Unknotting the brain: Some simple approaches to modeling neural circuits and their dynamics

Jay Pina
York University Department of Physics and Astronomy
Centre for Vision Research

Work with:
Joel Zylberberg (York), Bard Ermentrout (Univ. Pittsburgh)
Blake Richards (McGill), Colleen Gillon (Univ. Toronto), Tim Henley (York)
Yoshua Bengio (Montreal), Timothy Lillicrap (Google DeepMind)
Allen Institute for Brain Sciences

WARNING:

Please be aware that some example images are shown that may cause seizures in individuals with pattern sensitive epilepsy, and visual discomfort in others. Do not proceed with viewing this presentation if these are a concern for you.

This applies to slides/pages **32 and 37**

Brains are intricate networks of vast numbers of neurons

- Brains are highly inhomogeneous, densely interconnected networks of electrically active neurons

- Mammalian brains have
  - anywhere from $\sim 3 \times 10^7$ (naked mole rat) to as many as $\sim 3 \times 10^{11}$ neurons (African elephant)
  - ~100 to 1,000 different neuronal types
  - dozens of distinct anatomical regions
    - which can themselves have subareas

- Each individual neuron is in turn comprised of many different components and connects to ~1,000 to 10,000 other neurons

- How do we begin to model, let alone understand, such systems?


Science Photo Library – KTSDESIGN/Getty Images

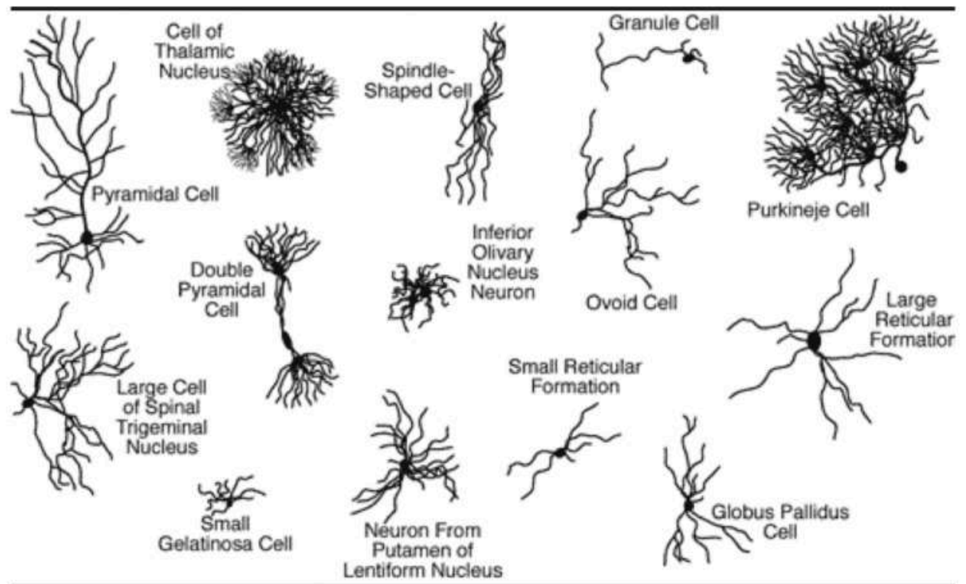

https://www.math.univ-toulouse.fr/~gfaye/CIMI/lecturesDSalort.pdf

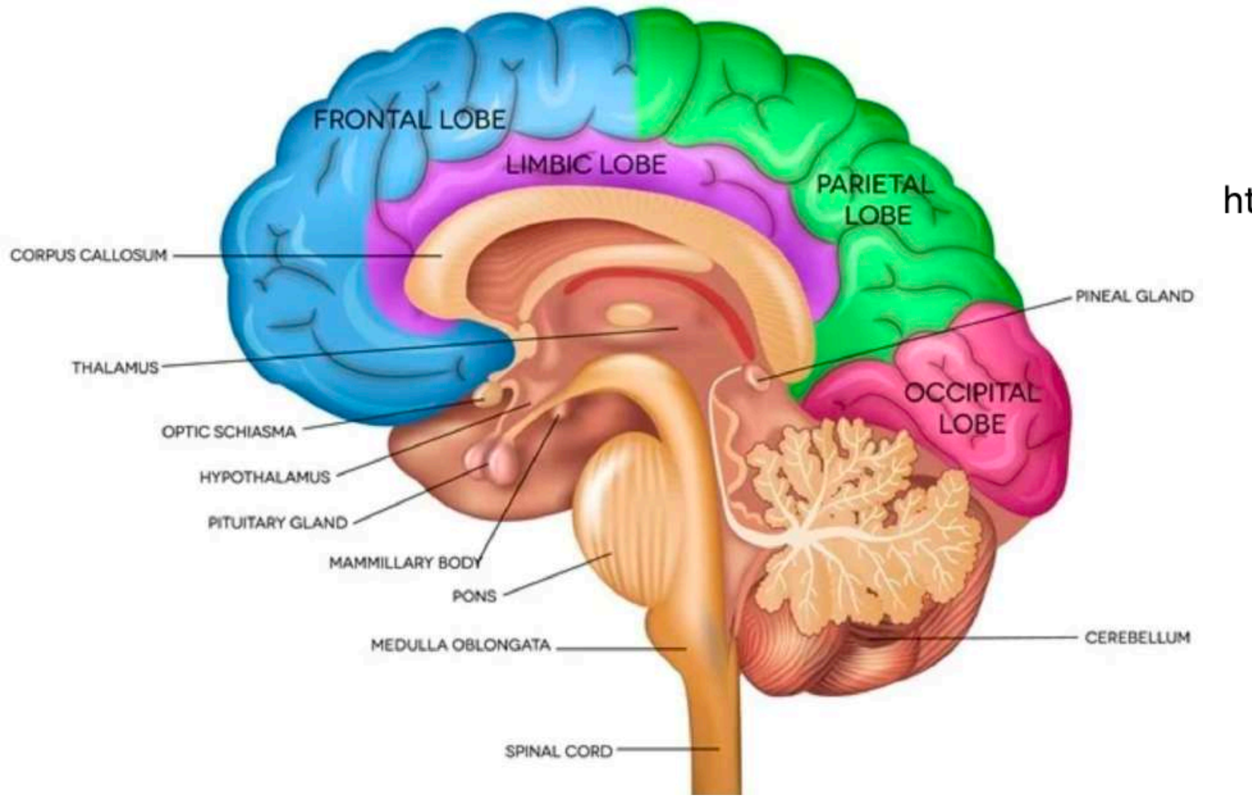

Image Credit: Tefi/Shutterstock.com

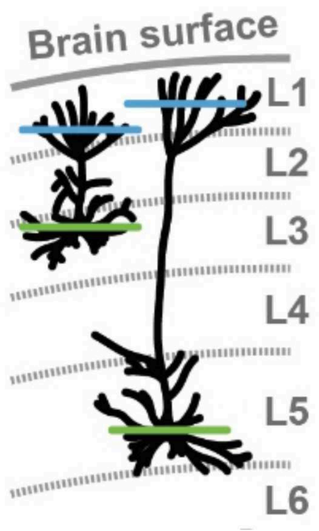

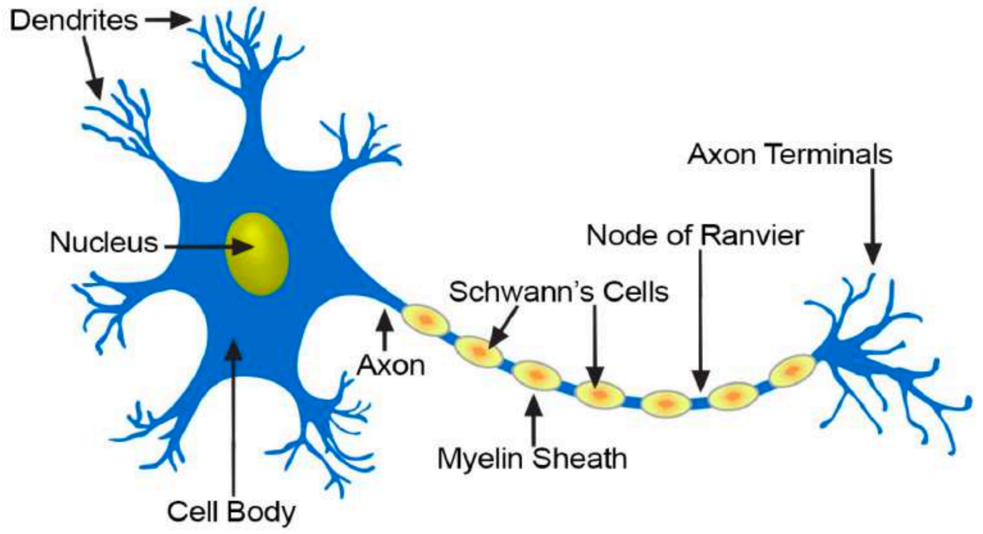

https://training.seer.cancer.gov/anatomy/nervous/tissue.html

Do we really need all of this complexity?

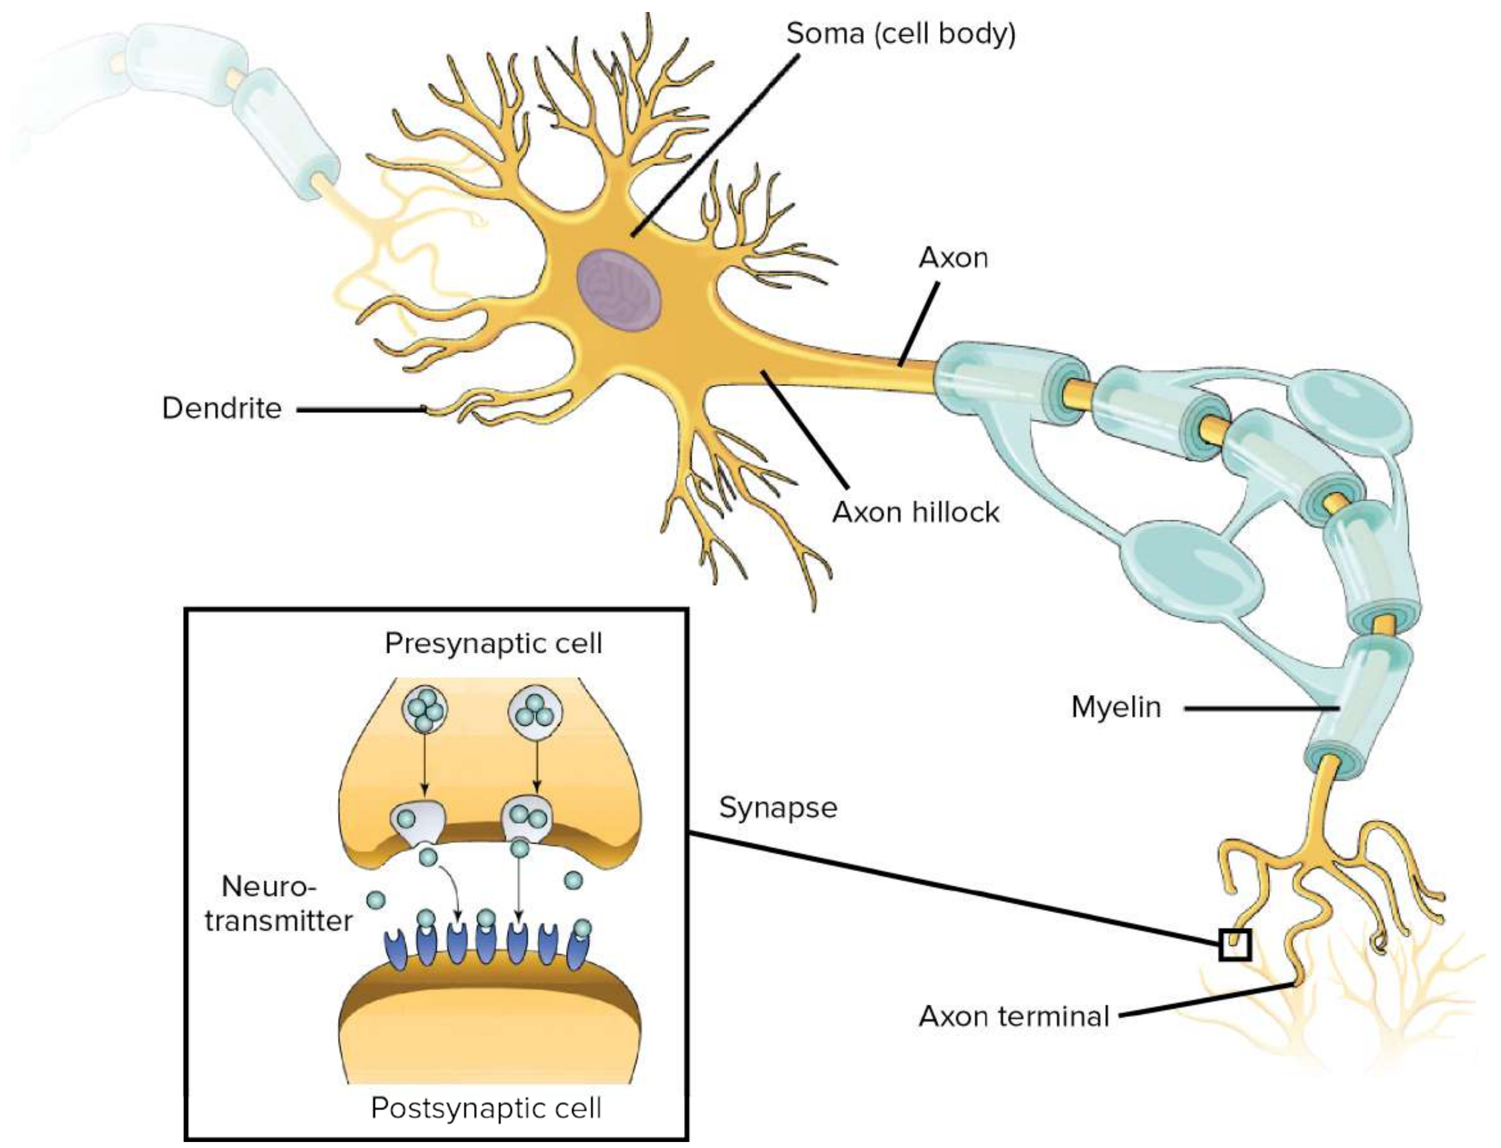

Image modified from "Neurons and glial cells: Figure 2" and "Synapse," by OpenStax College, Biology (CC BY 3.0).

How experimentalists look at neurons

One way: pretend the brain is in fact homogeneous with very simple neurons and see how far you can go!

How theorists look at neurons

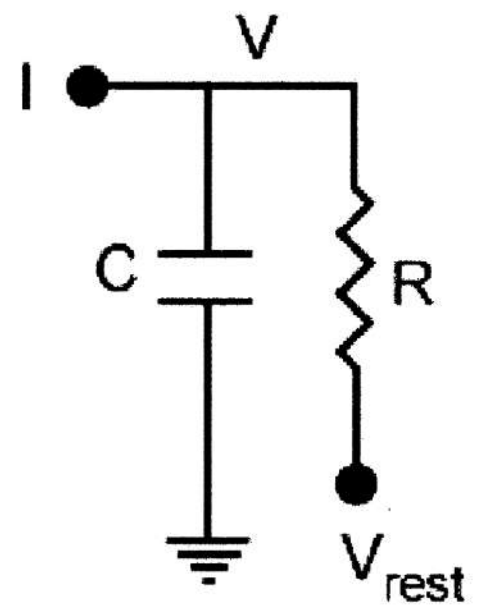

Abbott, 1999

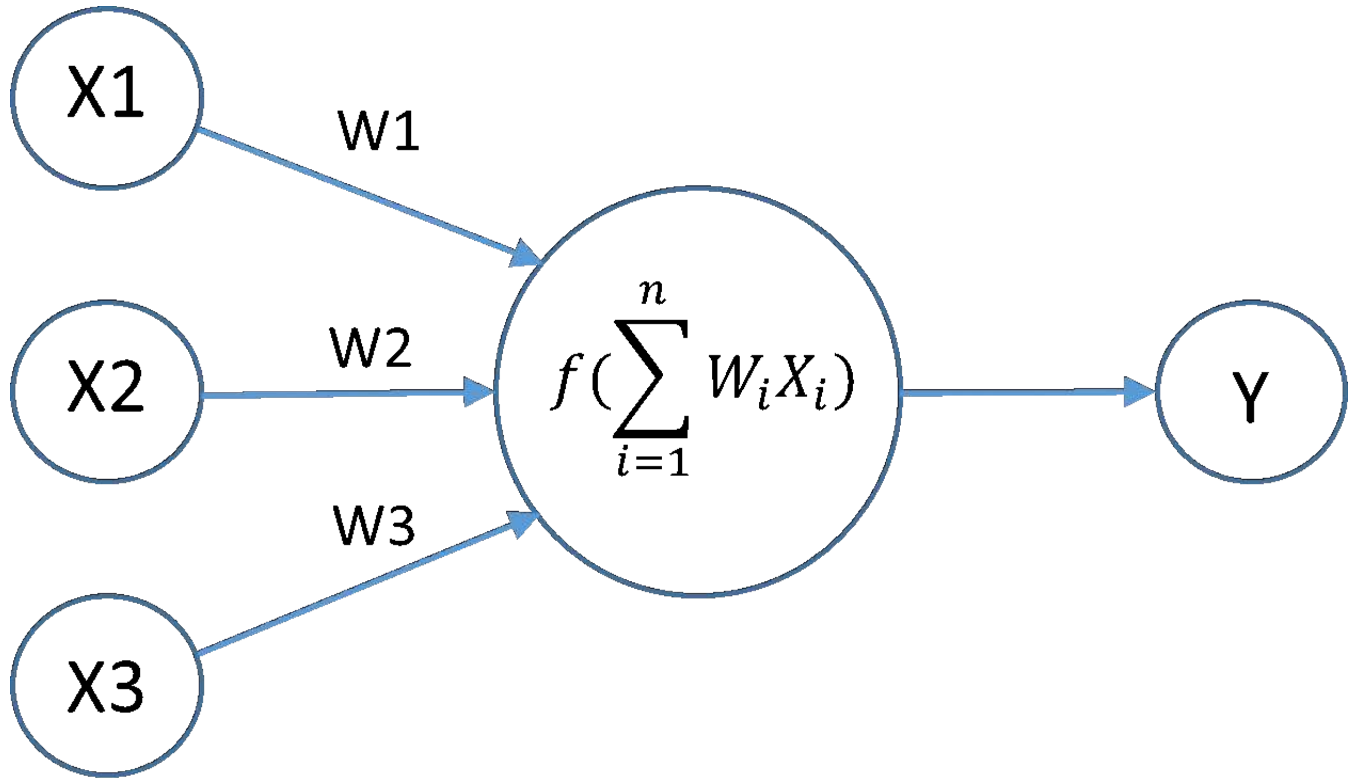

https://medium.com/chingu/neuron-explained-using-simple-algebra-example-b18f5e280845

Neurons communicate via propagating electrical signals

- 18th cent: Galvani and his wife observed that frog's legs contracted when stimulated by electricity (led to the first battery by Volta and to *Frankenstein* by Shelley!)

- 19th cent: discovery of cells, voltage across cell membrane, action potential (speed determined by von Helmholtz, who also studied vision)

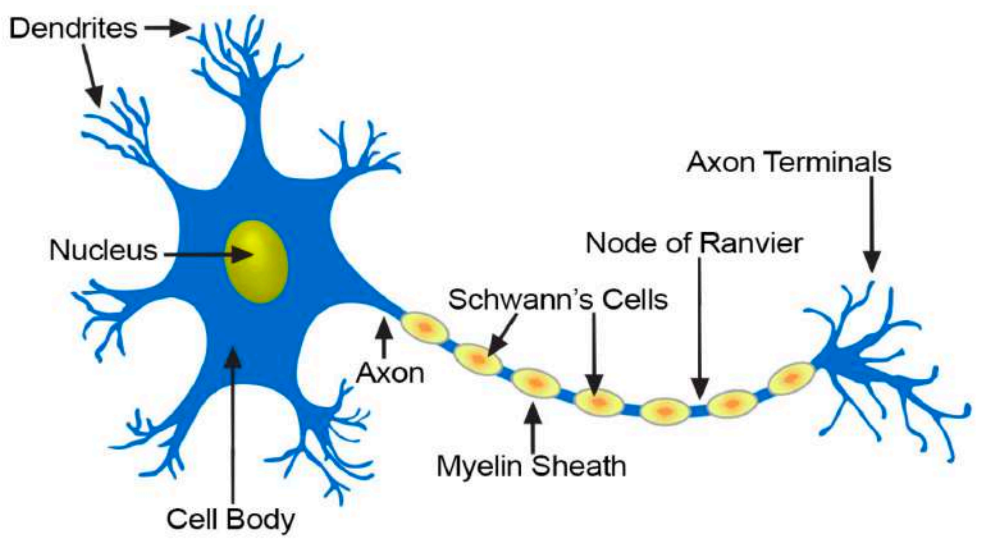

https://training.seer.cancer.gov/anatomy/nervous/tissue.html

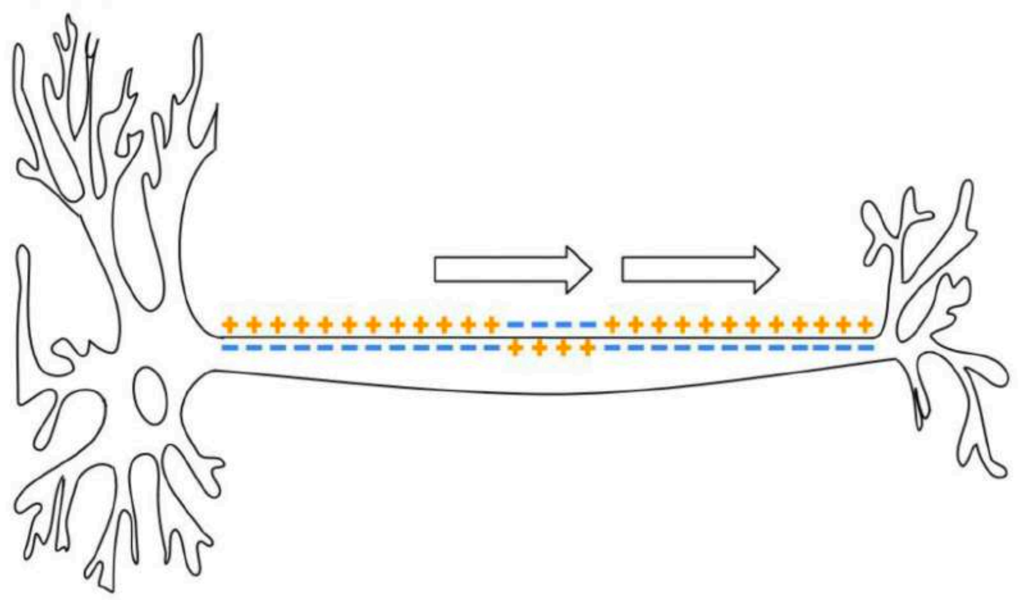
https://en.wikipedia.org/wiki/Action_potential#/media/File:Action_Potential.gif

Action Potential

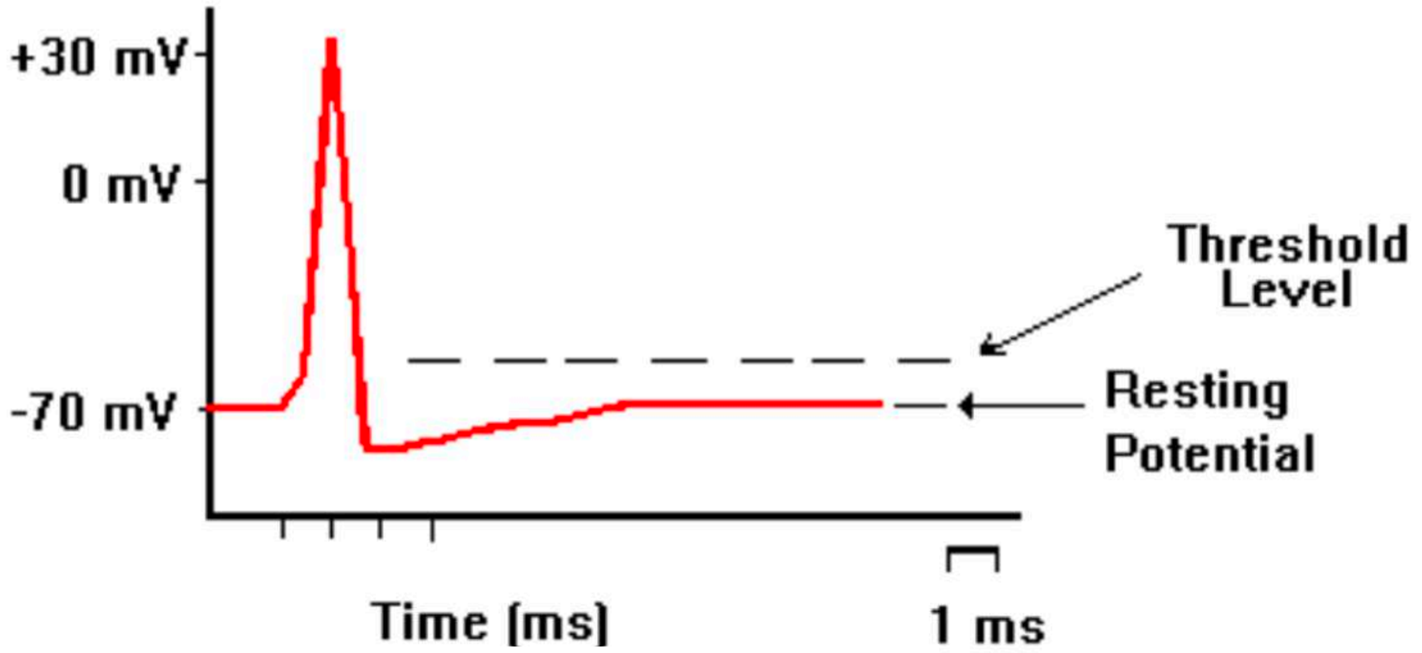

faculty.washington.edu/chudler/ap.html

Eric H. Chudler

Initial neuronal models included dynamic circuit and static feedforward models

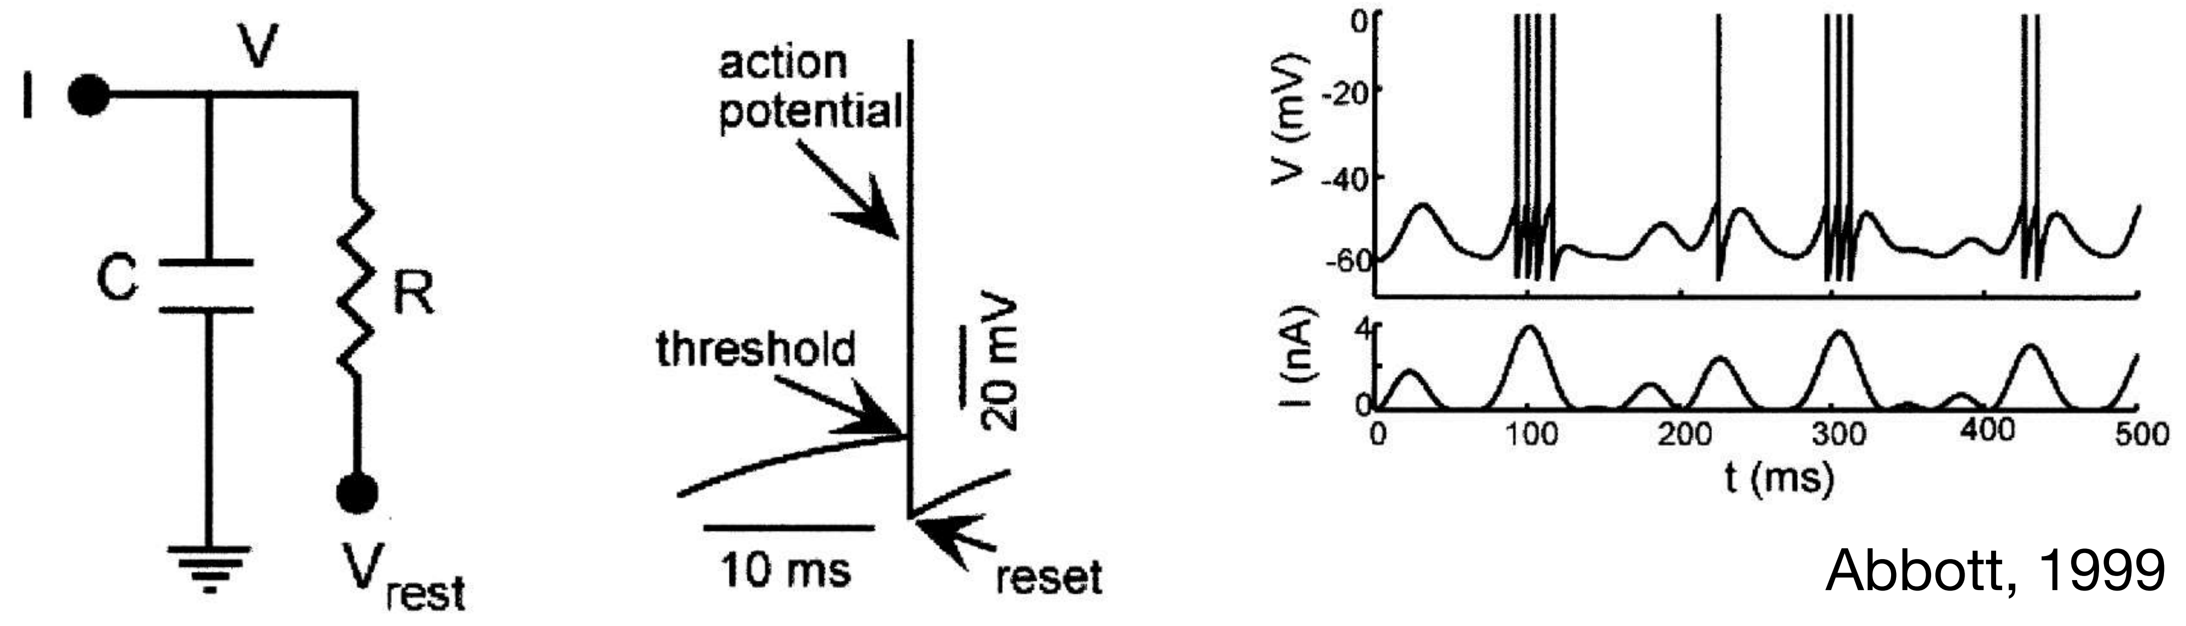

Abbott, 1999

- 1907: first circuit model of action potentials by Lapicque

$$\tau V(t) = -V(t) + I(t)$$

- 1943: first model of neuronal computations by McCullough and Pitts

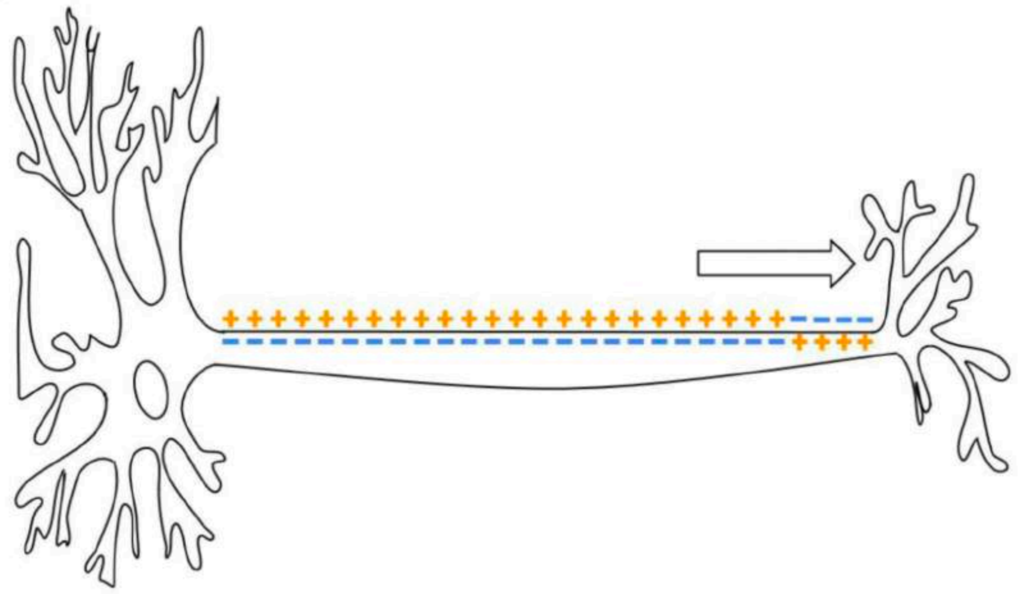

https://en.wikipedia.org/wiki/Action_potential#/media/File:Action_Potential.gif

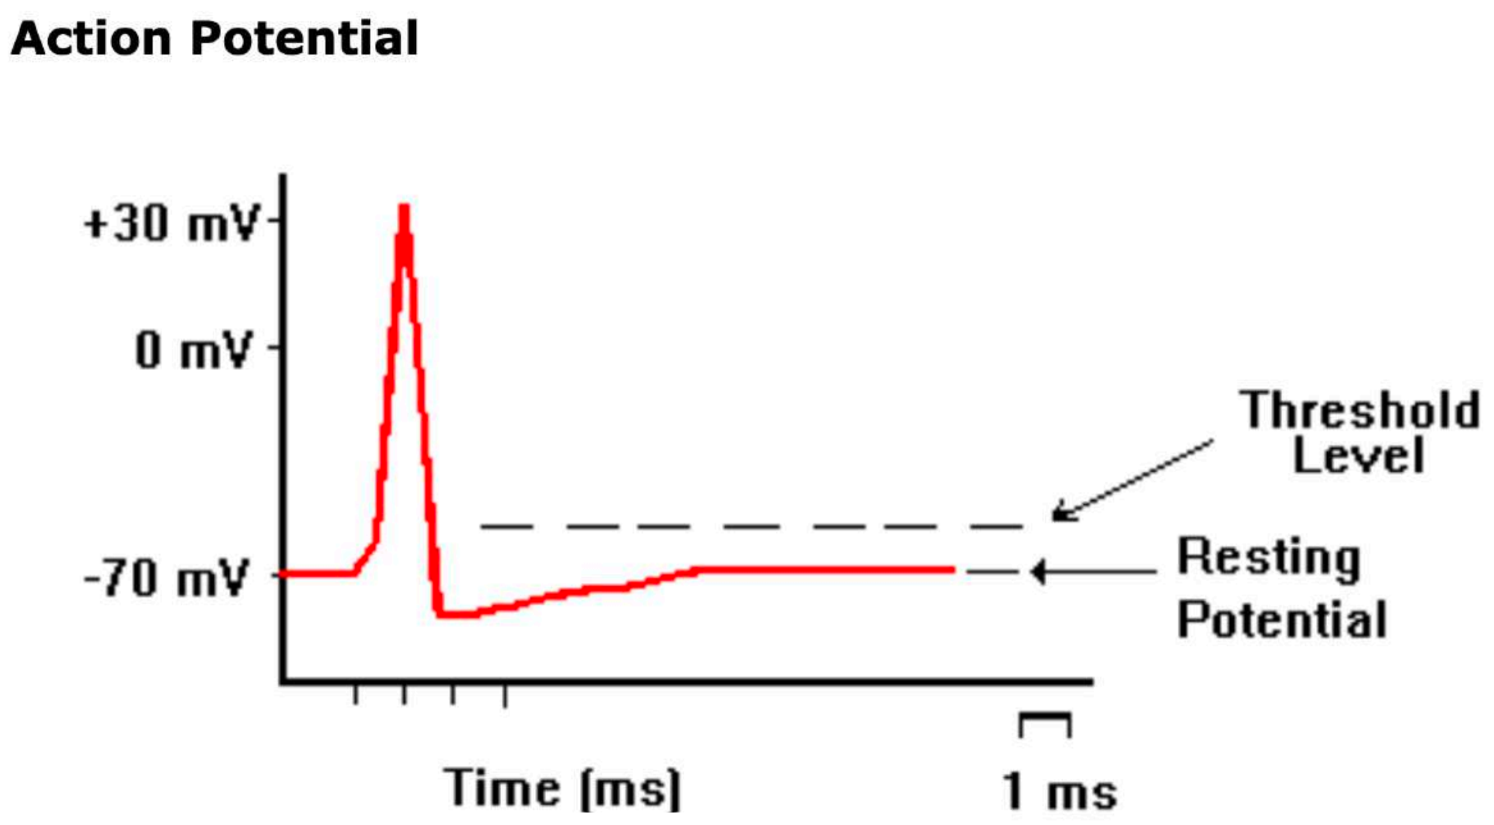

faculty.washington.edu/chudler/ap.html

Eric H. Chudler

Initial neuronal models included dynamic circuit and static feedforward models

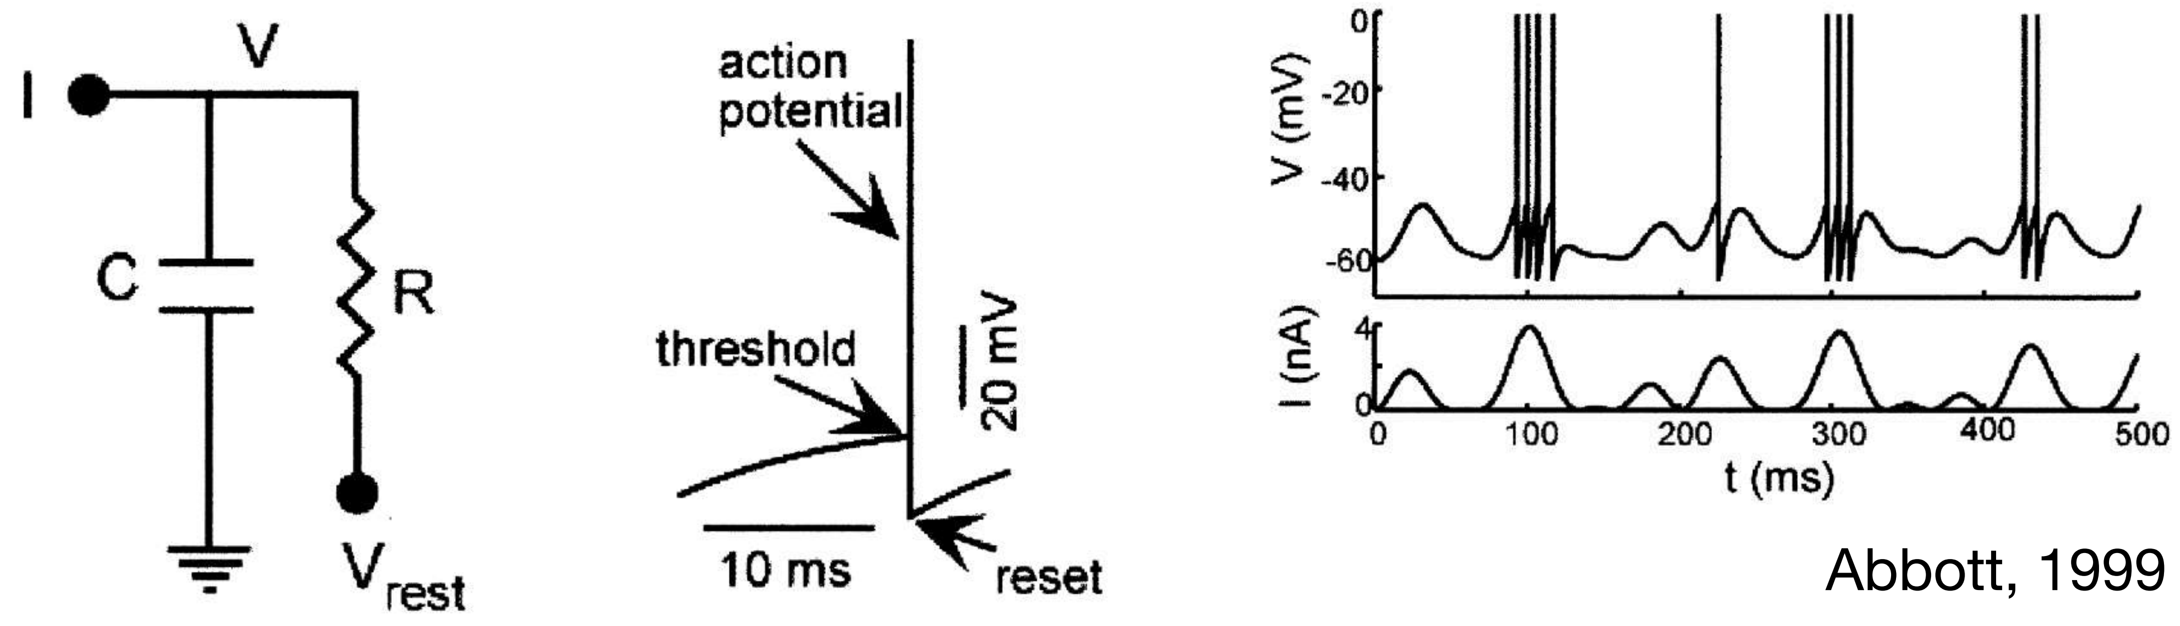

Abbott, 1999

- 1907: first circuit model of action potentials by Lapicque

$$\tau V(t) = -V(t) + I(t)$$

- 1943: first model of neuronal computations by McCullough and Pitts

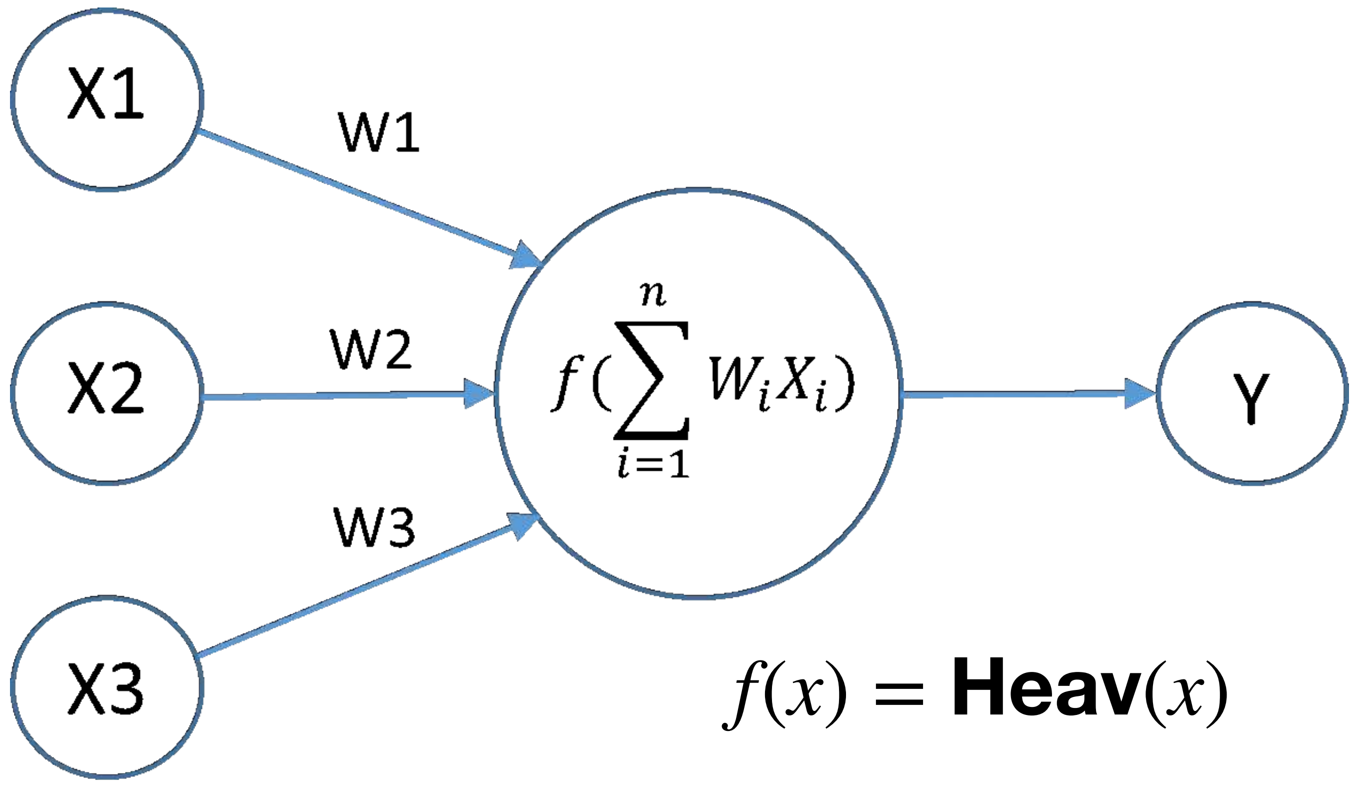

$Y \in \{0,1\}$

$f(x) = \textbf{Heav}(x)$

$$W1 = W2 = W3 = \ldots = WN$$

(fixed)

Initial neuronal models included dynamic circuit and static feedforward models

- 1907: first circuit model of action potentials by Lapicque
- 1943: first model of neuronal computations by McCullough and Pitts
- 1952: Hodgkin-Huxley model (1963: Nobel)
  - Much more accurate model of action potential

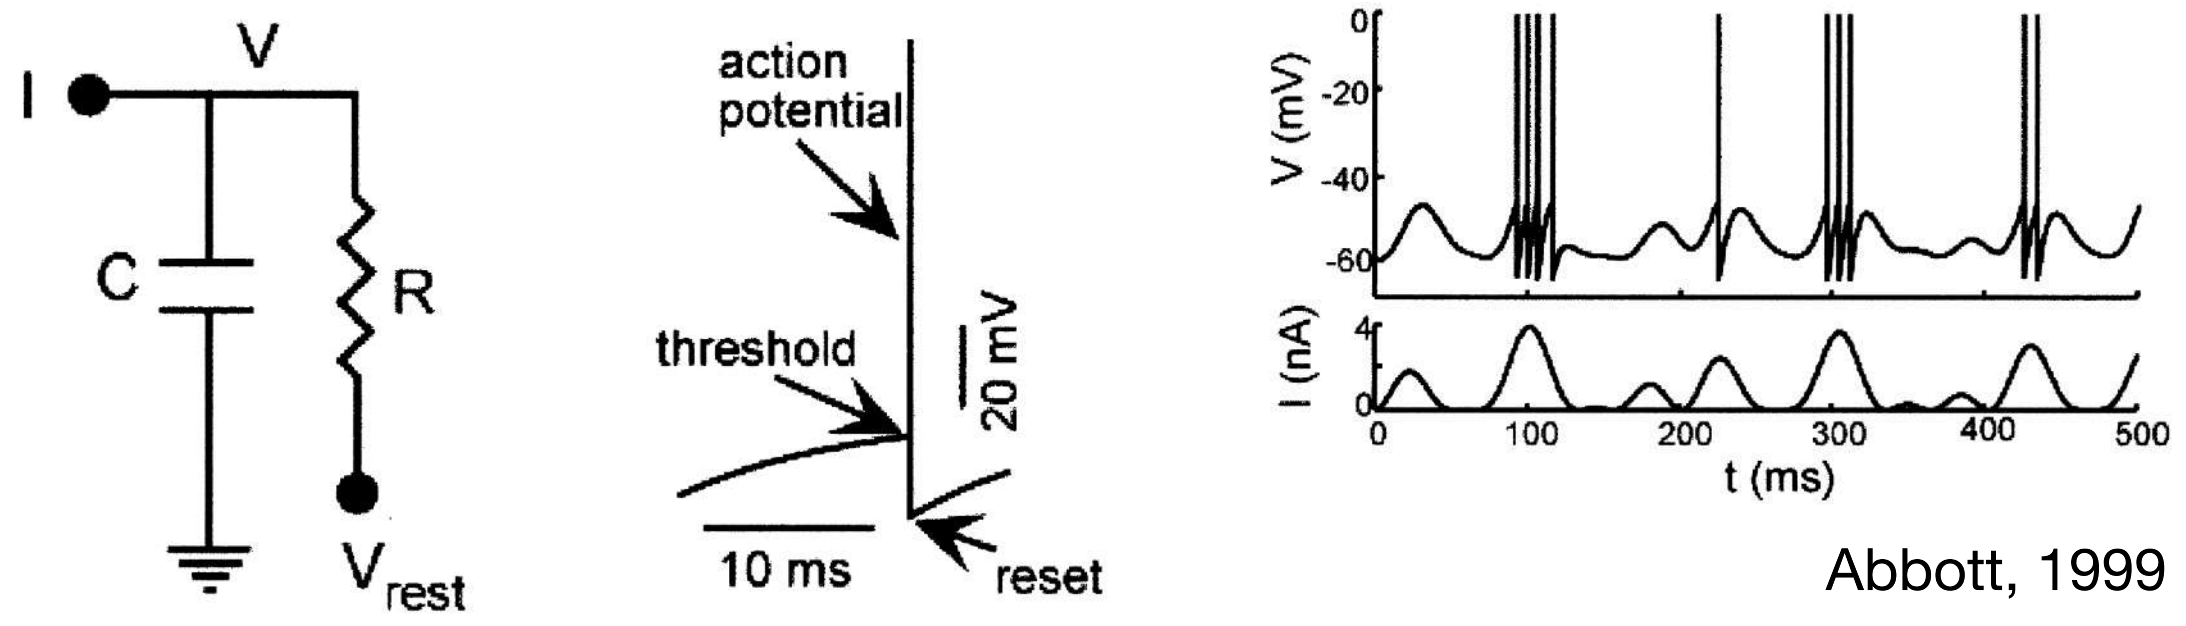

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

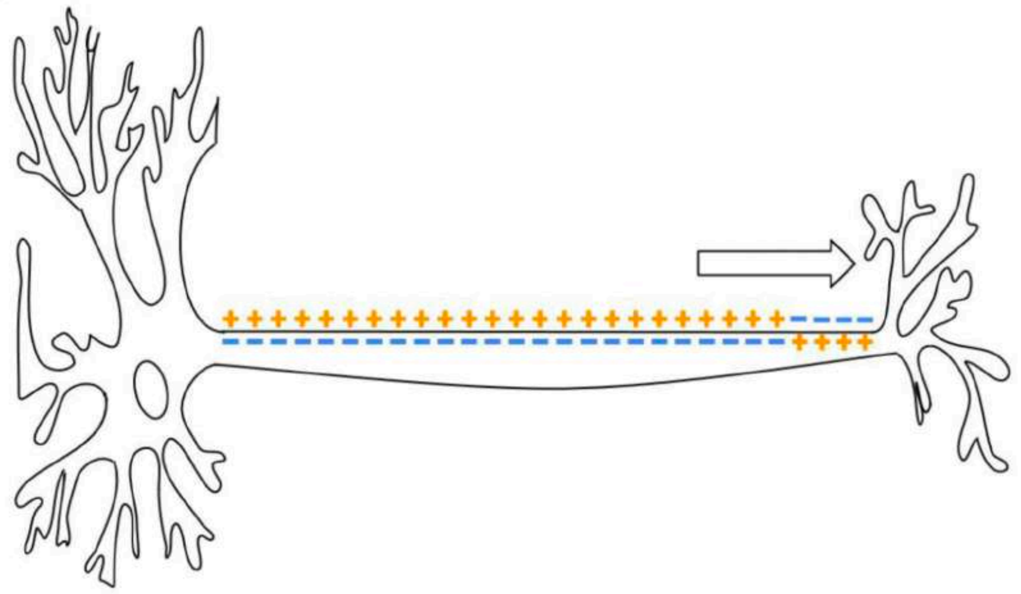

https://en.wikipedia.org/wiki/Action_potential#/media/File:Action_Potential.gif

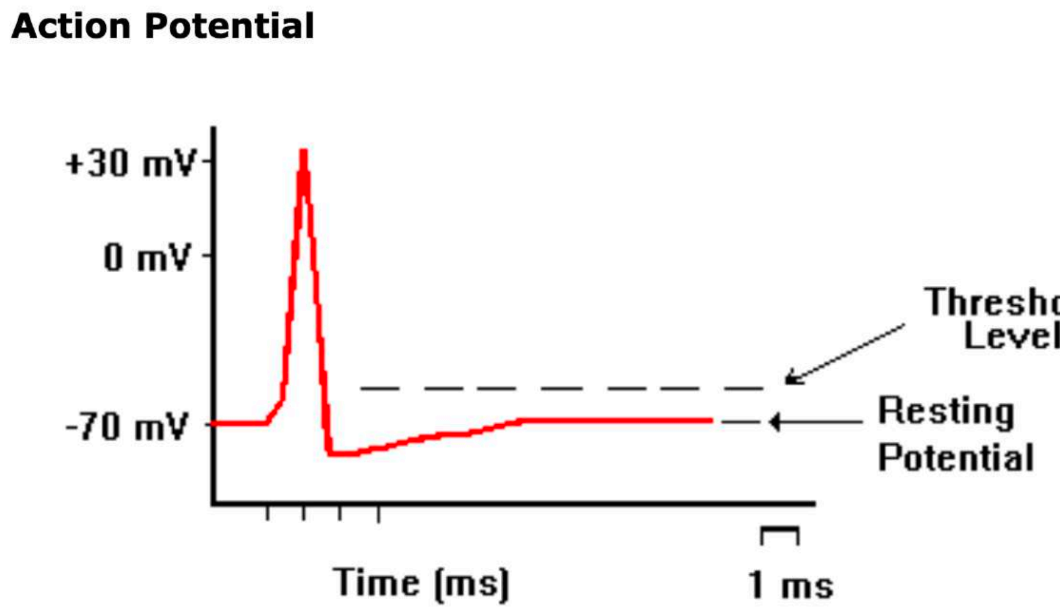

faculty.washington.edu/chudler/ap.html

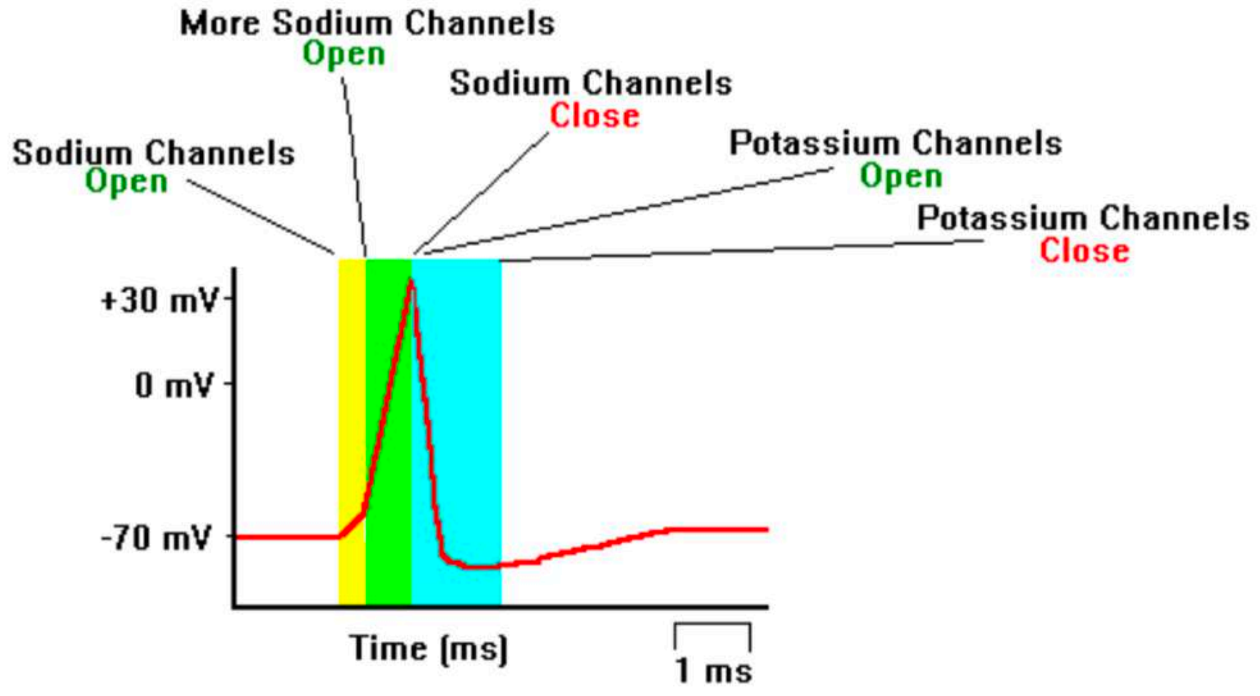

Eric H. Chudler

Initial neuronal models included dynamic circuit and static feedforward models

- 1907: first circuit model of action potentials by Lapicque
- 1943: first model of neuronal computations by McCullough and Pitts
- 1952: Hodgkin-Huxley model (1963: Nobel)
  - Much more accurate model of action potential
  - Sodium and potassium conductances modeled through nonlinear equations

