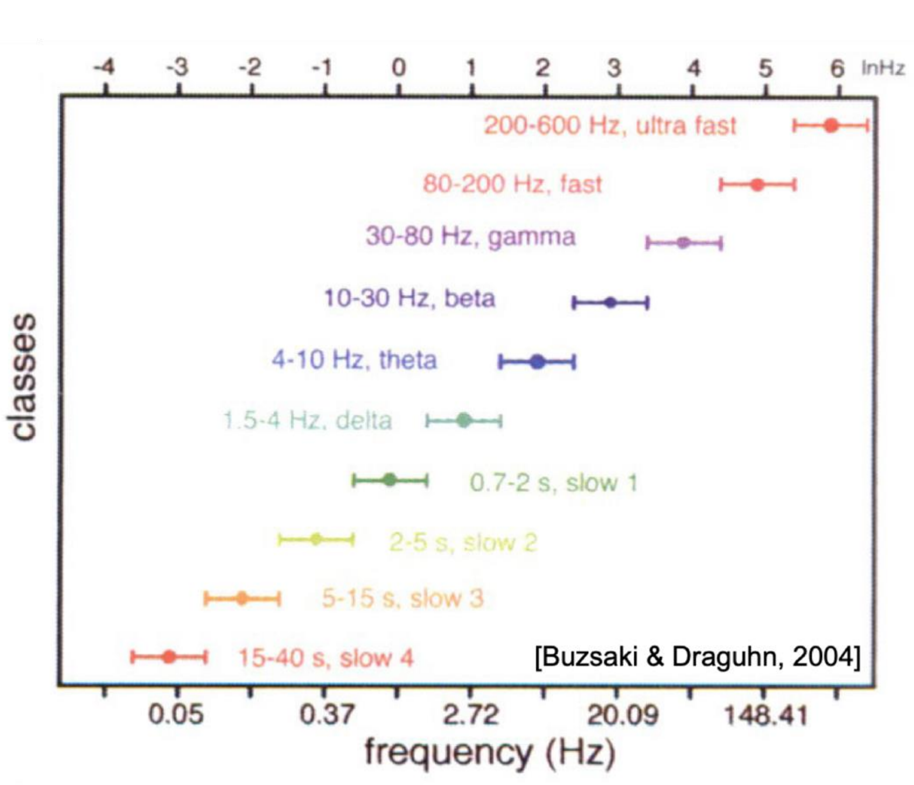

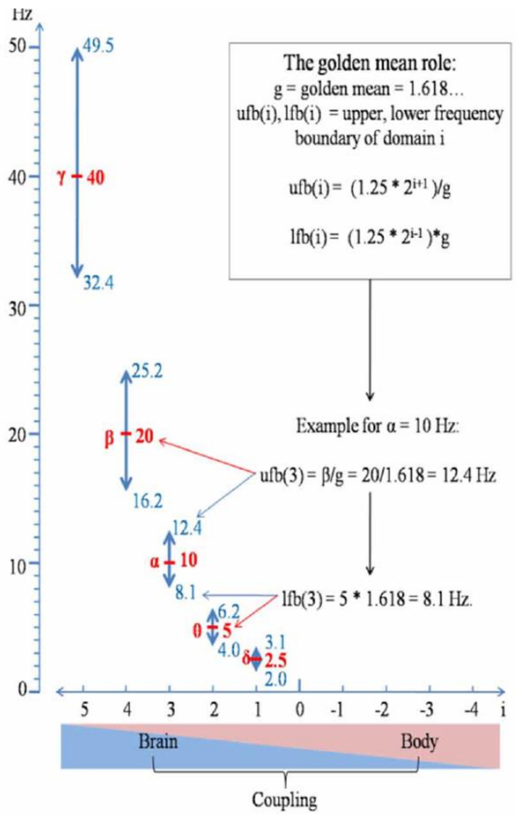

[Klimesch, 2013]