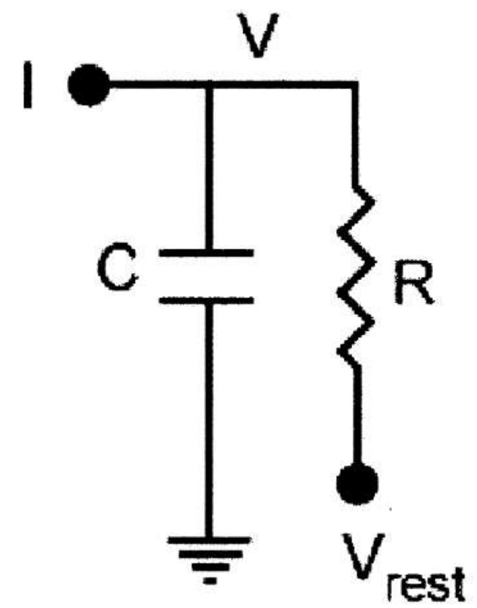

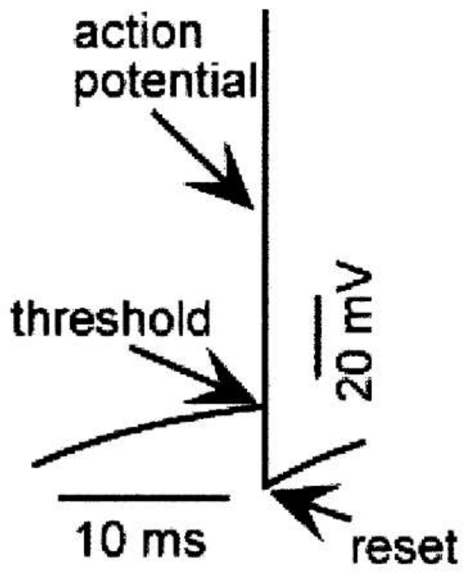

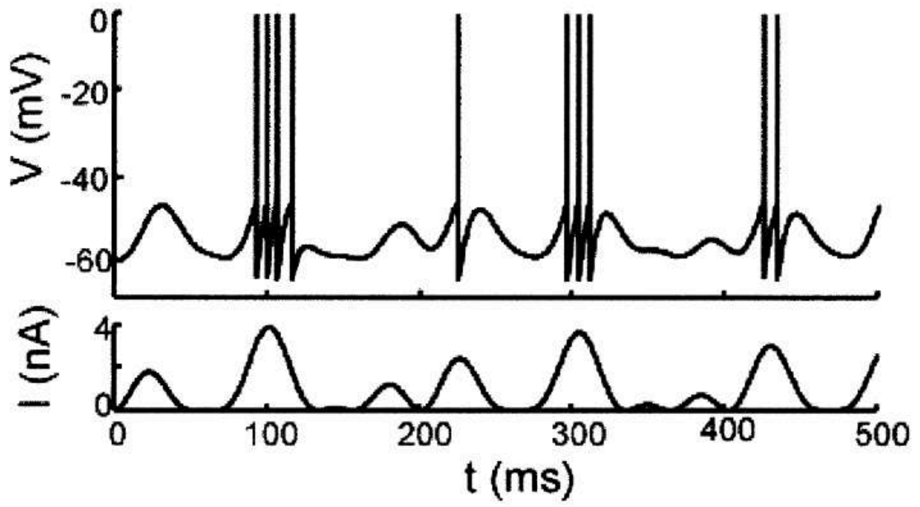

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

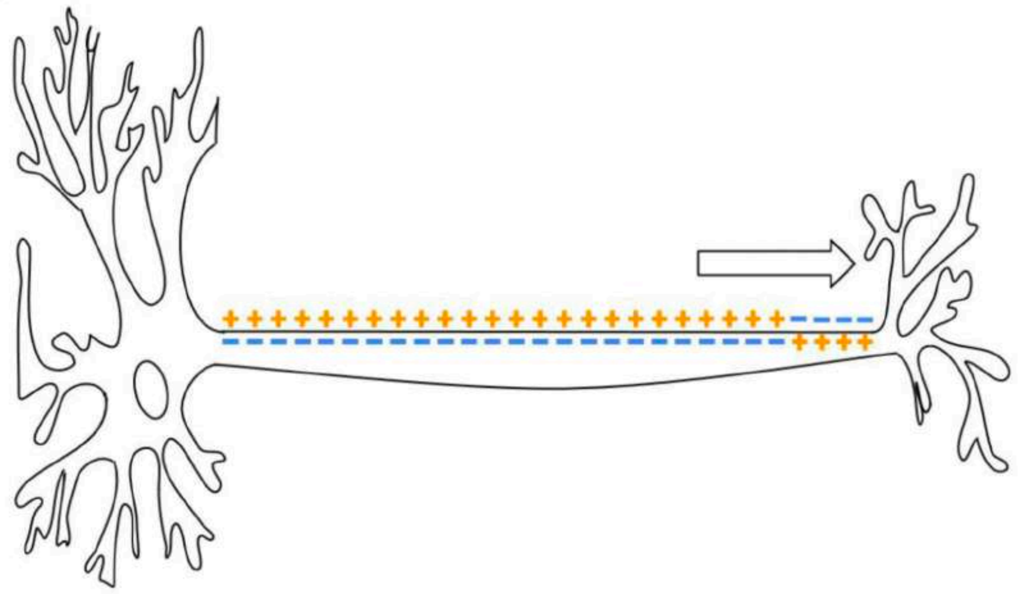

https://en.wikipedia.org/wiki/Action_potential#/media/File:Action_Potential.gif

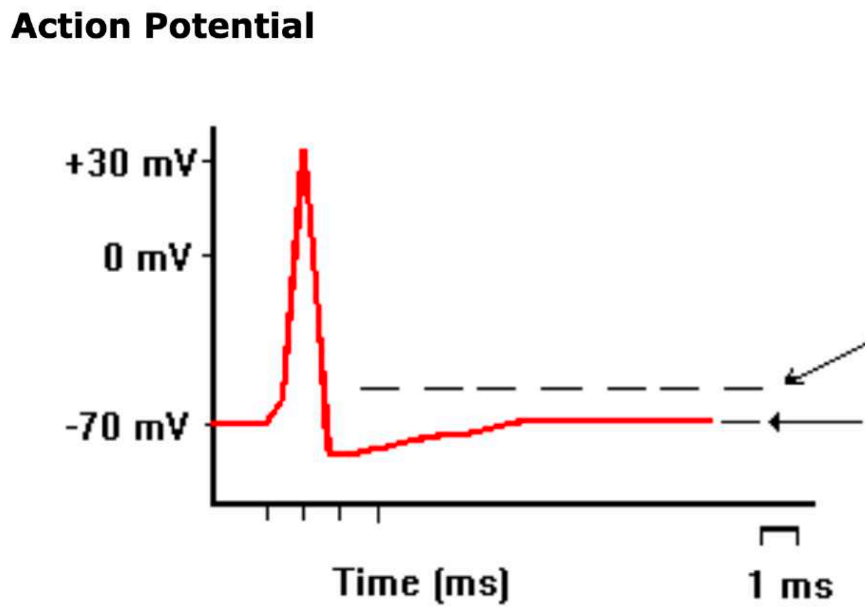

faculty.washington.edu/chudler/ap.html

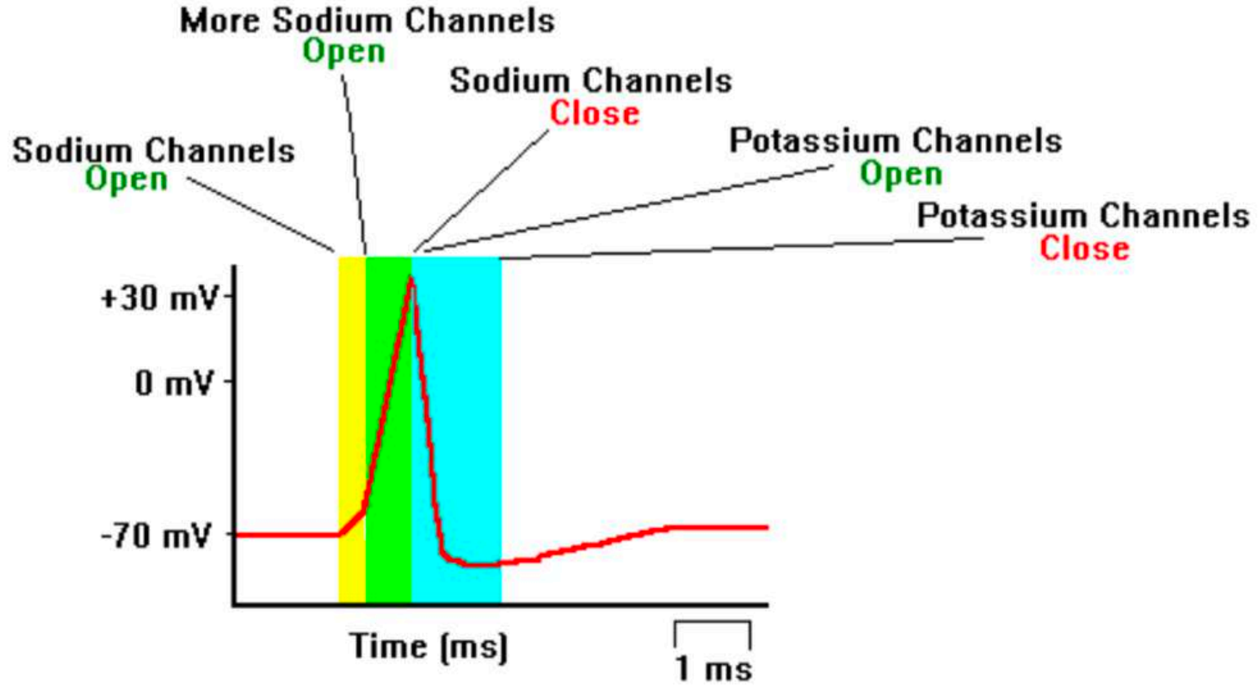

Eric H. Chudler

Initial neuronal models included dynamic circuit and static feedforward models

- 1907: first circuit model of action potentials by Lapicque
- 1943: first model of neuronal computations by McCullough and Pitts
- 1952: Hodgkin-Huxley model (1963: Nobel)
  - Much more accurate model of action potential
  - Sodium and potassium conductances modeled through nonlinear equations
  - System of 4 ODEs

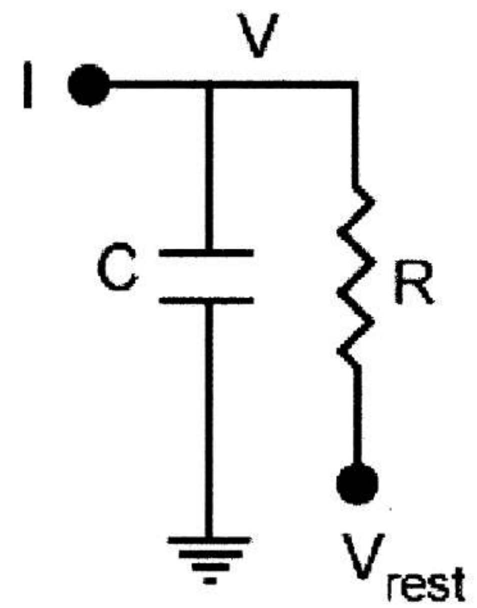

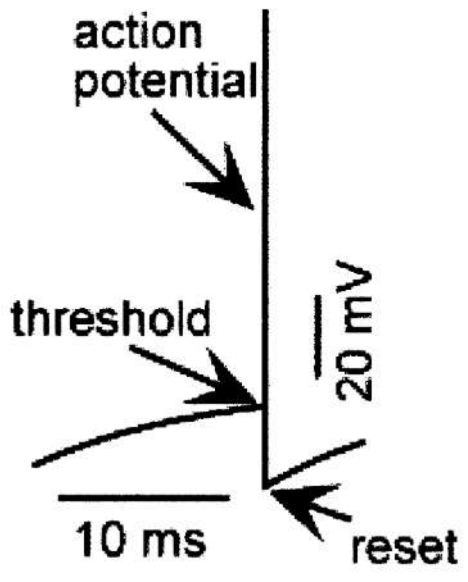

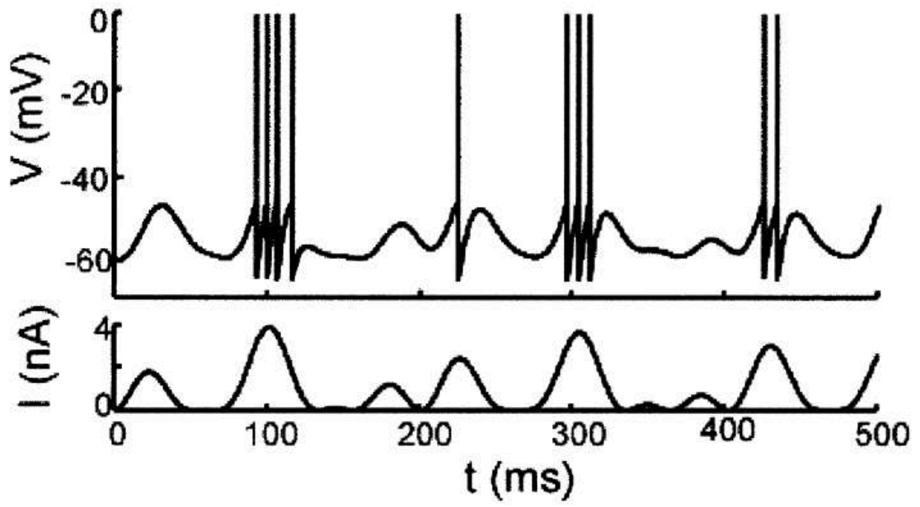

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

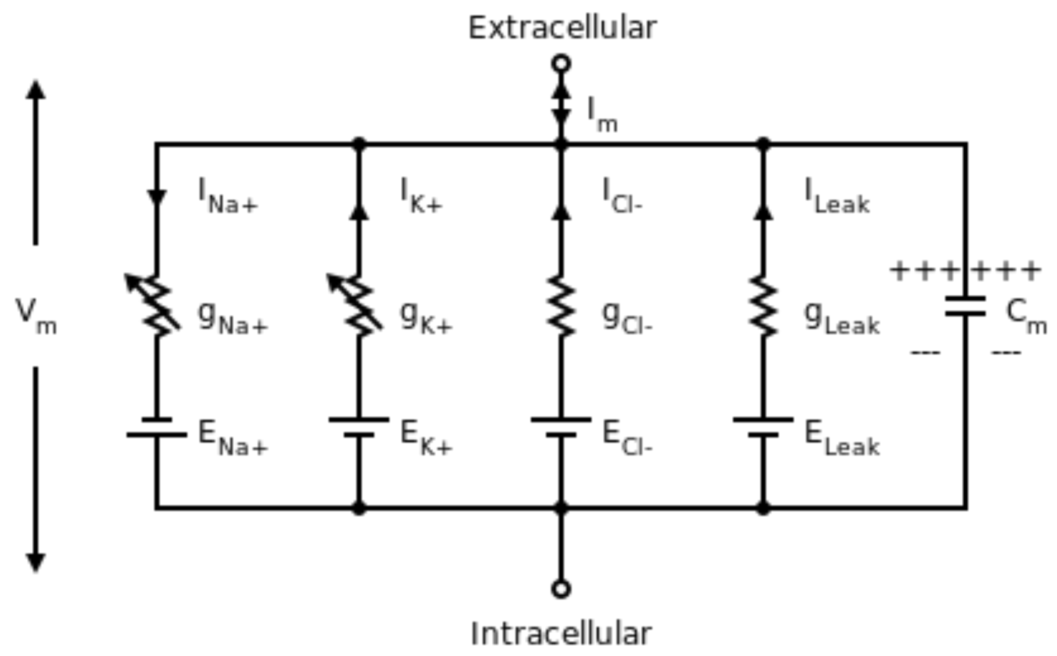

https://upload.wikimedia.org/wikipedia/commons/thumb/1/1f/MembraneCircuit.svg/336px-MembraneCircuit.svg

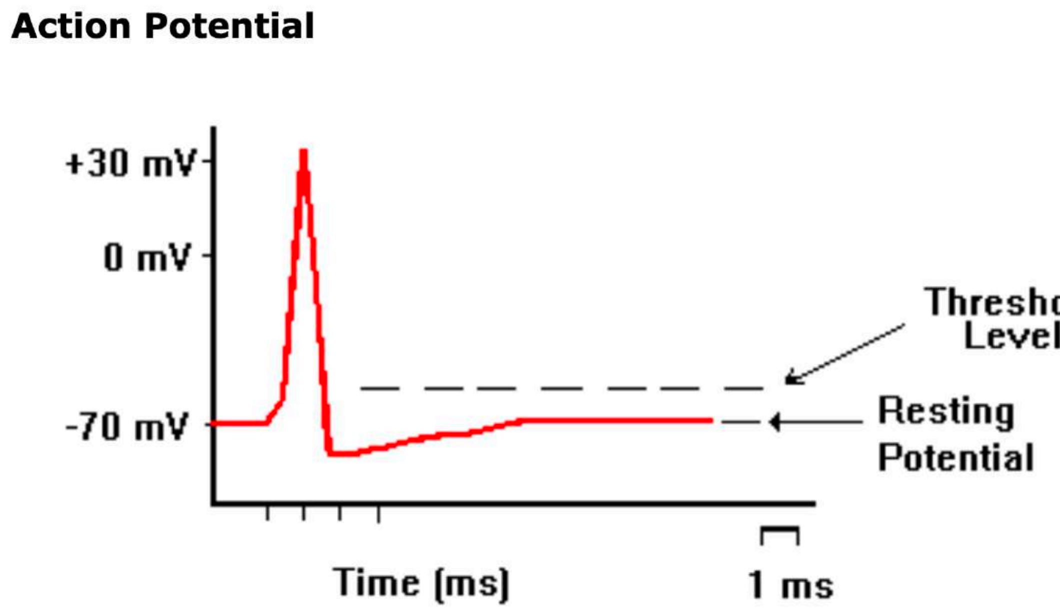

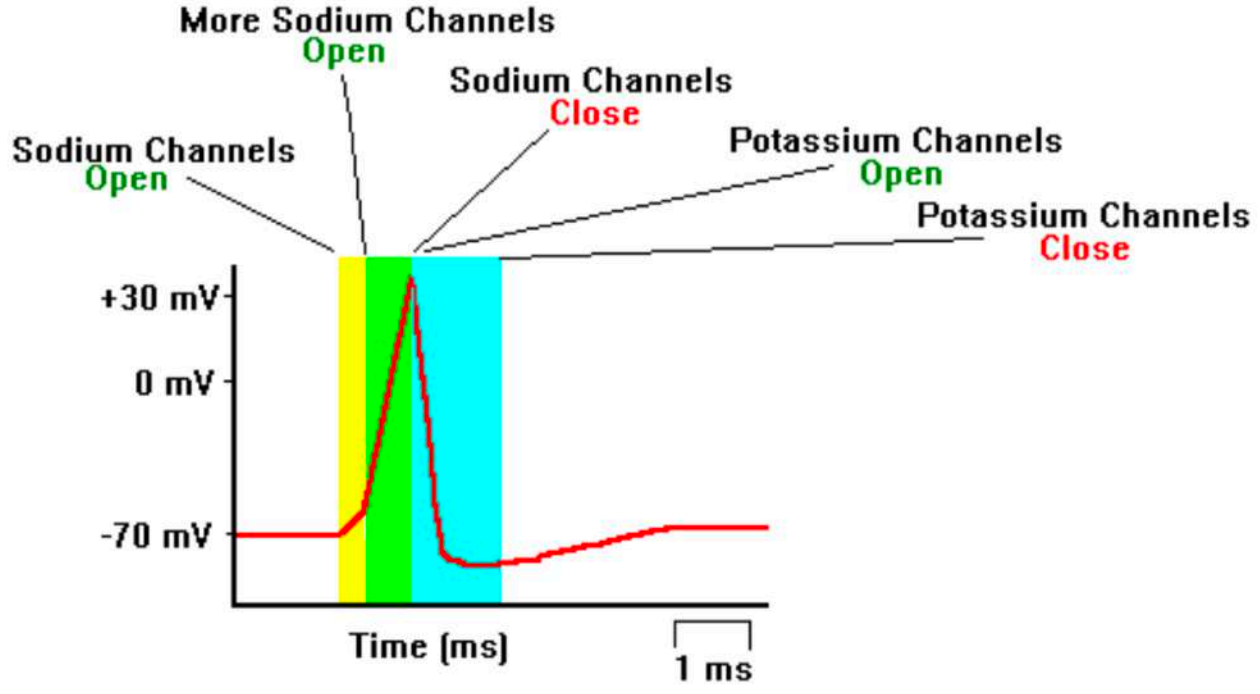

faculty.washington.edu/chudler/ap.html

Eric H. Chudler

Initial neuronal models included dynamic circuit and static feedforward models

$$C\frac{dV}{dt} = I - g_{Na}m^3h(V - E_{Na}) - g_K n^4(V - E_K) - g_L(V - E_L)$$
$$\frac{dm}{dt} = a_m(V)(1-m) - b_m(V)m$$
$$\frac{dh}{dt} = a_h(V)(1-h) - b_h(V)h$$
$$\frac{dn}{dt} = a_n(V)(1-n) - b_n(V)n$$
$$a_m(V) = 0.1(V+40)/(1-\exp(-(V+40)/10))$$
$$b_m(V) = 4\exp(-(V+65)/18)$$
$$a_h(V) = 0.07\exp(-(V+65)/20)$$
$$b_h(V) = 1/(1+\exp(-(V+35)/10))$$
$$a_n(V) = 0.01(V+55)/(1-\exp(-(V+55)/10))$$
$$b_n(V) = 0.125\exp(-(V+65)/80)$$

Hodgkin-Huxley equations

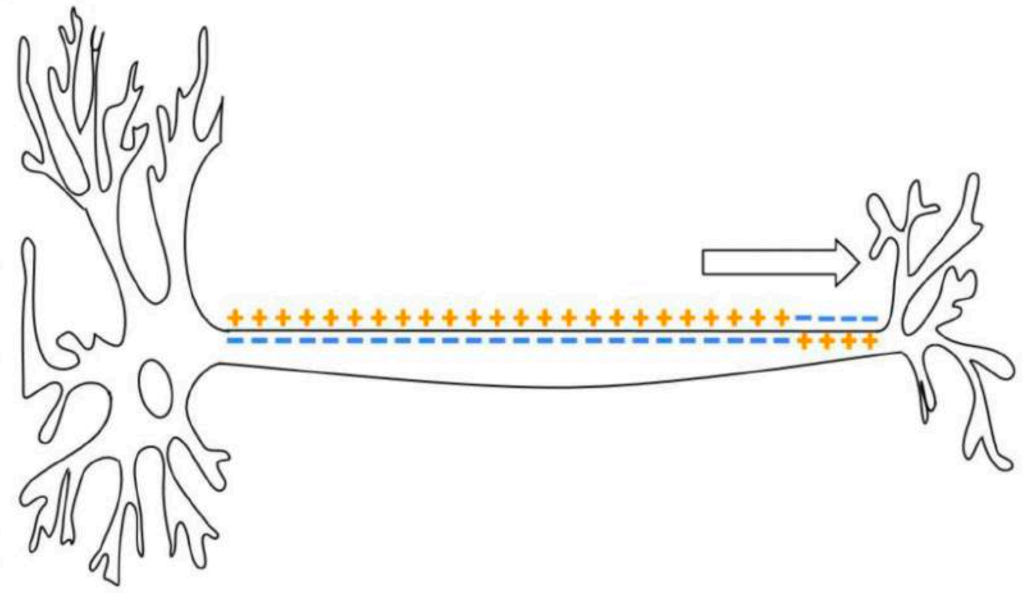

https://en.wikipedia.org/wiki/Action_potential#/media/File:Action_Potential.gif

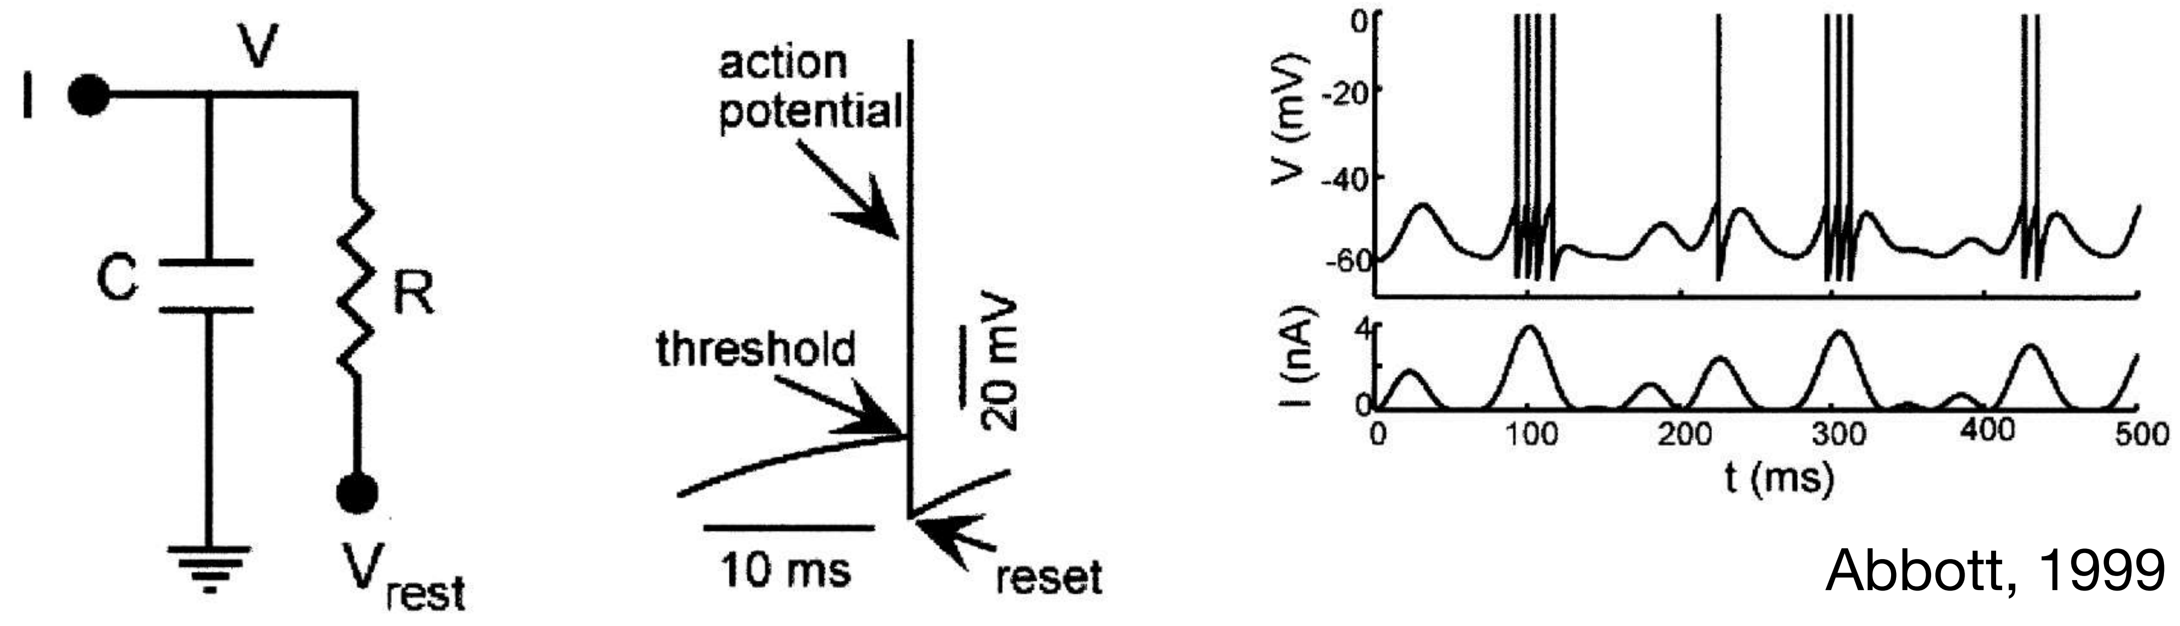

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

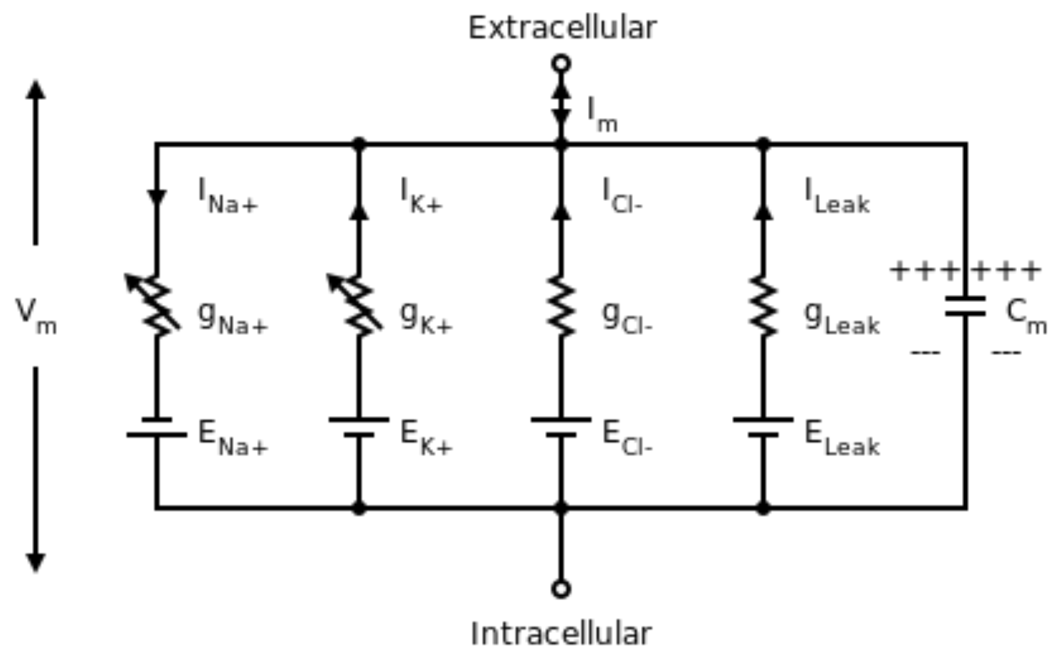

https://upload.wikimedia.org/wikipedia/commons/thumb/1/1f/MembraneCircuit.svg/336px-MembraneCircuit.svg

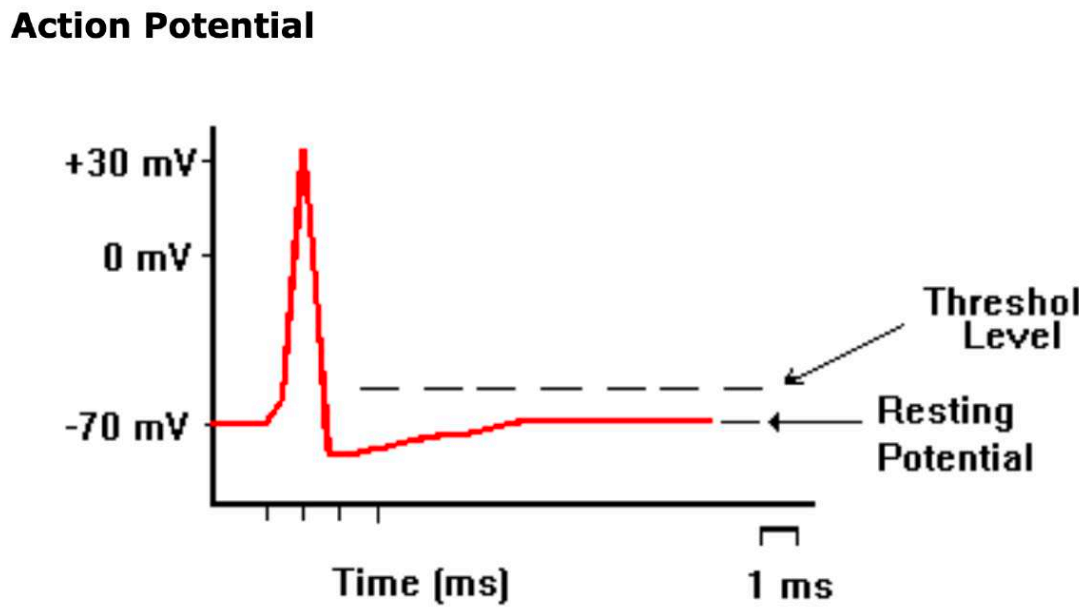

faculty.washington.edu/chudler/ap.html

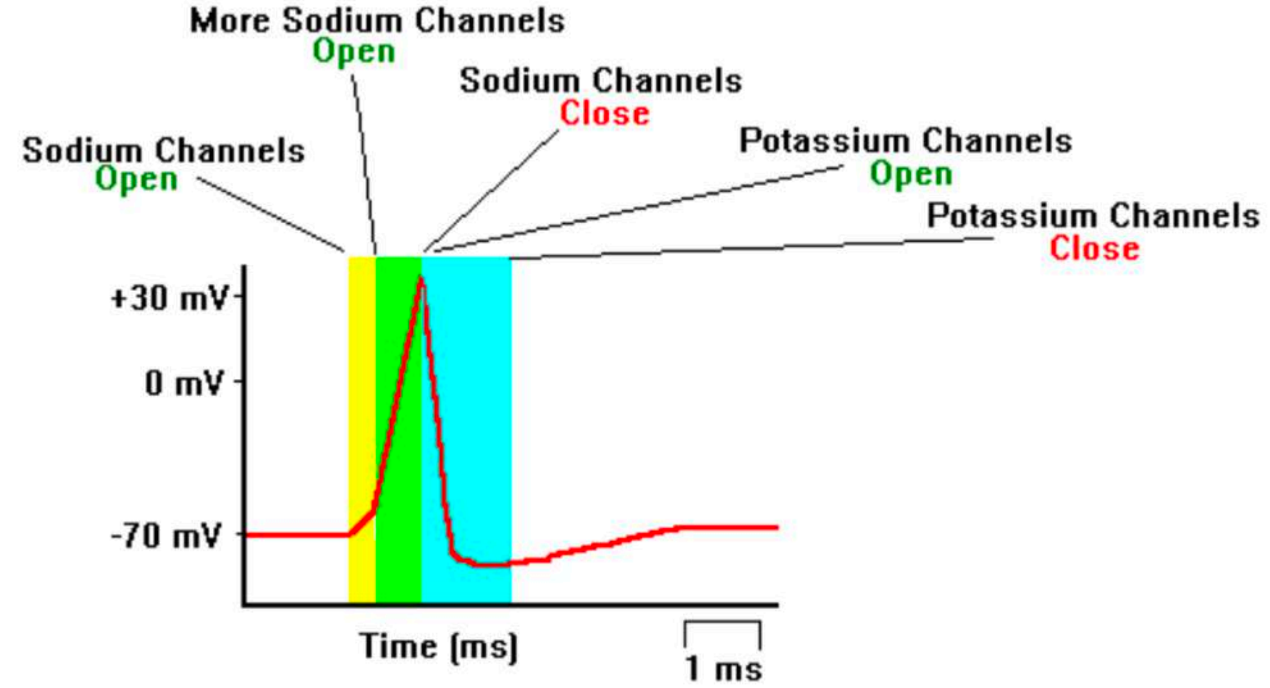

Eric H. Chudler

Mean-field models allow for tractable equations that capture large-scale dynamics

- 1907: first circuit model of action potentials by Lapicque

- 1943: first model of neuronal computations by McCullough and Pitts

- 1952: Hodgkin-Huxley model (1963: Nobel)

- 1956: Neural fields by Beurle

    - 1972, 1973: Wilson-Cowan equations include inhibition

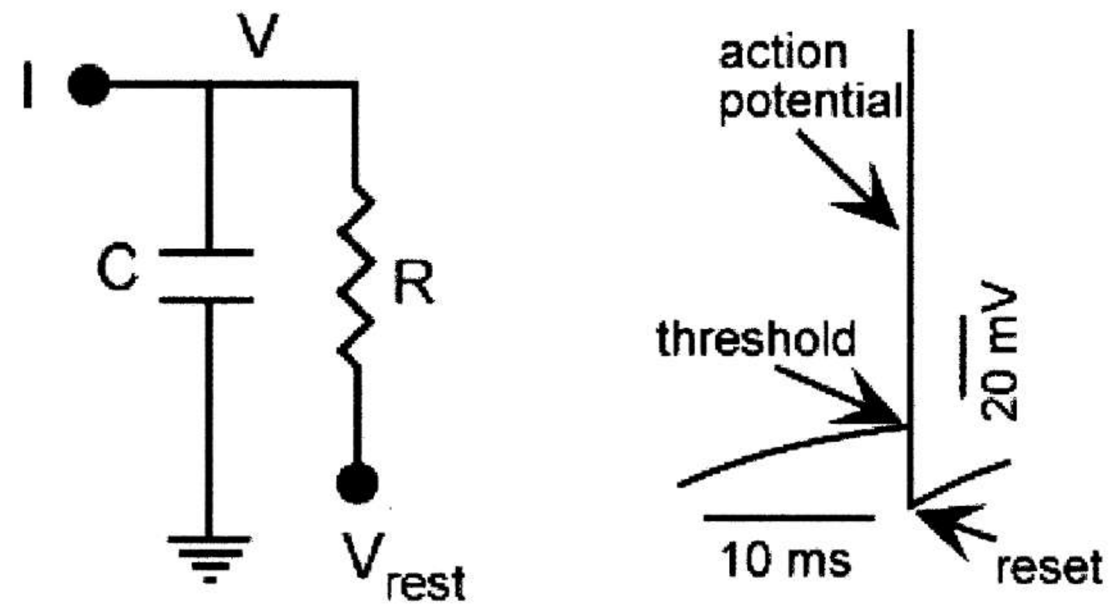

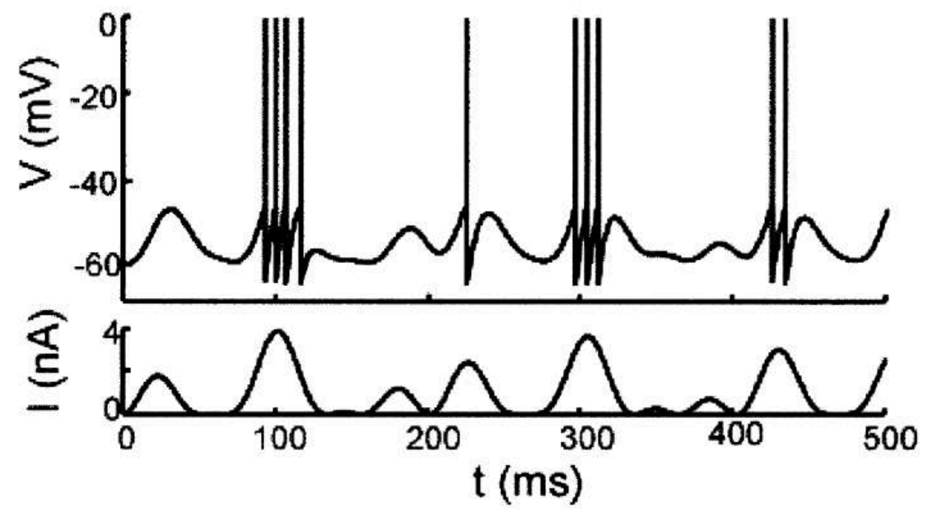

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

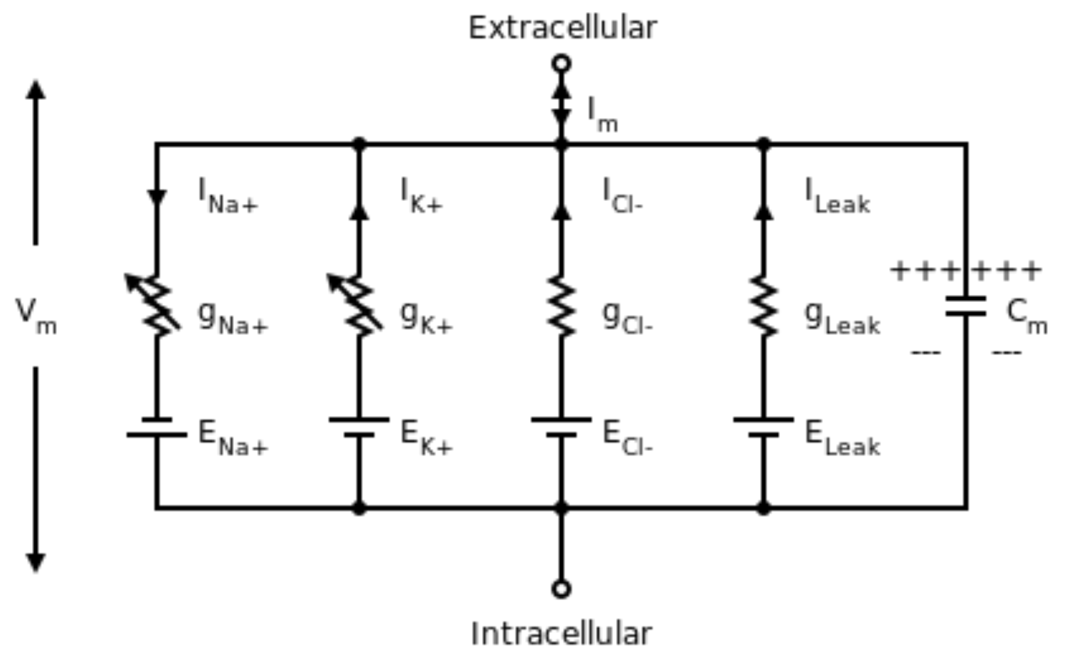

https://upload.wikimedia.org/wikipedia/commons/thumb/1/1f/MembraneCircuit.svg/336px-MembraneCircuit.svg

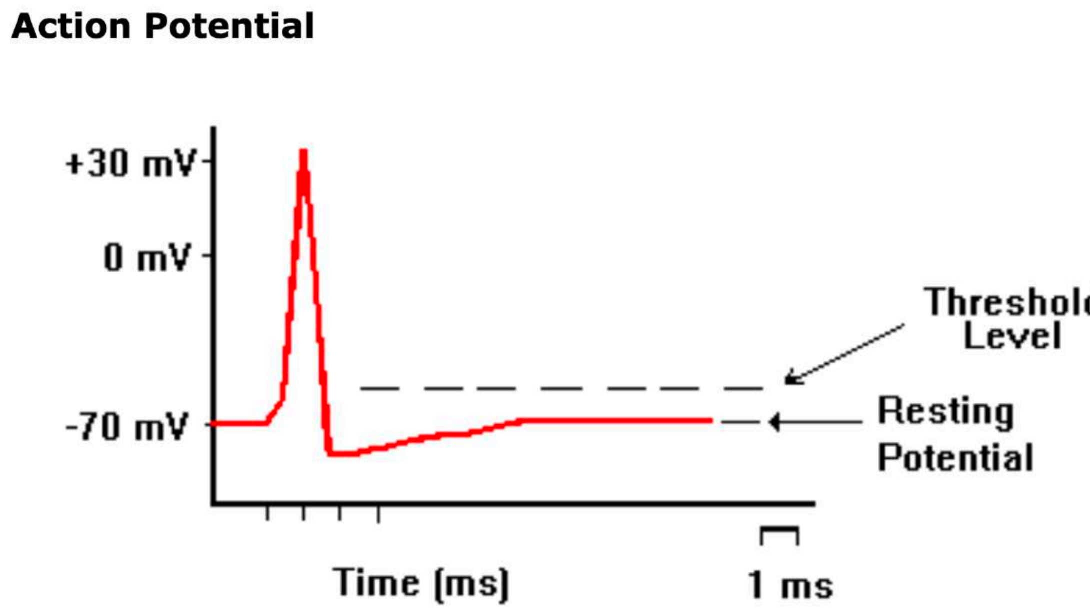

faculty.washington.edu/chudler/ap.html

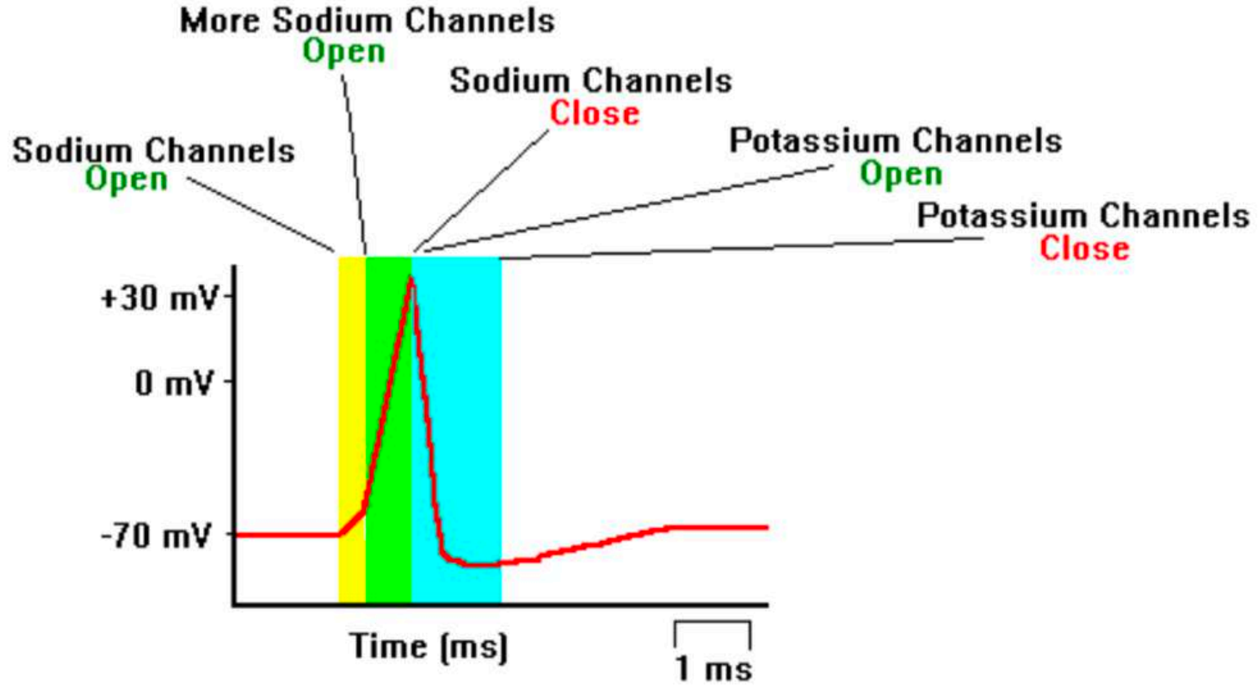

Eric H. Chudler

Mean-field models allow for tractable equations that capture large-scale dynamics

- 1907: first circuit model of action potentials by Lapicque
- 1943: first model of neuronal computations by McCullough and Pitts
- 1952: Hodgkin-Huxley model (1963: Nobel)
- 1956: Neural fields by Beurle
  - 1972, 1973: Wilson-Cowan equations include inhibition

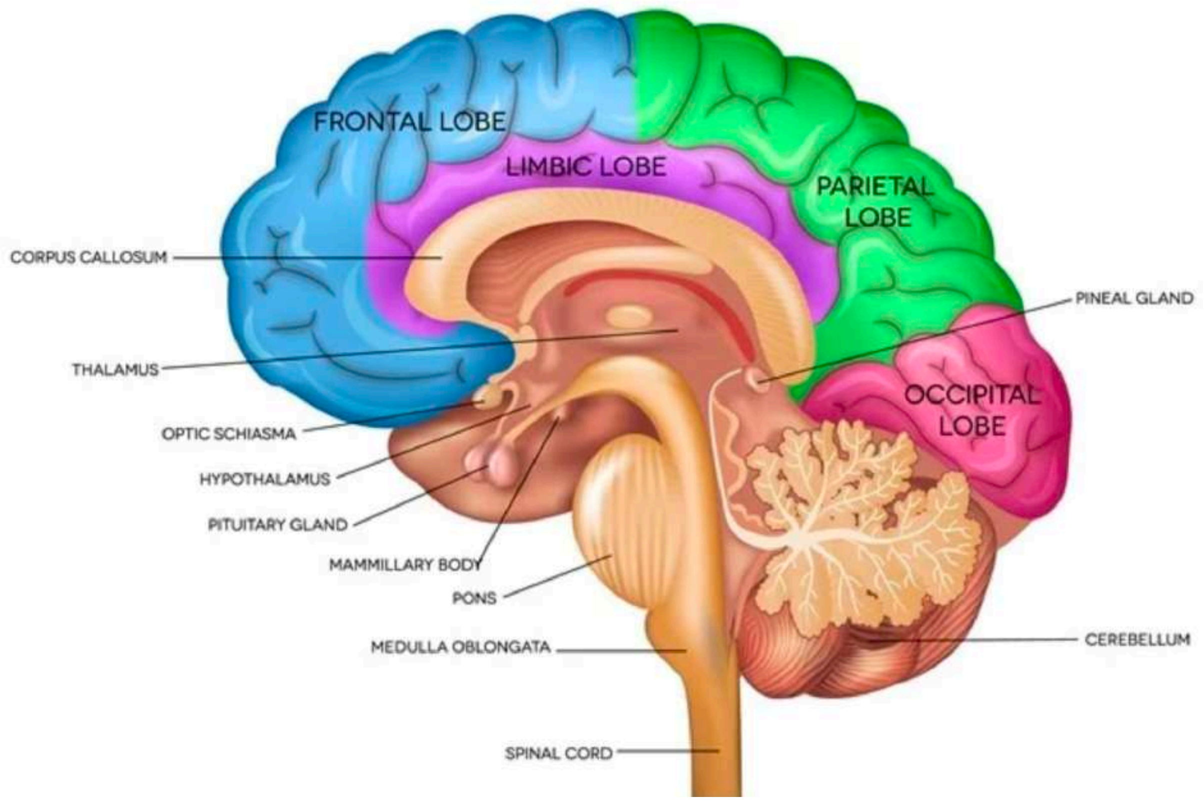

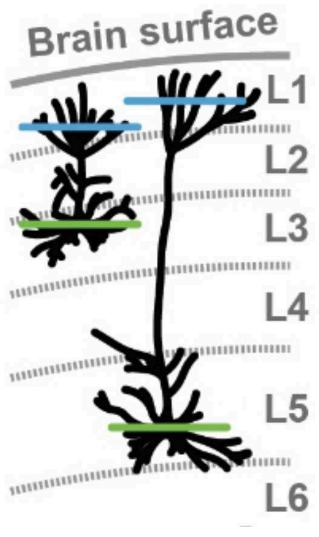

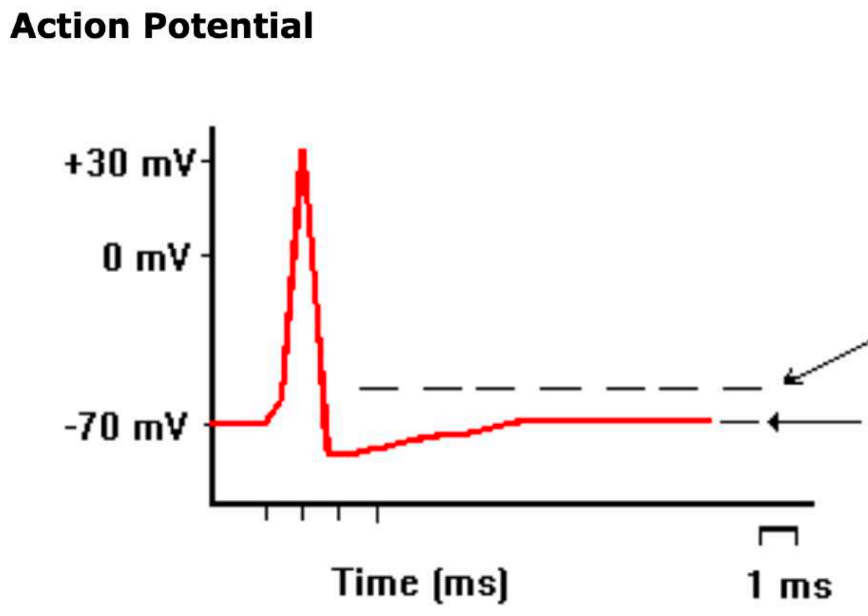

faculty.washington.edu/chudler/ap.html

https://upload.wikimedia.org/wikipedia/commons/thumb/1/1f/MembraneCircuit.svg/336px-MembraneCircuit.svg

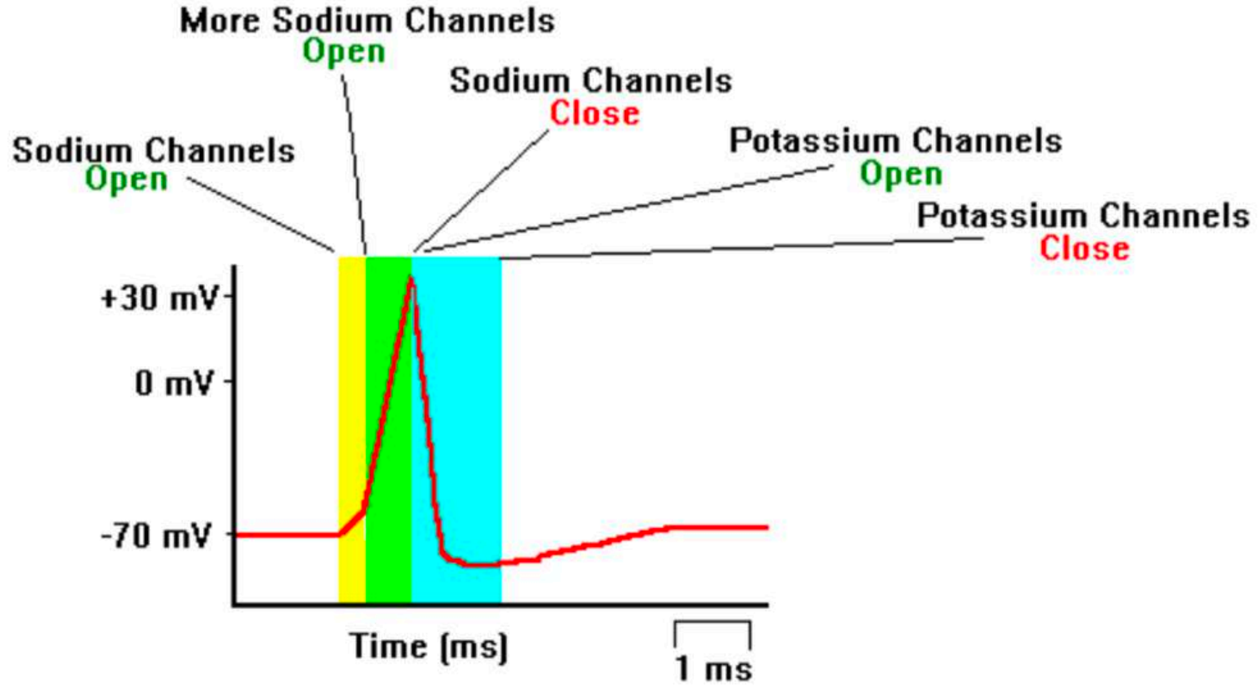

Eric H. Chudler

Neural field models approximate cortex as a continuum of neurons

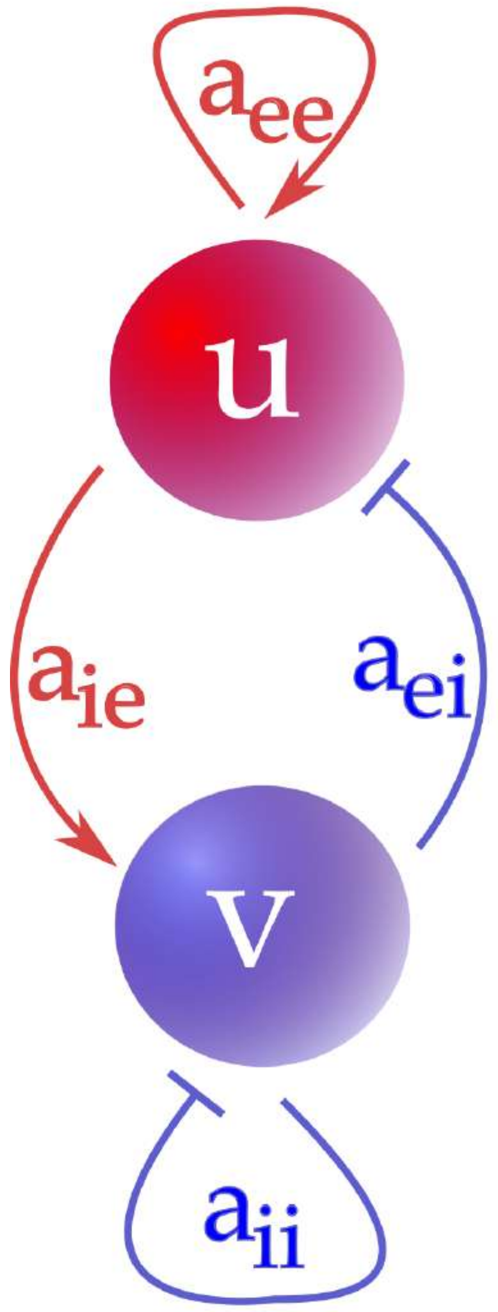

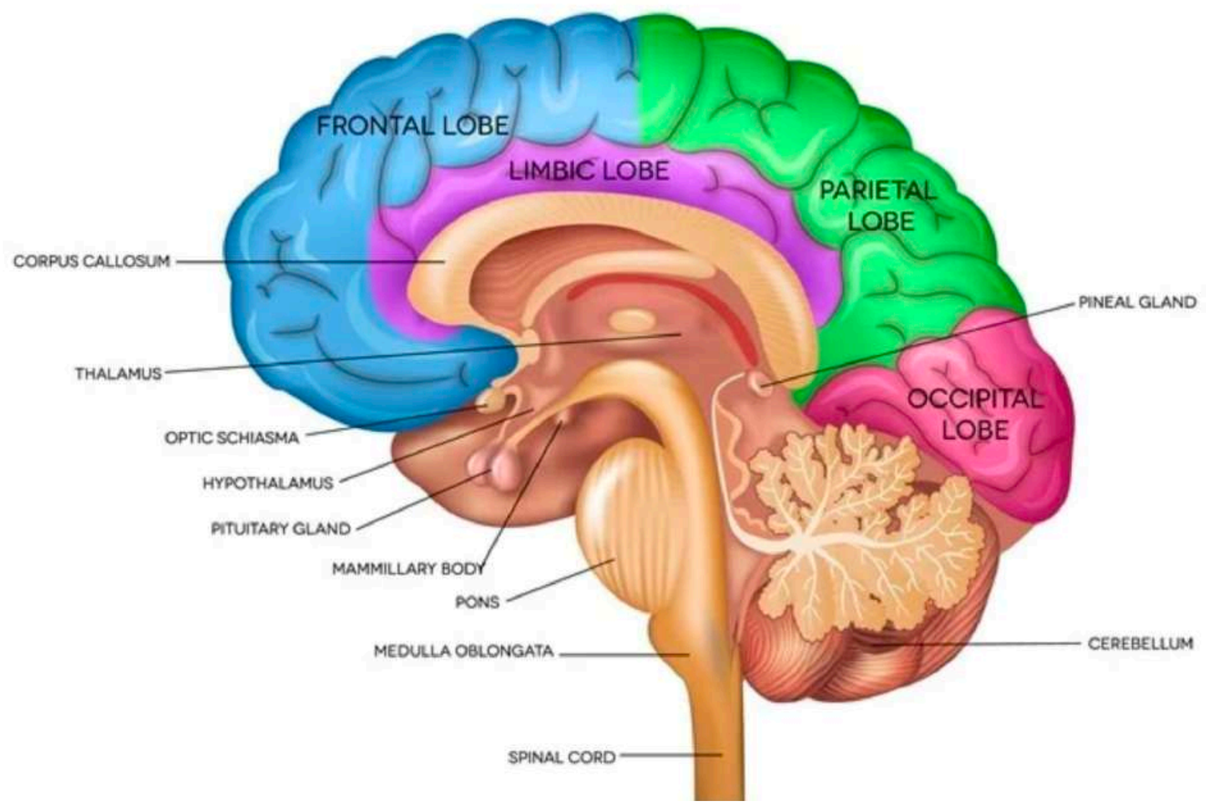

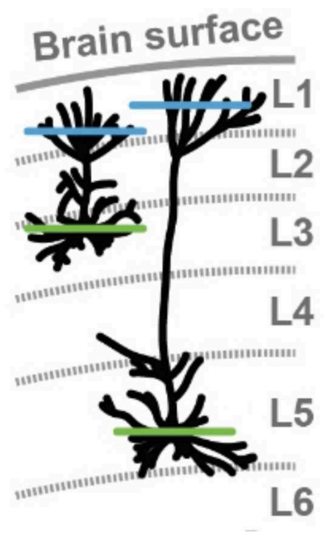

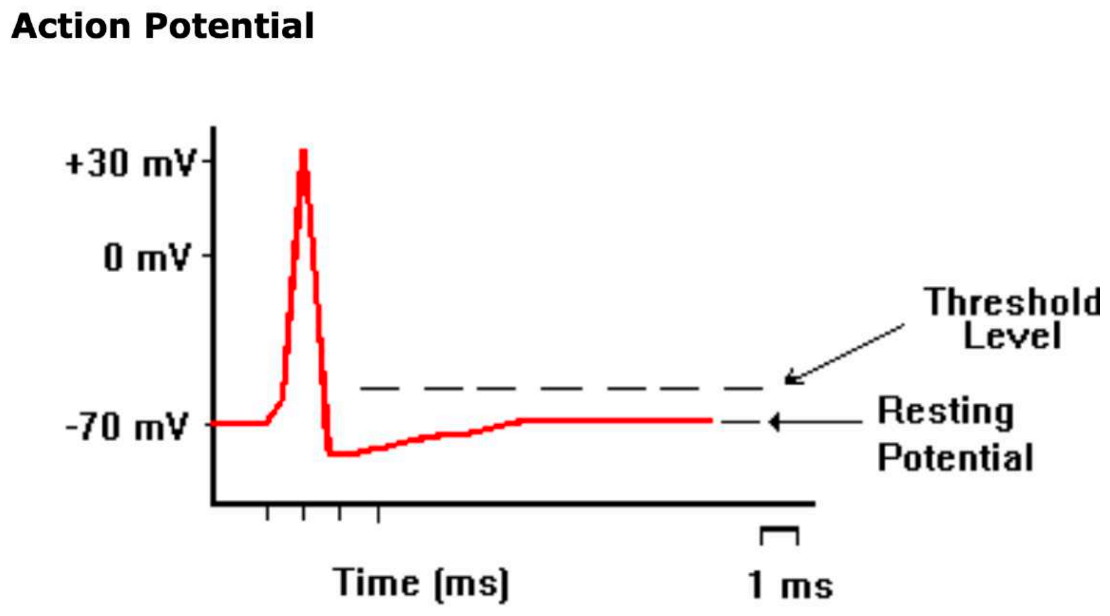

https://upload.wikimedia.org/wikipedia/commons/thumb/1/1f/MembraneCircuit.svg/336px-MembraneCircuit.svg

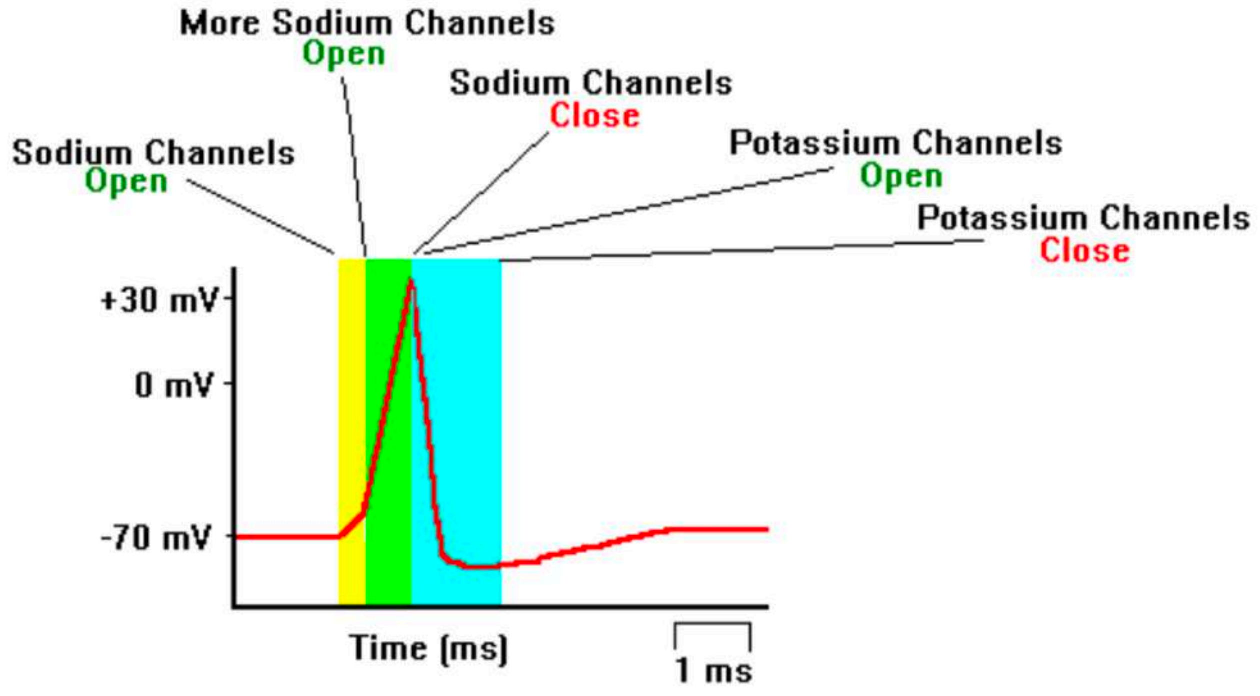

u - excitatory
v - inhibitory

faculty.washington.edu/chudler/ap.html

Eric H. Chudler

Neural field models approximate cortex as a continuum of neurons

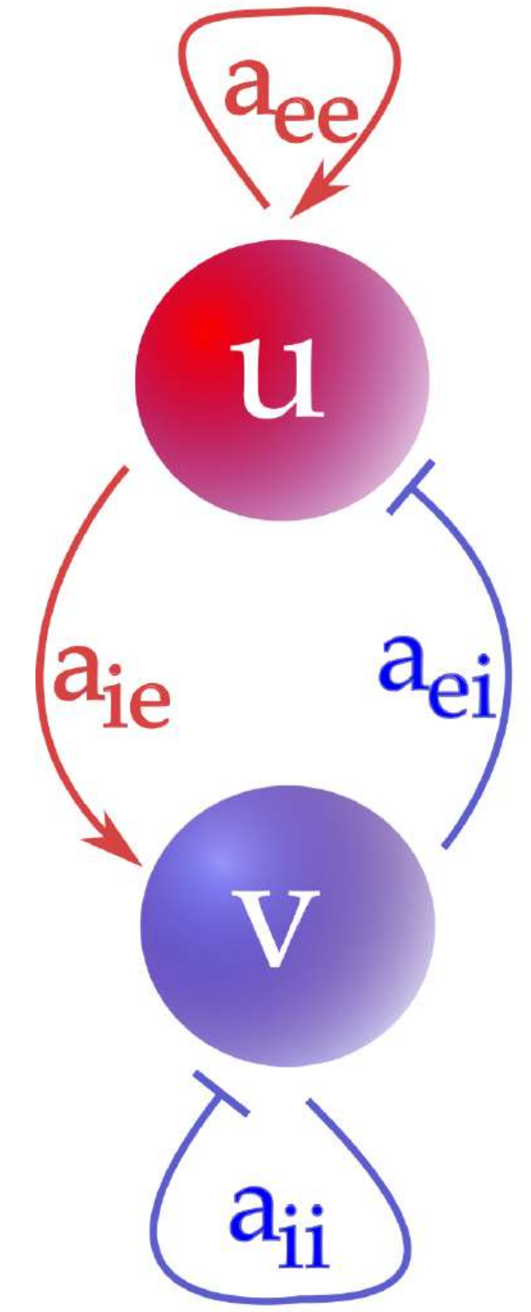

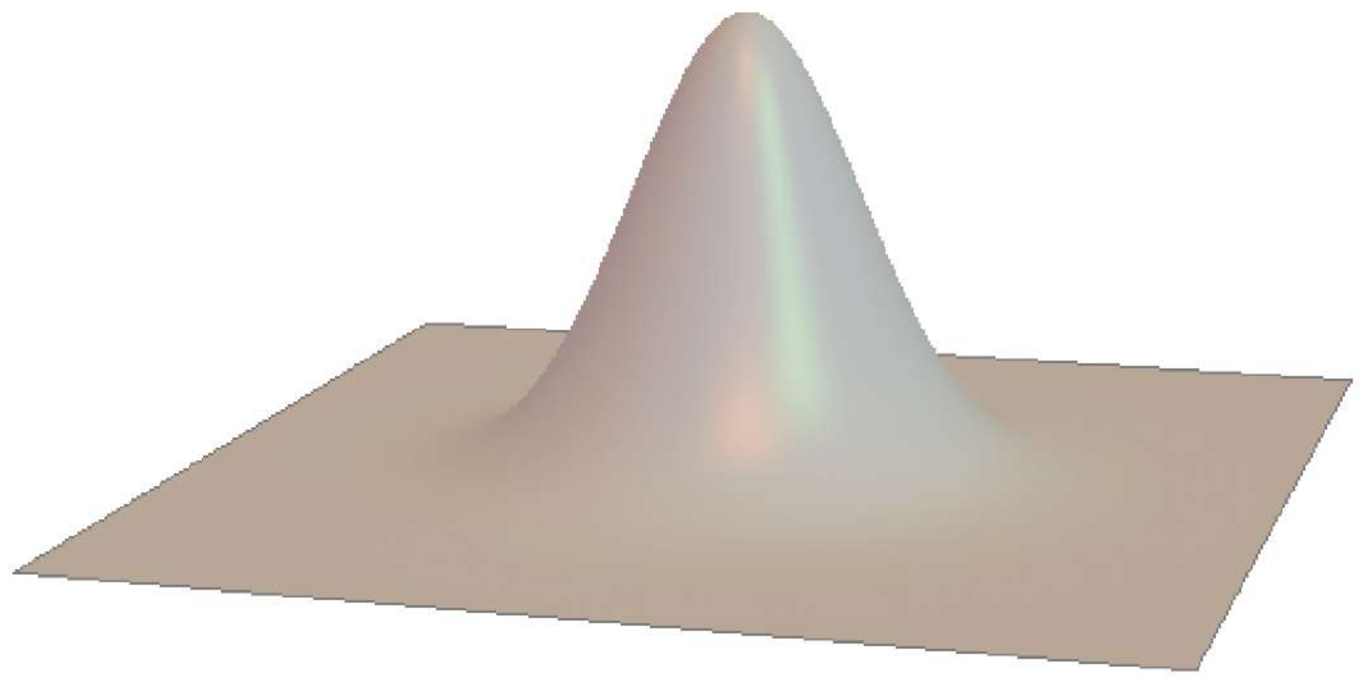

$$J_{\alpha\beta}(x) = a_{\alpha\beta}K_\beta(x)$$

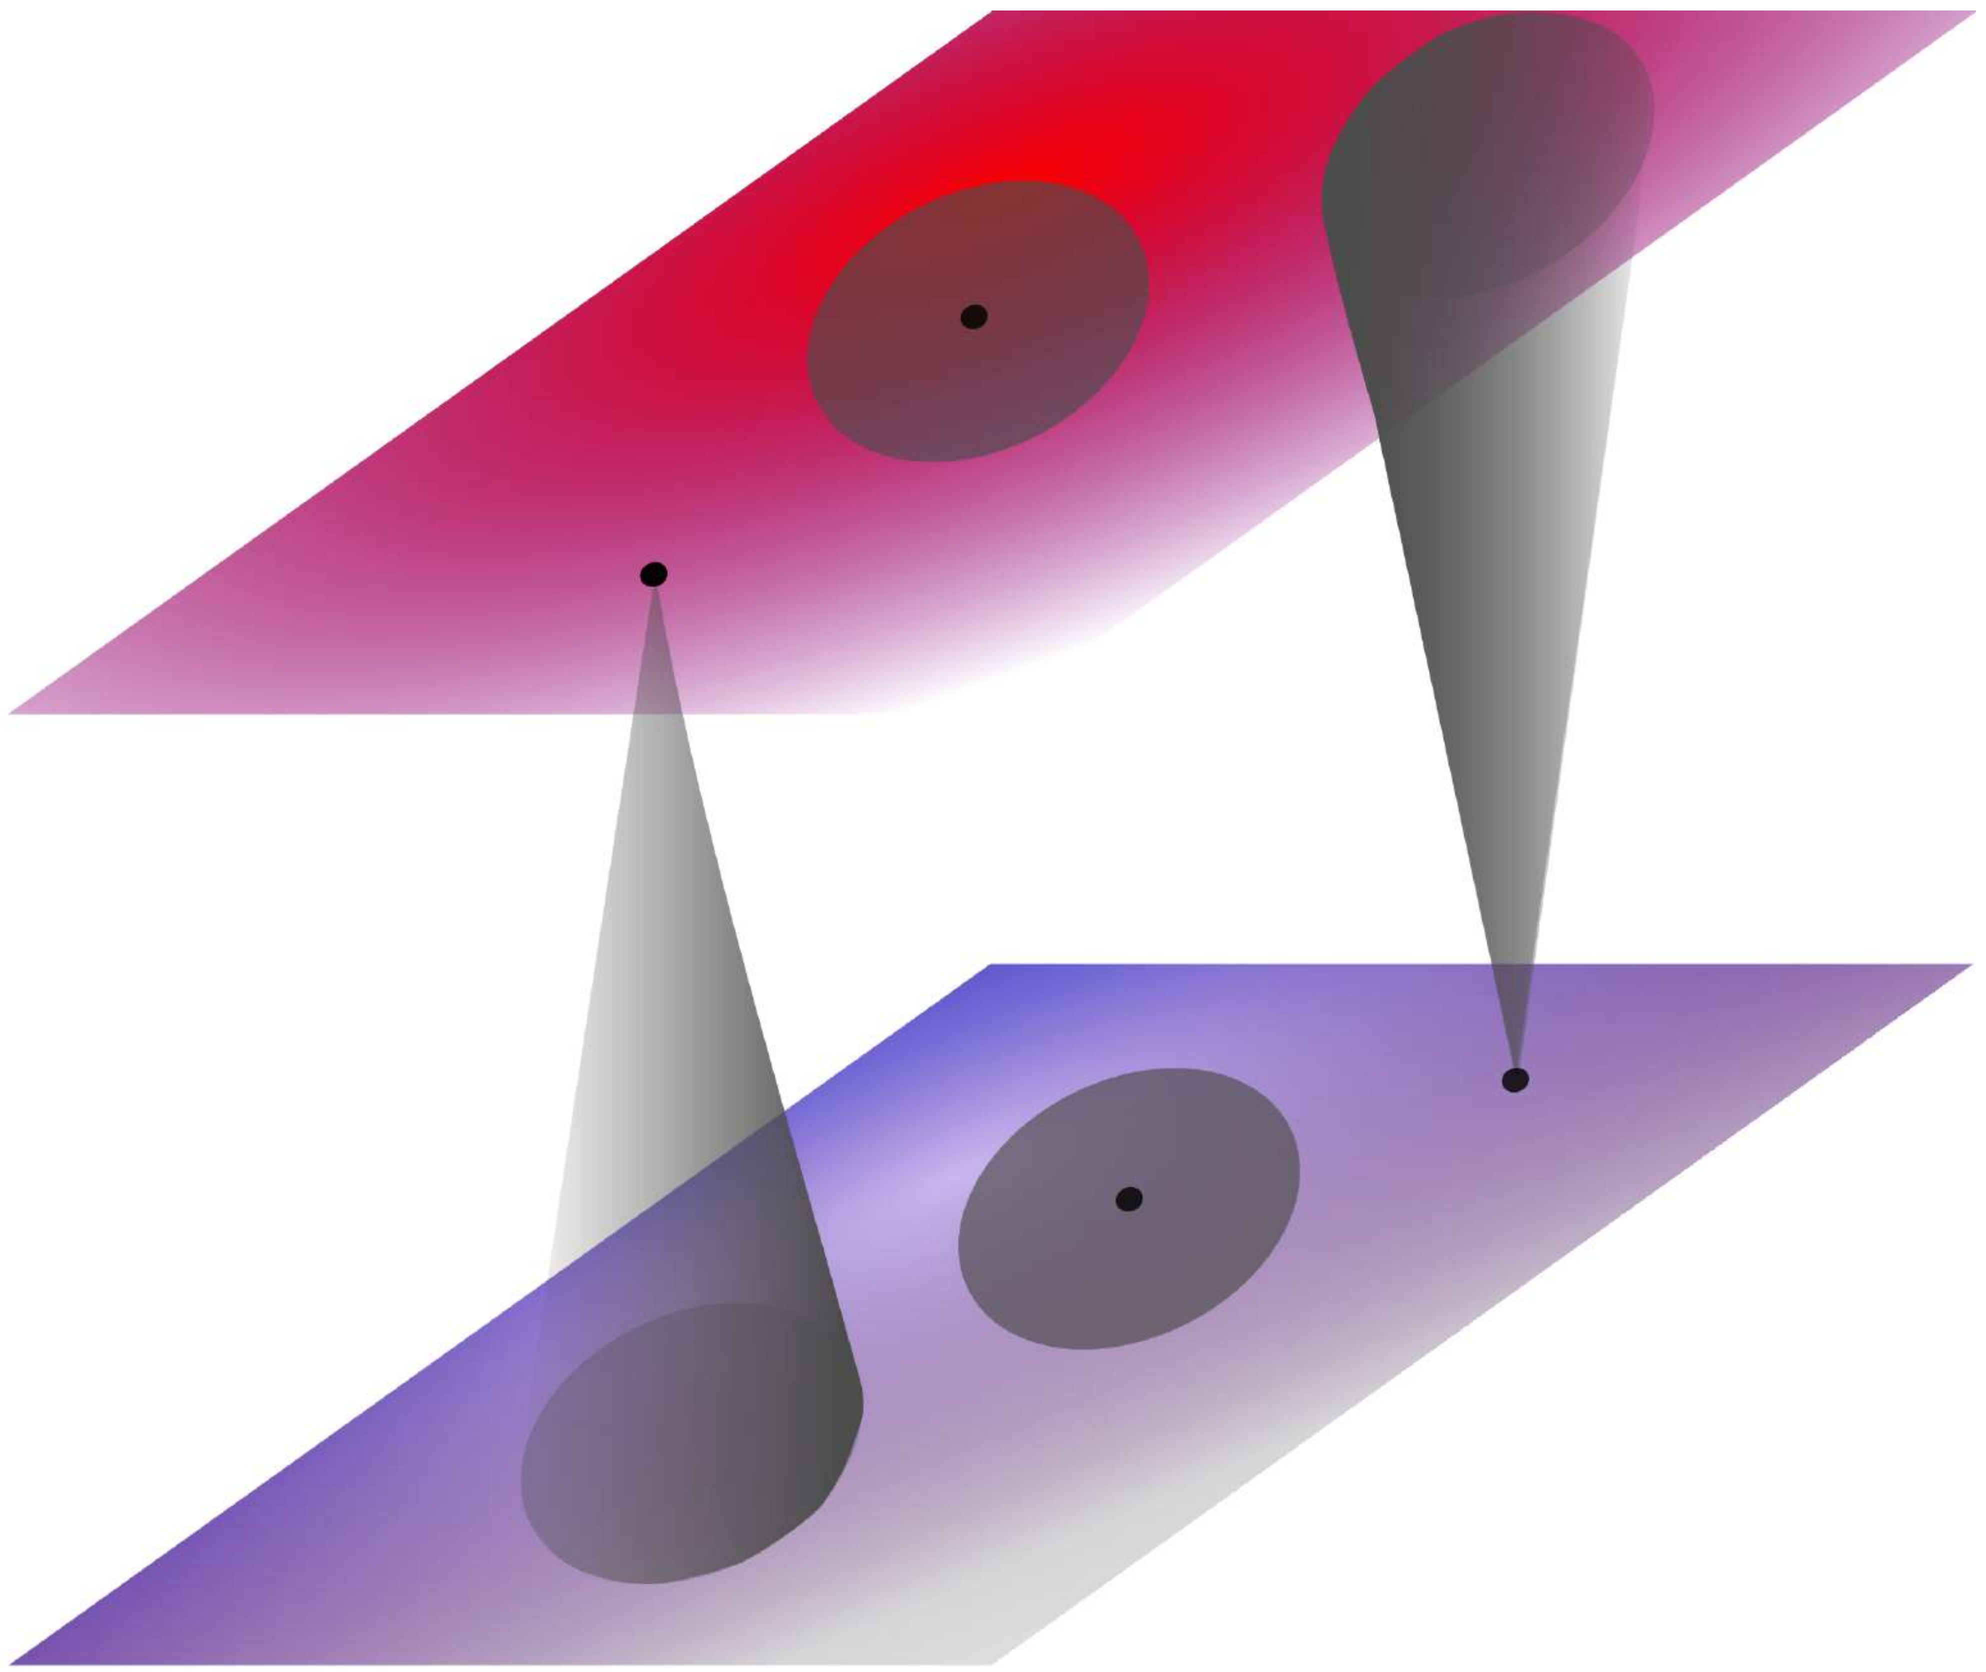

$$\frac{\partial u(x,t)}{\partial t} = -u(x,t) + f_e(J_{ee}(x) * u(x,t) - J_{ei}(x) * v(x,t))$$

$$\tau\frac{\partial v(x,t)}{\partial t} = -v(x,t) + f_i(J_{ie}(x) * u(x,t) - J_{ii}(x) * v(x,t))$$

$$f_{e,i}(u) = \frac{1}{1+\exp(-4(u-\theta_{e,i}))}$$

u - excitatory
v - inhibitory
* - spatial convolution

But… what about those static feedforward models?

- 1943: McCullough-Pitts model
- 1949: Hebbian learning - synaptic weights change to allow for learning
- 1958: perceptron by Rosenblatt
    - Mark I perceptron machine (compare: Blue/Human Brain Project)
- 1969: Need more than one layer of “neurons” (XOR - Minsky and Papert)

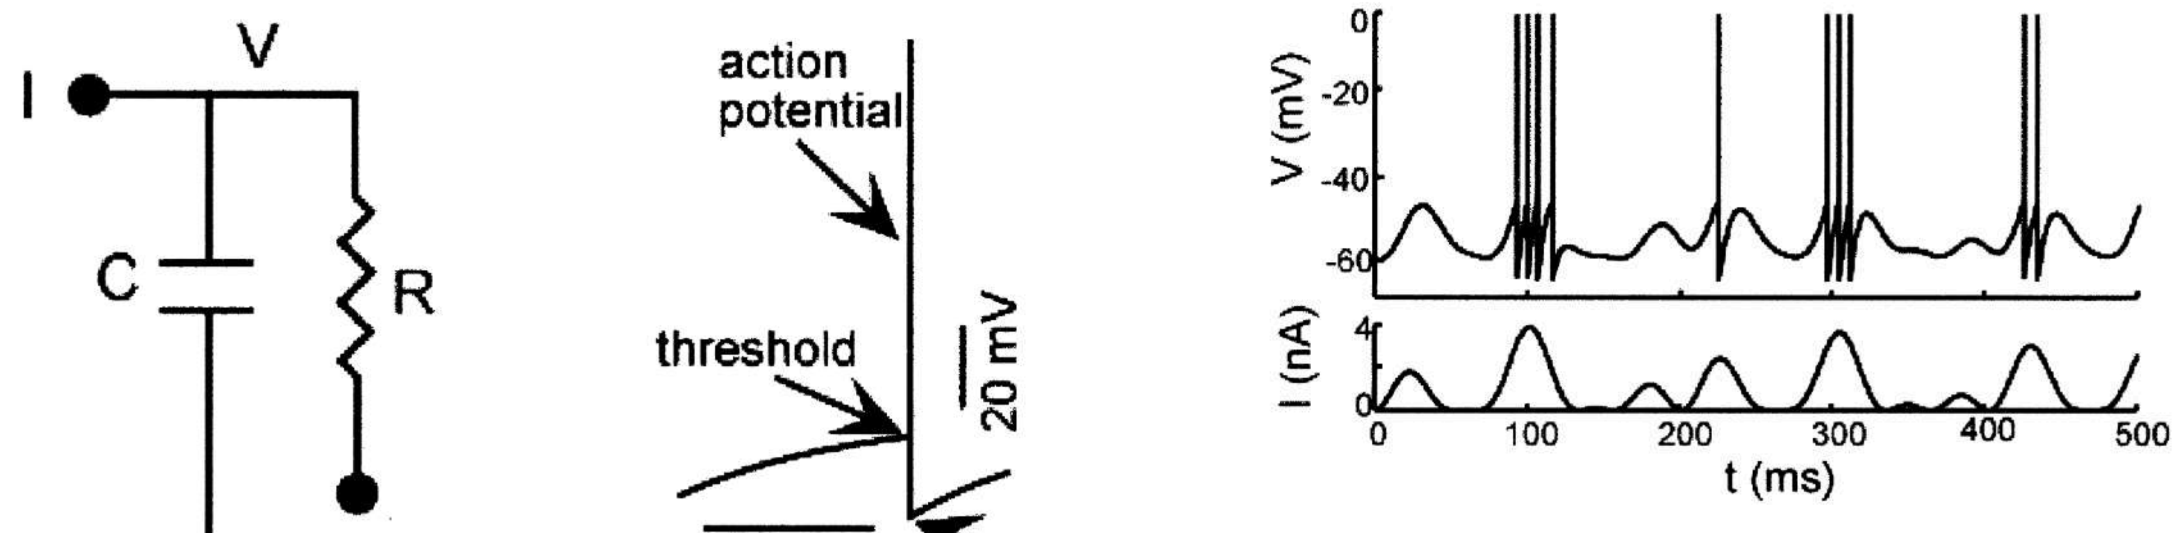

Abbott, 1999

$$\tau V(t) = -V(t) + I(t)$$

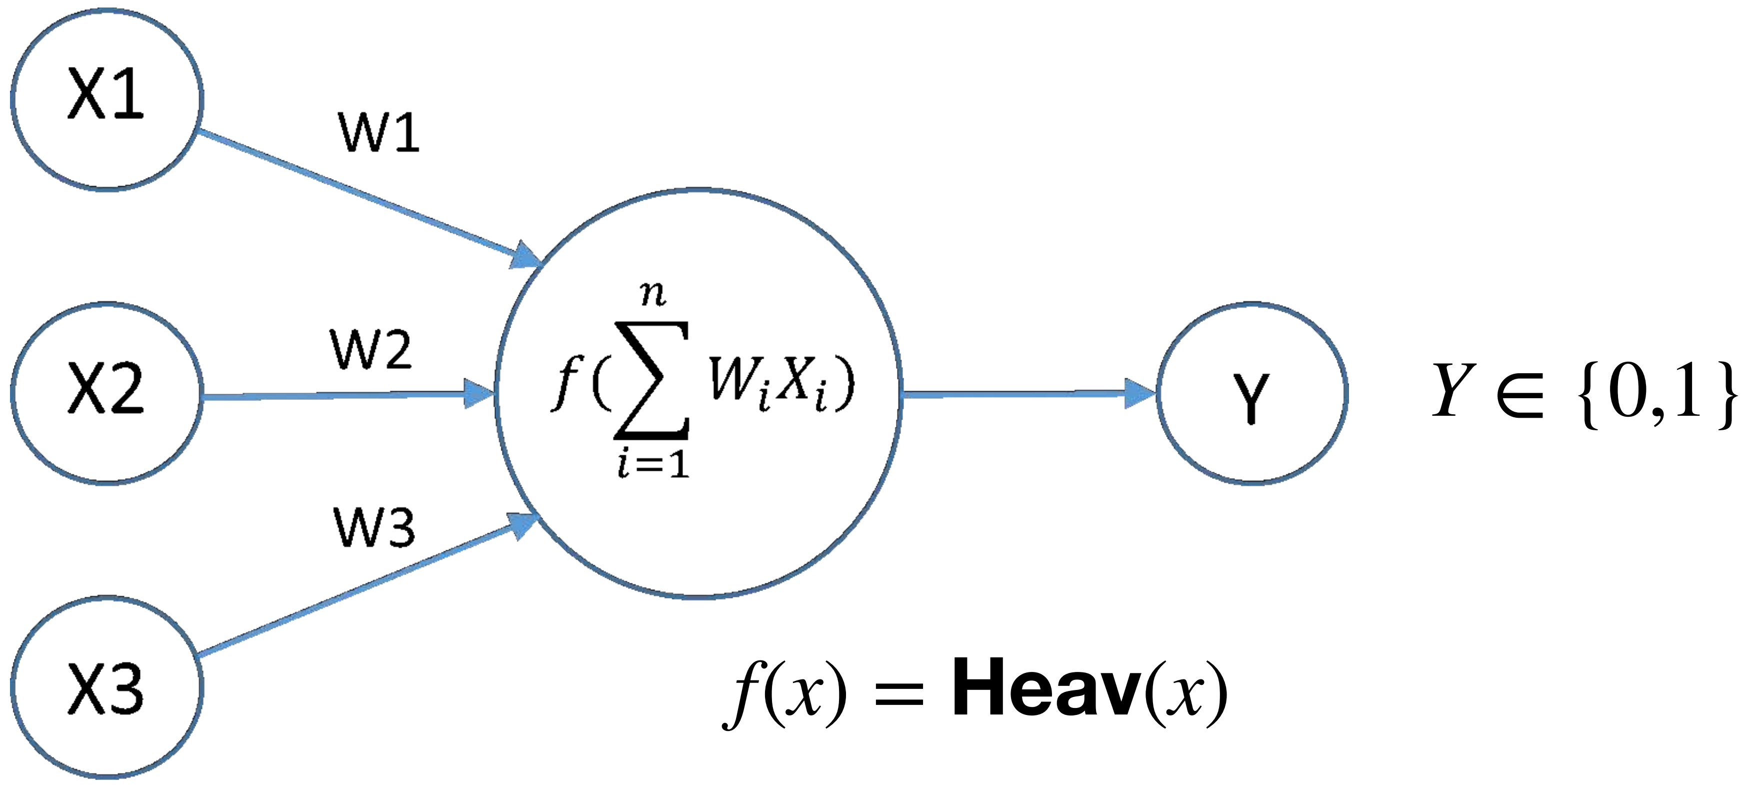

$$W1 = W2 = W3 = \ldots = WN$$

(fixed)

But… what about those static feedforward models?

- 1943: McCullough-Pitts model
- 1949: Hebbian learning - synaptic weights change to allow for learning
- 1958: perceptron by Rosenblatt
    - Mark I perceptron machine (compare: Blue/Human Brain Project)
- 1969: Need more than one layer of “neurons” (XOR - Minsky and Papert)

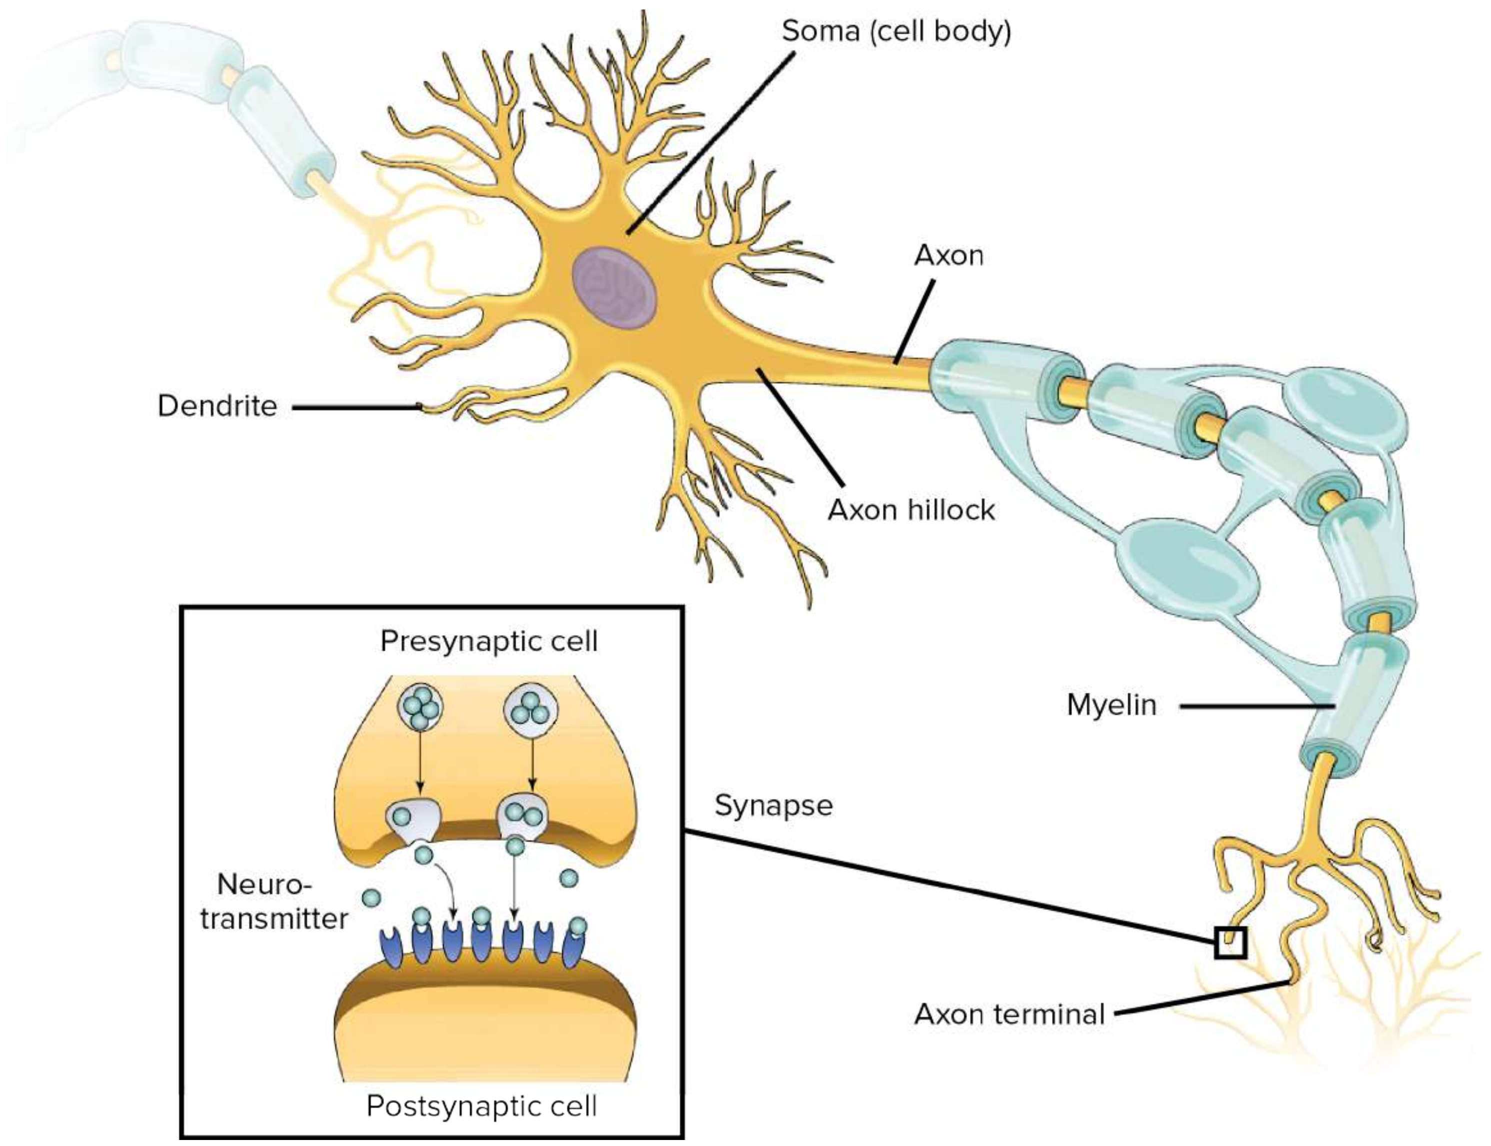

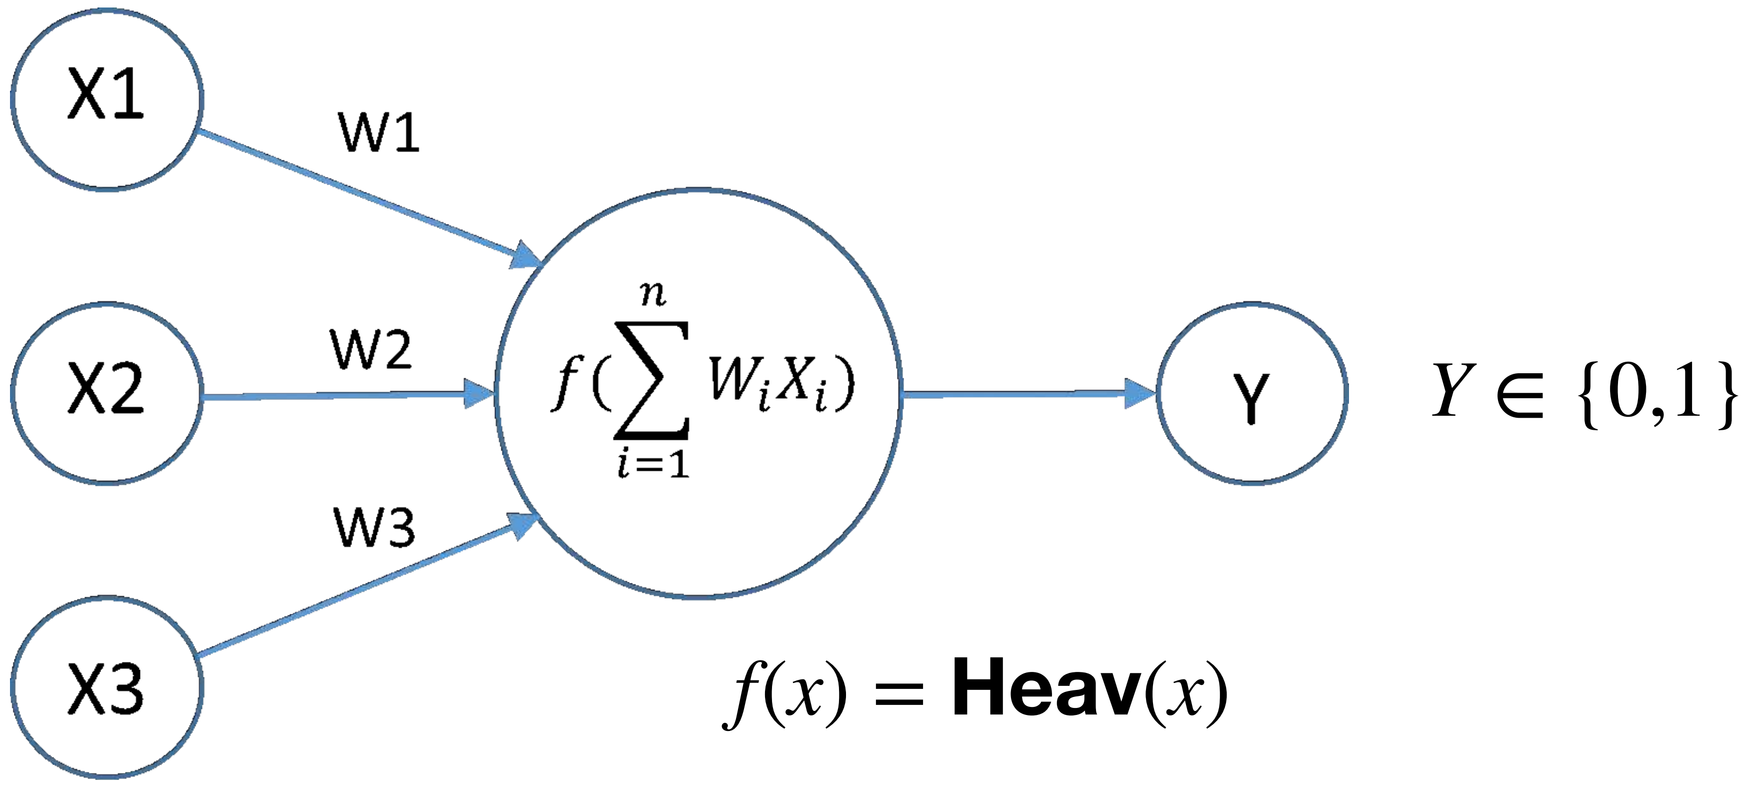

$$W1 = W2 = W3 = \ldots = WN$$

(fixed)

But… what about those static feedforward models?

- 1943: McCullough-Pitts model
- 1949: Hebbian learning - synaptic weights change to allow for learning
- 1958: perceptron by Rosenblatt
  - Mark I perceptron machine (compare: Blue/Human Brain Project)
- 1969: Need more than one layer of “neurons” (XOR - Minsky and Papert)

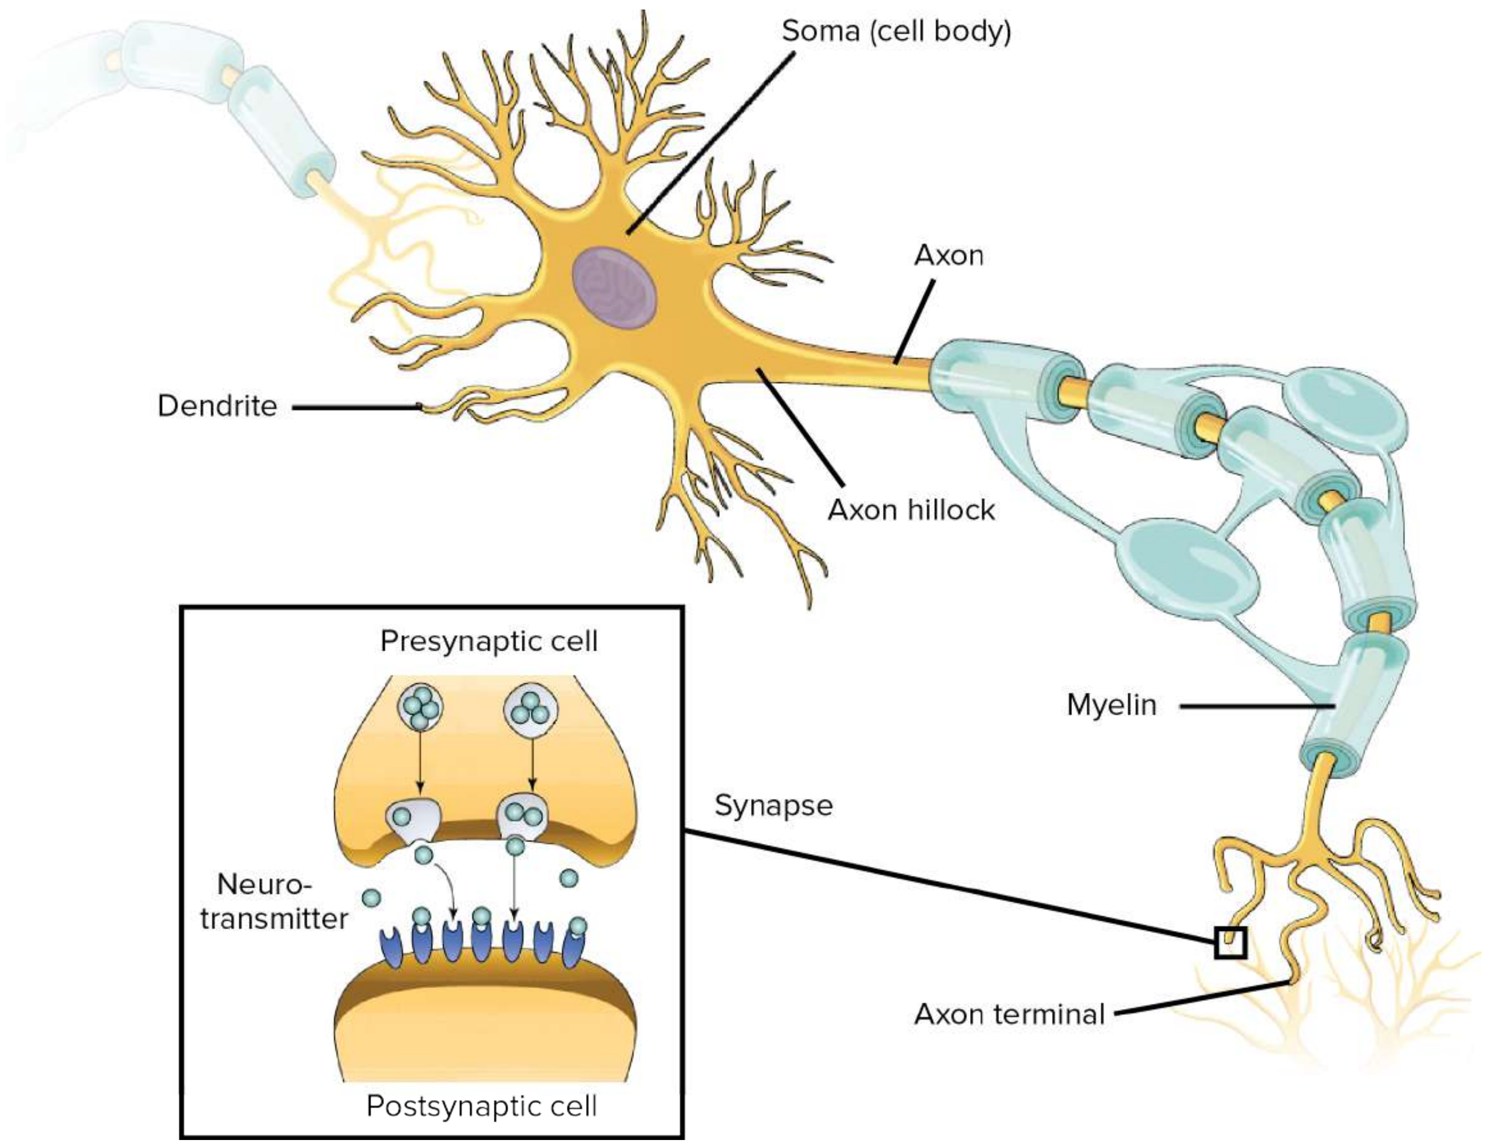

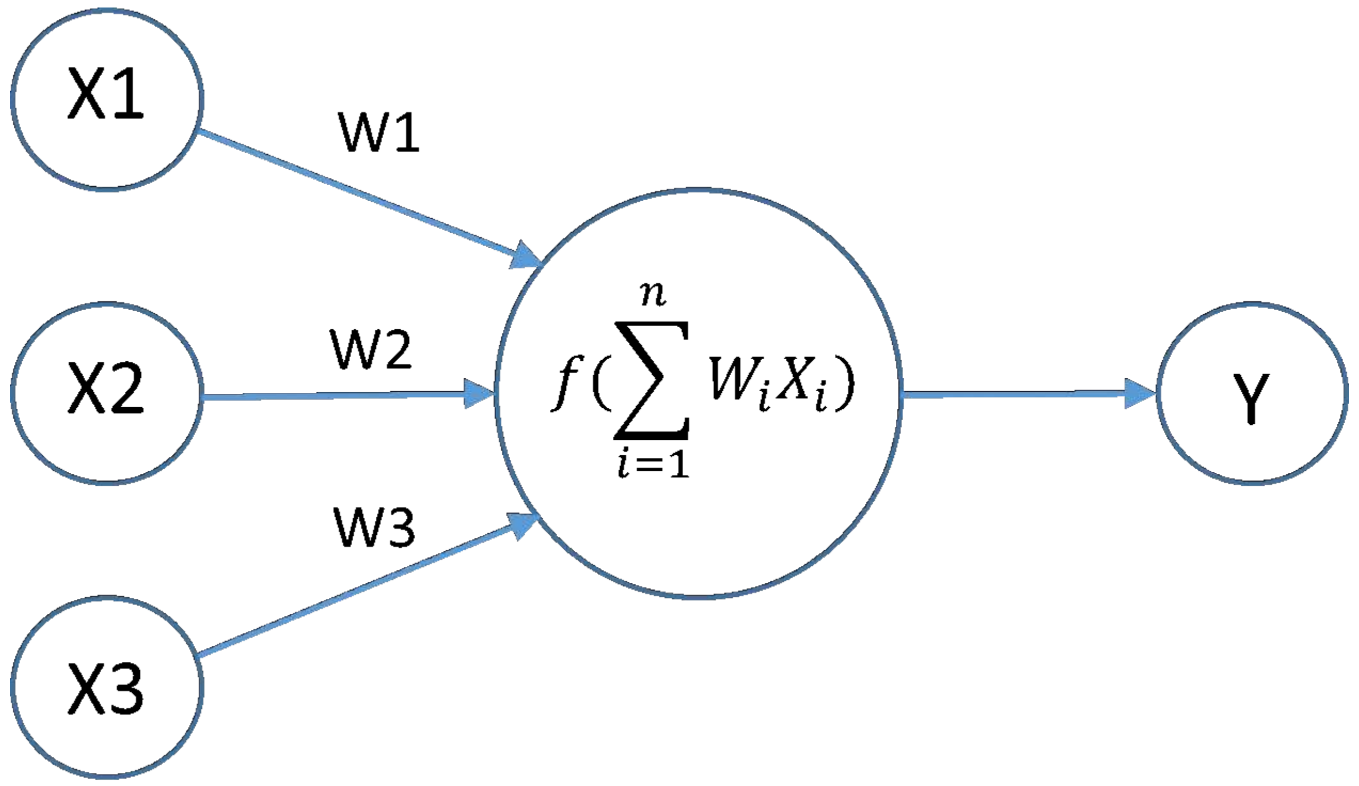

But… what about those static feedforward models?