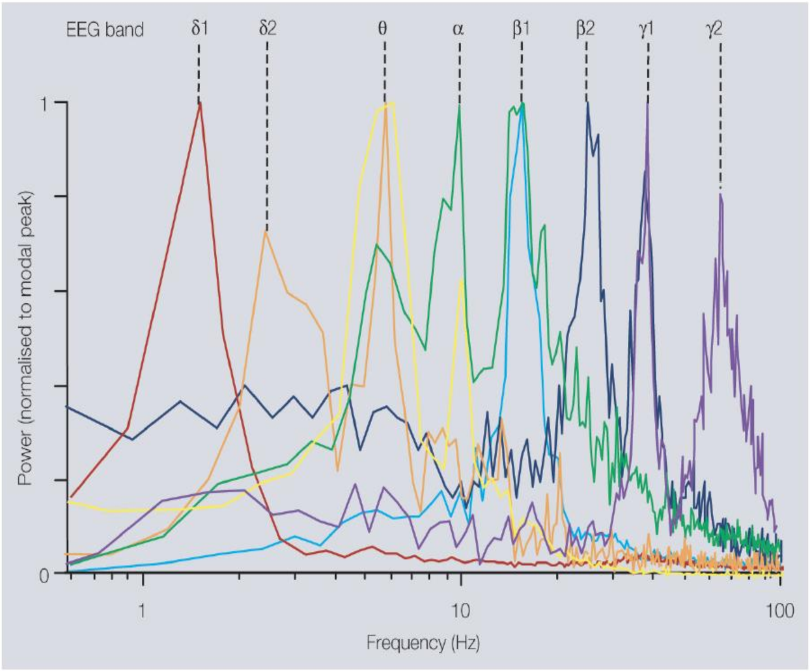

[Roopun et al., 2004]

Q: What rhythms in the RNN?