- 1943: McCullough-Pitts model
- 1949: Hebbian learning - synaptic weights change to allow for learning
- 1958: perceptron by Rosenblatt
    - Mark I perceptron machine (compare: Blue/Human Brain Project)
- 1969: Need more than one layer of "neurons" (XOR - Minsky and Papert)

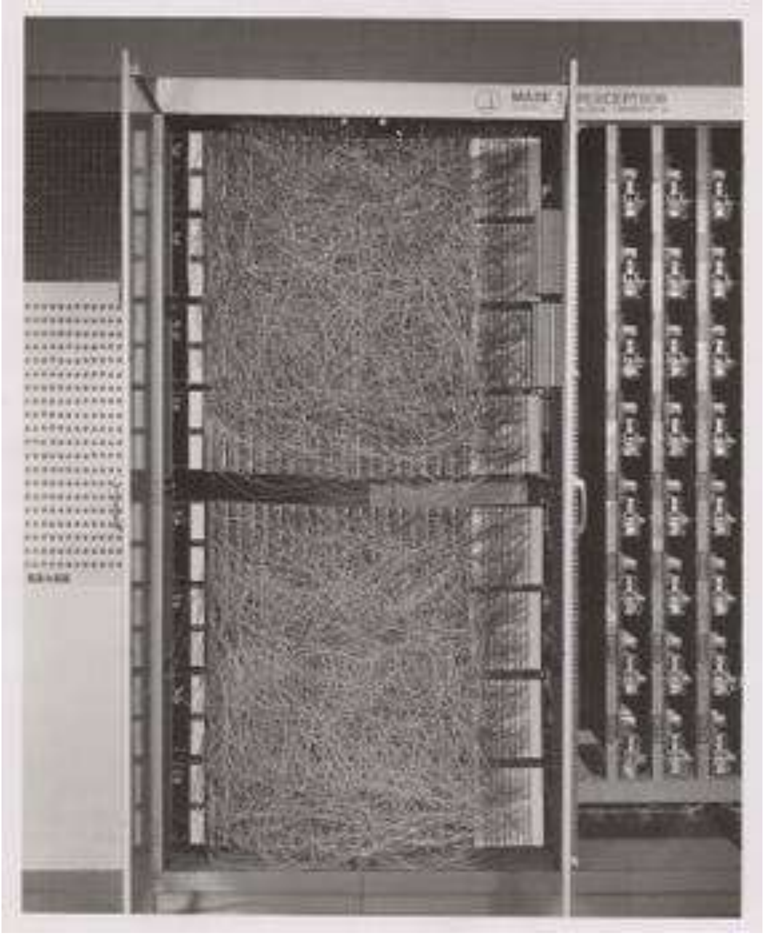

https://en.wikipedia.org/wiki/Perceptron

Mark I Perceptron displayed at the Smithsonian museum

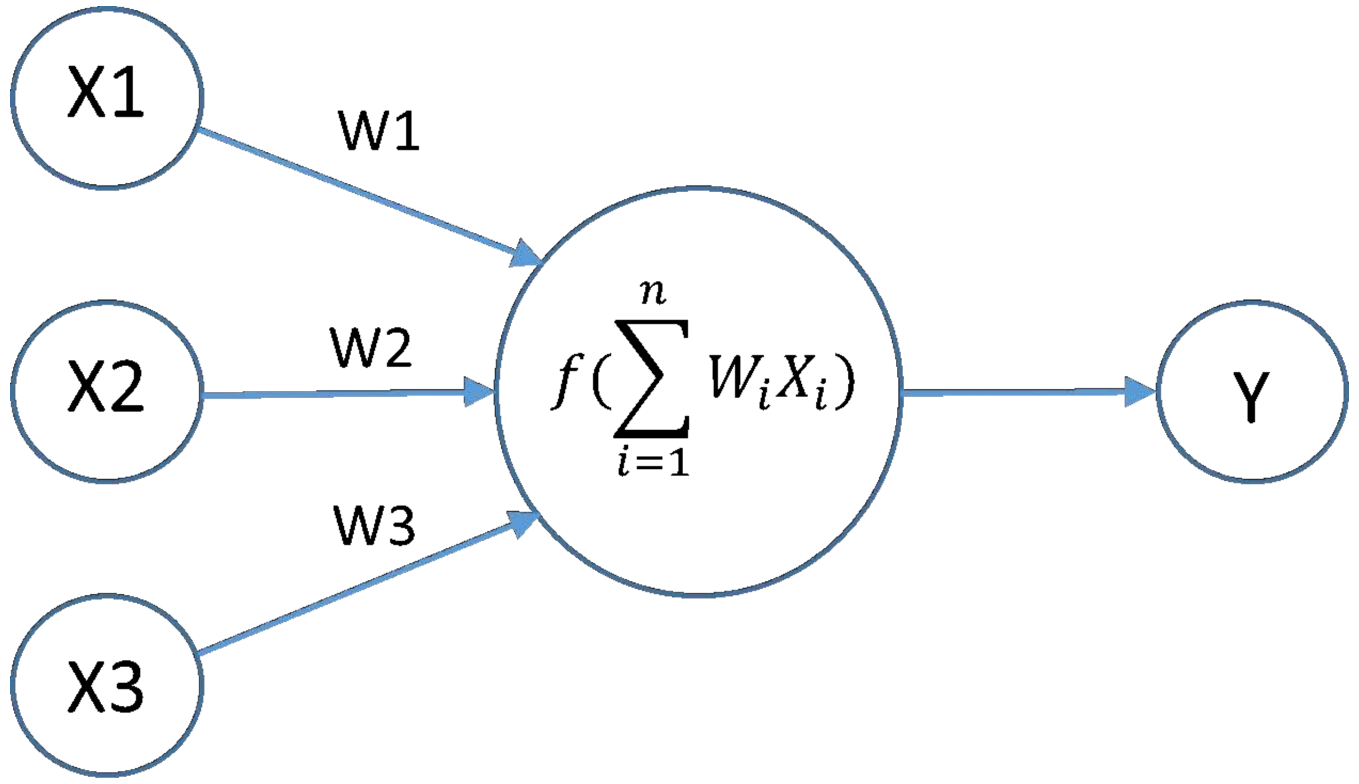

Multilayered ("deep") perceptron networks prove more practical in applications

- 1967: Supervised learning on "deep" feedforward networks (multilayer perceptrons) by Ivakhnenko and Lapa

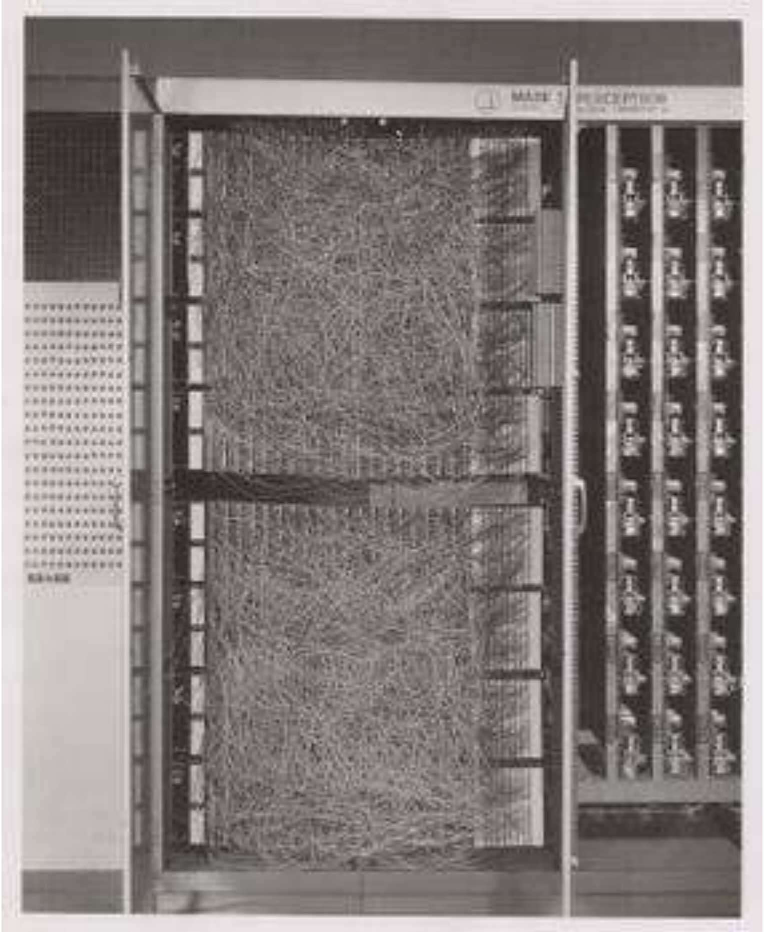

https://en.wikipedia.org/wiki/Perceptron

Mark I Perceptron displayed at the Smithsonian museum

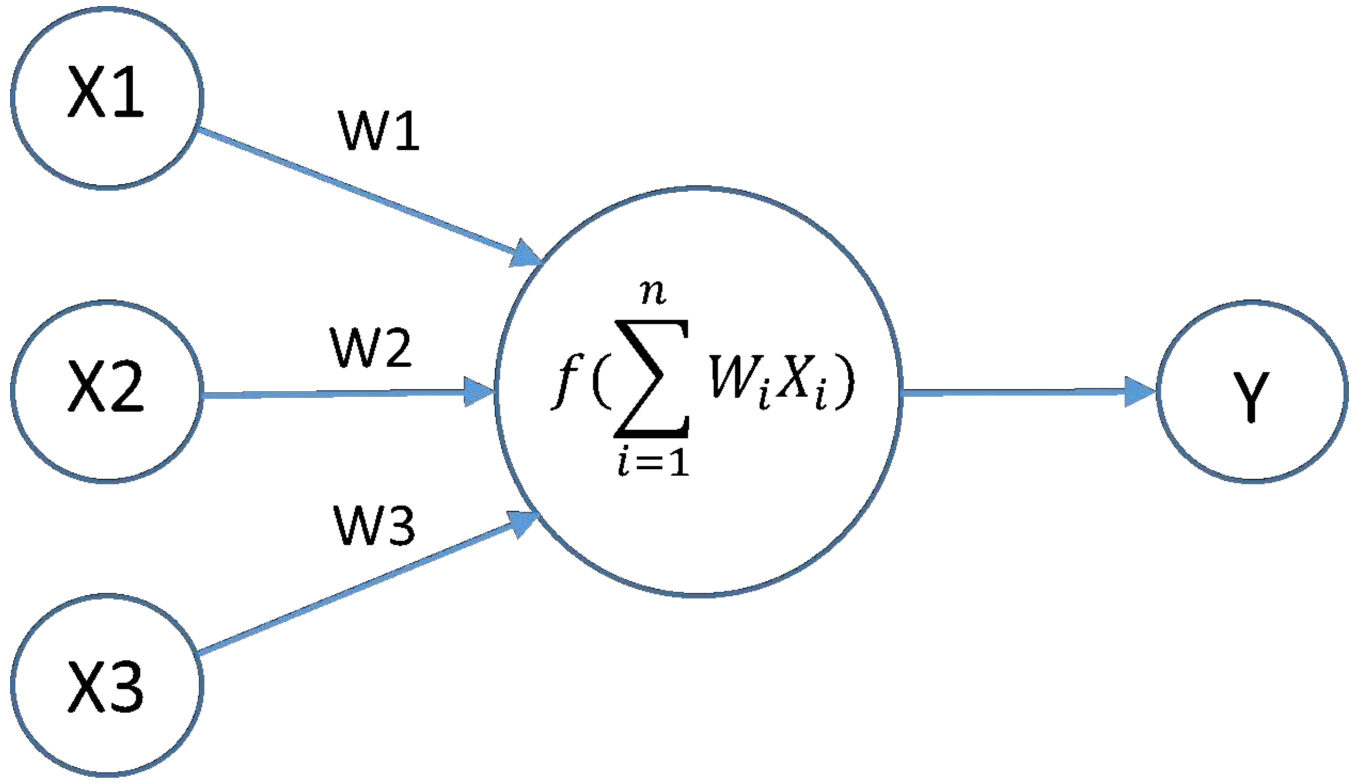

Multilayered ("deep") perceptron networks prove more practical in applications

- 1967: Supervised learning on "deep" feedforward networks (multilayer perceptrons) by Ivakhnenko and Lapa

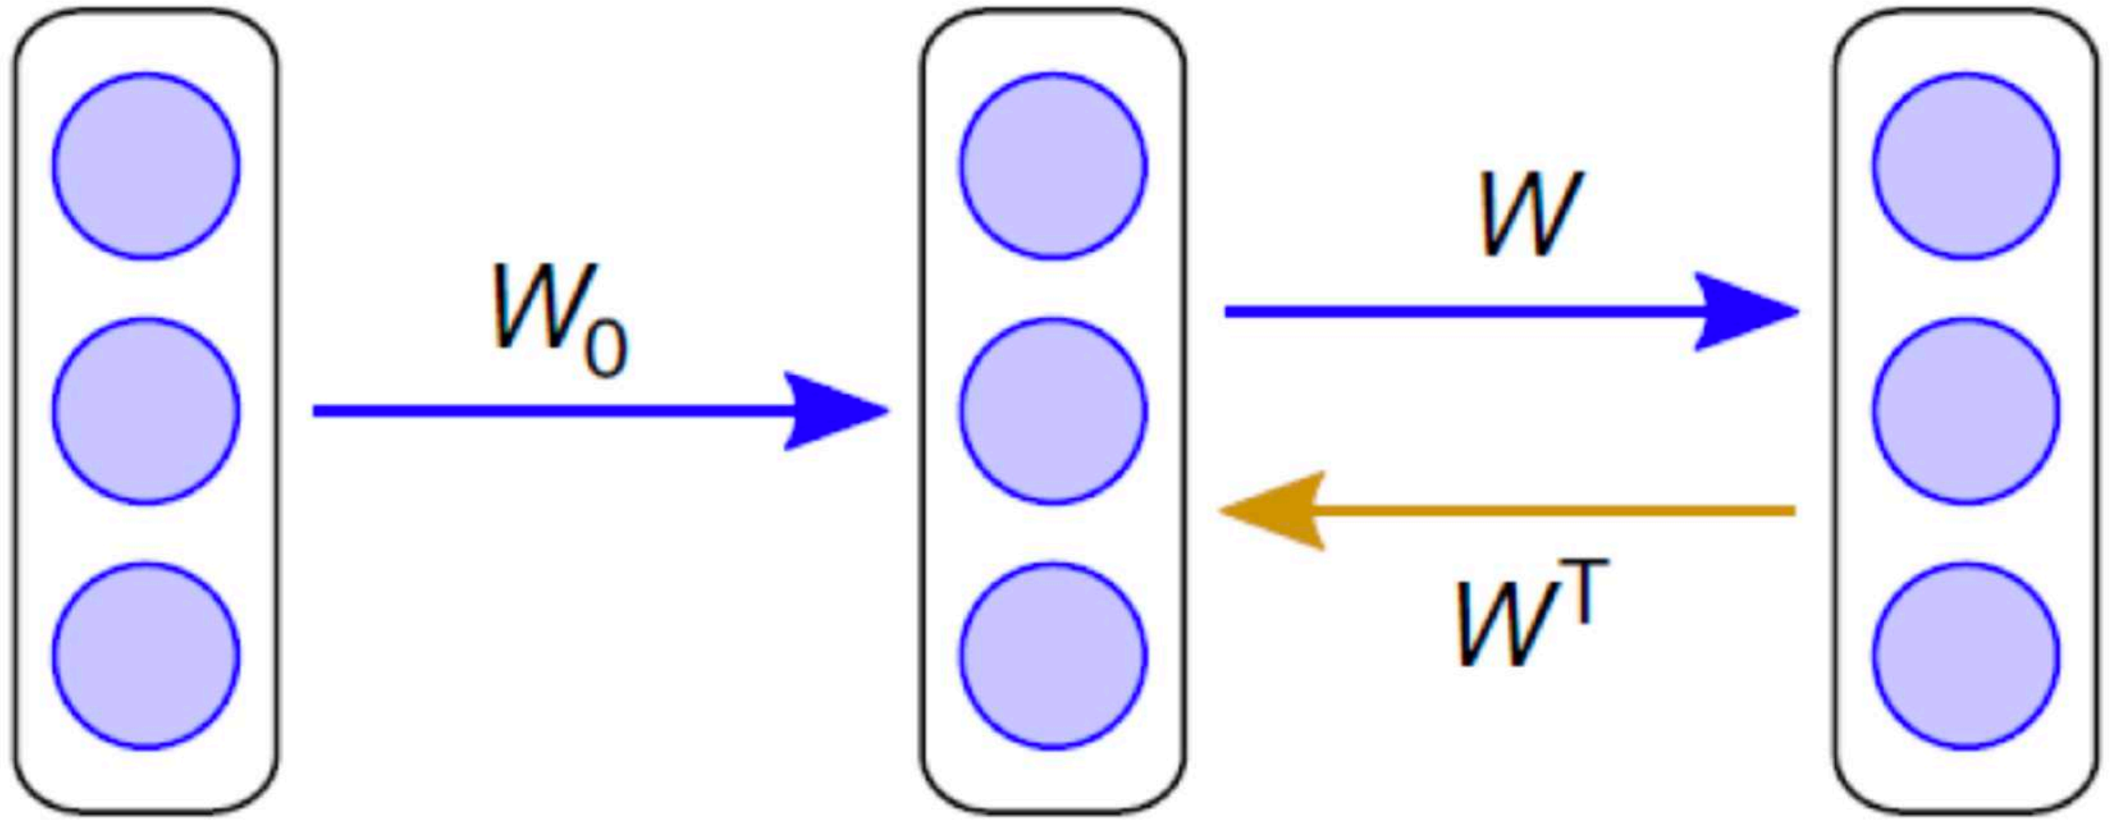

Lillicrap…Akerman, 2016

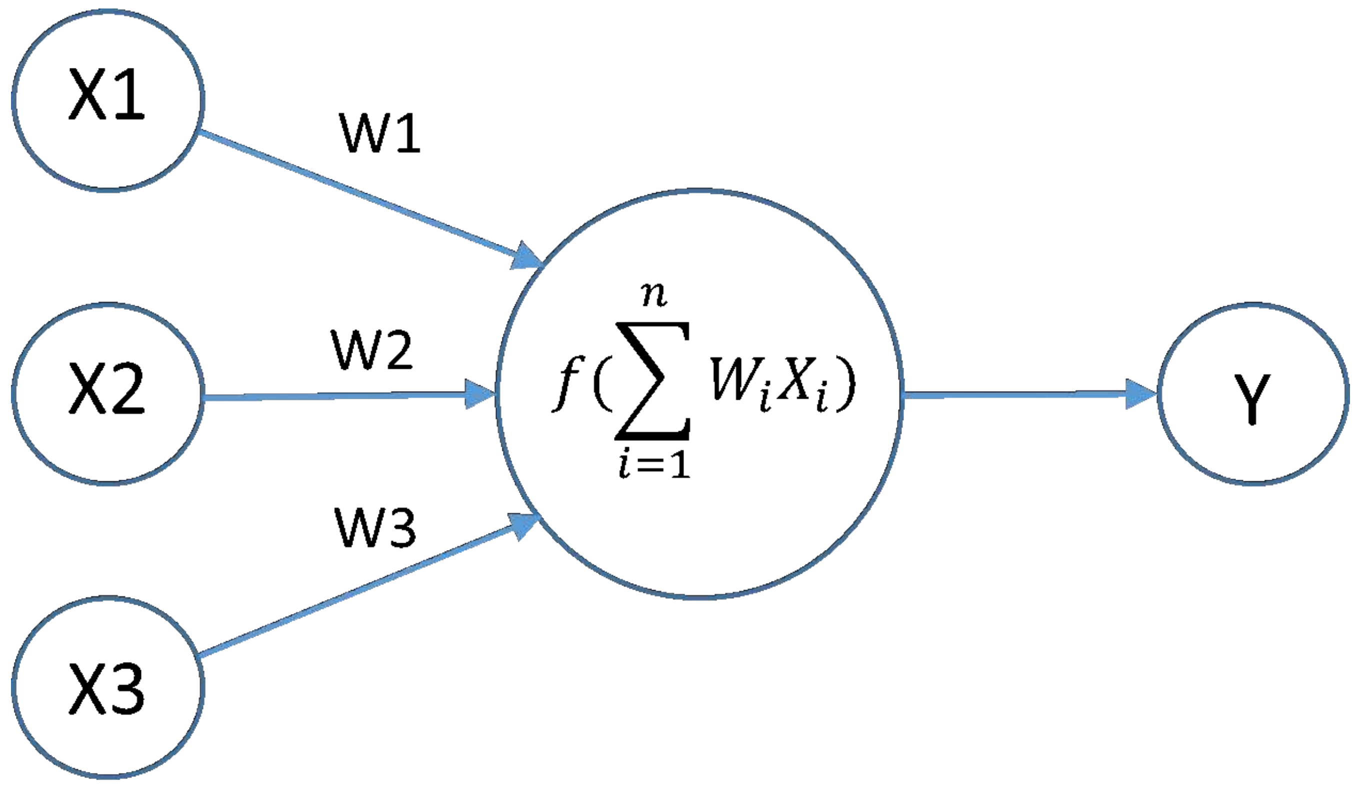

Multilayered ("deep") perceptron networks prove more practical in applications

- 1967: Supervised learning on "deep" feedforward networks (multilayer perceptrons) by Ivakhnenko and Lapa

- 1986: Backpropagation to update weights by Rumelhart, Hinton & Williams

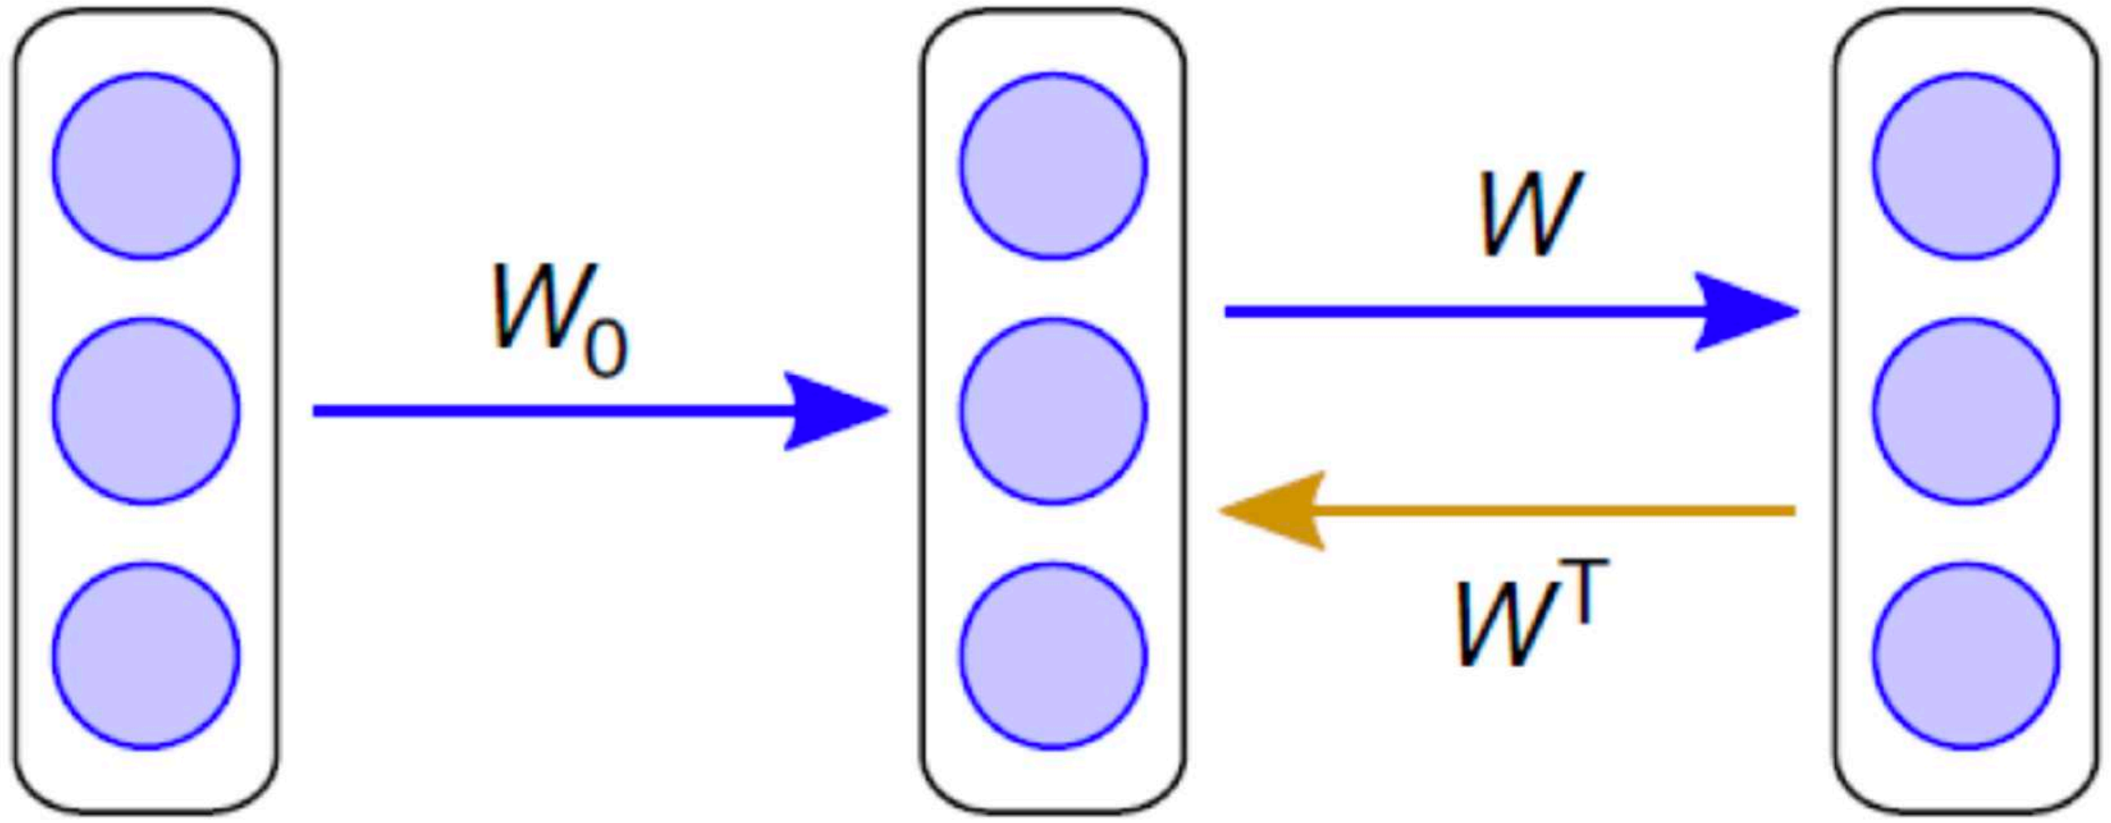

Lillicrap…Akerman, 2016

Backprop: $\Delta W \propto \frac{\partial L}{\partial W} = -\mathbf{e} \cdot \mathbf{h}$

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \mathbf{h}}\right)\left(\frac{\partial \mathbf{h}}{\partial W_0}\right) = -W^T \mathbf{e} \cdot \mathbf{x}$$

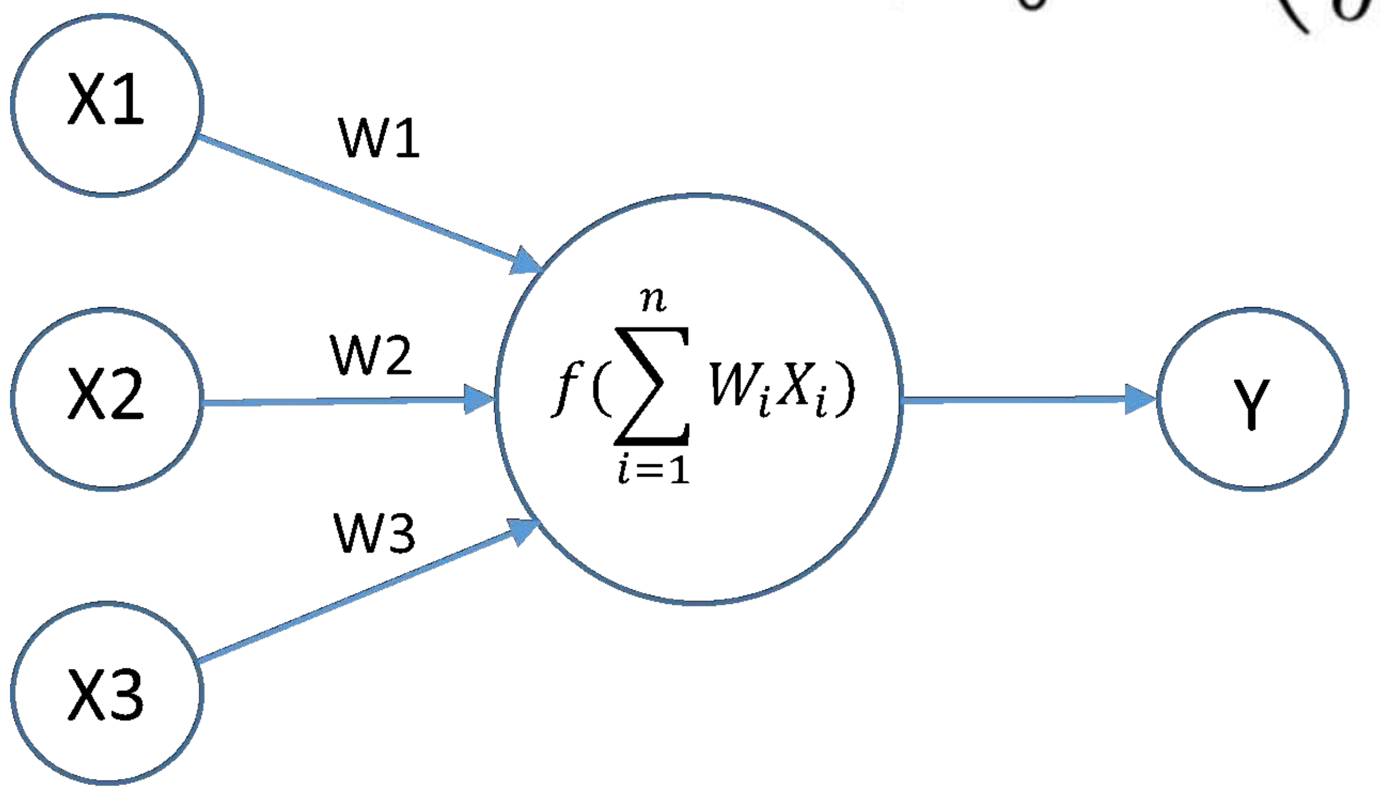

Convolutional networks borrowed further from neurobiological findings

- 1967: Supervised learning on "deep" feedforward networks (multilayer perceptrons) by Ivakhnenko and Lapa

- 1986: Backpropagation to update weights by Rumelhart, Hinton & Williams

- 1989: Convolutional neural networks with backpropagation by LeCun

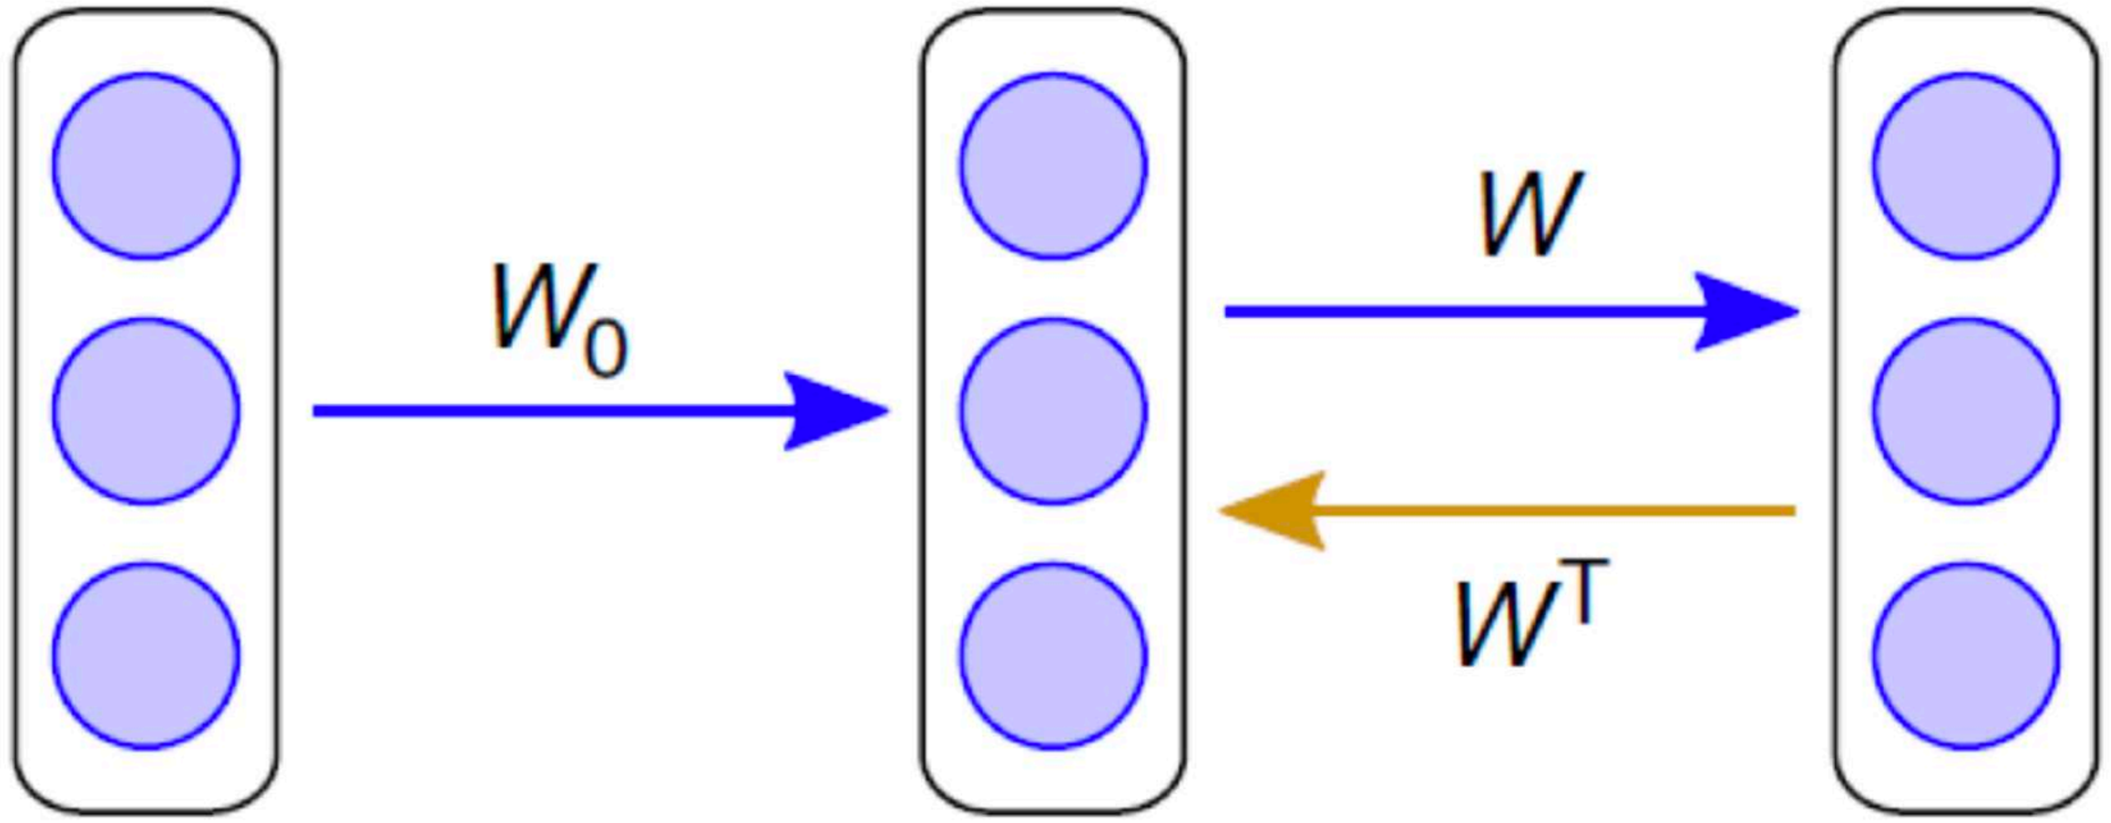

Lillicrap…Akerman, 2016

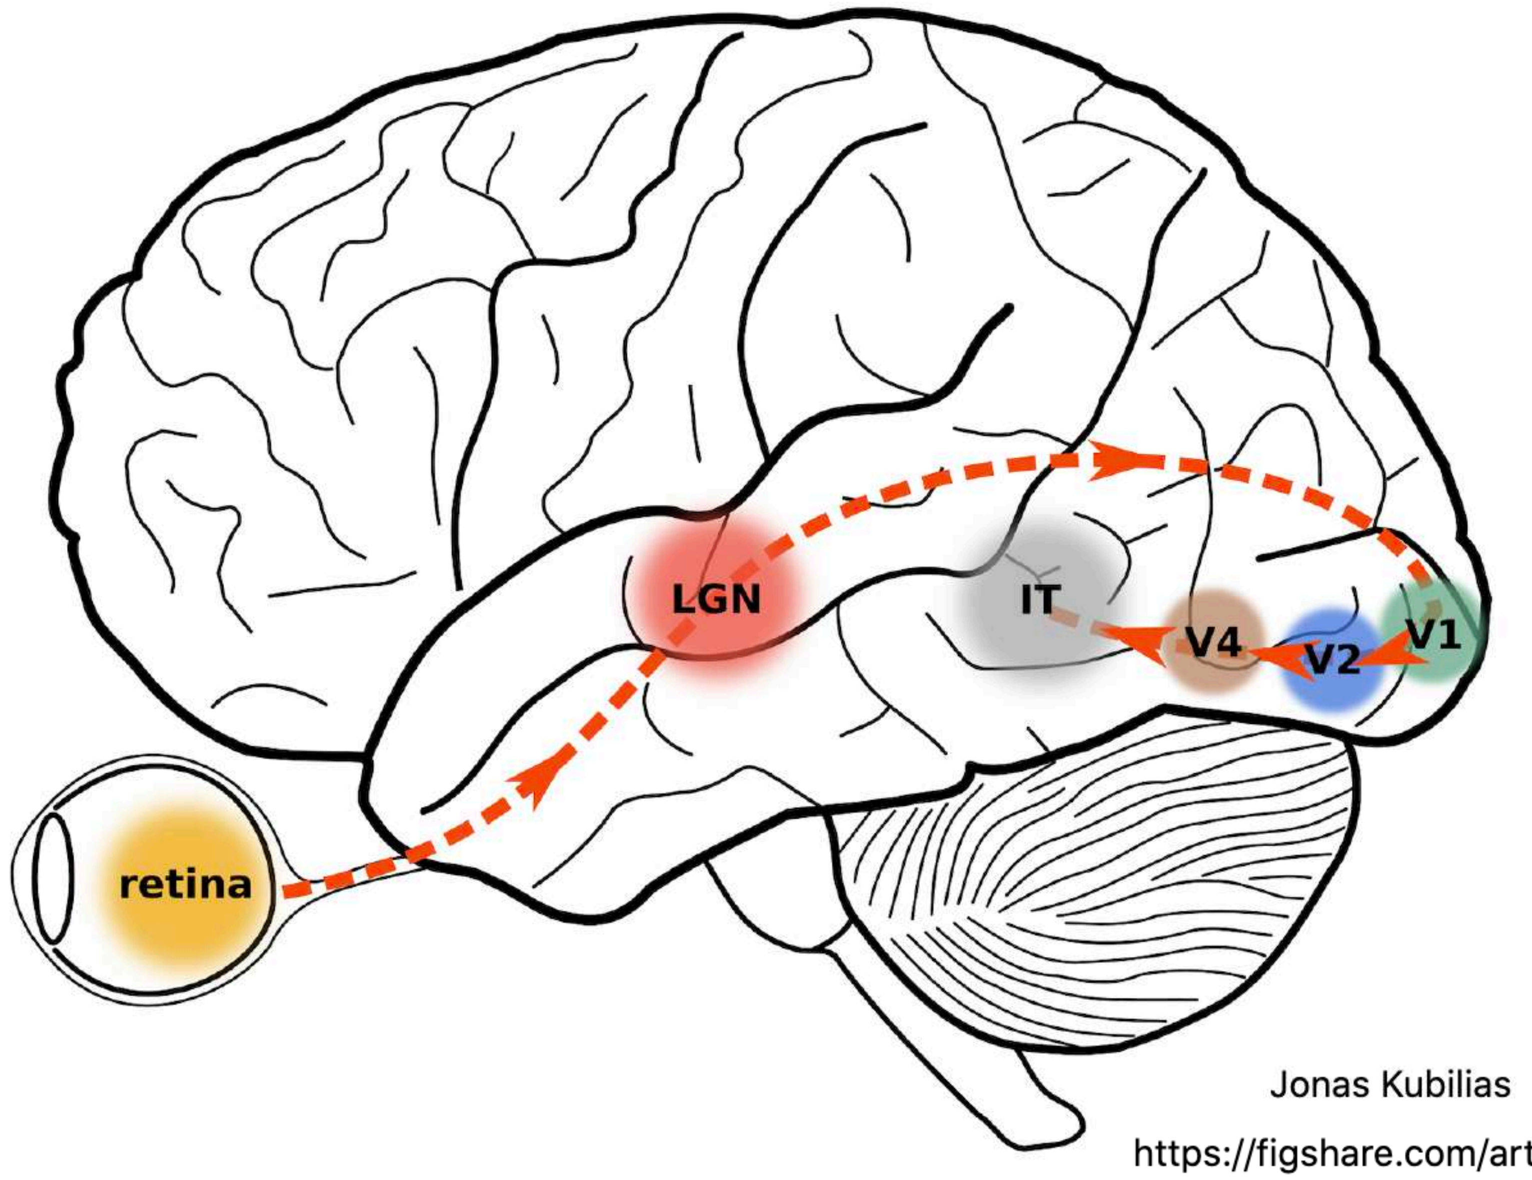

Jonas Kubilias

https://figshare.com/articles/figure/Ventral_visual_stream/106794

Convolutional networks borrowed further from neurobiological findings

- 1967: Supervised learning on “deep” feedforward networks (multilayer perceptrons) by Ivakhnenko and Lapa

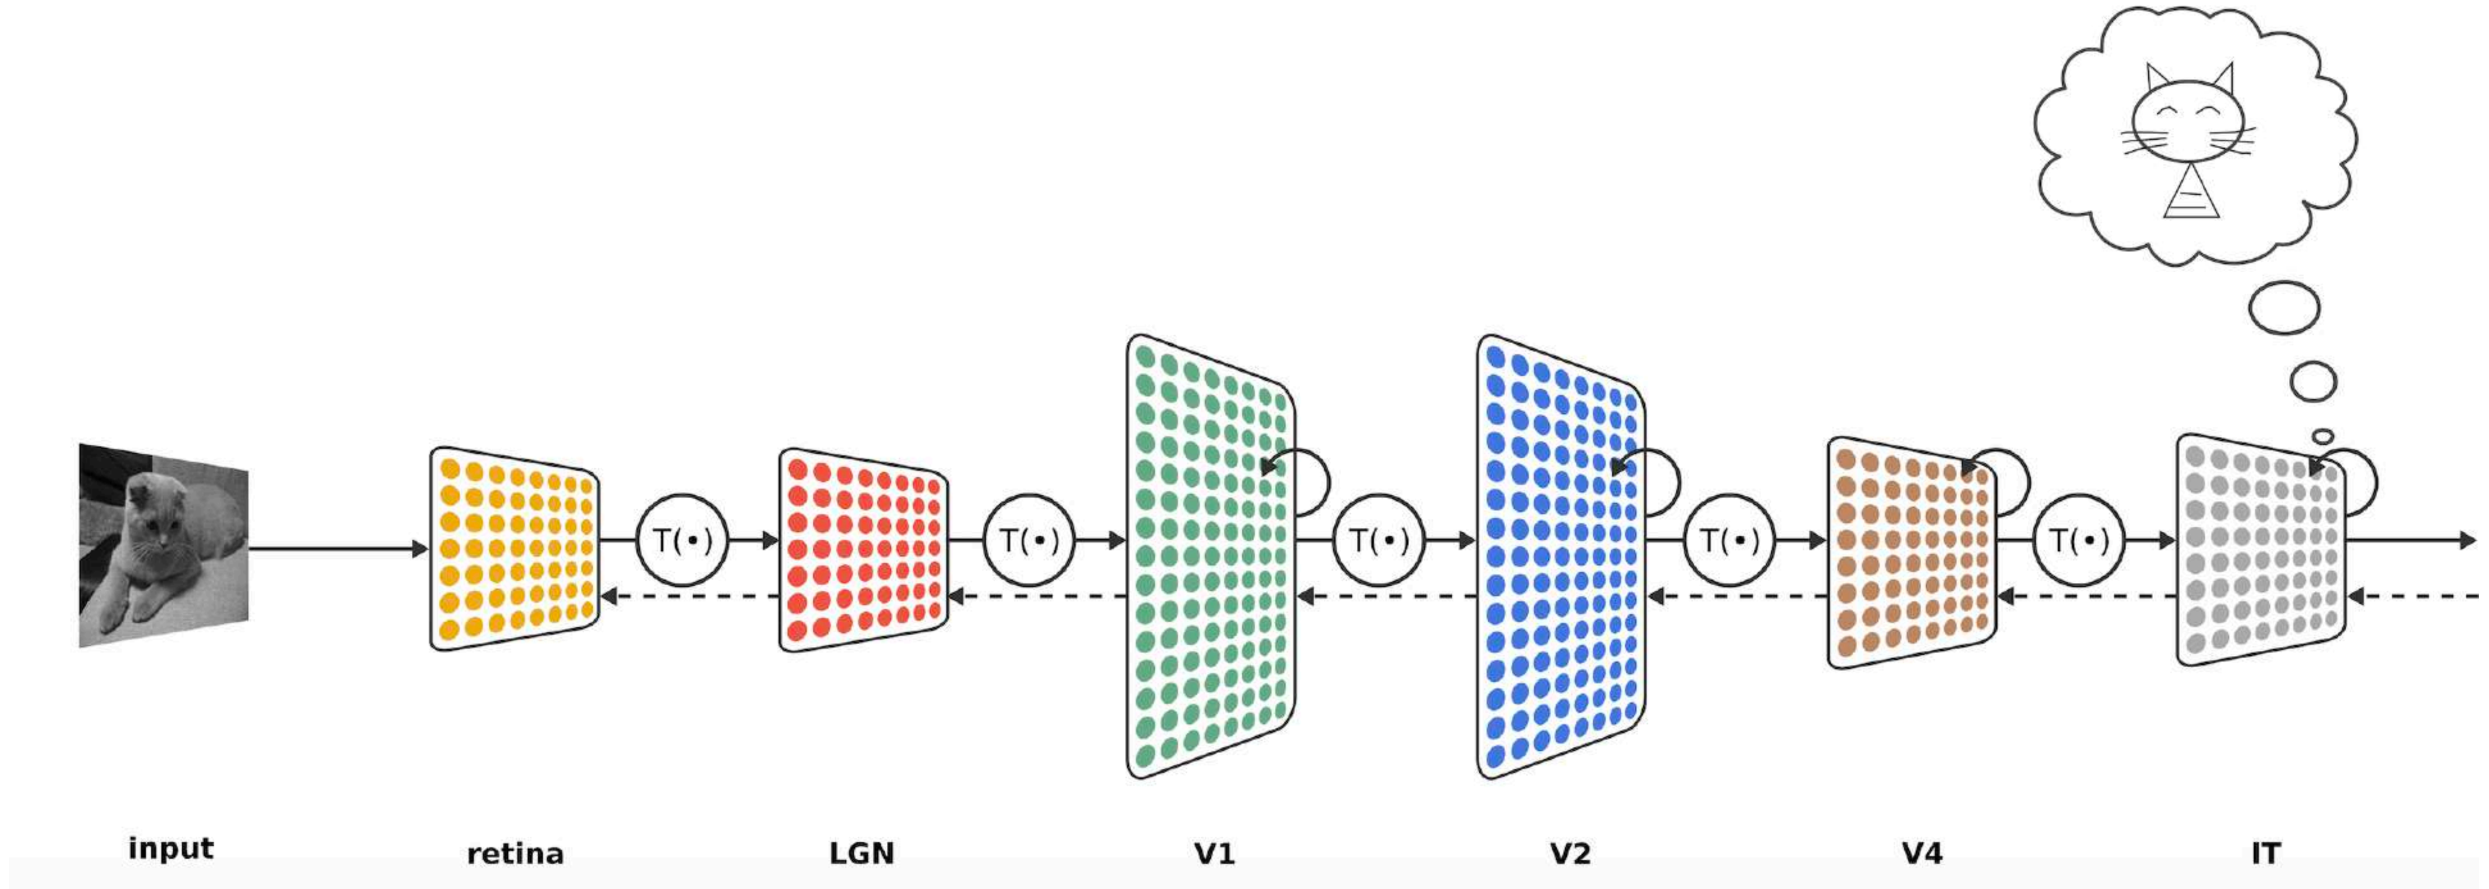

- 1986: Backpropagation to update weights by Rumelhart, Hinton & Williams

- 1989: Convolutional neural networks with backpropagation by LeCun

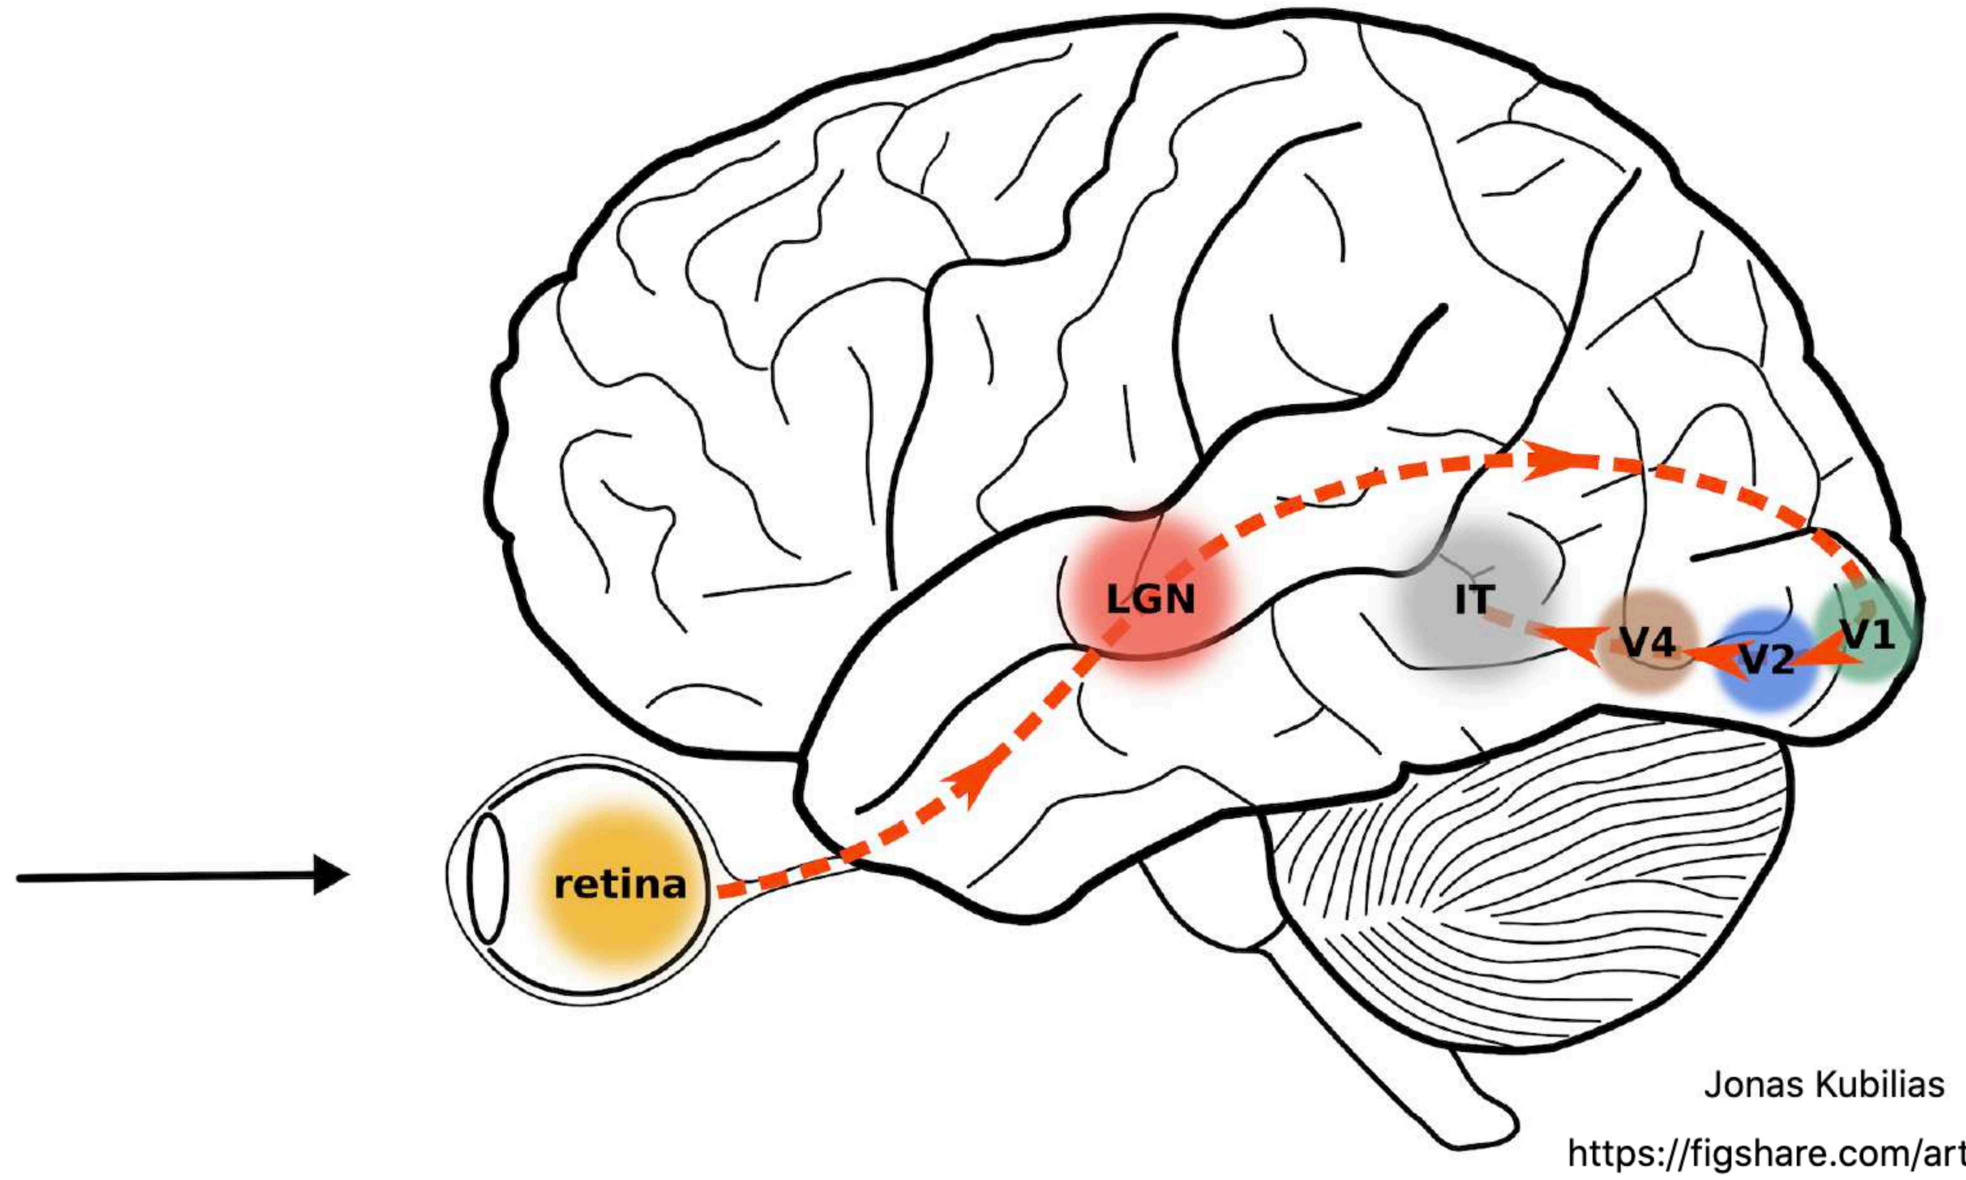

Jonas Kubilias

https://figshare.com/articles/figure/Ventral_visual_stream/106794

Deep networks can be made to be much more biologically realistic

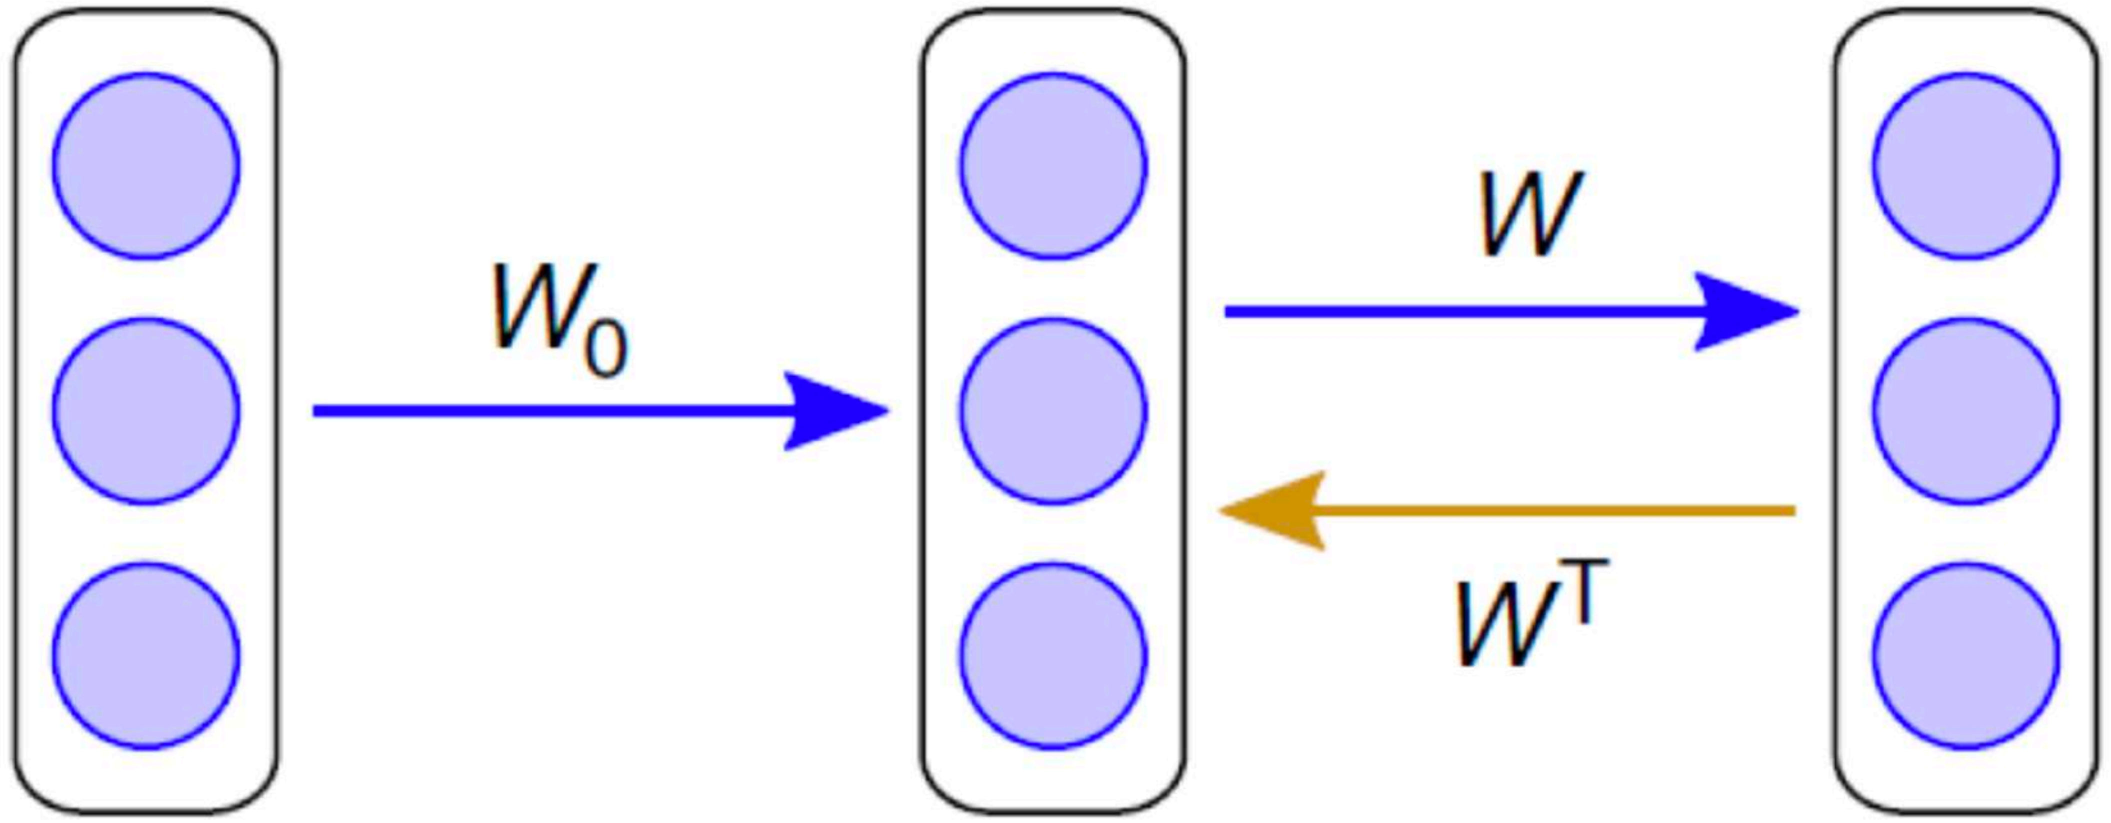

Lillicrap…Akerman, 2016

- Realistic?

$$\text{Backprop:}\ \Delta W \propto \frac{\partial L}{\partial W} = -\mathbf{e} \cdot \mathbf{h}$$

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \mathbf{h}}\right)\left(\frac{\partial \mathbf{h}}{\partial W_0}\right) = -W^T \mathbf{e} \cdot \mathbf{x}$$

Deep networks can be made to be much more biologically realistic

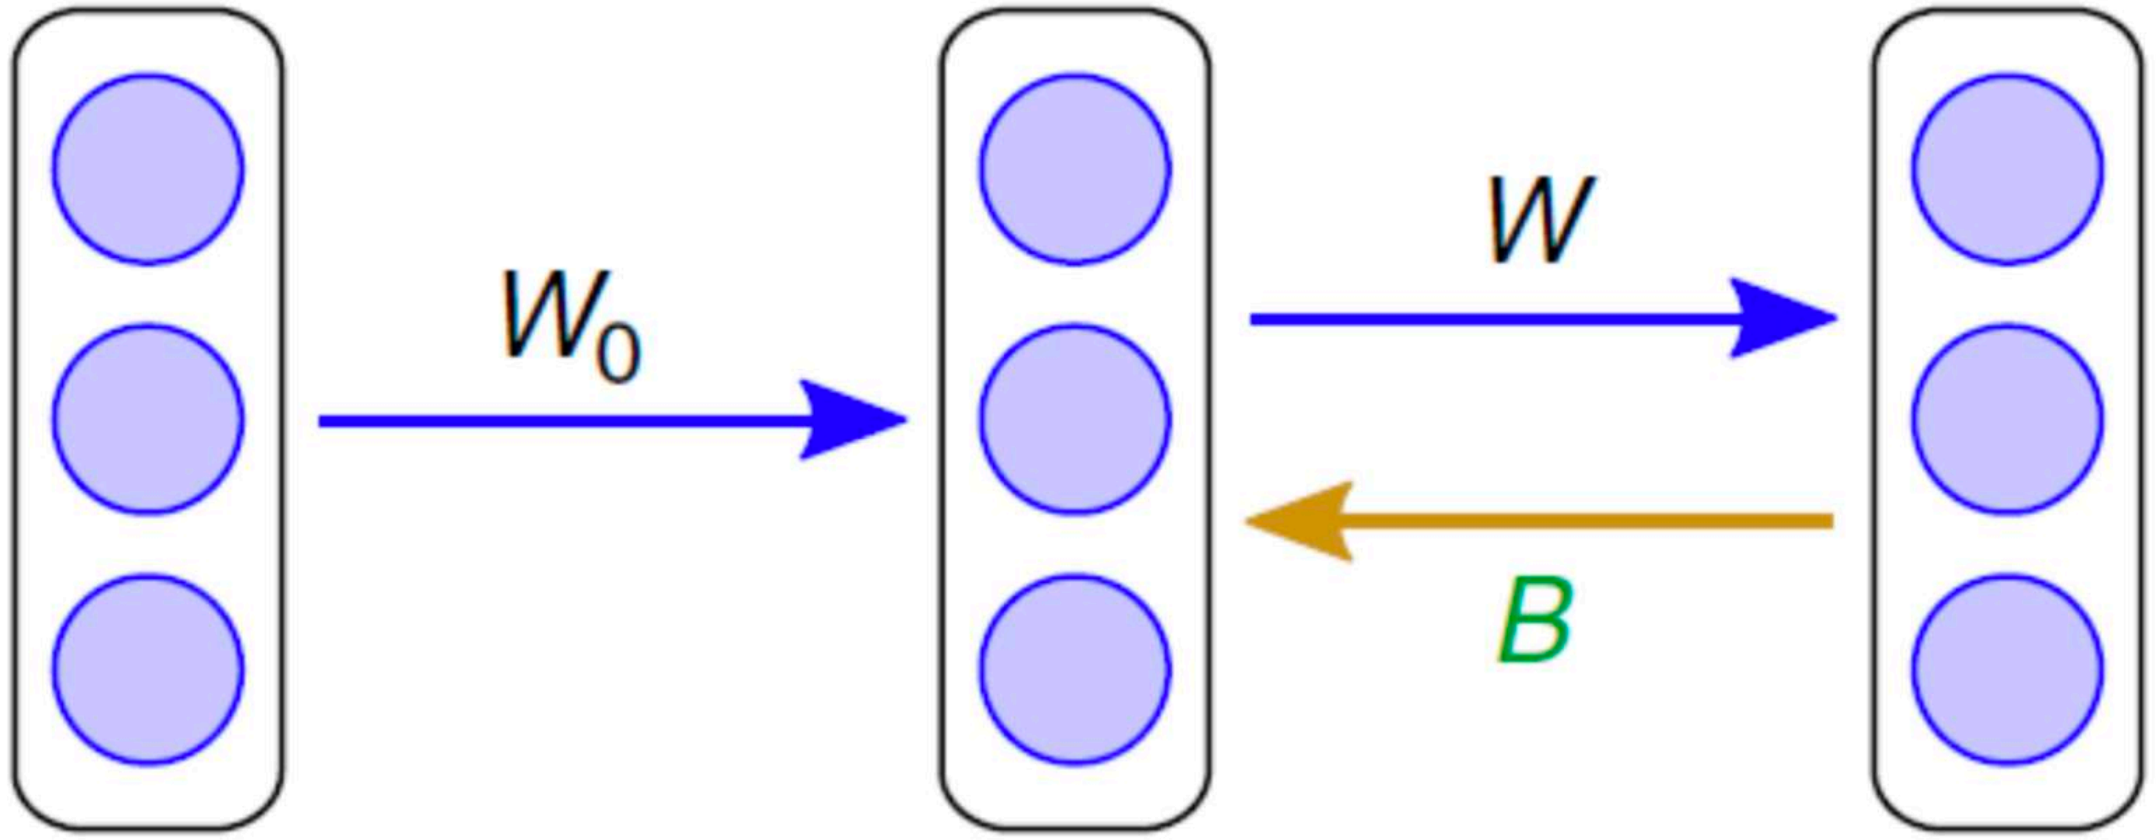

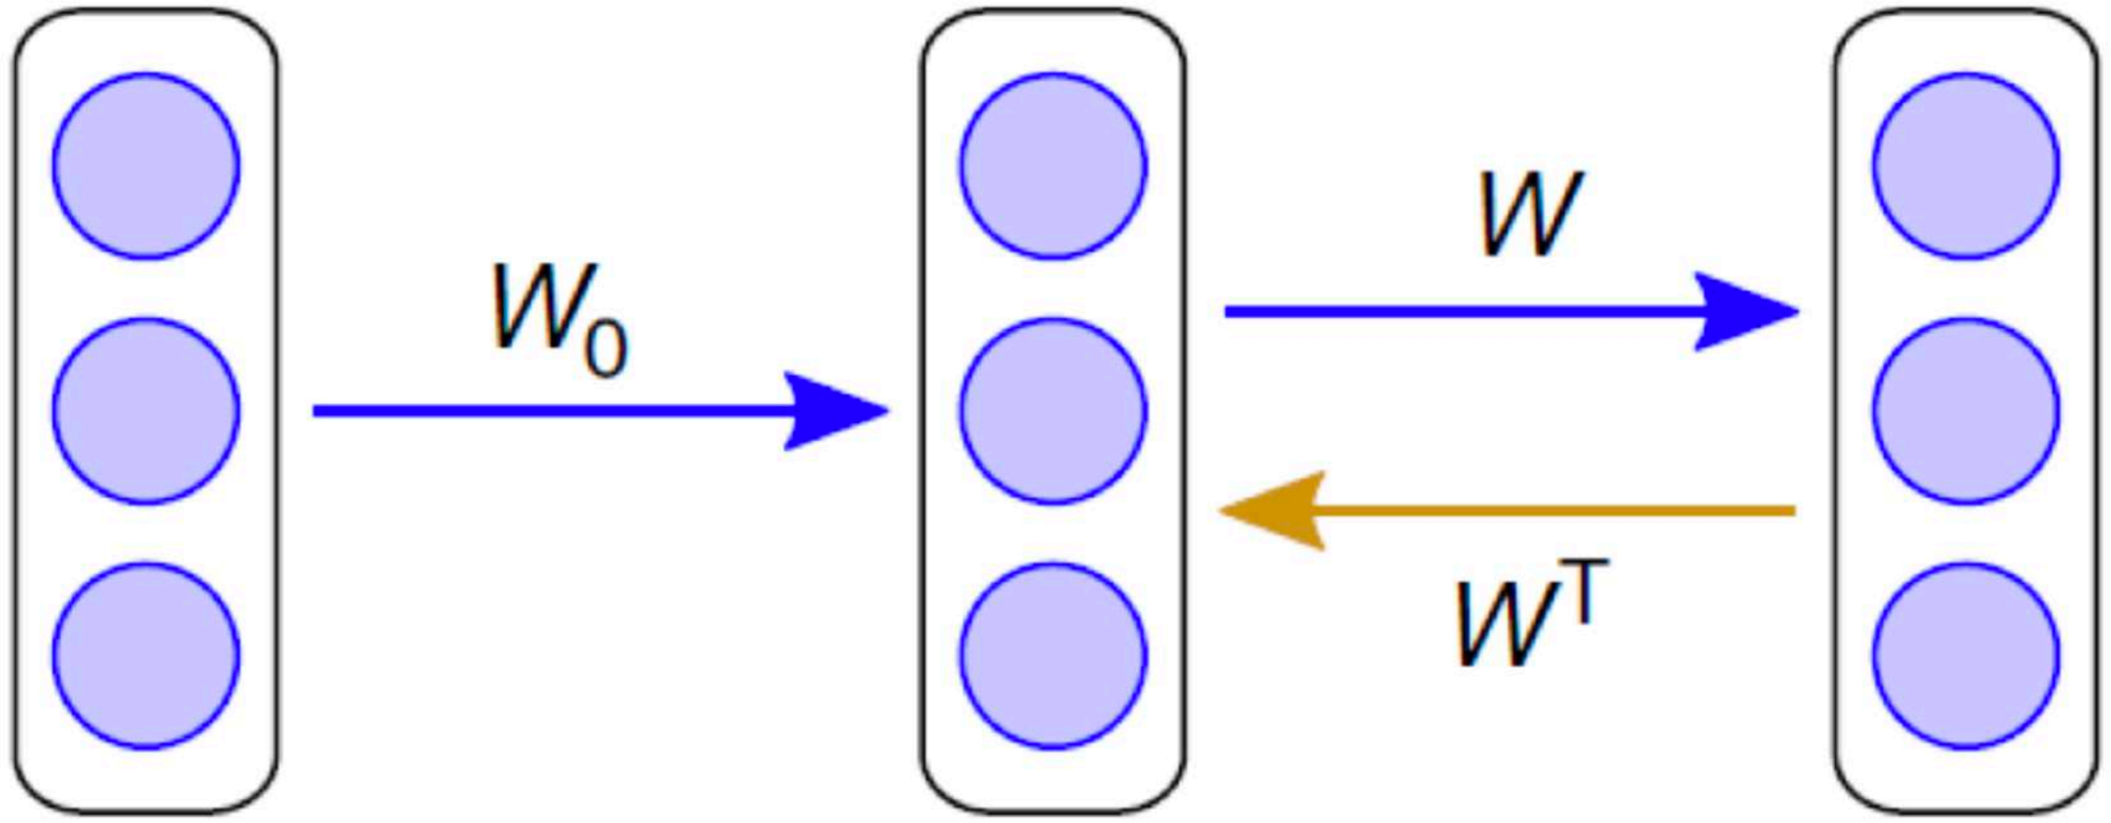

Backprop: $\Delta W \propto \frac{\partial L}{\partial W} = -\boldsymbol{e} \cdot \boldsymbol{h}$

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \boldsymbol{h}}\right)\left(\frac{\partial \boldsymbol{h}}{\partial W_0}\right) = -W^T \boldsymbol{e} \cdot \boldsymbol{x}$$

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \boldsymbol{h}}\right)\left(\frac{\partial \boldsymbol{h}}{\partial W_0}\right) = -W^T \boldsymbol{e} \cdot \boldsymbol{x}$$

Feedback alignment: $\Delta W_0 \propto B\boldsymbol{e} \cdot \boldsymbol{x}$, where B is a random matrix (with entries uniformly drawn from, e.g., [-0.5, 0.5])

Lillicrap…Akerman, 2016

Deep networks can be made to be much more biologically realistic

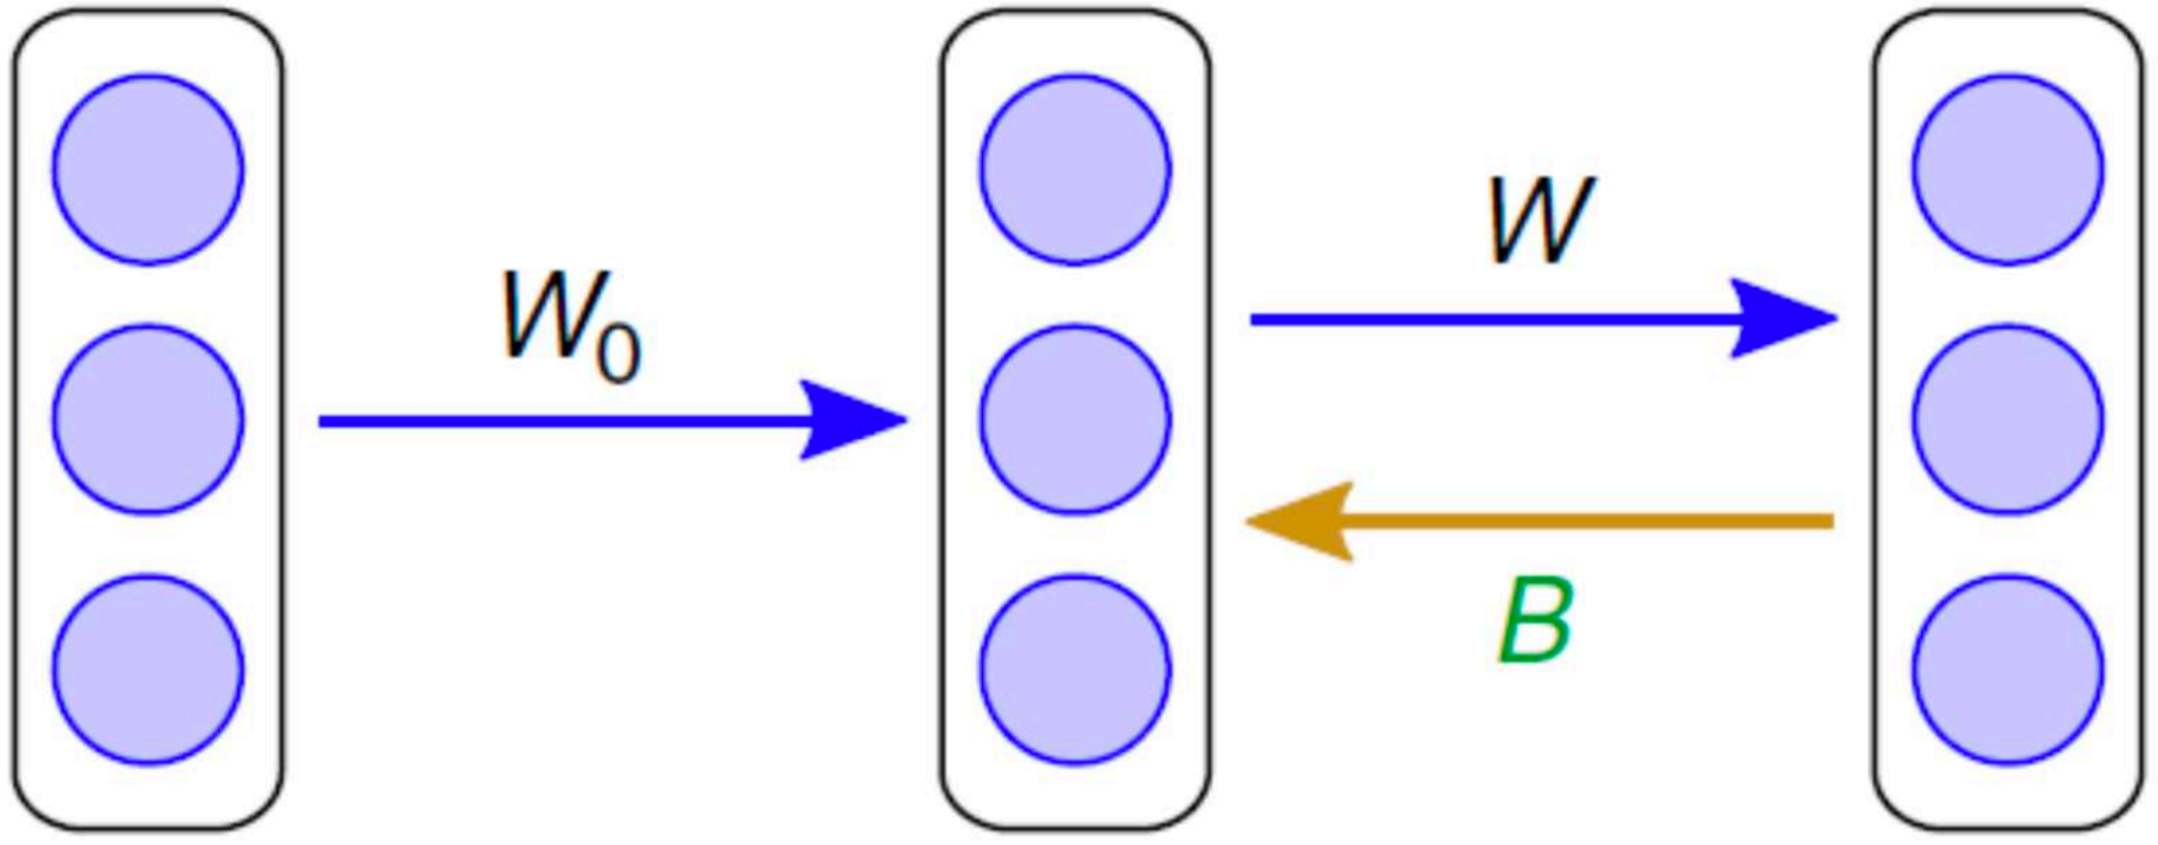

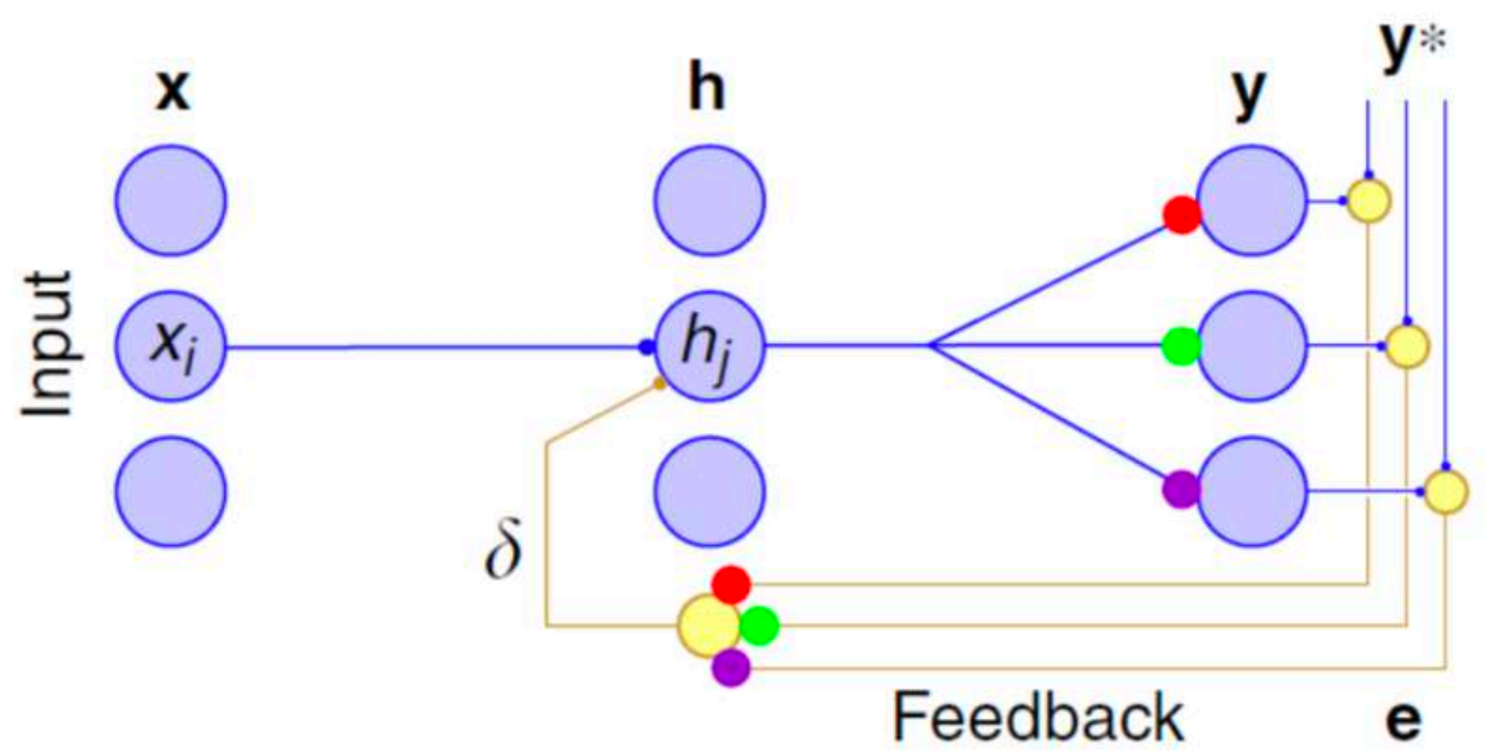

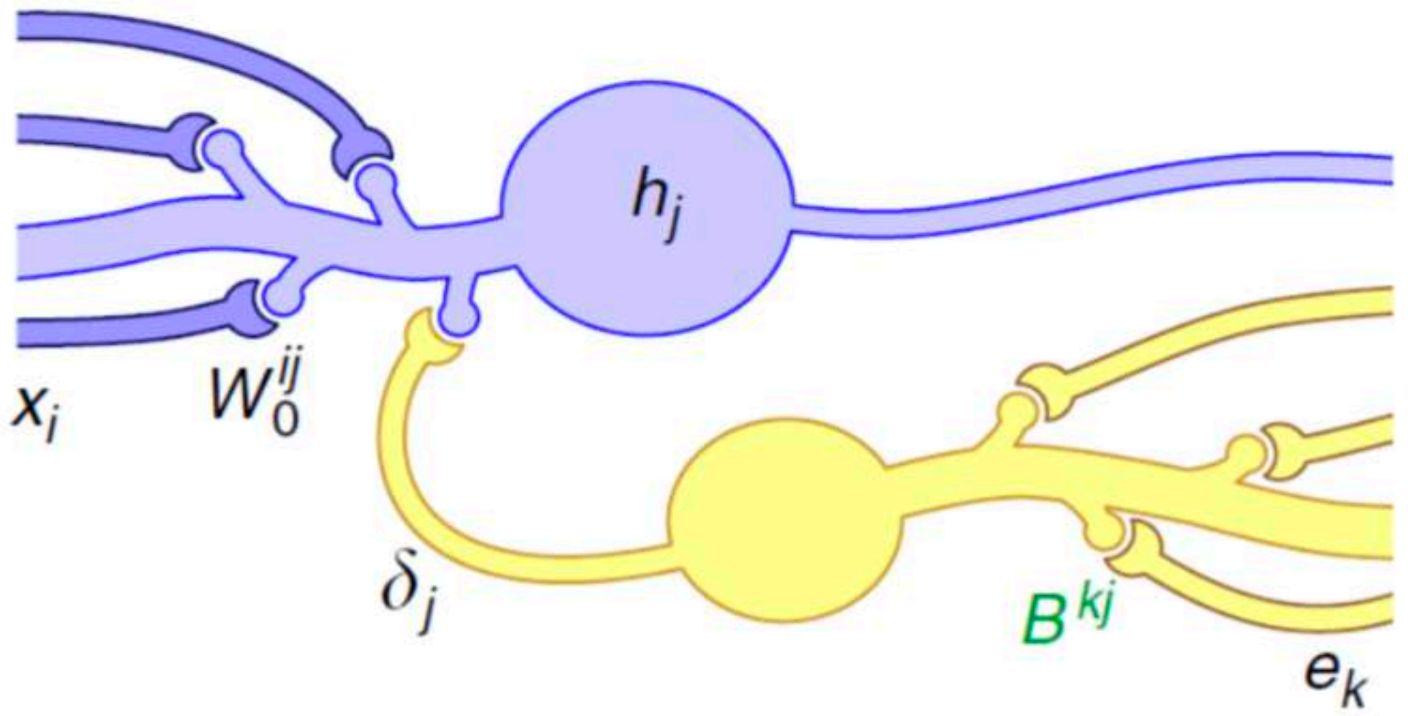

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \boldsymbol{h}}\right)\left(\frac{\partial \boldsymbol{h}}{\partial W_0}\right) = -W^T \boldsymbol{e} \cdot \boldsymbol{x}$$

Feedback alignment: $\Delta W_0 \propto B\boldsymbol{e} \cdot \boldsymbol{x}$, where B is a random matrix (with entries uniformly drawn from, e.g., [-0.5, 0.5])

Lillicrap…Akerman, 2016

Deep networks can be made to be much more biologically realistic

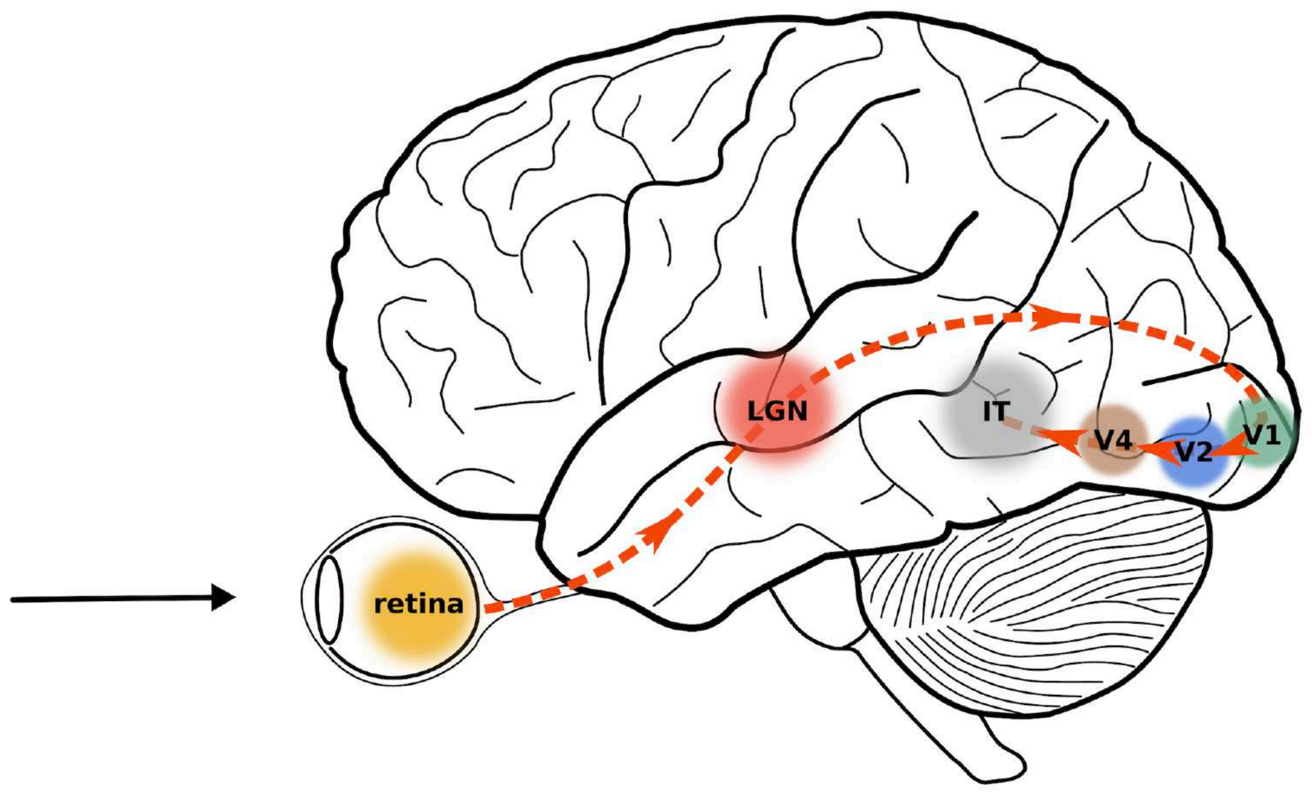

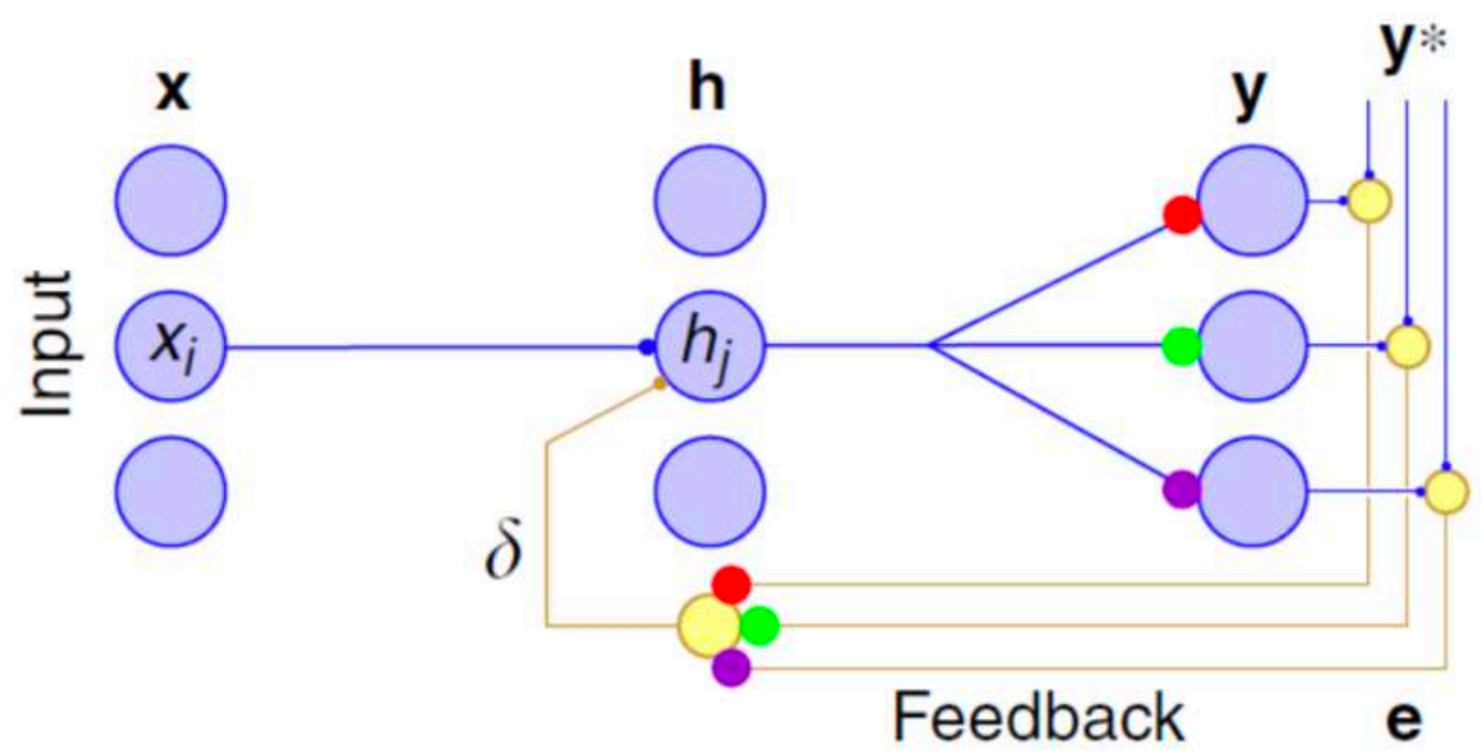

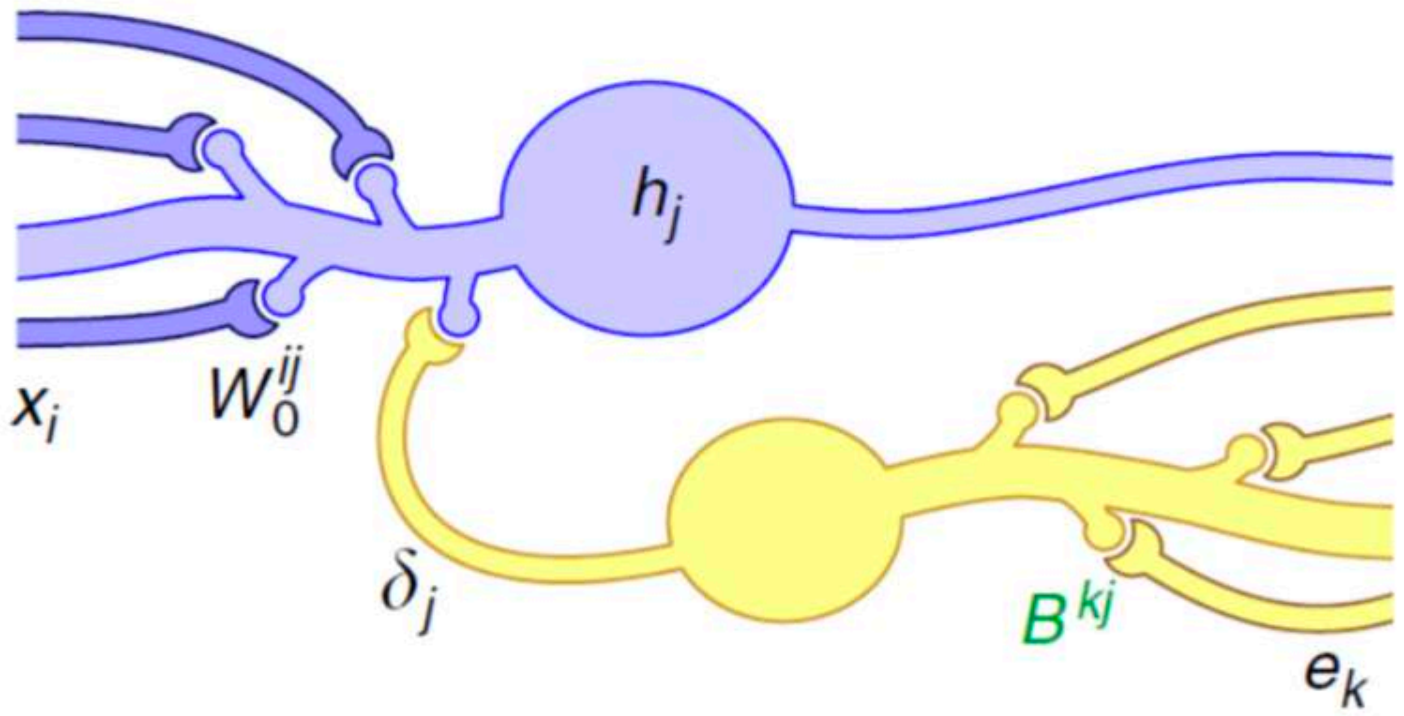

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \boldsymbol{h}}\right)\left(\frac{\partial \boldsymbol{h}}{\partial W_0}\right) = -W^T \boldsymbol{e} \cdot \boldsymbol{x}$$

Feedback alignment: $\Delta W_0 \propto B\boldsymbol{e} \cdot \boldsymbol{x}$, where B is a random matrix (with entries uniformly drawn from, e.g., [-0.5, 0.5])

Lillicrap…Akerman, 2016

Deep networks can be made to be much more biologically realistic

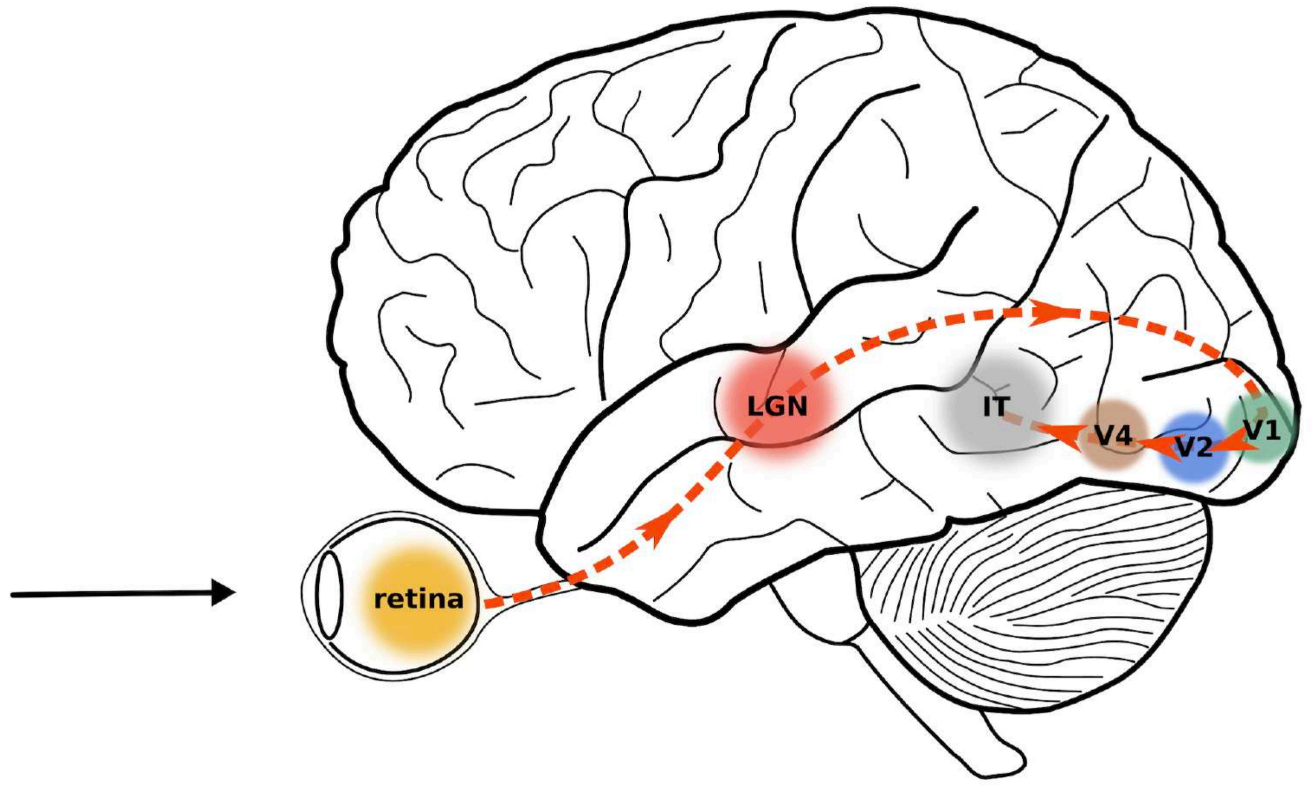

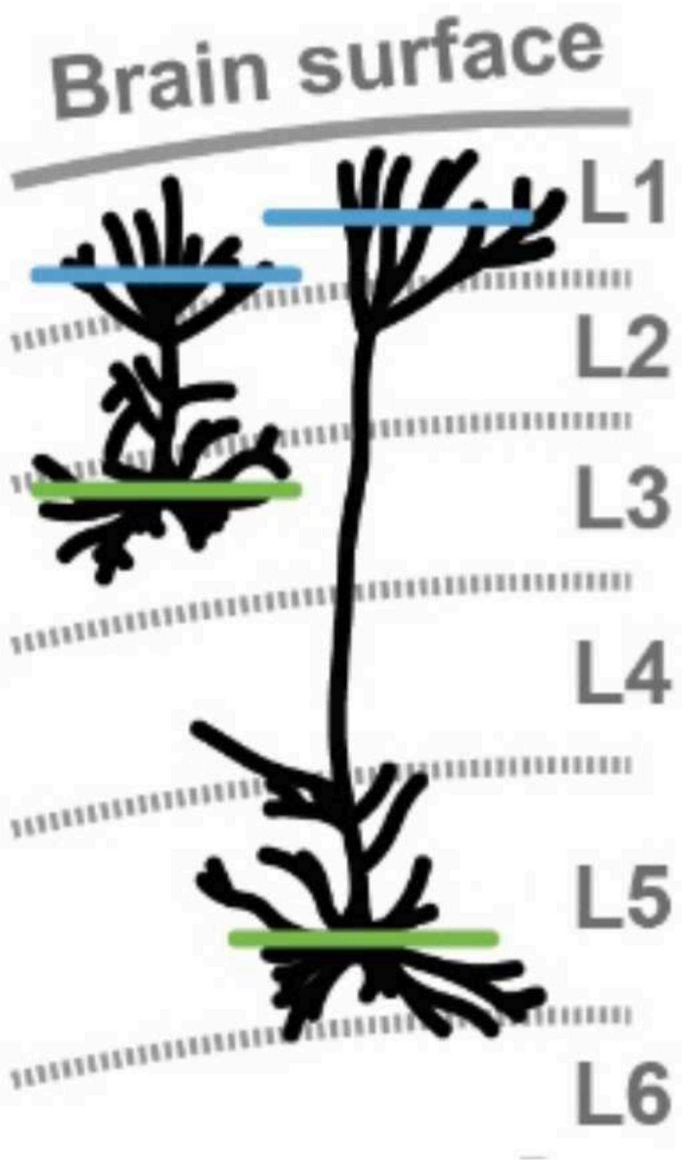

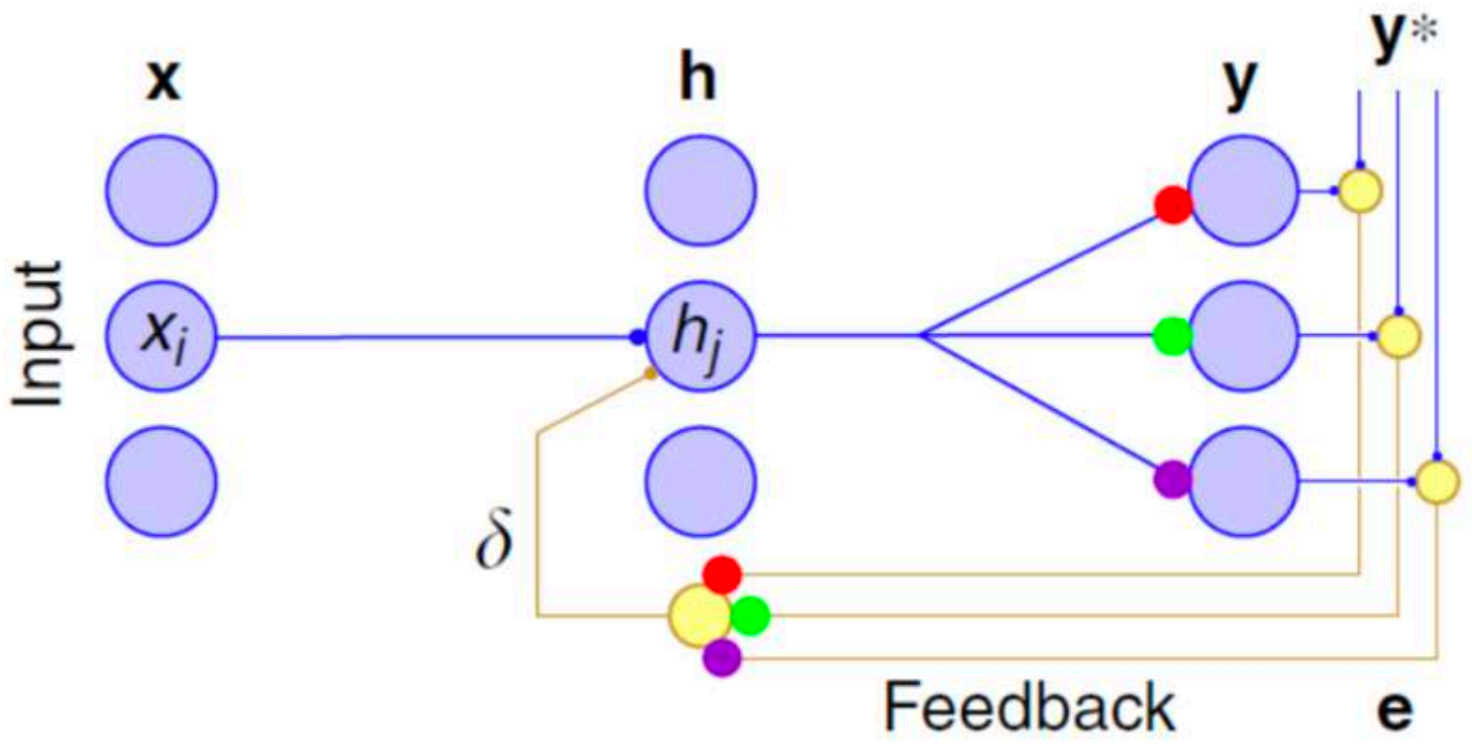

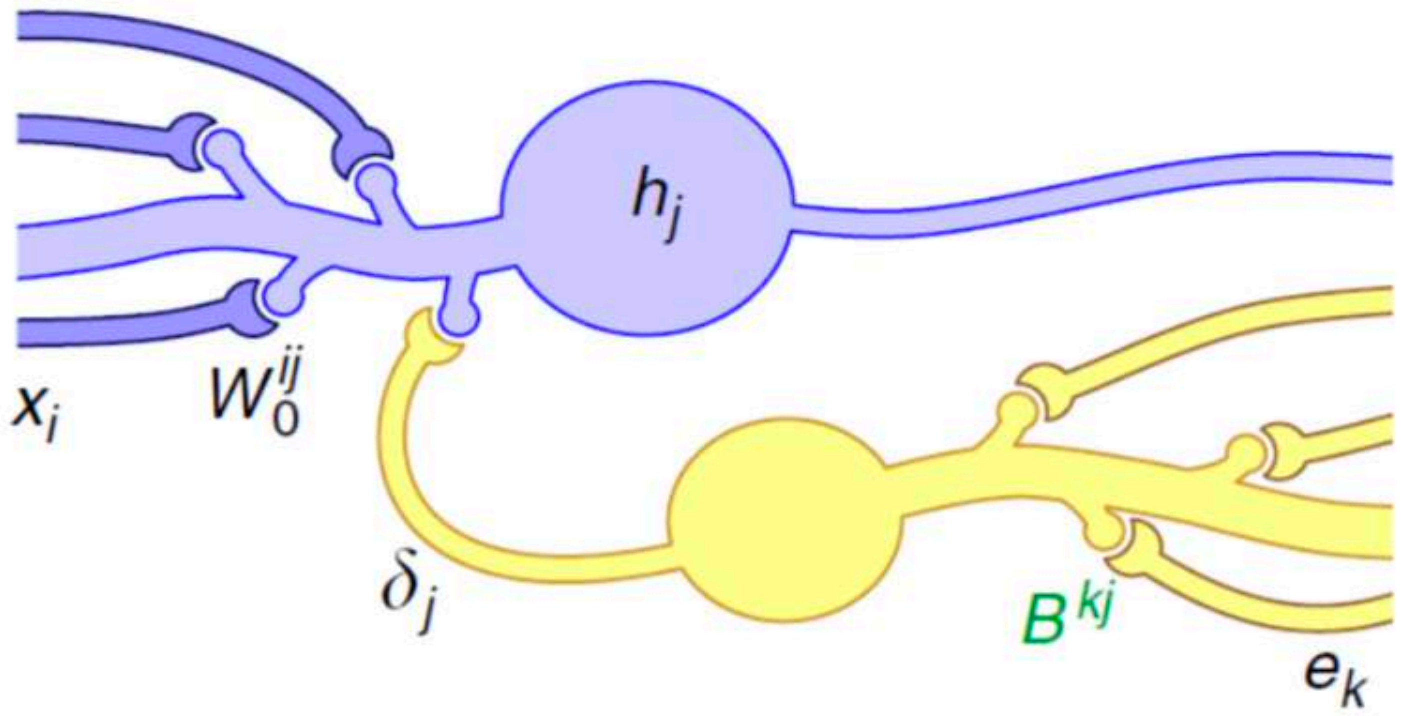

$$\Delta W_0 \propto \frac{\partial L}{\partial W_0} = \left(\frac{\partial L}{\partial \boldsymbol{h}}\right)\left(\frac{\partial \boldsymbol{h}}{\partial W_0}\right) = -W^T \boldsymbol{e} \cdot \boldsymbol{x}$$