Golden rhythms

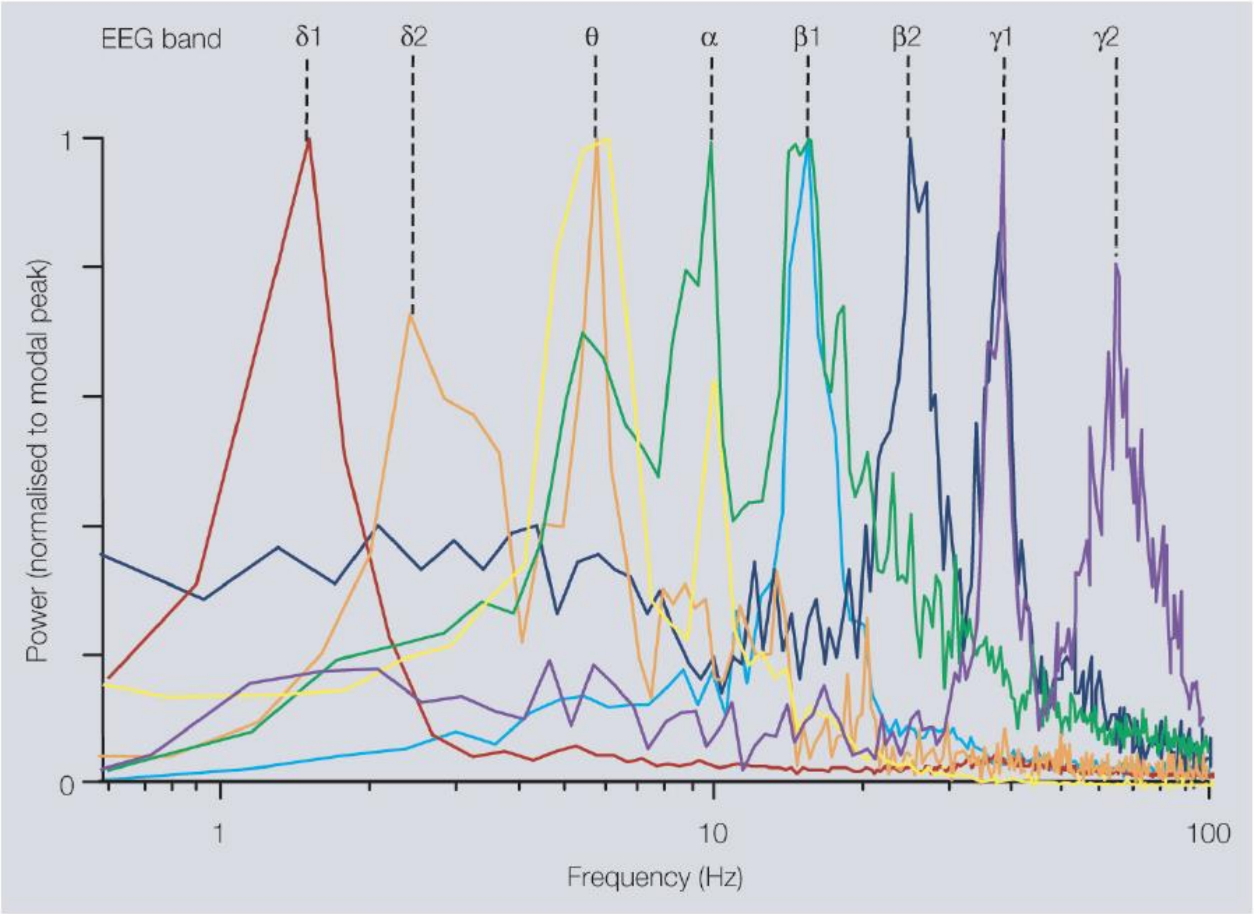

[Roopun et al., 2004]

f_k = center frequency of band k

$$\frac{f_k}{f_{k-1}} = \phi$$

"in neocortex the ratio of adjacent frequency bands is approximately phi" [Roopun et al., 2004]