Feedback alignment: $\Delta W_0 \propto B\boldsymbol{e} \cdot \boldsymbol{x}$, where B is a random matrix (with entries uniformly drawn from, e.g., [-0.5, 0.5])

Lillicrap…Akerman, 2016

Spatial resonance: Certain static visual stimuli can cause seizures and discomfort

- (A) Periodic patterns with certain spatial frequencies can cause epileptic seizures
- (B) In those without epilepsy, the same stimuli can cause headaches, illusions, and general aversion and discomfort
- (C) More complicated images composed of many wavenumbers can also cause discomfort

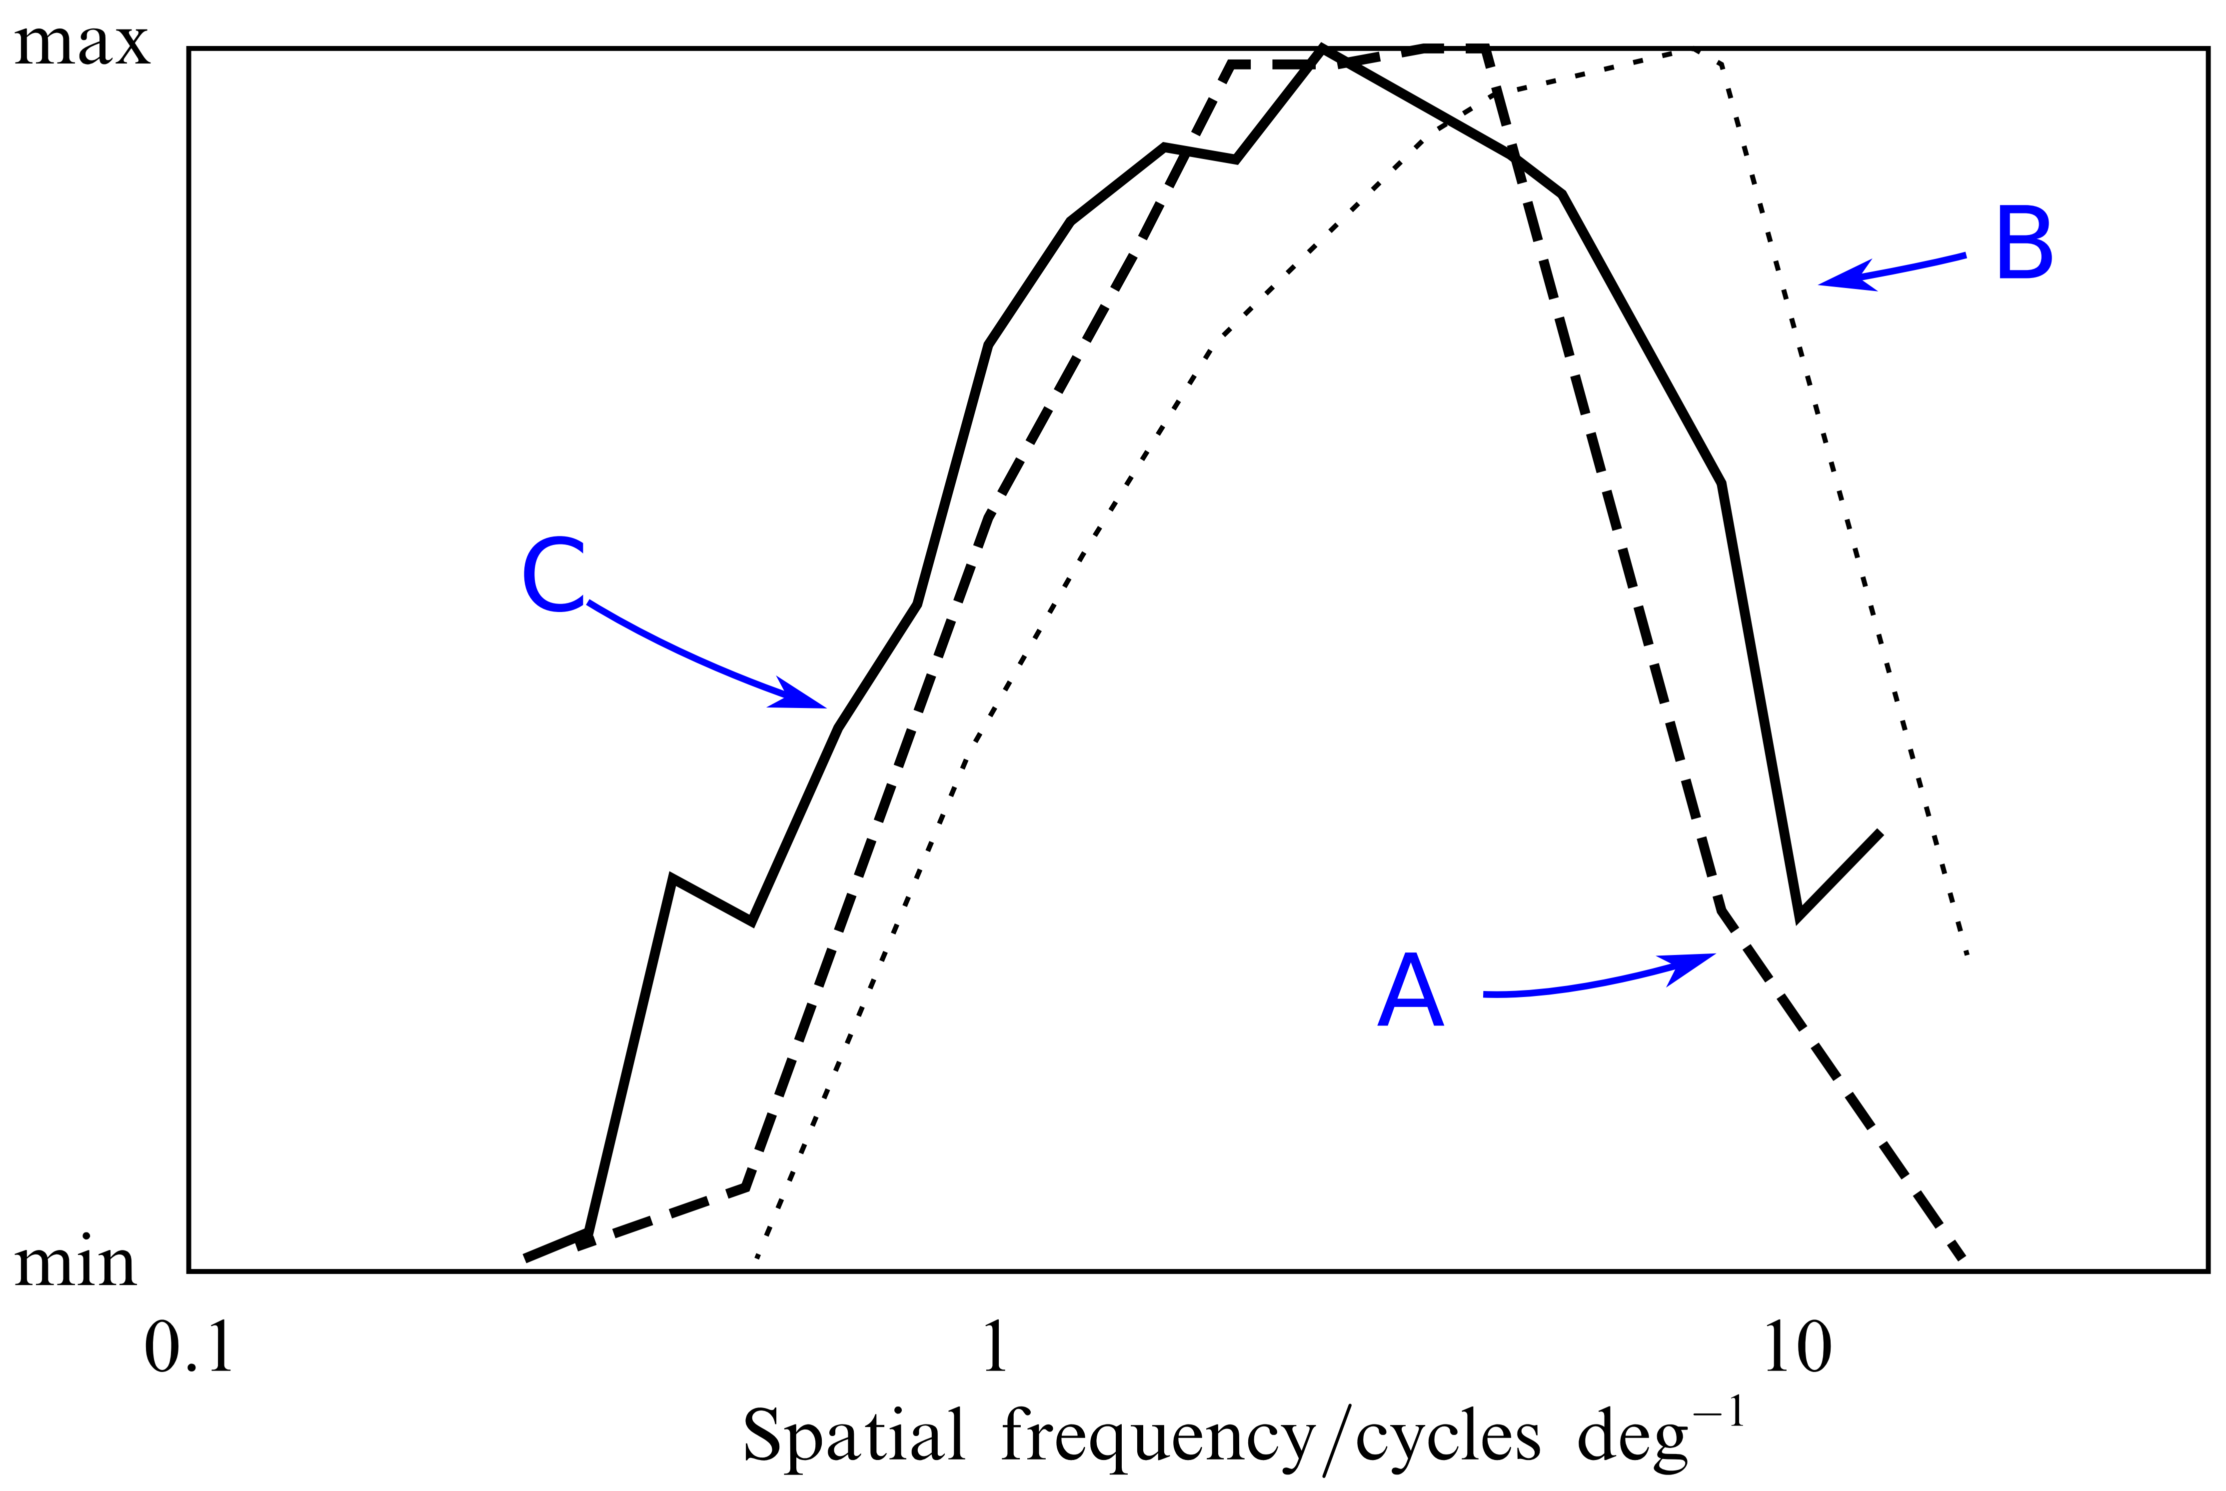

Fernandez & Wilkins, 2008

Oscillatory patterns are observed in response to the stimuli

If you have a visual epilepsy, please look away for the next slide, as an example stimulus will be shown

Oscillatory patterns are observed in response to the stimuli

1

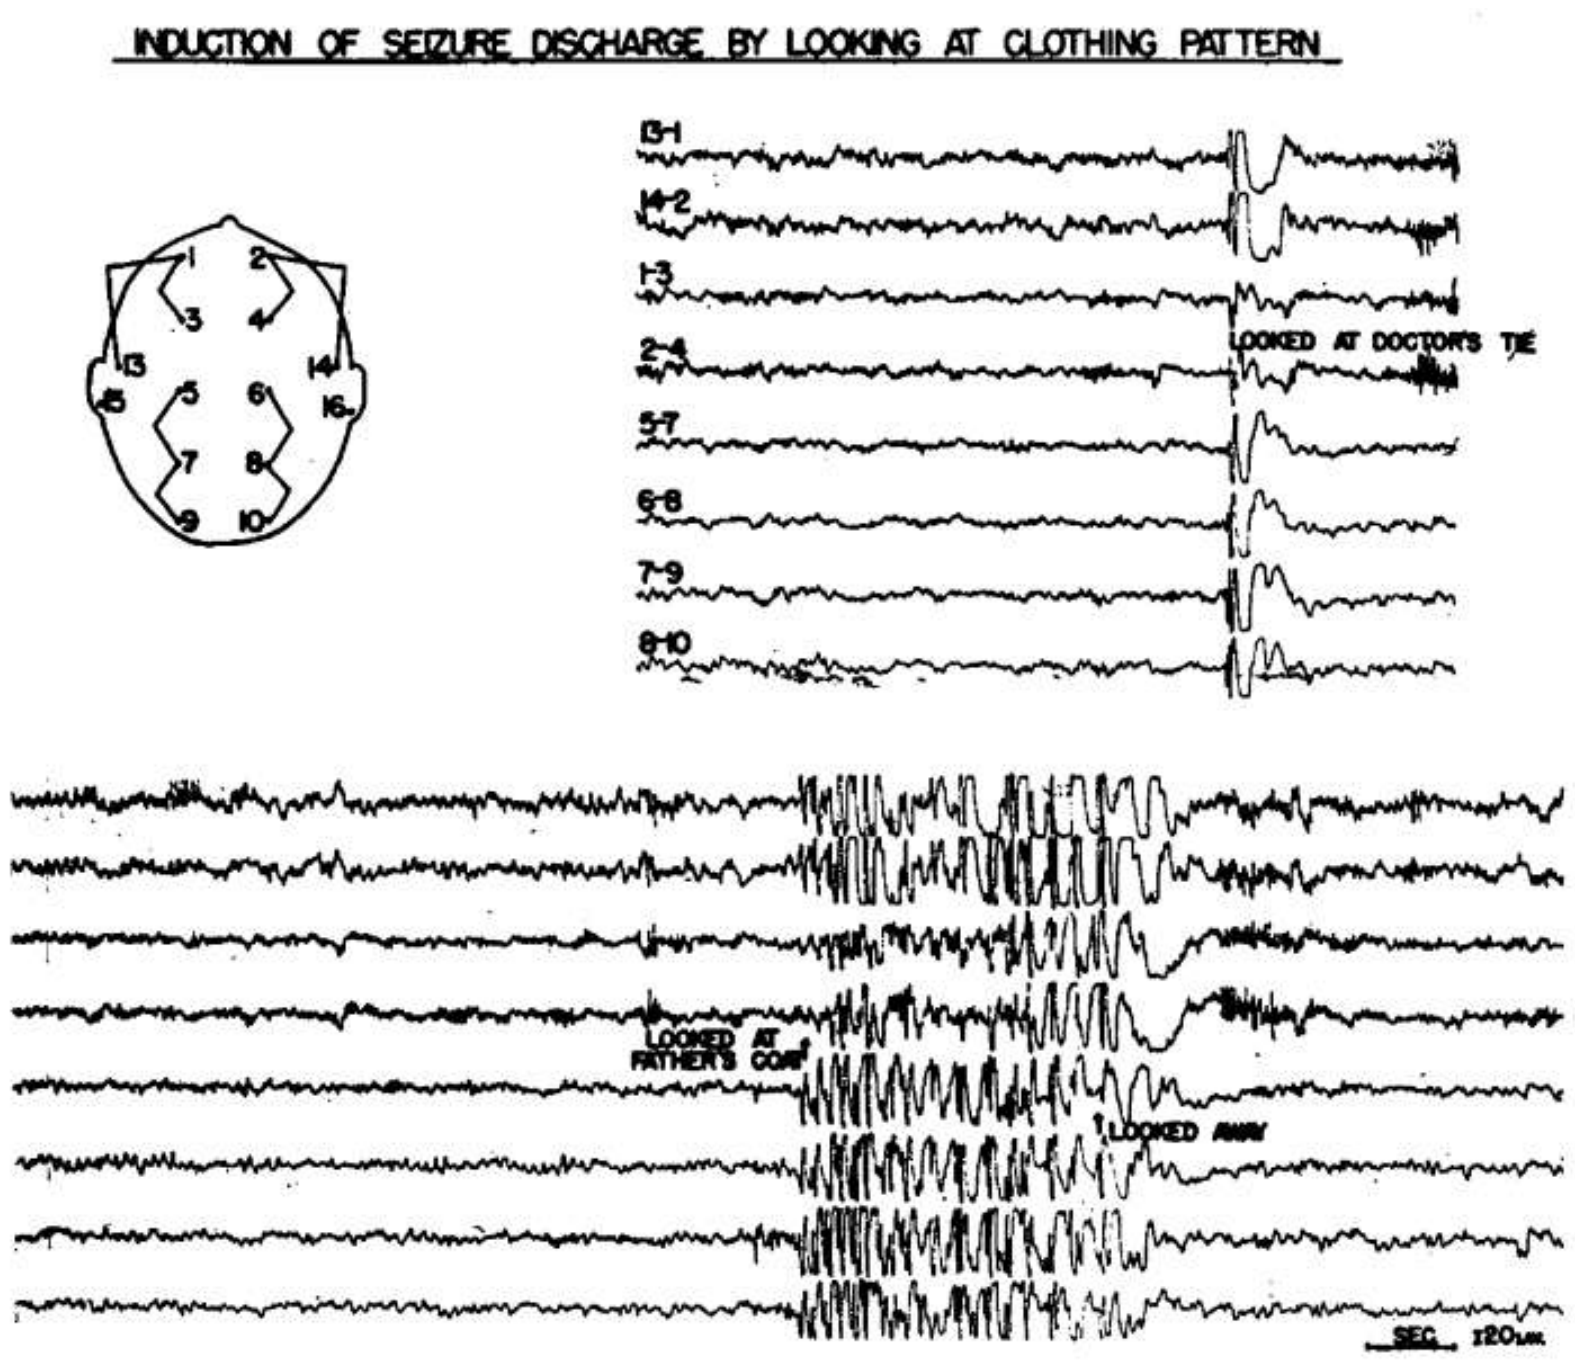

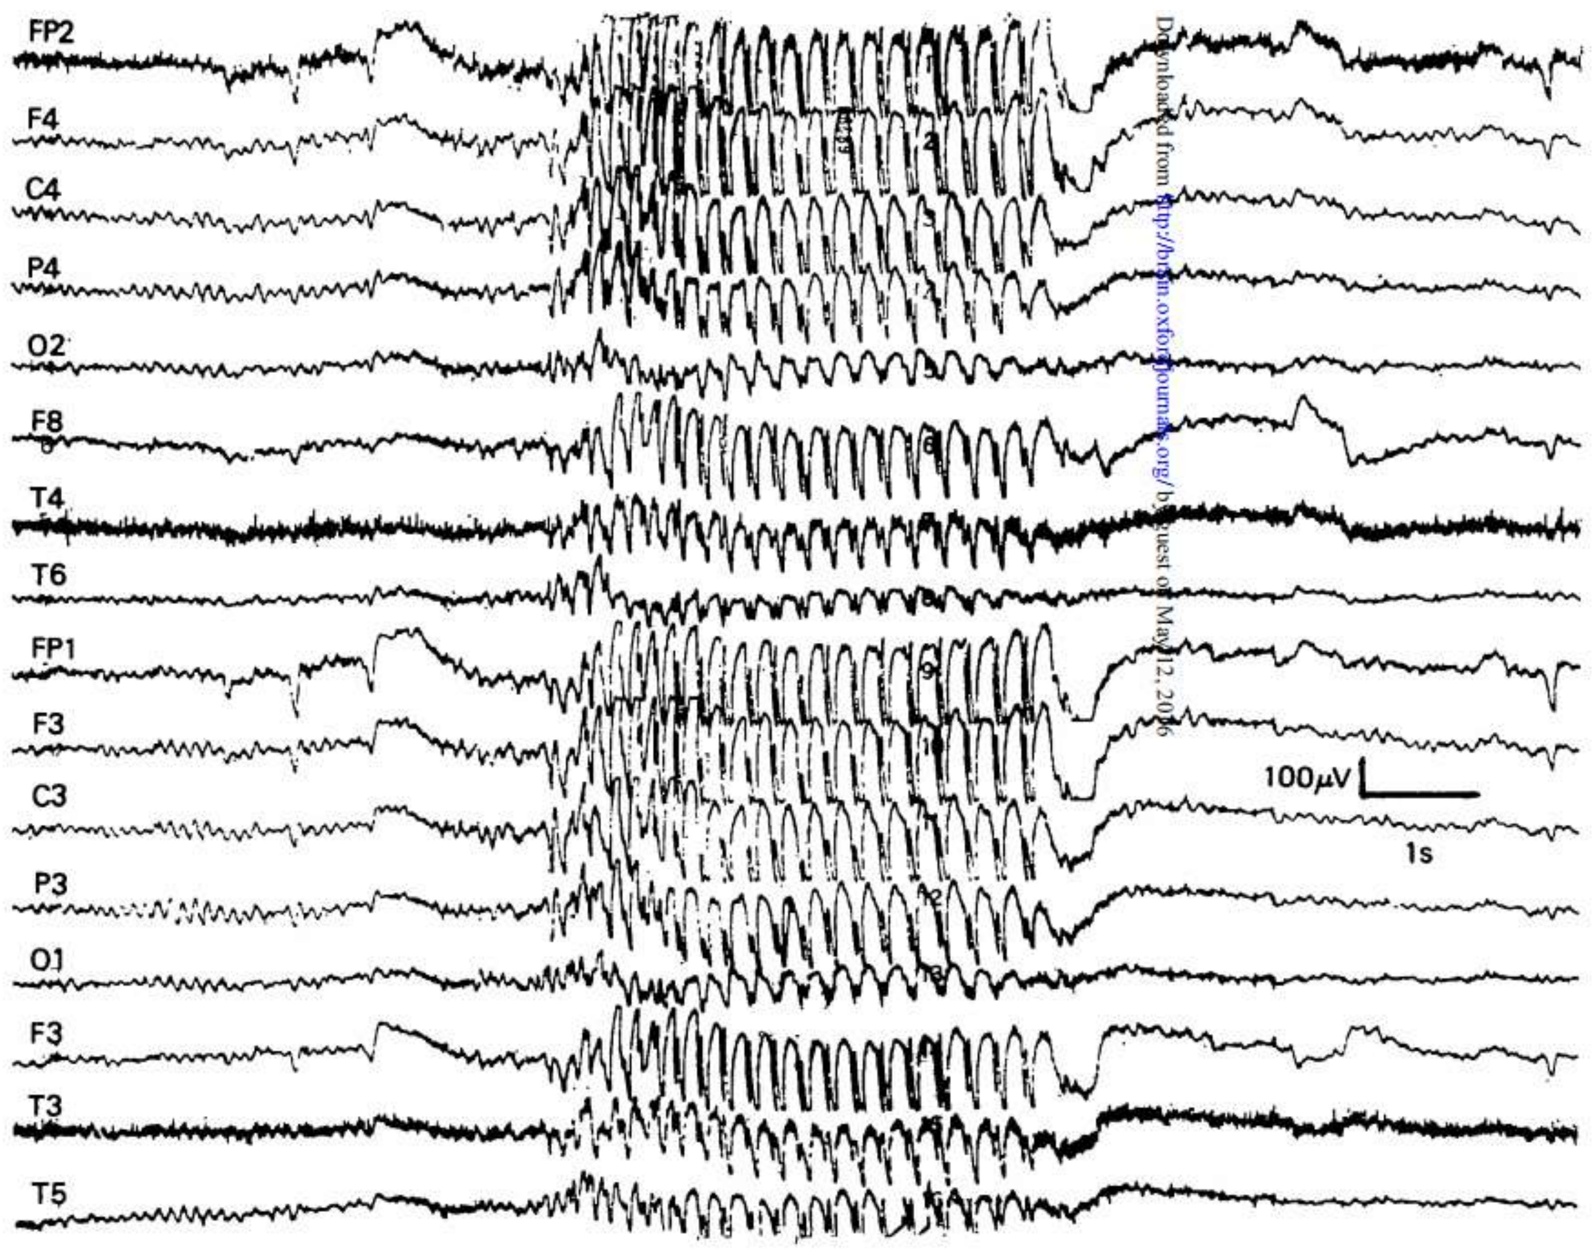

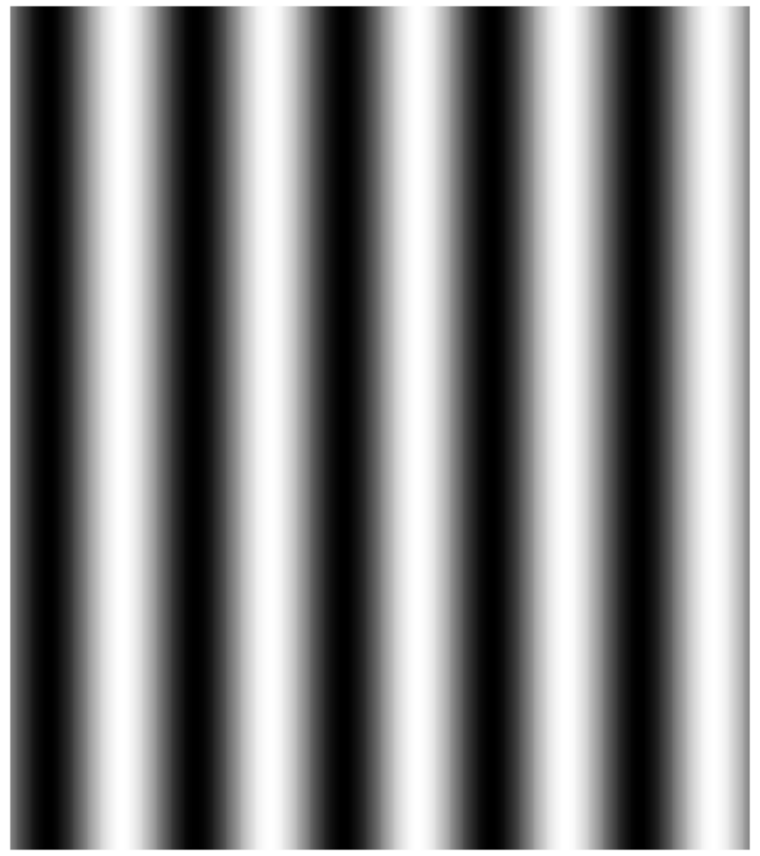

- Striped patterns, such as sine- and square-wave gratings trigger such seizures
- Screen doors, copper mesh, corduroy could all trigger epileptiform activity

[1] Bickford & Keith, 1953

Oscillatory patterns are observed in response to the stimuli

1

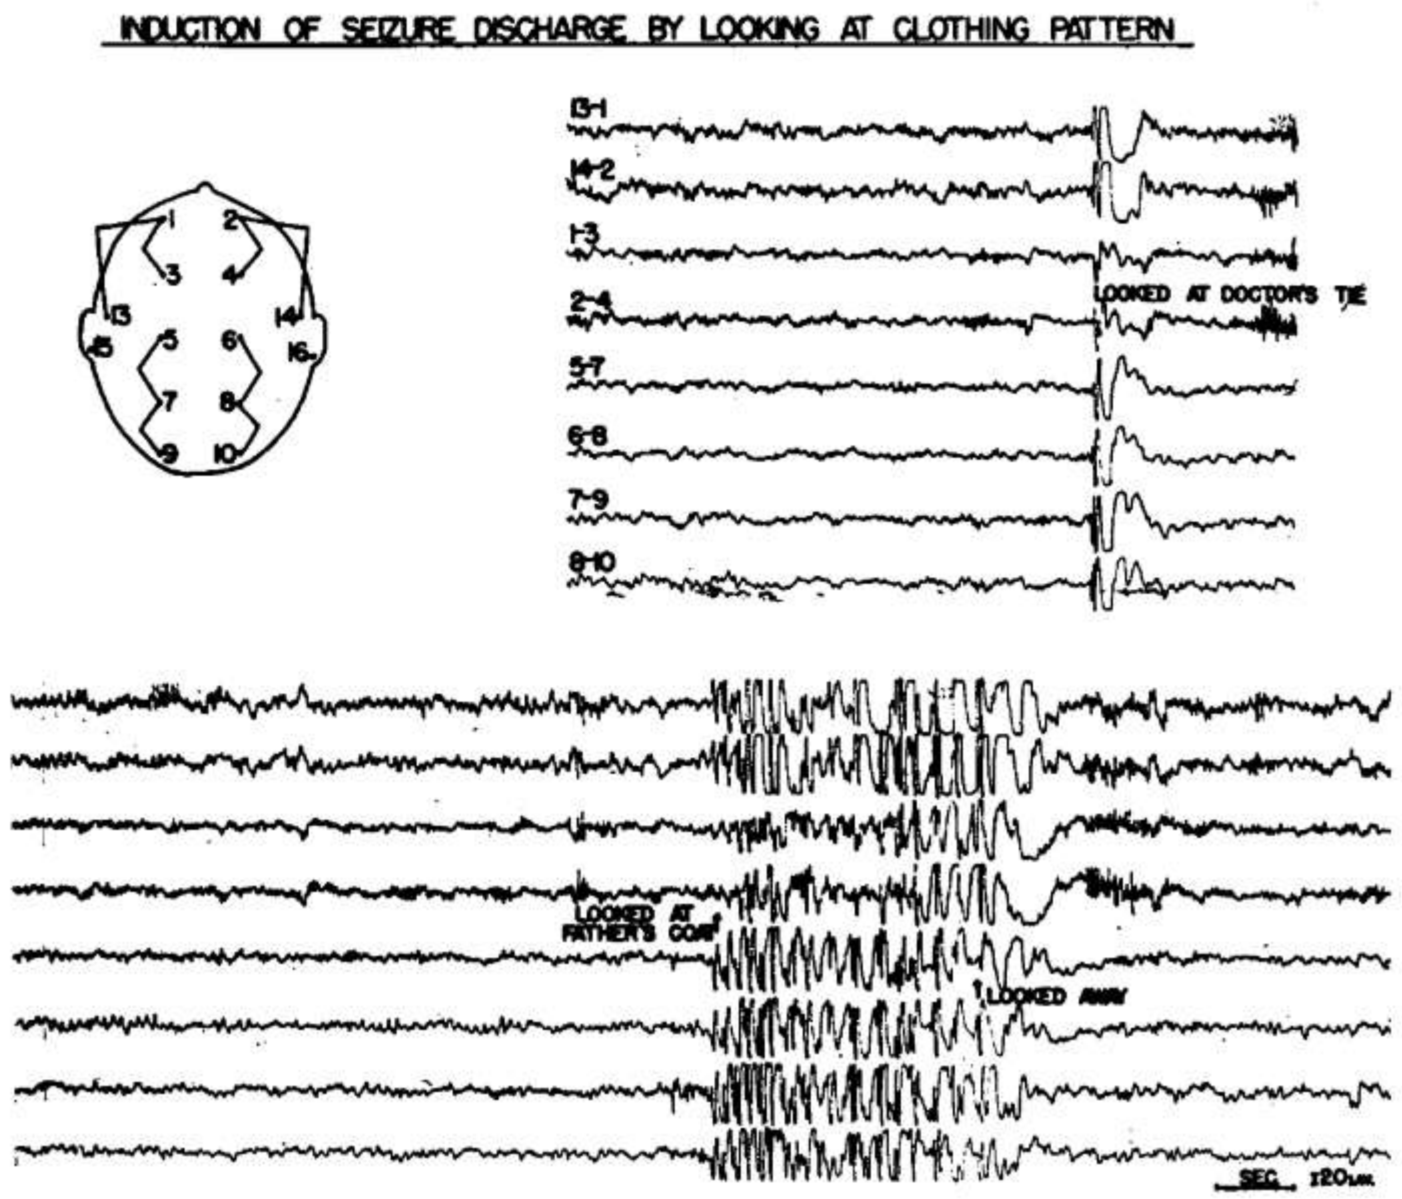

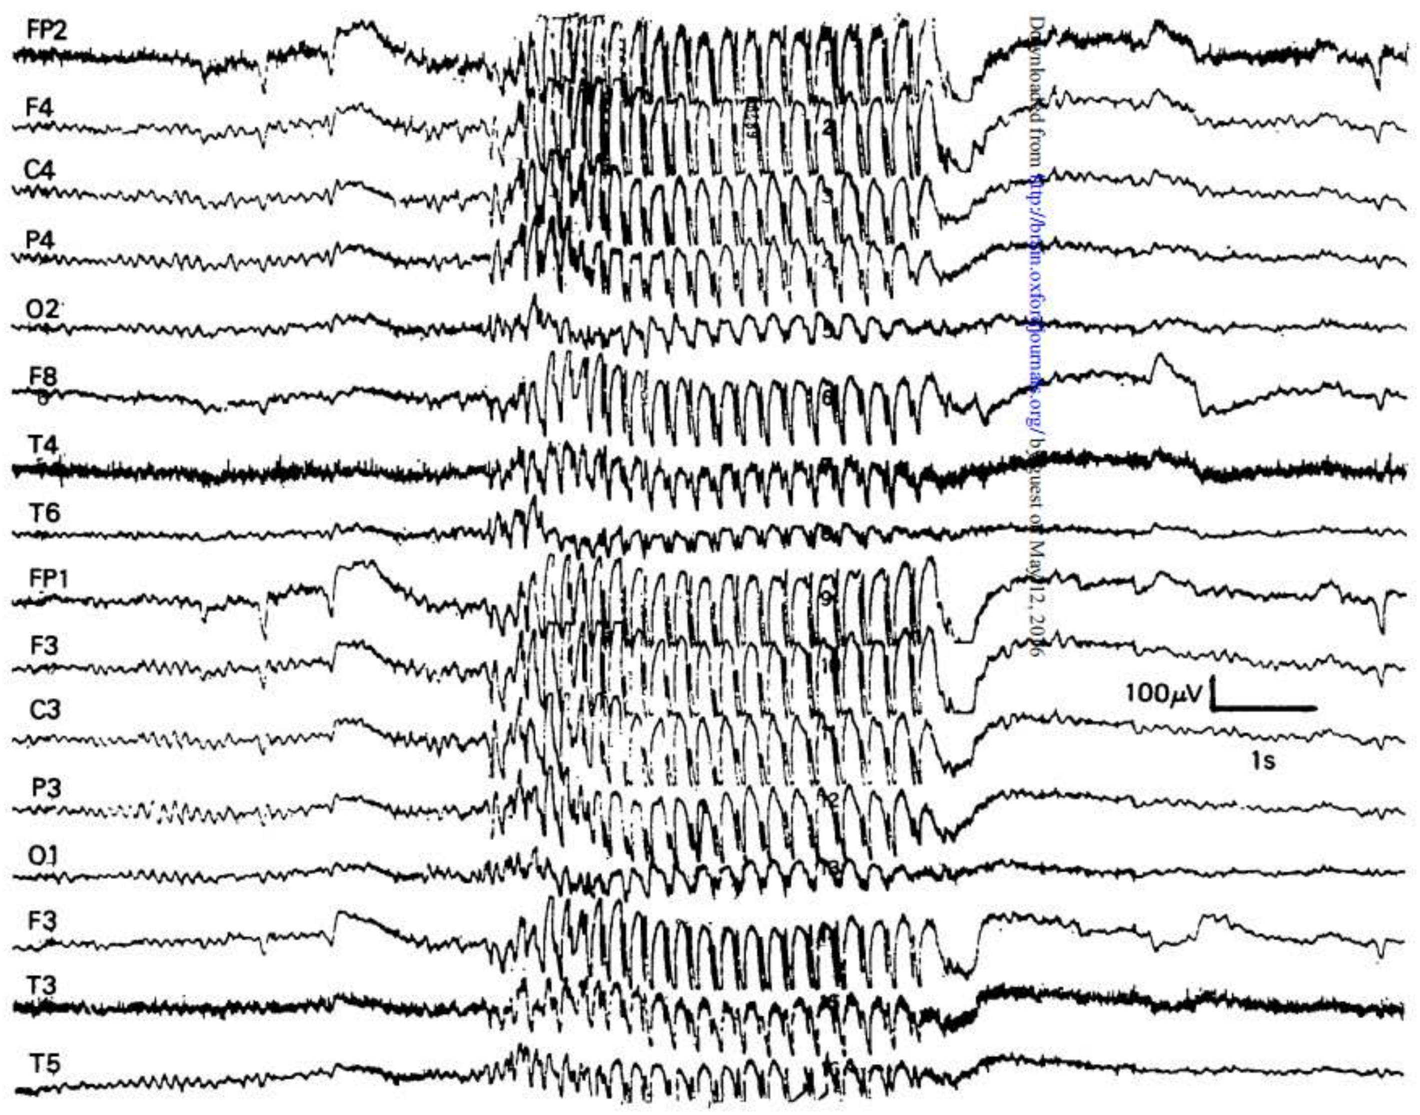

2

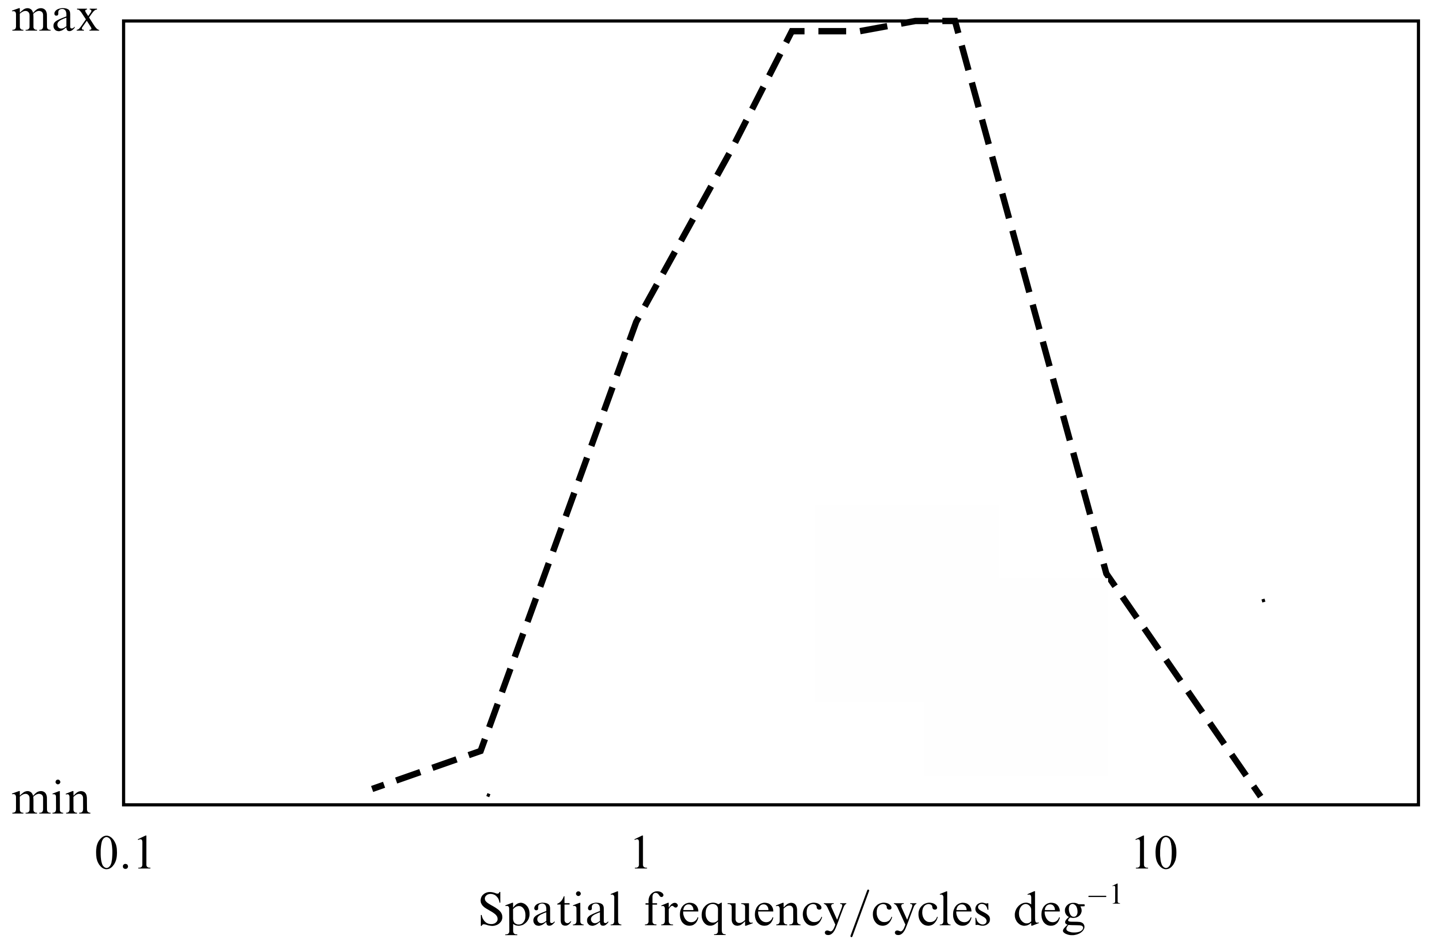

- Spatial frequencies esp. near 2-4 cpd induce the seizures
- Similar oscillatory activity observed in the case of visual discomfort

1 Bickford & Keith, 1953
2 Fernandez & Wilkins, 2008

Oscillatory patterns are observed in response to the stimuli

1

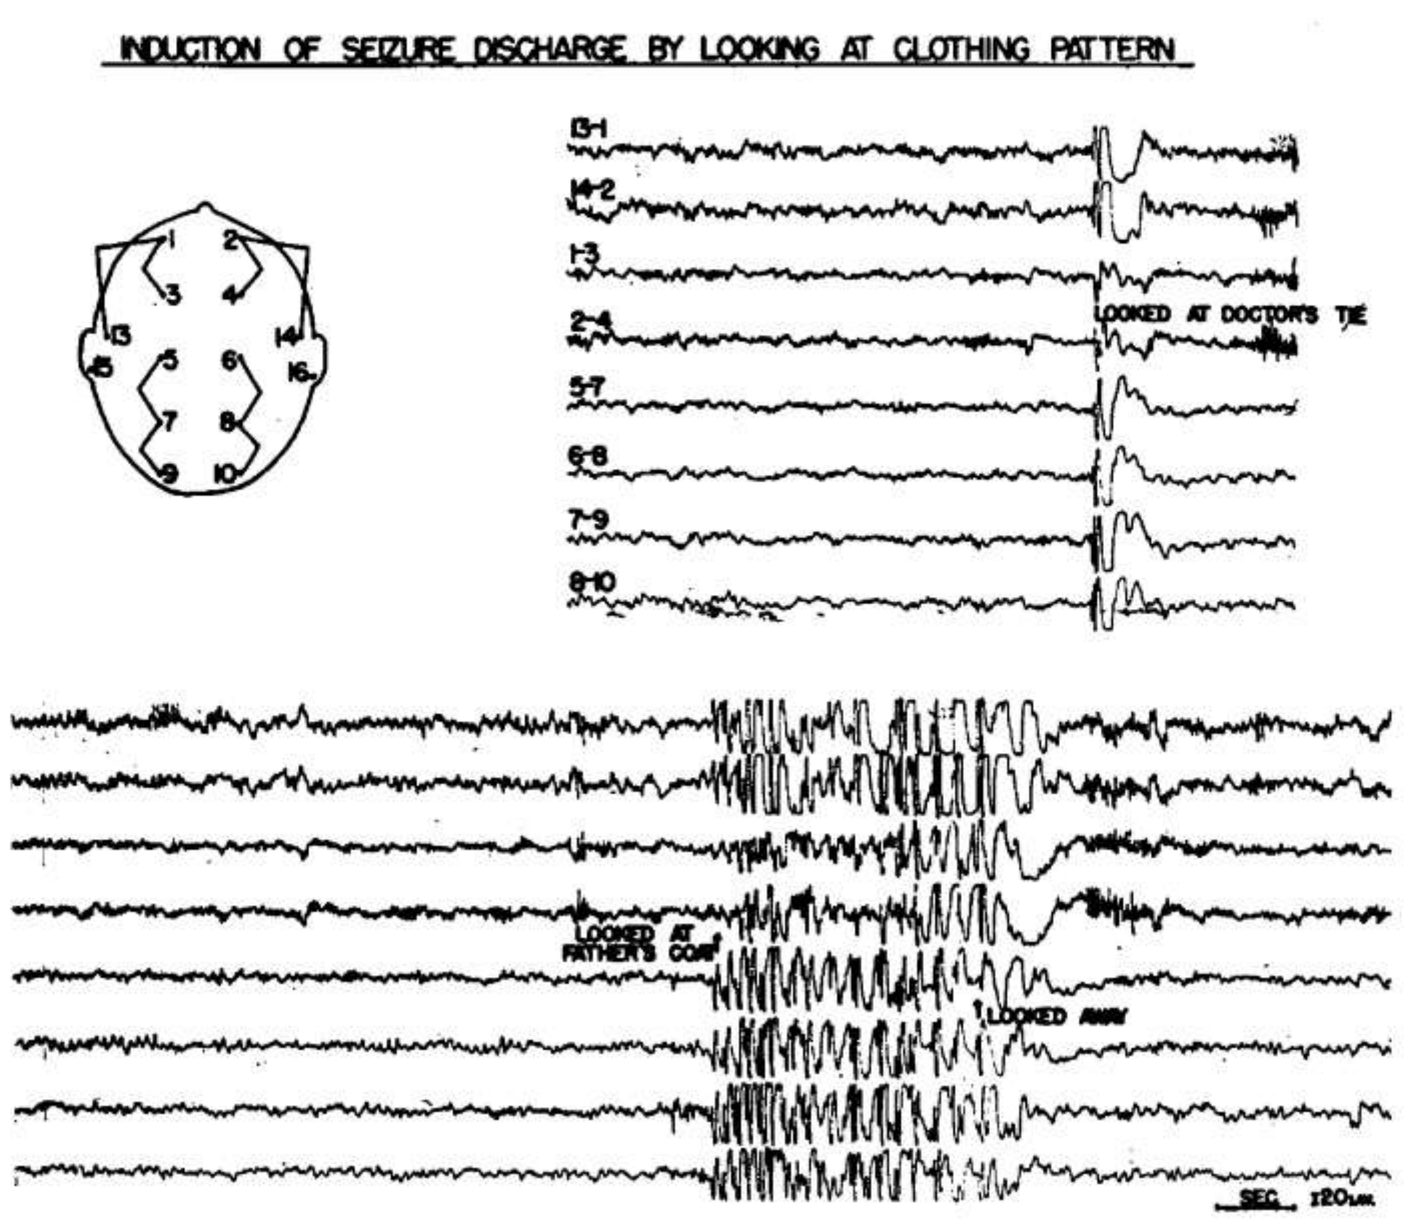

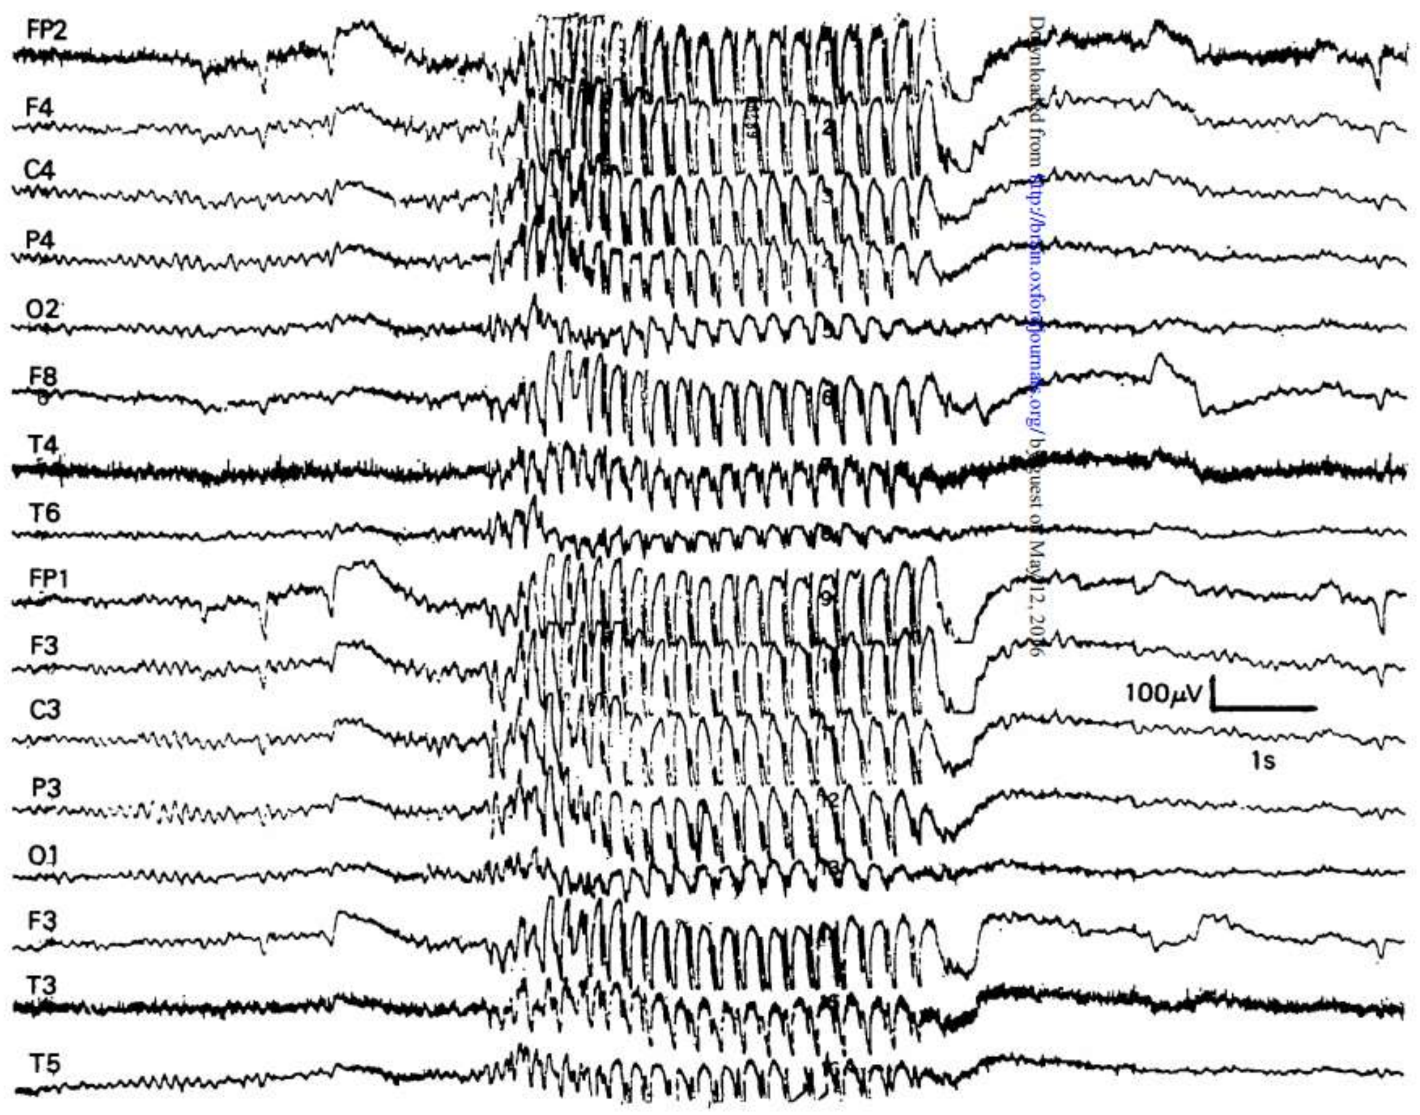

- Large-scale activity suggests a mean-field approach

- Spatial frequencies esp. near 2-4 cpd induce the seizures

- Similar oscillatory activity observed in the case of visual discomfort

[1] Bickford & Keith, 1953
[2] Fernandez & Wilkins, 2008

Neural fields provide a natural starting point

- Large-scale dynamic activity suggests a population-level mean-field approach such as neural fields

$$f_{e,i}(u) = \frac{1}{1 + \exp(-4(u - \theta_{e,i}))}$$

$$J_{\alpha\beta}(x) = a_{\alpha\beta} K_{\beta}(x)$$

$$\frac{\partial u(x,t)}{\partial t} = -u(x,t) + f_e(J_{ee}(x) * u(x,t) - J_{ei}(x) * v(x,t))$$

$$\tau \frac{\partial v(x,t)}{\partial t} = -v(x,t) + f_i(J_{ie}(x) * u(x,t) - J_{ii}(x) * v(x,t))$$