Golden ratio, ϕ

$$\frac{a+b}{a} = \frac{a}{b} = \phi = \frac{1+\sqrt{5}}{2} = 1.618\ldots$$

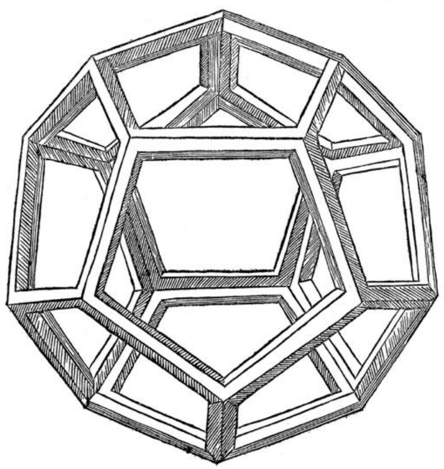

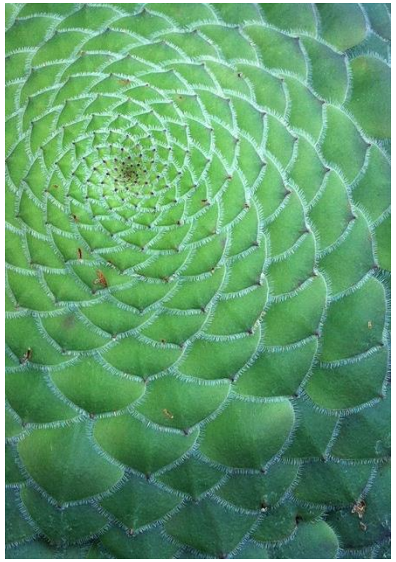

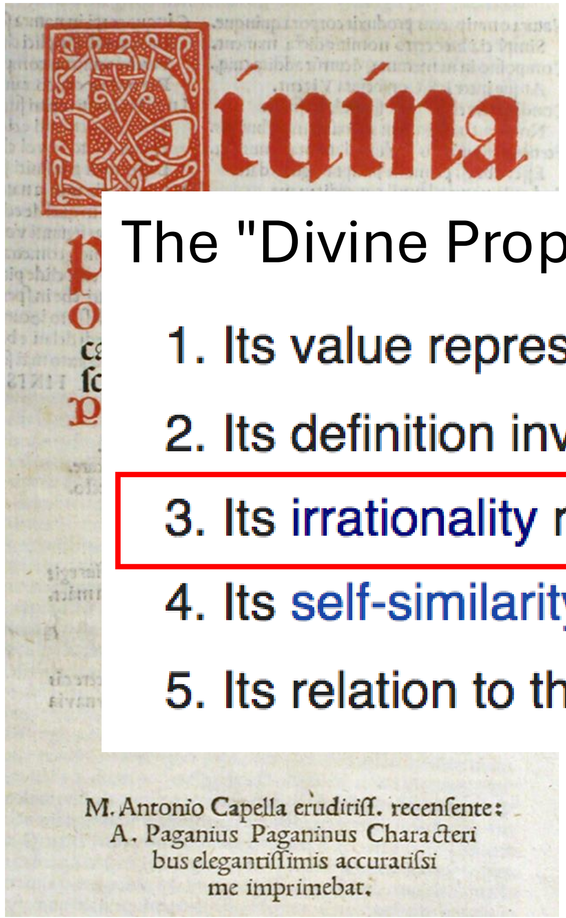

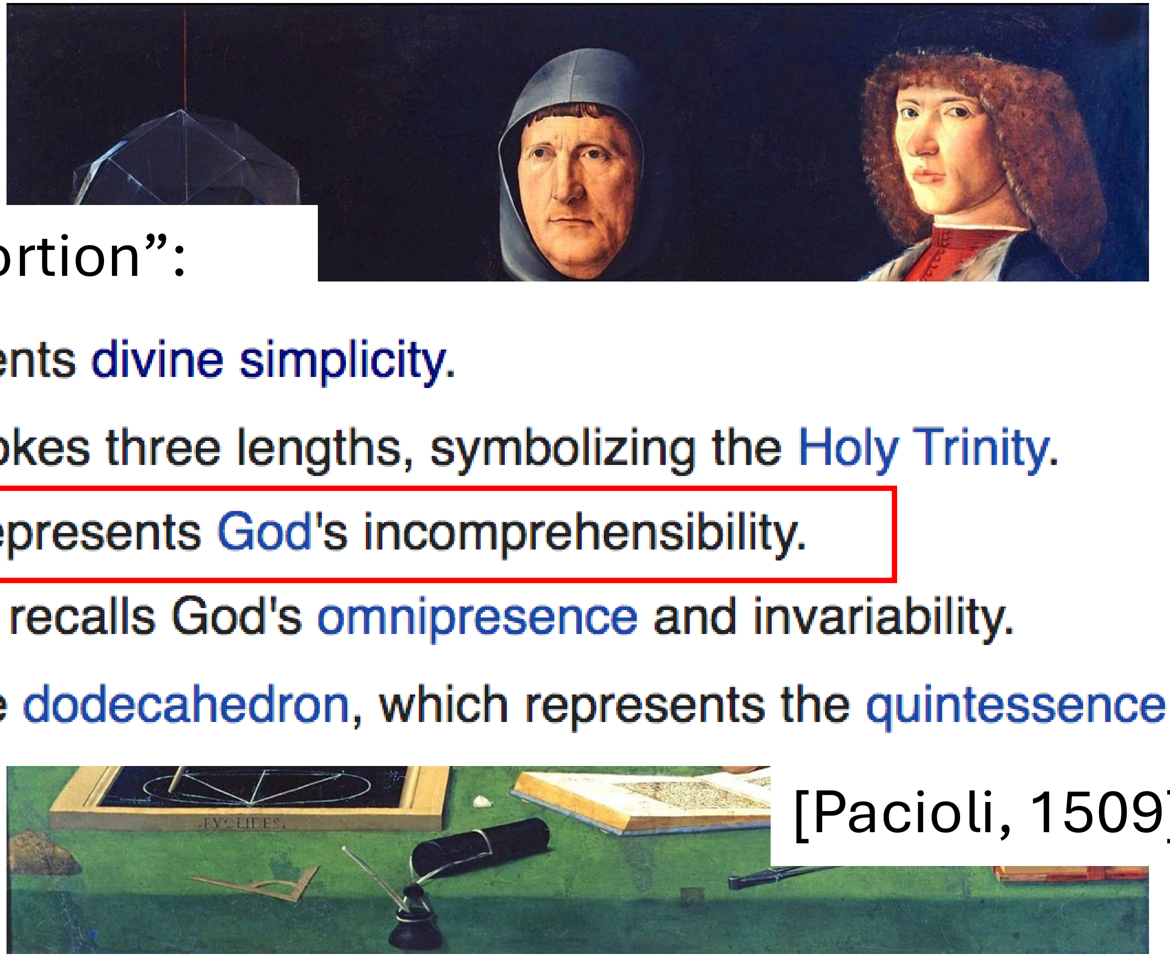

The "Divine Proportion”:

1. Its value represents divine simplicity.
2. Its definition invokes three lengths, symbolizing the Holy Trinity.
3. Its irrationality represents God's incomprehensibility.
4. Its self-similarity recalls God's omnipresence and invariability.
5. Its relation to the dodecahedron, which represents the quintessence

[Pacioli, 1509]

Golden rhythms

<u>Best example</u>: beta1 (15 Hz), beta2 (25 Hz), and gamma (40 Hz)

$$\frac{40\ \text{Hz}}{25\ \text{Hz}} \approx \frac{25\ \text{Hz}}{15\ \text{Hz}} \approx 1.6 \approx \phi$$

Note: $15\ \text{Hz} + 25\ \text{Hz} = 40\ \text{Hz}$

$65\ \text{ms} = 40\ \text{ms} + 25\ \text{ms}$

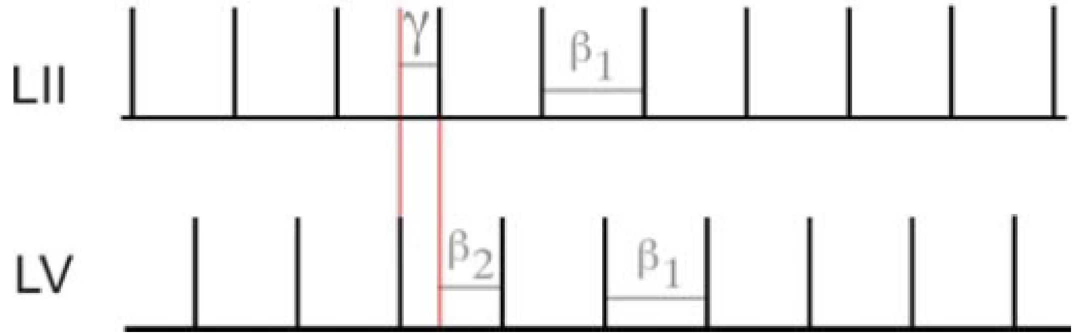

Rhythm generation through period concatenation

[Roopun et al., 2008a; Roopun et al., 2008b; Kramer et al., 2008; Kramer 2022]

Golden rhythms evidence (?)