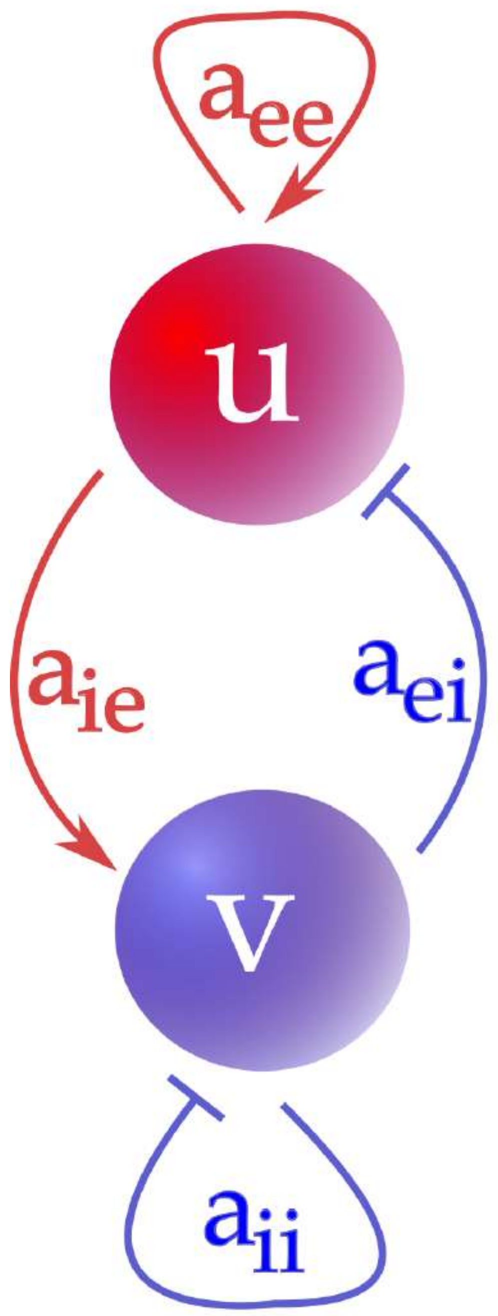

u - excitatory
v - inhibitory

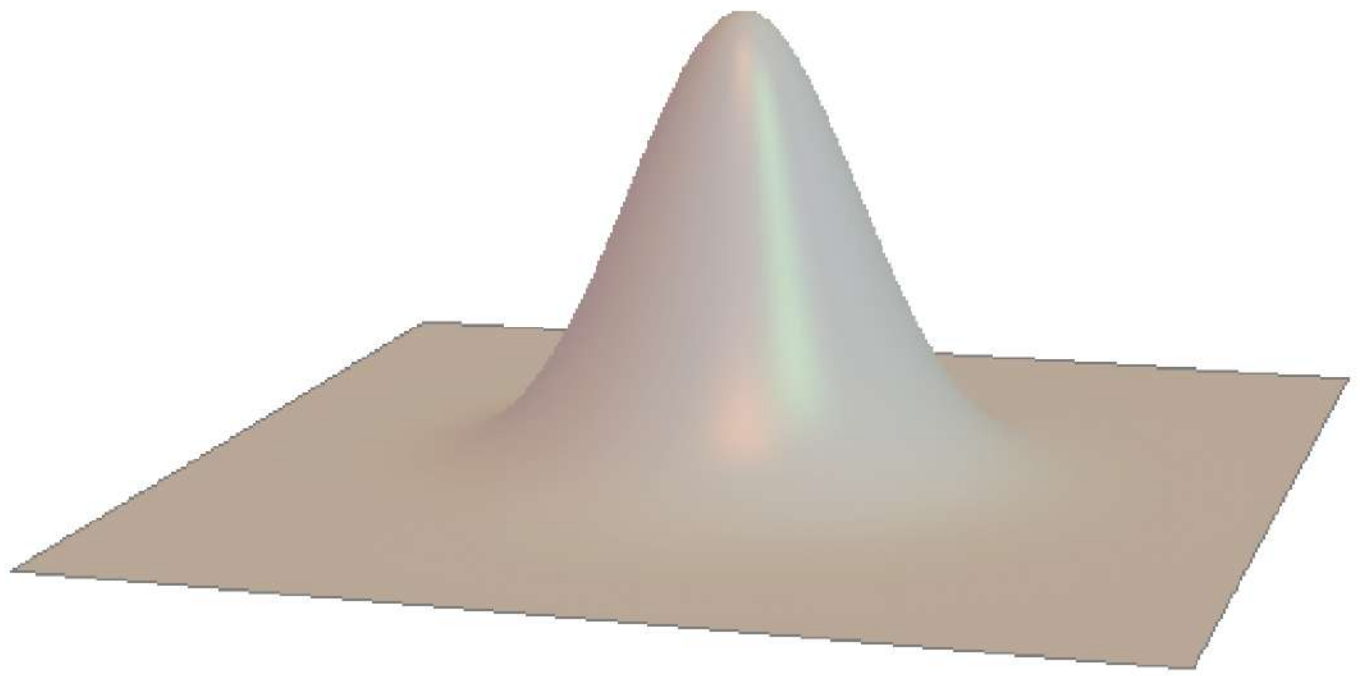

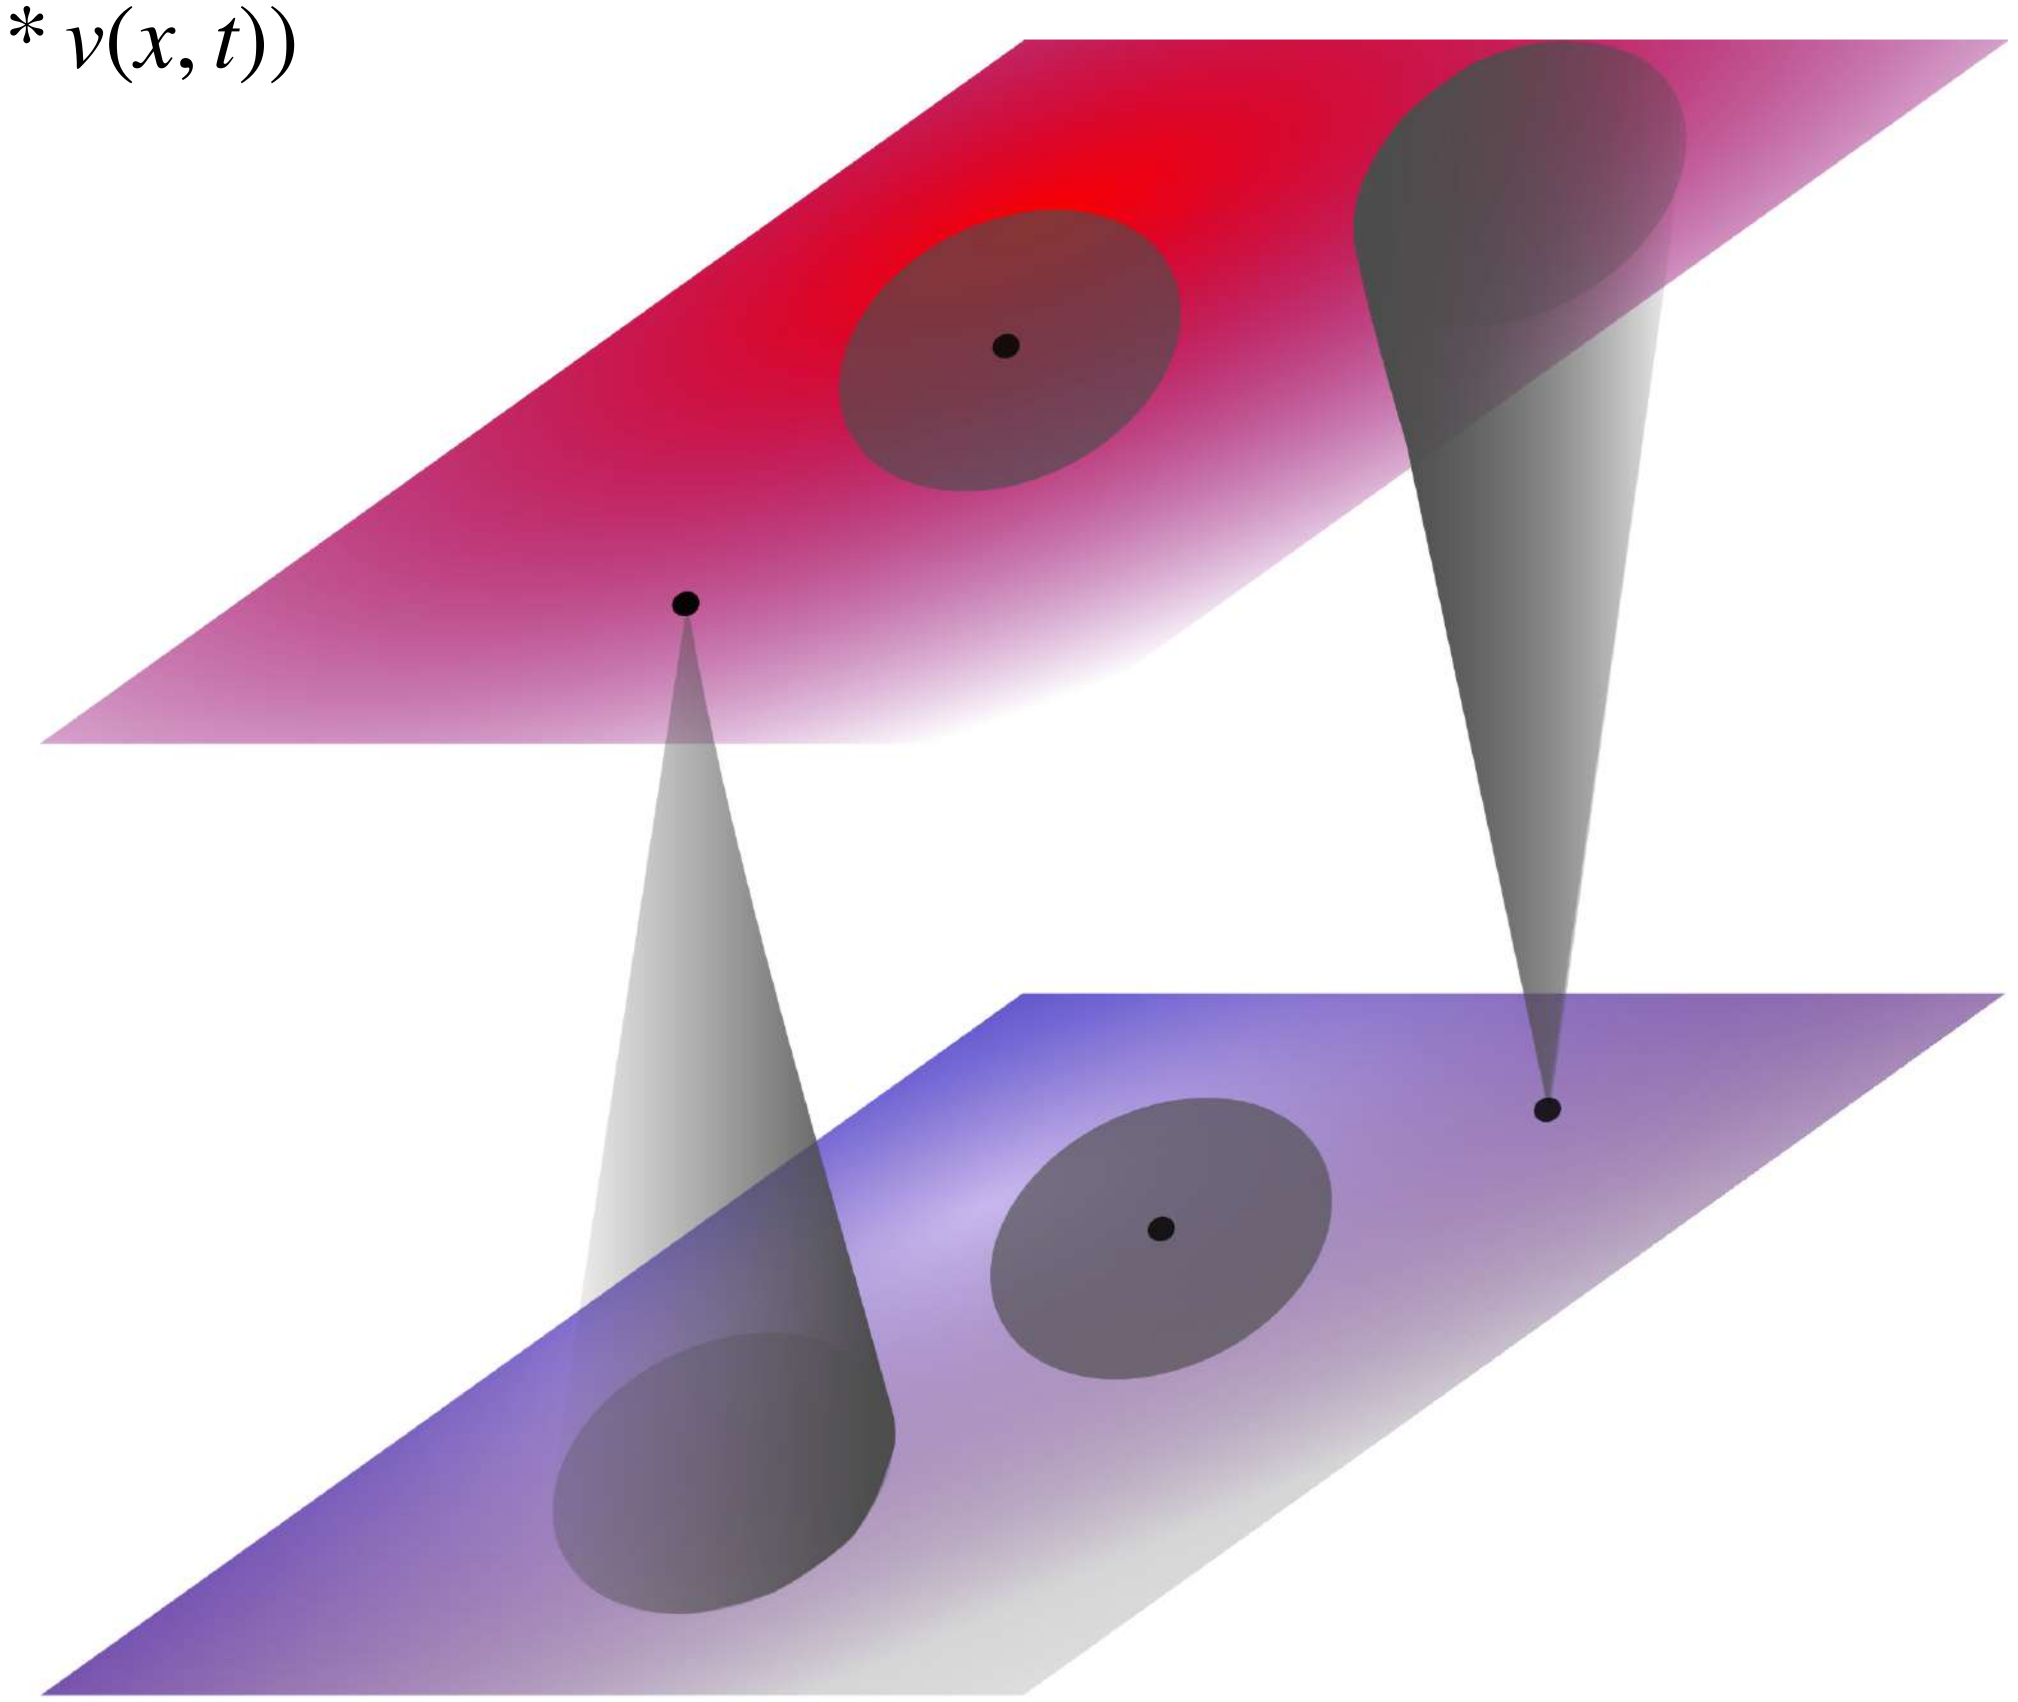

Neural fields provide a natural starting point

- Large-scale dynamic activity suggests a population-level mean-field approach such as neural fields

$$\frac{\partial u(x,t)}{\partial t} = -u(x,t) + f_e(J_{ee}(x) * u(x,t) - J_{ei}(x) * v(x,t) + q\ S(x;k))$$

$$\tau\ \frac{\partial v(x,t)}{\partial t} = -v(x,t) + f_i(J_{ie}(x) * u(x,t) - J_{ii}(x) * v(x,t) + q \cdot r\ S(x;k))$$

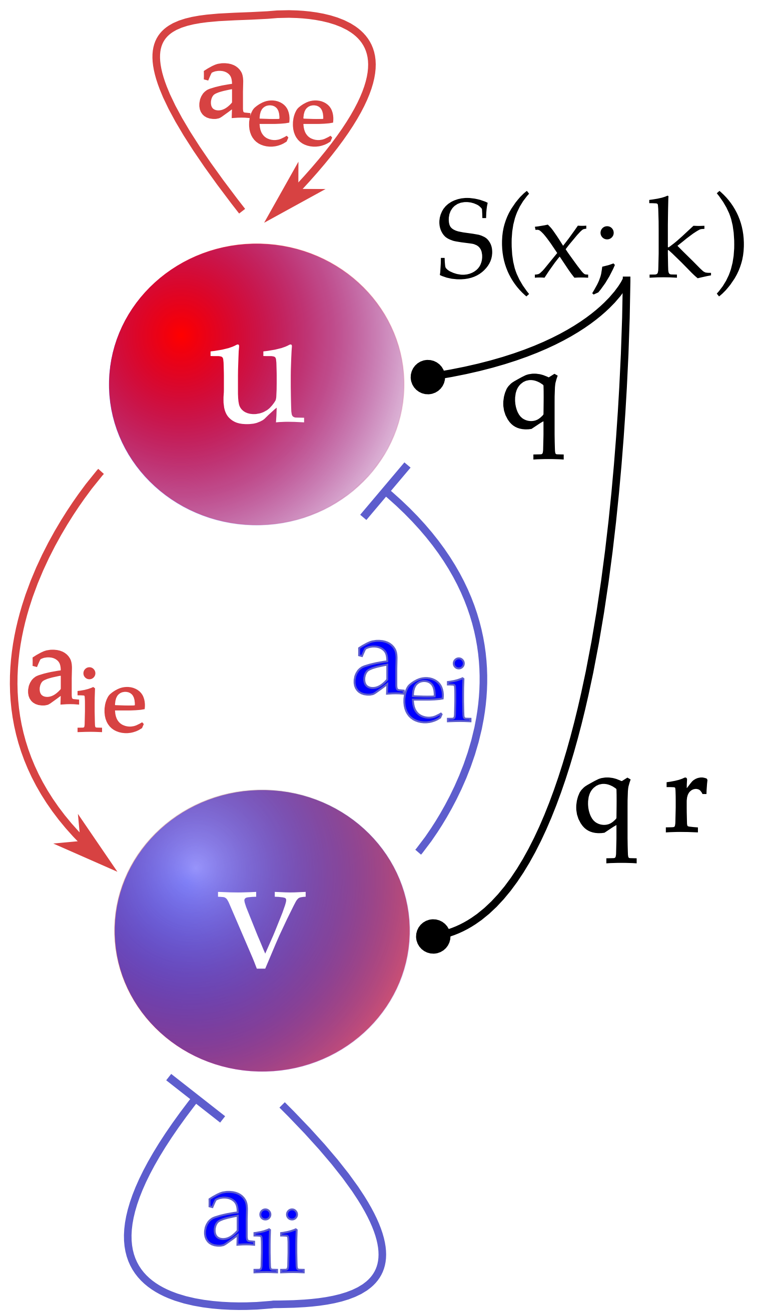

u - excitatory
v - inhibitory

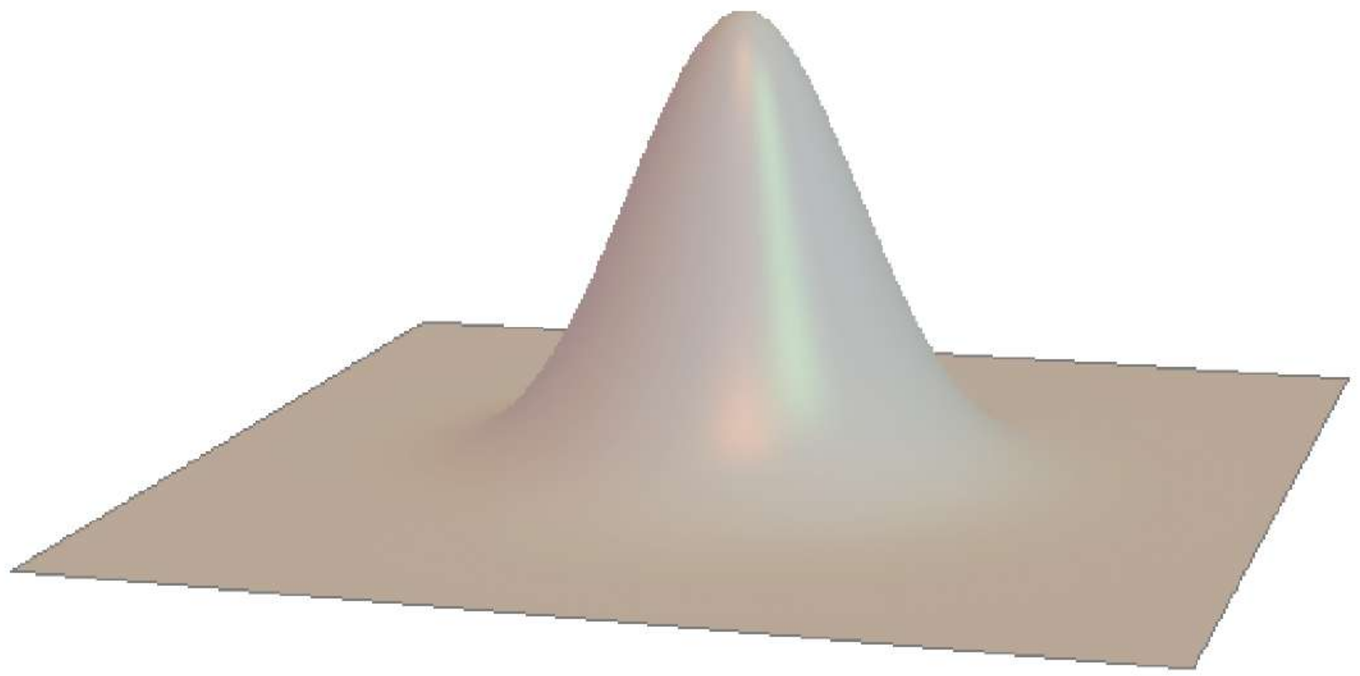

$$S(x;k) = \cos\left(\frac{2\pi kx}{N}\right)$$

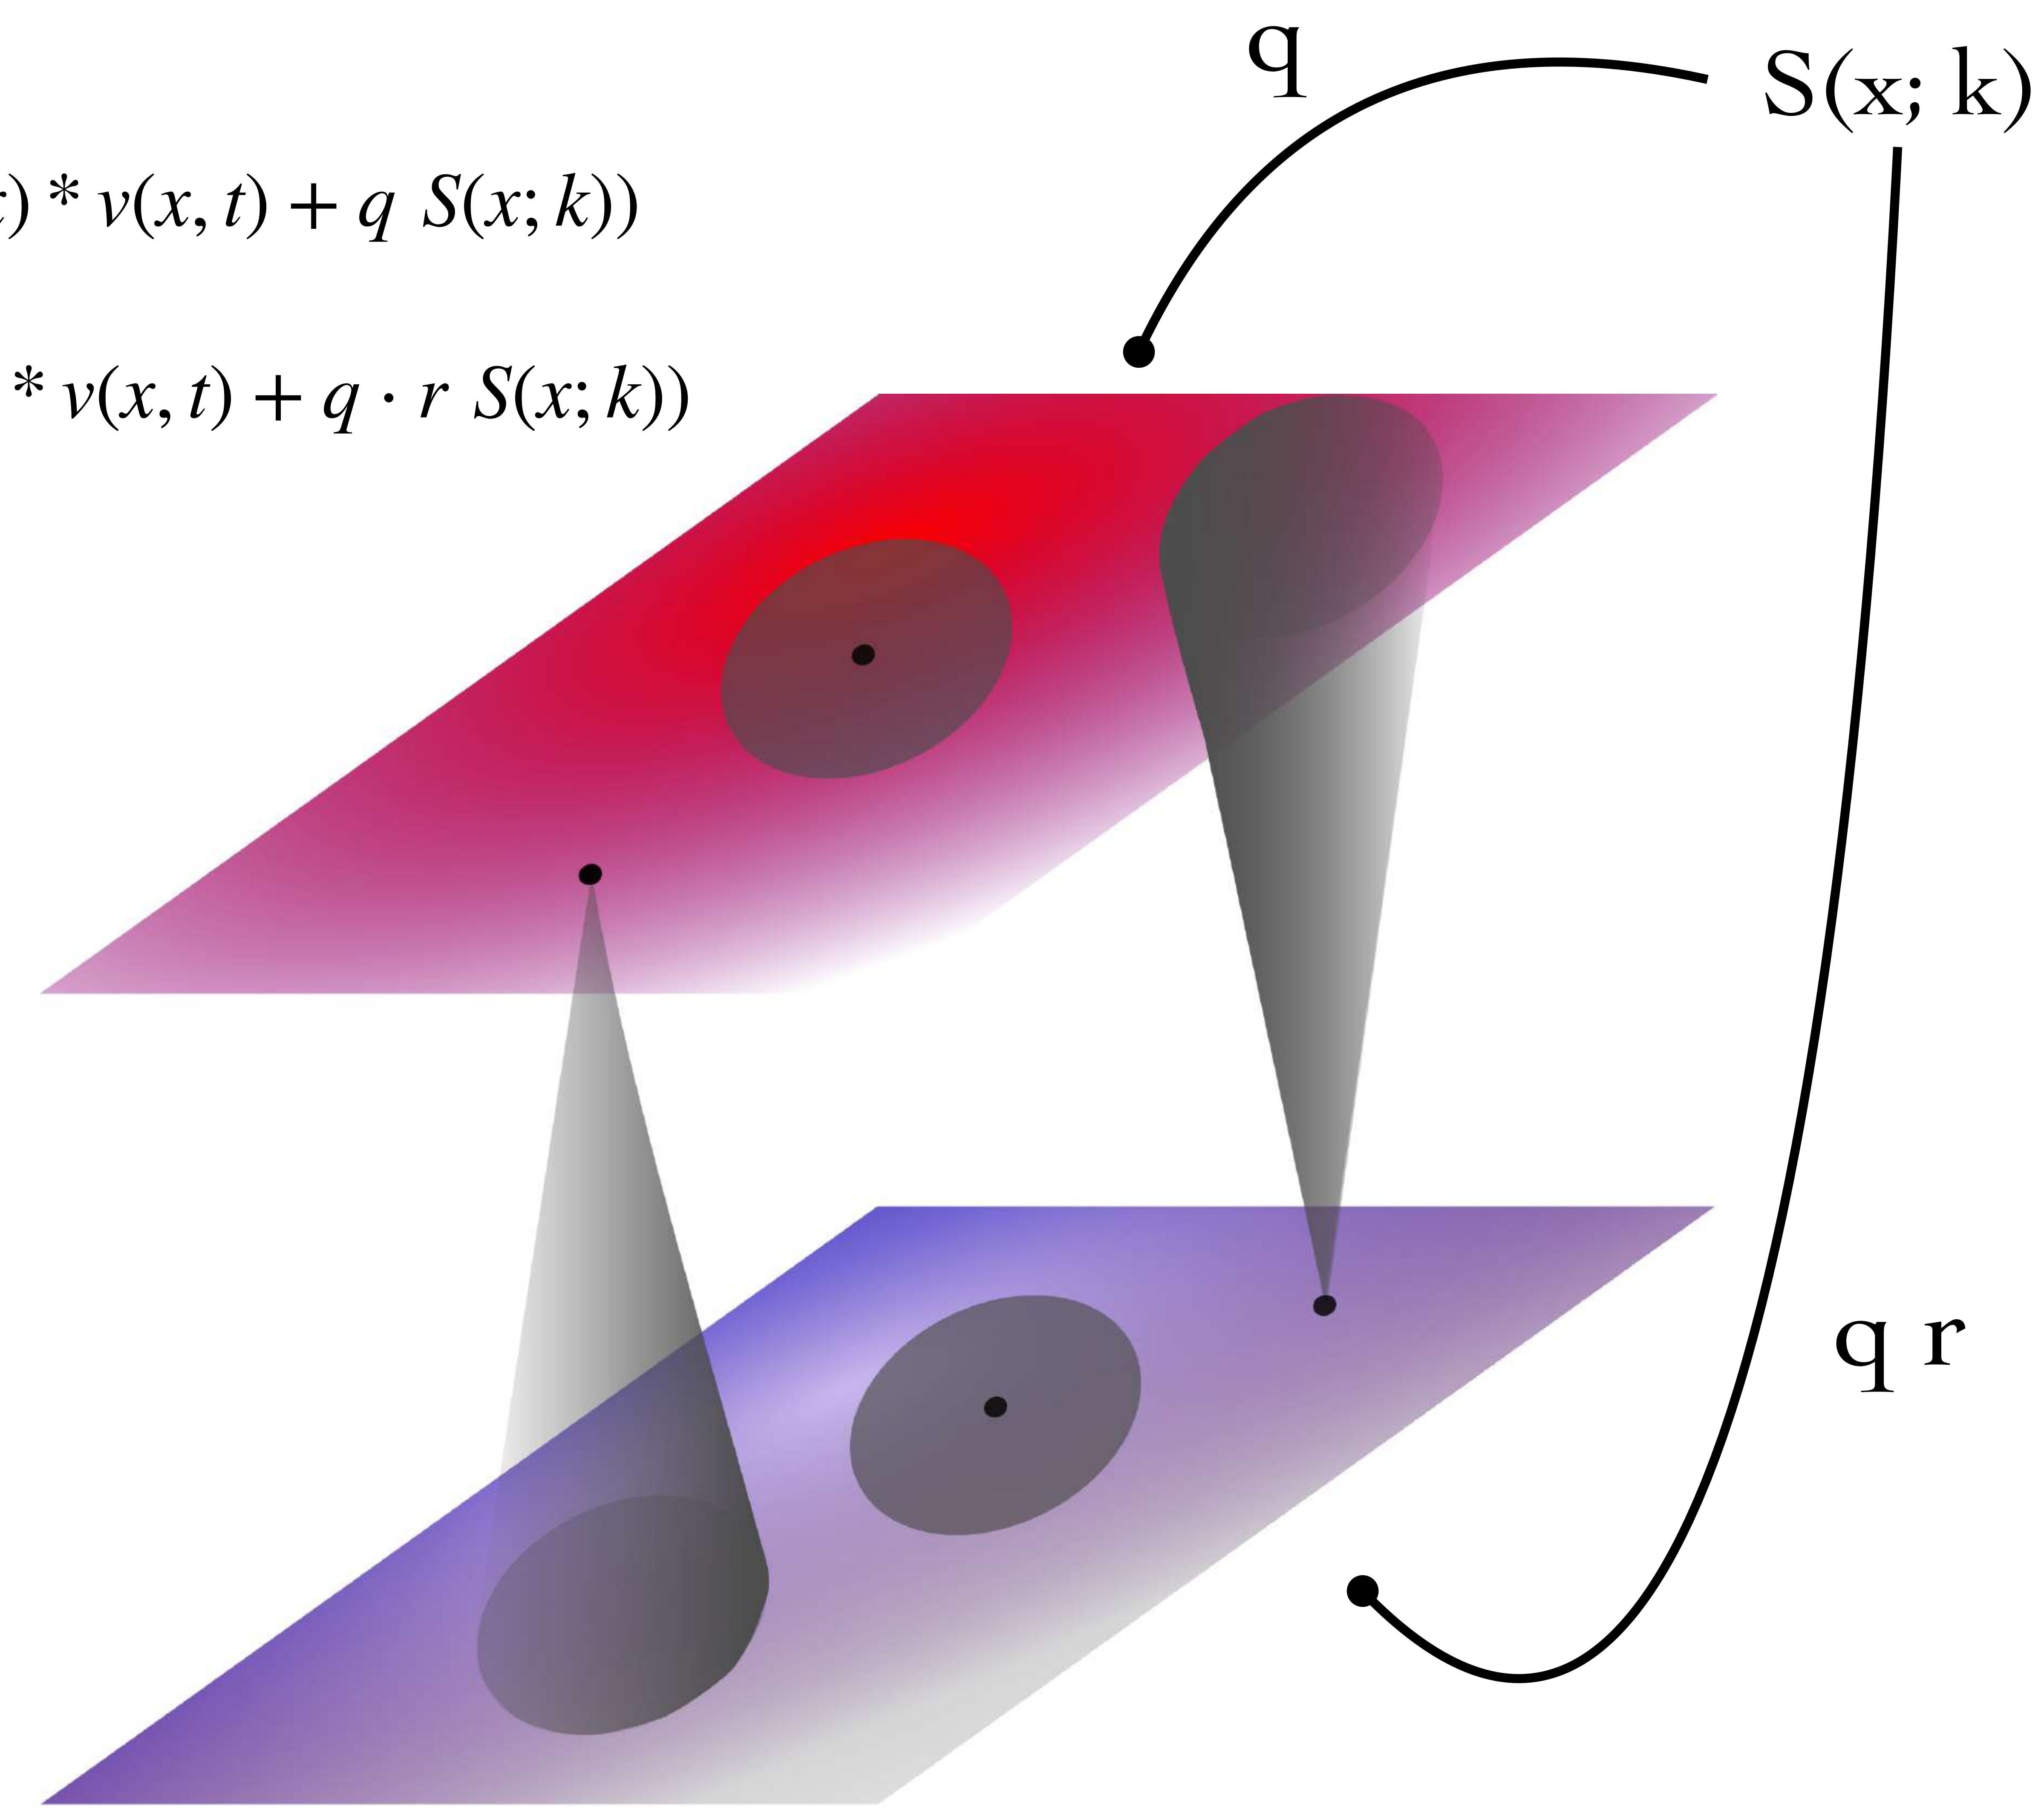

Neural fields provide a natural starting point

- Large-scale dynamic activity suggests a population-level mean-field approach such as neural fields

$$\frac{\partial u(x,t)}{\partial t} = -u(x,t) + f_e(J_{ee}(x) * u(x,t) - J_{ei}(x) * v(x,t) + q\ S(x;k))$$

$$\tau\frac{\partial v(x,t)}{\partial t} = -v(x,t) + f_i(J_{ie}(x) * u(x,t) - J_{ii}(x) * v(x,t) + q \cdot r\ S(x;k))$$

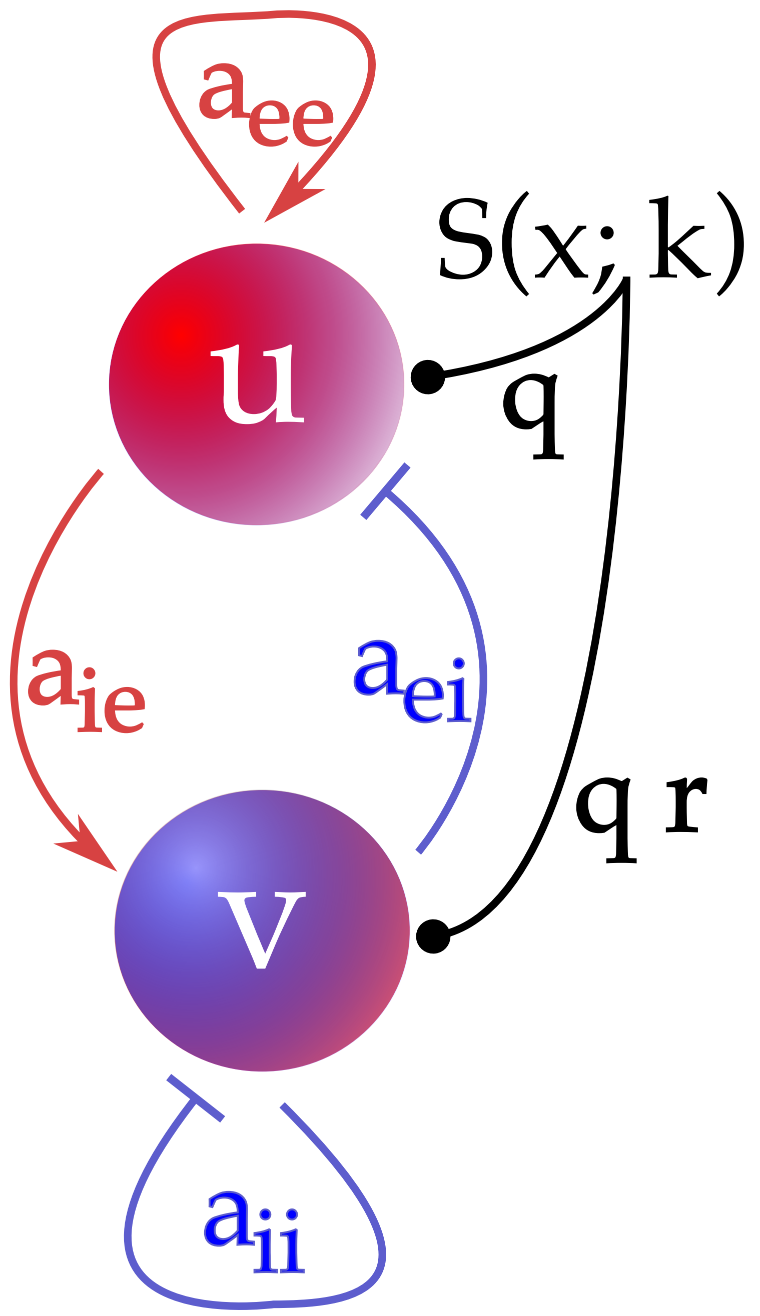

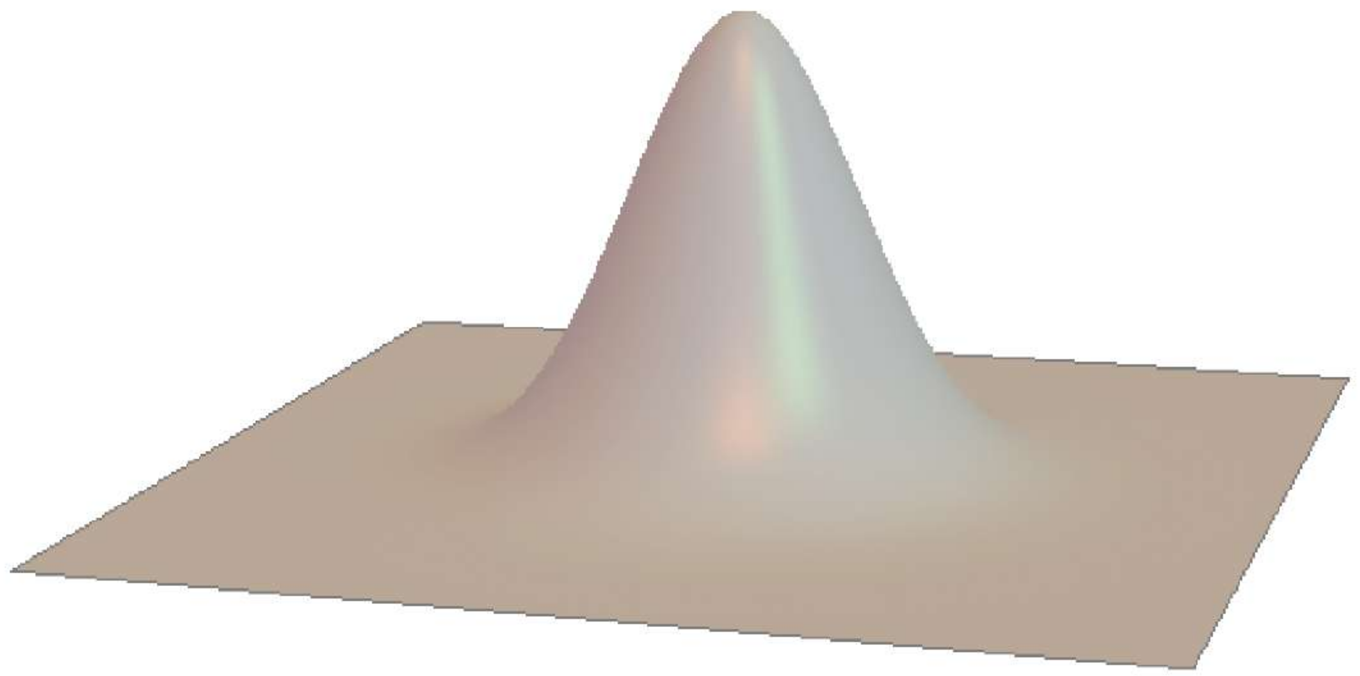

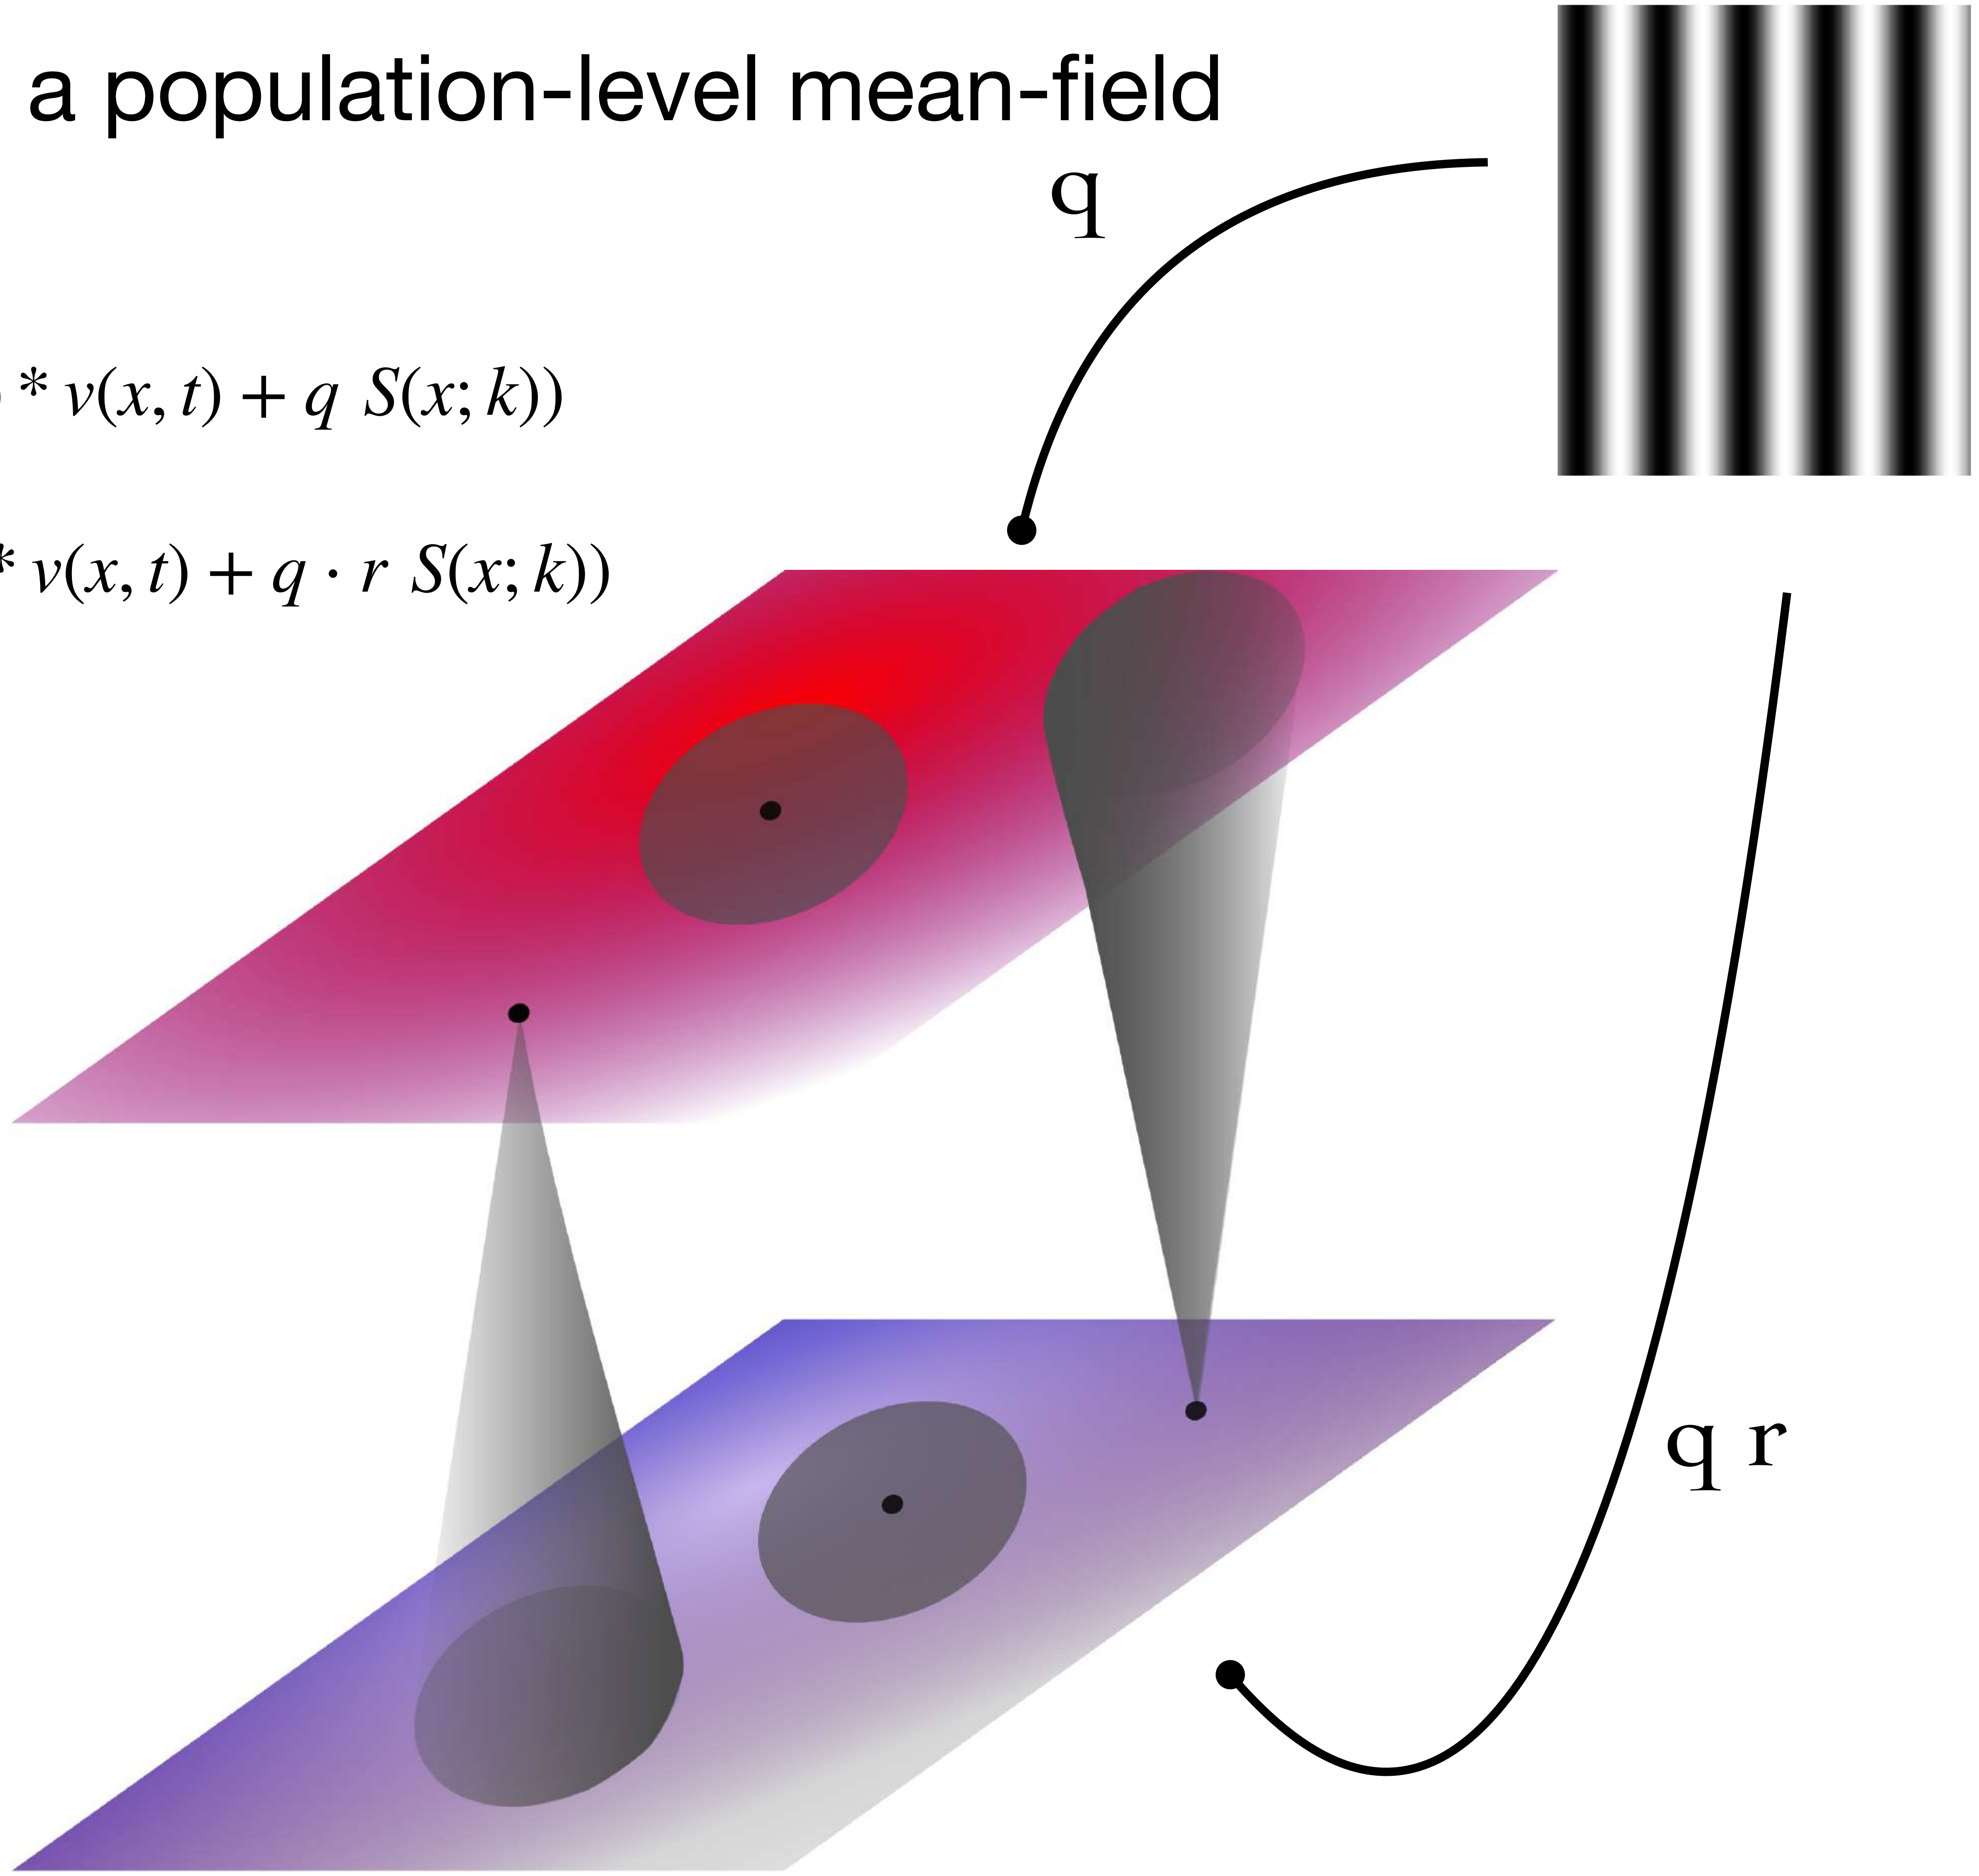

u - excitatory
v - inhibitory

Resonant oscillations suggest Turing-Hopf bifurcation

- Hopf bifurcation: changing a parameter results in the appearance of oscillations