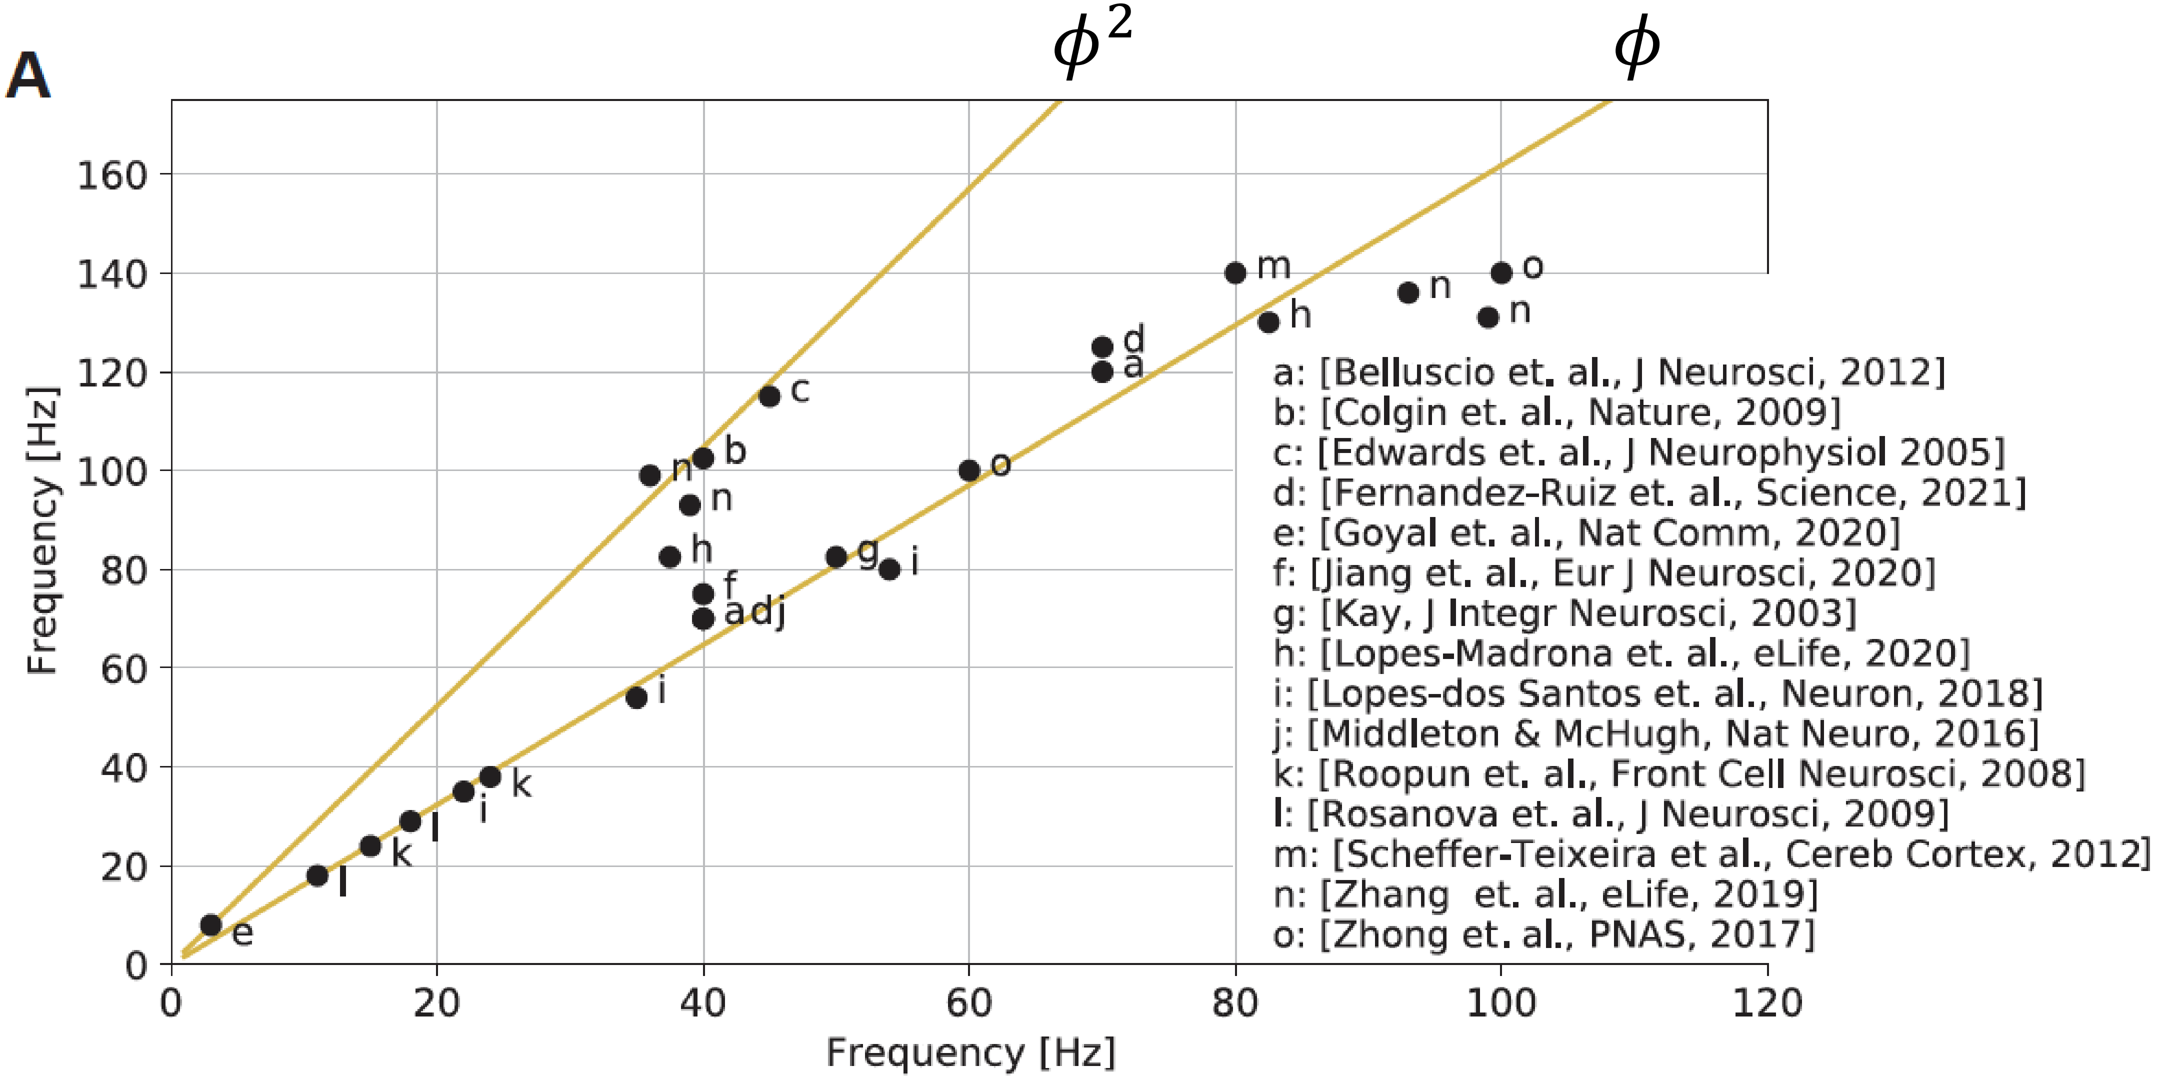

Golden rhythms evidence (?)

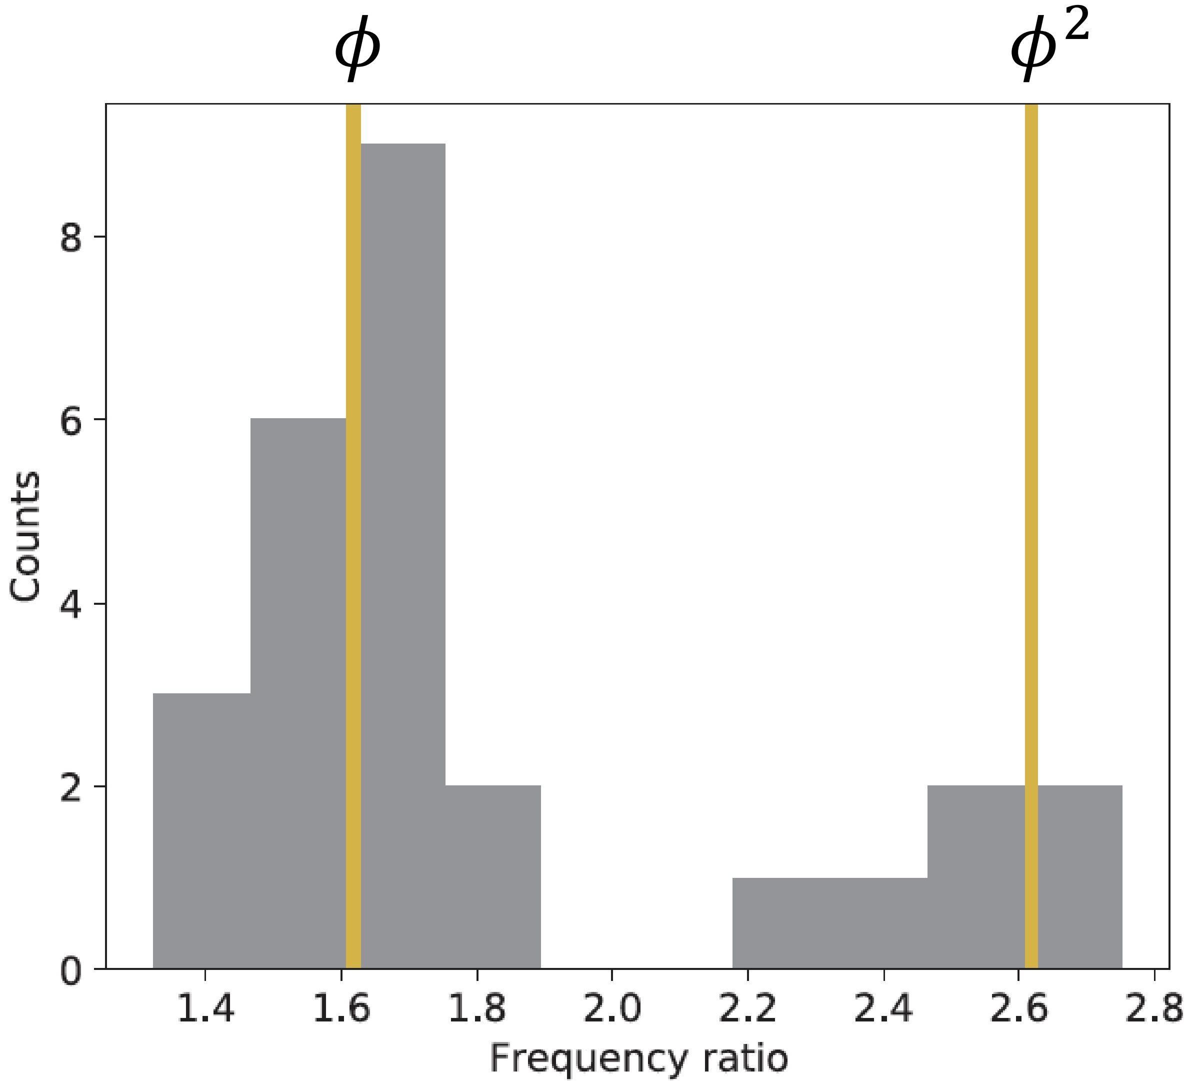

My favorite idea

Q. What is the fundamental timescale for life on earth?

A. 23 hrs, 56 min, sidereal period.

1/(23 hrs, 56 min) = 1/86160 s = 1.16e-5 Hz

Multiply by ϕ (again and again) to establish $$\frac{f_{k+1}}{f_k} = \phi$$

Power	T [s]	f [Hz]	Power	T [s]	f [Hz]	Power	T [s]	f [Hz]	Label
0	86160	1.16E-05	12	268	0.004	24	0.83	1.20	Slow 1
1	53250	1.88E-05	13	165	0.006	25	0.51	2	Delta
2	32910	3.04E-05	14	102	0.010	26	0.32	3	Delta
3	20340	4.92E-05	15	63.2	0.016	27	0.20	5	Theta
4	12571	7.96E-05	16	39.0	0.026	28	0.12	8	Alpha
5	7769	1.29E-04	17	24.1	0.041	29	0.07	13	Beta1
6	4802	2.08E-04	18	14.9	0.067	30	0.05	22	Beta2
7	2968	3.37E-04	19	9.2	0.11	31	0.03	35	Low Gamma
8	1834	5.45E-04	20	5.7	0.18	32	0.02	57	Mid Gamma
9	1133	8.82E-04	21	3.5	0.28	33	0.01	91	High Gamma
10	701	0.001	22	2.2	0.46	34	0.01	148	Ripple
11	433	0.002	23	1.3	0.74	35	0.004	239	Fast Ripples

Q: Why are brain rhythms organized in specific frequency bands?

A: The sun ...

(My favorite) hypothesis

If intelligent life were to evolve on a planet like Earth,

… in a star system like our own,

… with neural physiology like our own,

… then rhythmic bands would exist with center frequencies that depend on the planet’s circadian cycle.

Summary

It’s fun to be a mathematician

… but, it’s important to constrain theories with real-world observations.

… do not take the previous ideas *too* seriously.

Q (big): Are rhythms in the biological brain advantageous?