- By adjusting the spatial profiles of the Gaussian kernel, the steady state of the system can be lost to oscillations with at a nonzero wavenumber, *m**

  (Turing-Hopf bifurcation)

- Then, *presumably*, the system will be more sensitive to stimuli with those wavenumbers

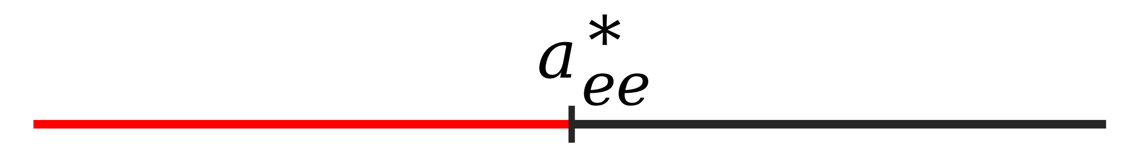

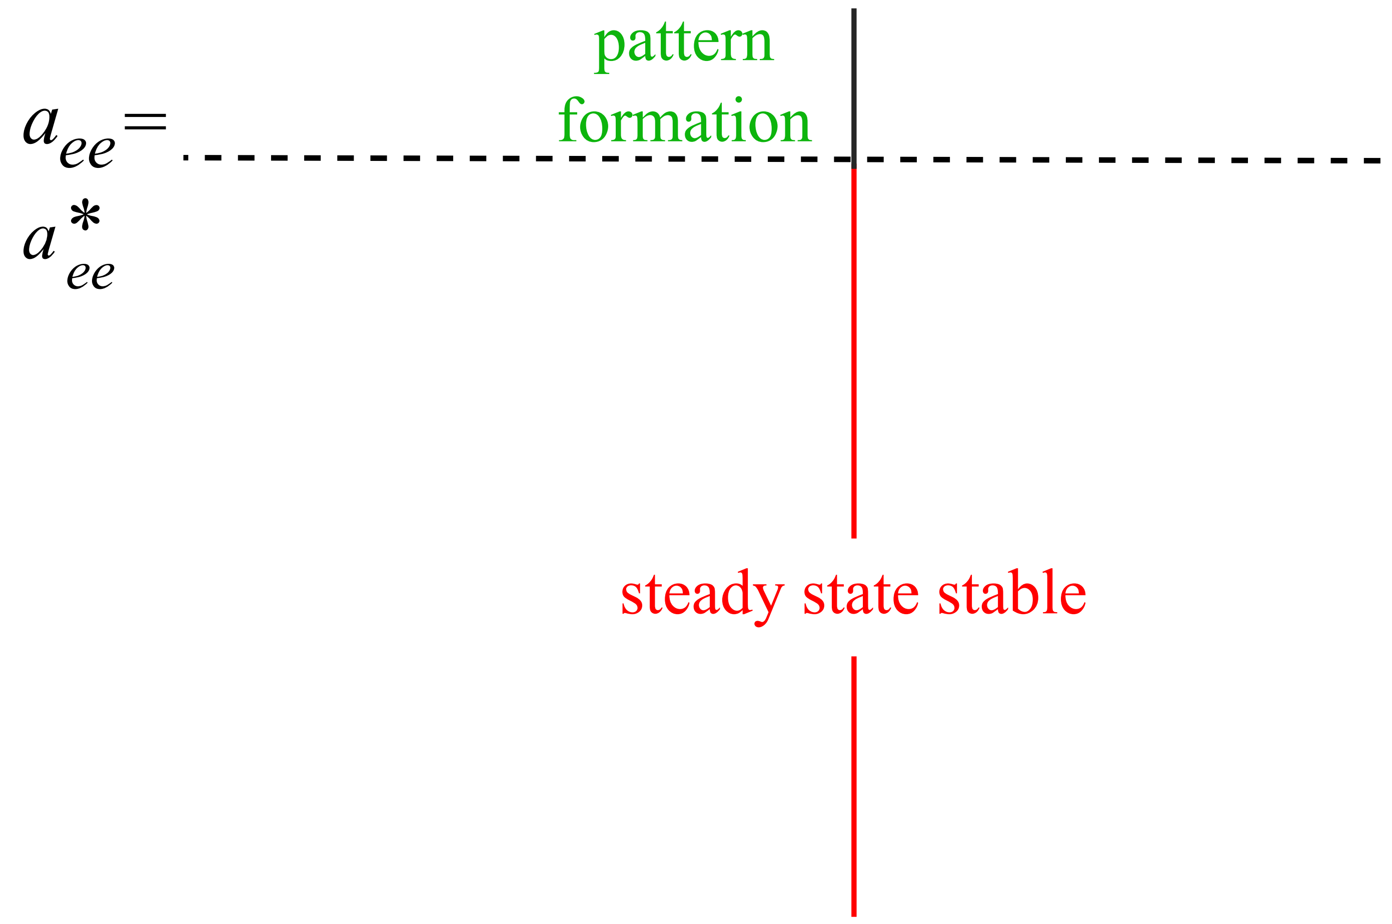

Resonant oscillations suggest Turing-Hopf bifurcation

- Hopf bifurcation: changing a parameter results in the appearance of oscillations
- By adjusting the spatial profiles of the Gaussian kernel, the steady state of the system can be lost to oscillations with at a nonzero wavenumber, m^* (Turing-Hopf bifurcation)
- Then, *presumably*, the system will be more sensitive to stimuli with those wavenumbers

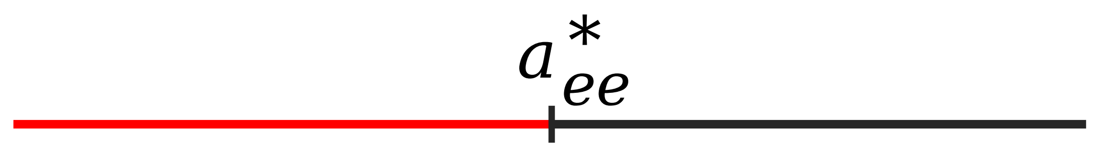

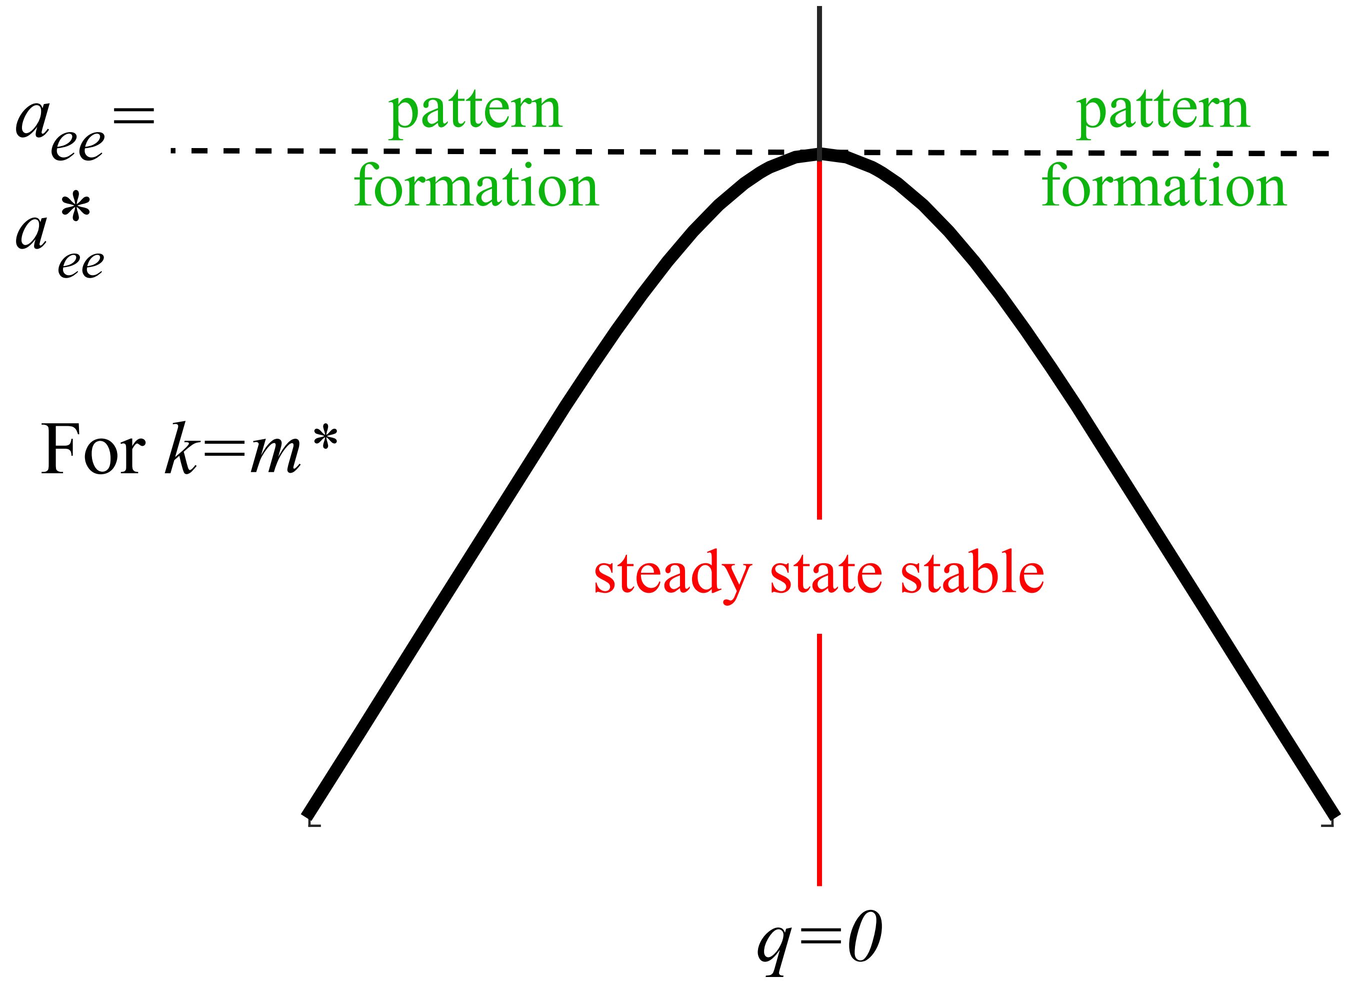

Resonant oscillations suggest Turing-Hopf bifurcation

- Linearize system, look for solutions that are periodic in space and time ($u, v \sim e^{i\mu t} e^{imx}$)

- End up with simple 2x2 linear system that will be a function of the wavenumber m

- Find when the eigenvalue is purely imaginary only at a nonzero wavenumber $m*$

- Since eigenvalues are given by $\lambda = T \pm \sqrt{T^2 - 4D}$ ($T =$ Trace, $D =$ Det), sufficient if
    - $T = 0$ at $m = m*$ and negative elsewhere
    - $D > 0$ everywhere
    - Then $\lambda = i\mu$ at $m*$

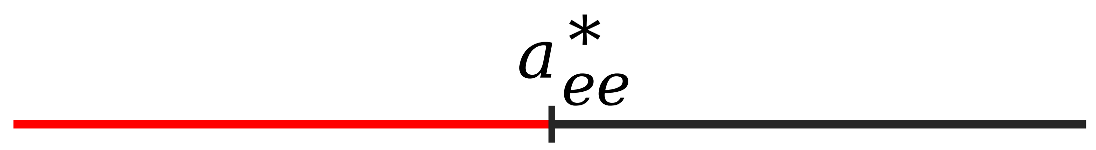

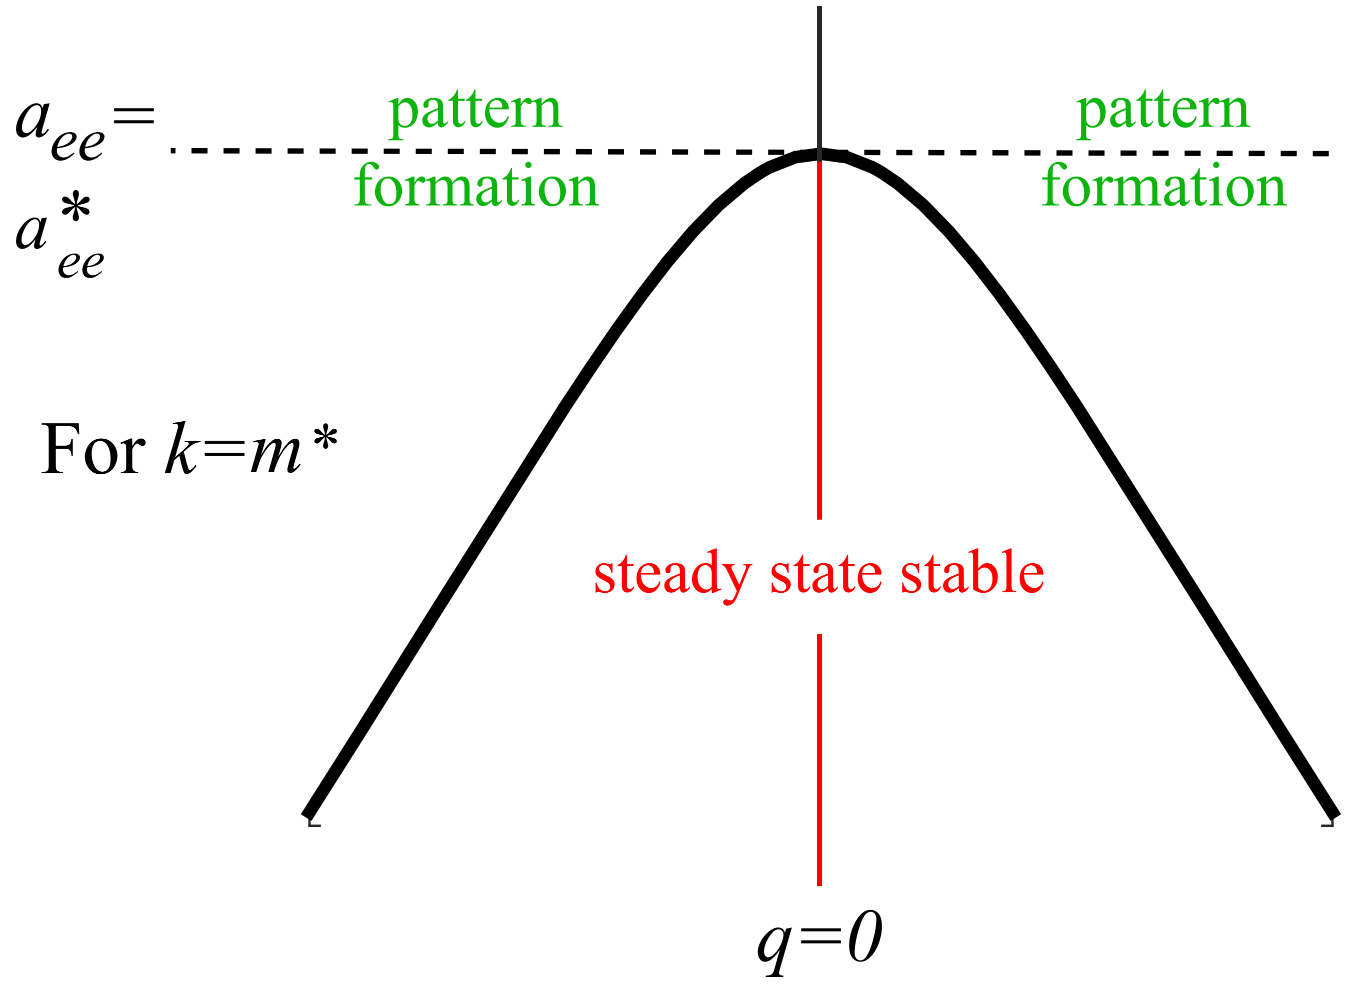

Resonant oscillations suggest Turing-Hopf bifurcation

- Linearize system, look for solutions that are periodic in space and time ($u, v \sim e^{i\mu t}e^{imx}$)
- End up with simple 2x2 linear system that will be a function of the wavenumber m
- Find when the eigenvalue is purely imaginary only at a nonzero wavenumber $m*$
- Since eigenvalues are given by $\lambda = T \pm \sqrt{T^2 - 4D}$ ($T = \text{Trace}, D = \text{Det}$), sufficient if
  - $T = 0$ at $m = m*$ and negative elsewhere
  - $D > 0$ everywhere
  - Then $\lambda = i\mu$ at $m*$

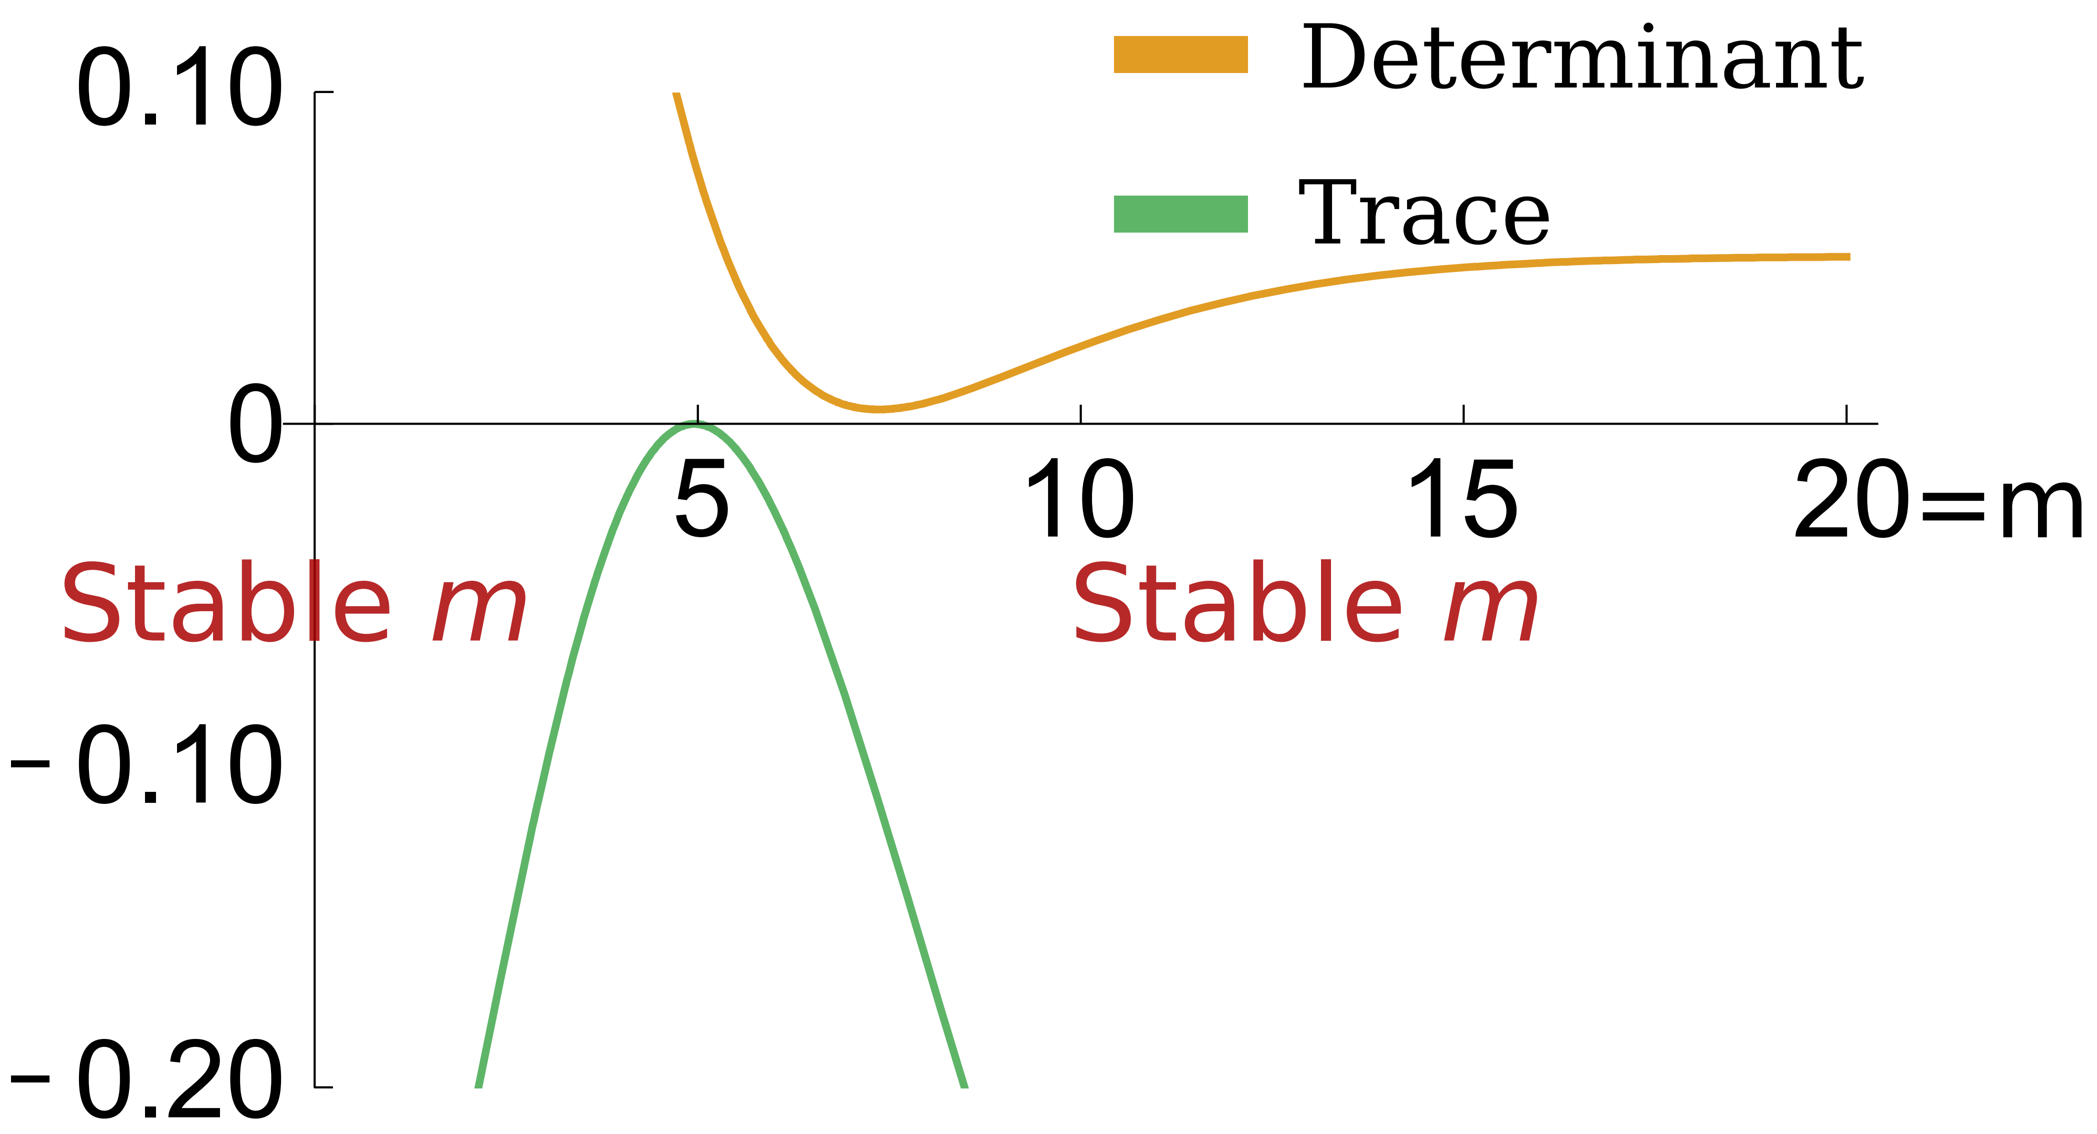

Resonant oscillations suggest Turing-Hopf bifurcation

- Linearize system, look for solutions that are periodic in space and time ($u, v \sim e^{i\mu t} e^{imx}$)
- End up with simple 2x2 linear system that will be a function of the wavenumber m
- Find when the eigenvalue is purely imaginary only at a nonzero wavenumber $m*$
- Since eigenvalues are given by $\lambda = T \pm \sqrt{T^2 - 4D}$ ($T = \text{Trace}, D = \text{Det}$), sufficient if
  - $T = 0$ at $m = m*$ and negative elsewhere
  - $D > 0$ everywhere
  - Then $\lambda = i\mu$ at $m*$

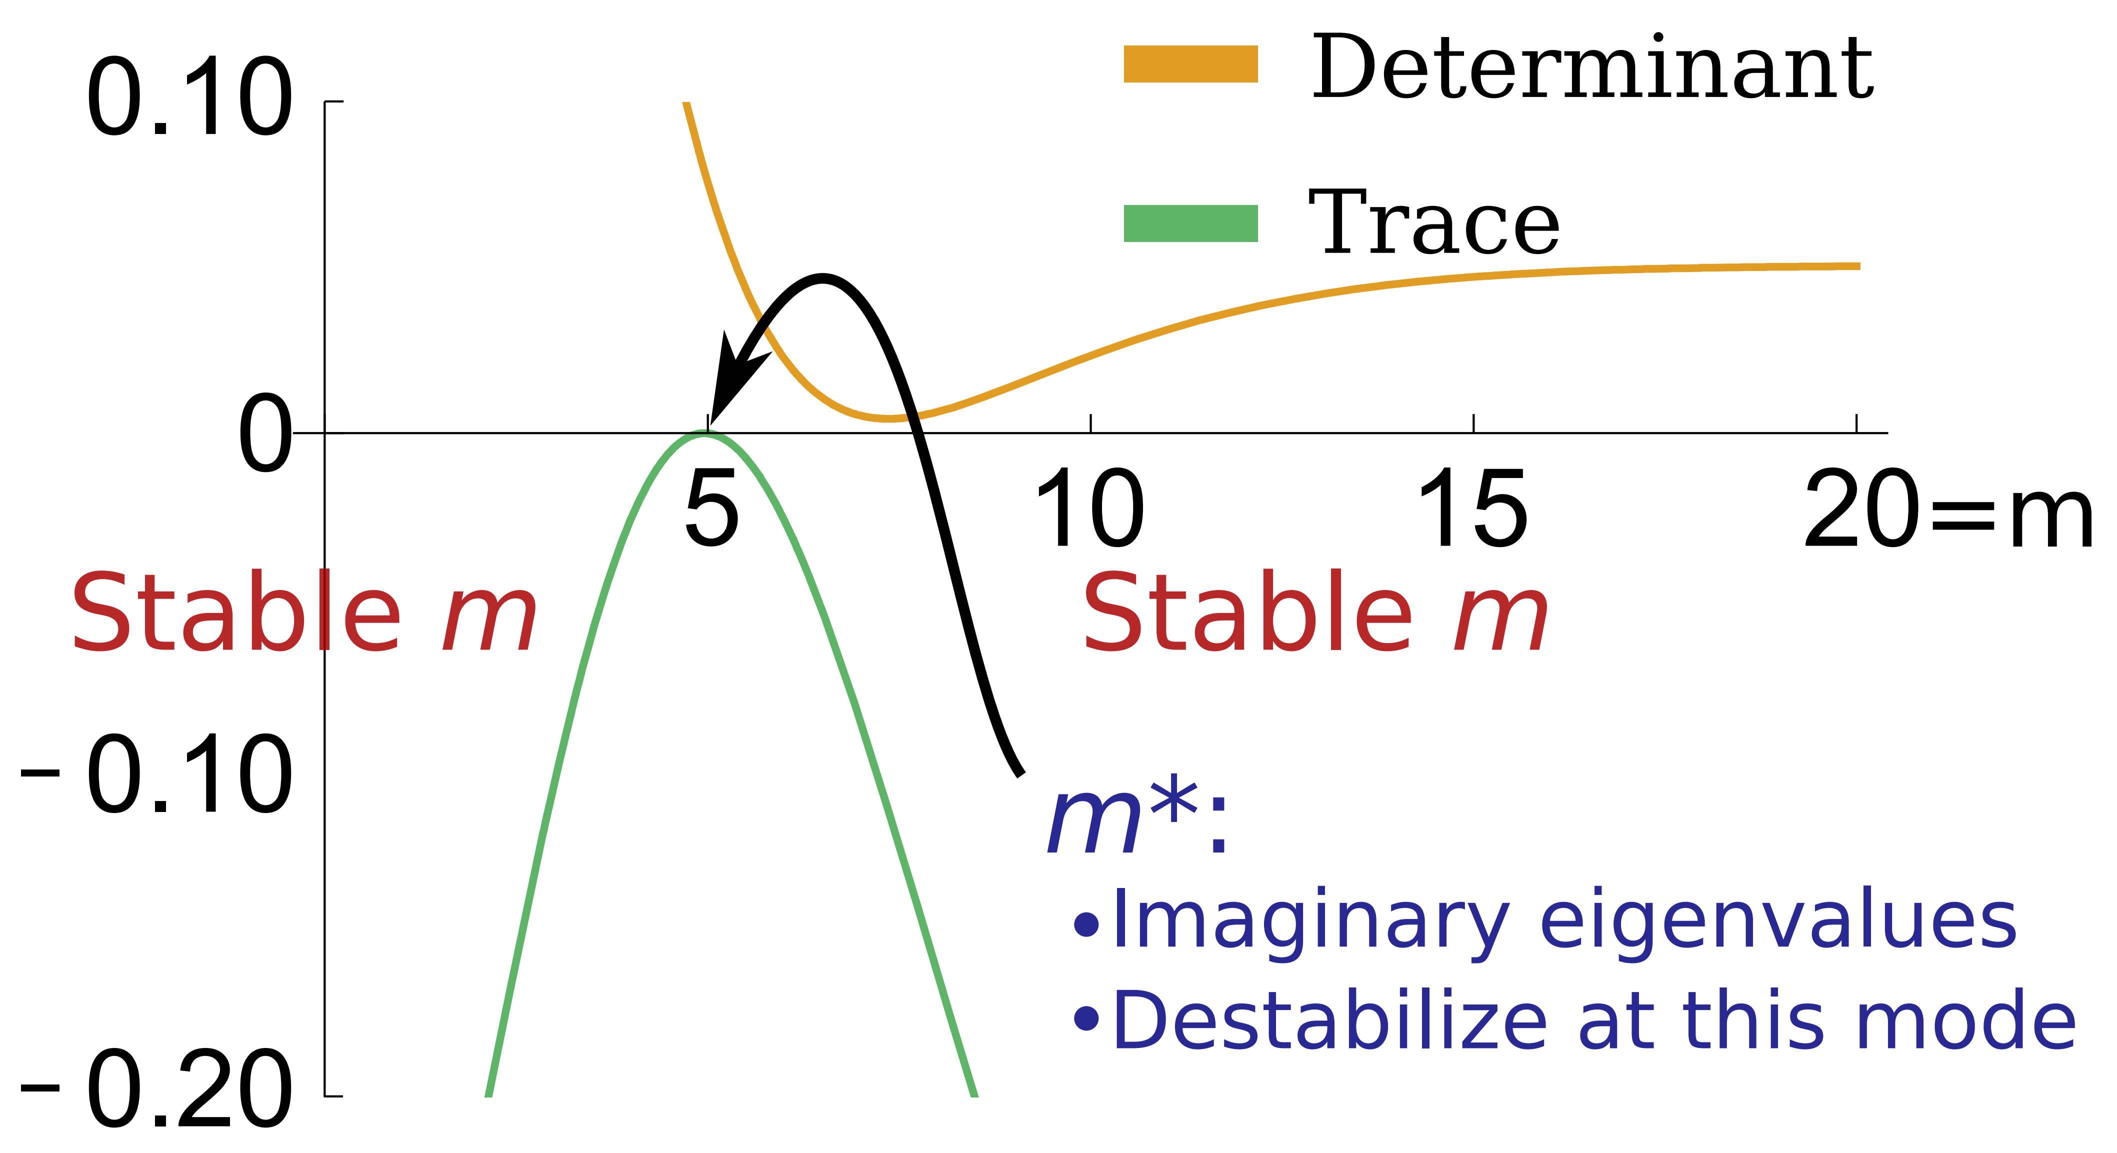

Poor-person's bifurcation diagram: simulate over a q, k range and probe variance

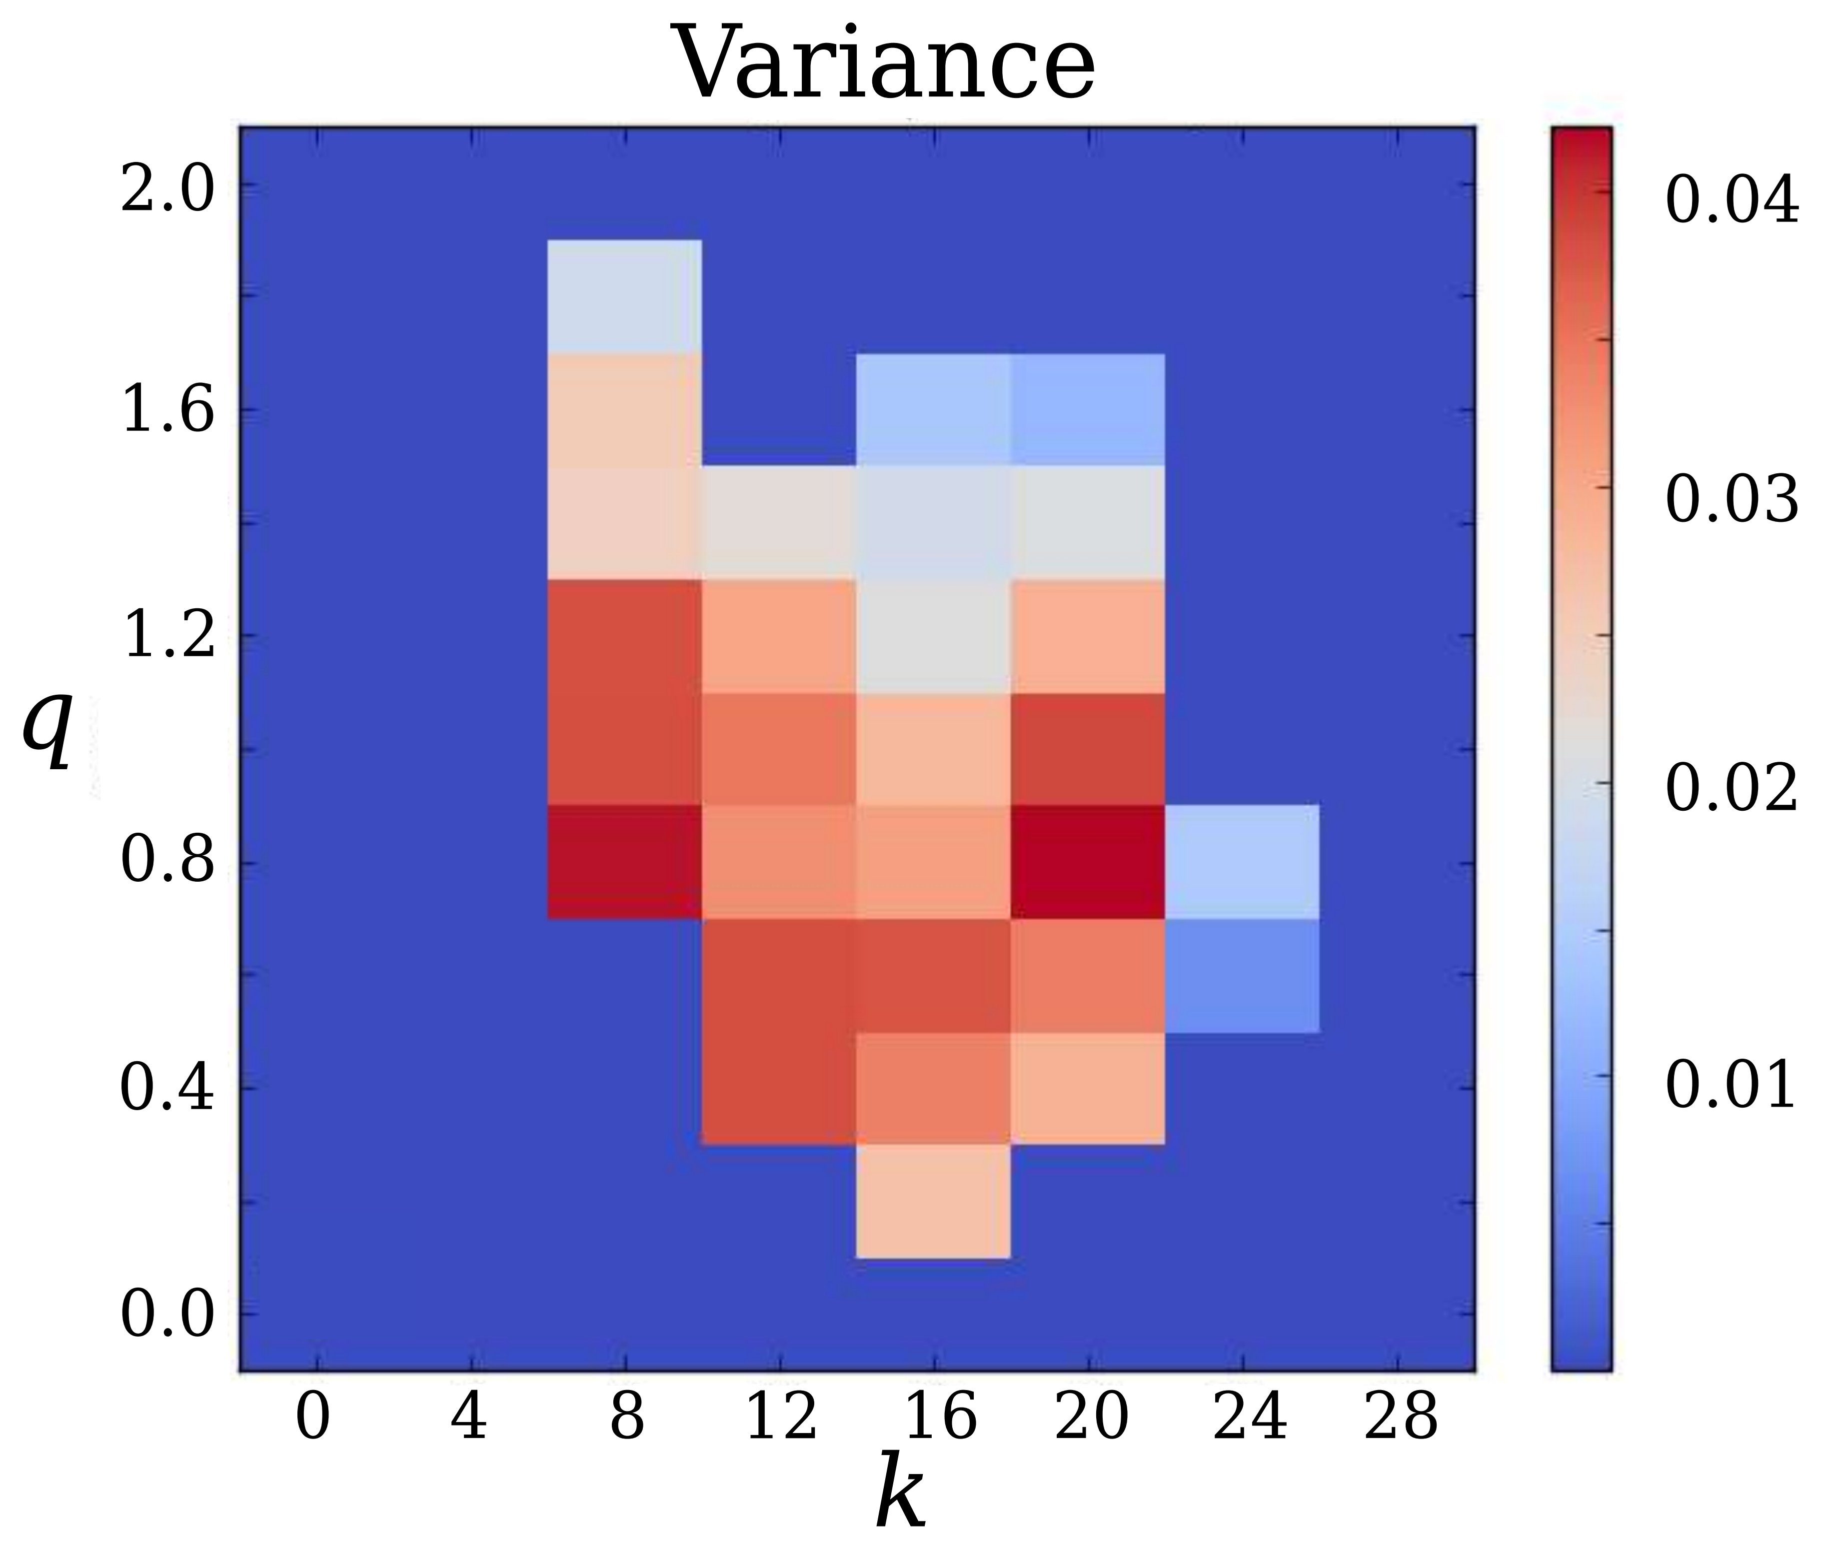

Poor-person's bifurcation diagram: simulate over a q, k range and probe variance

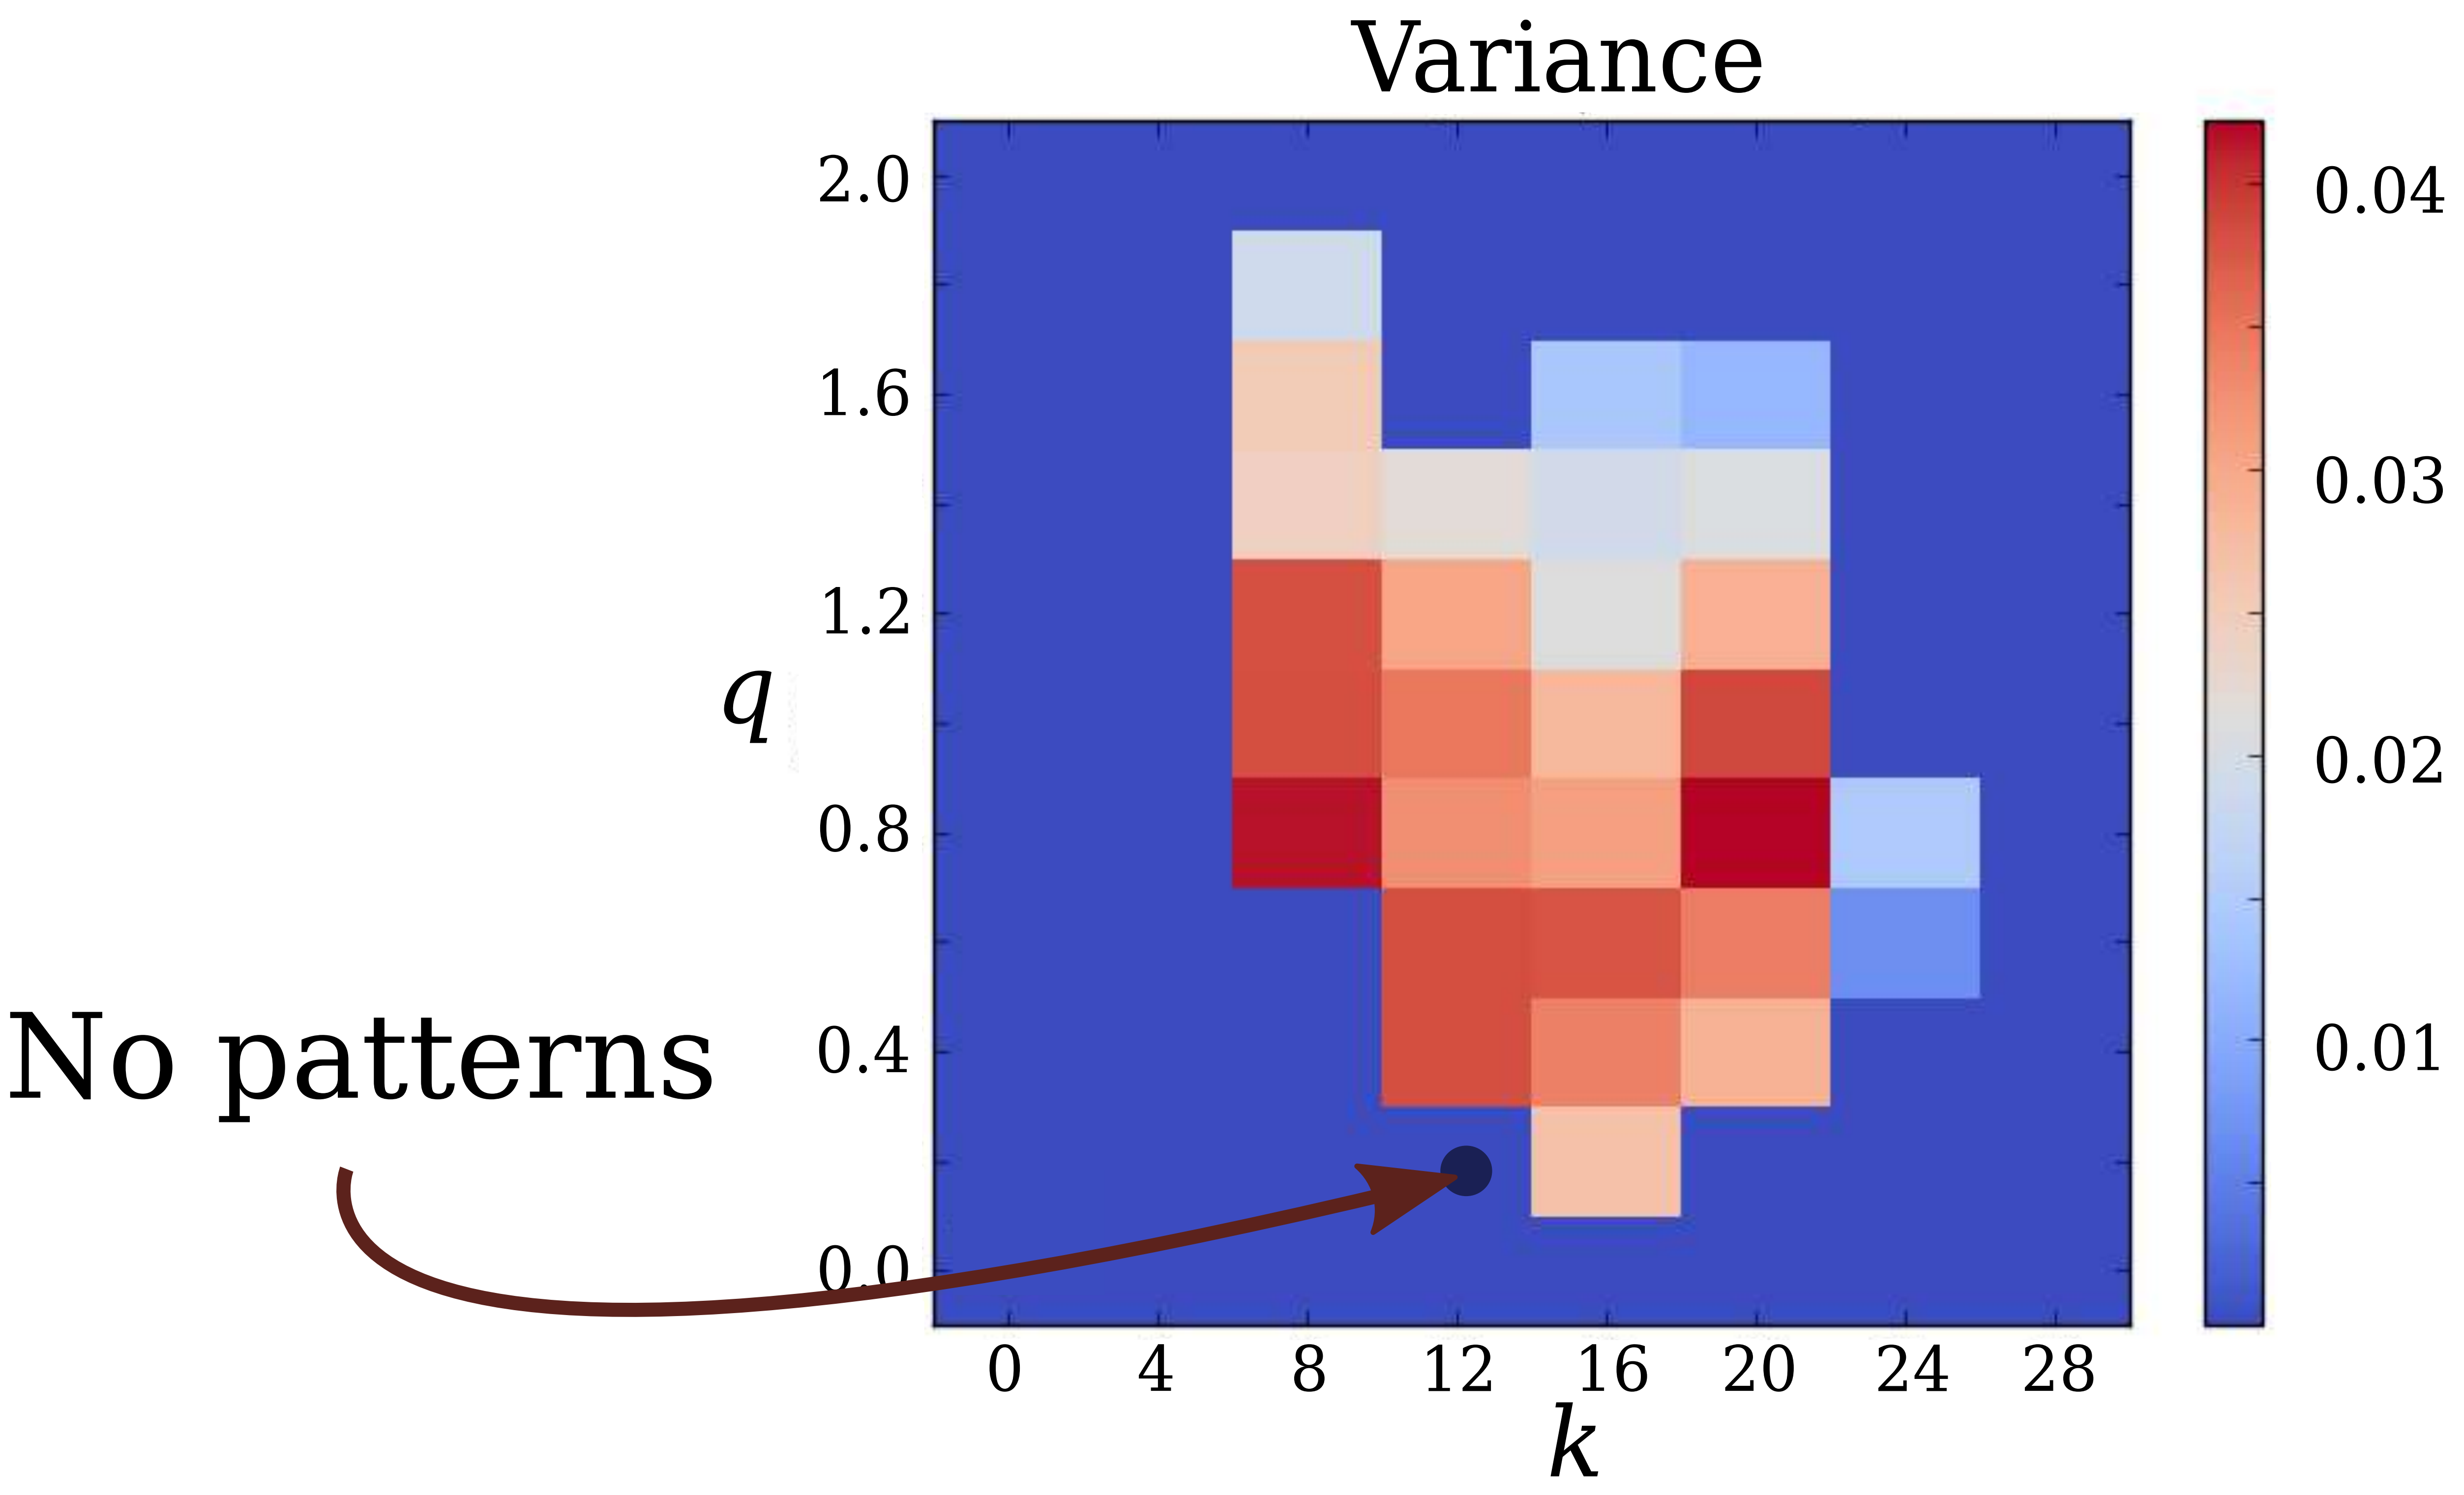

Poor-person's bifurcation diagram: simulate over a q, k range and probe variance

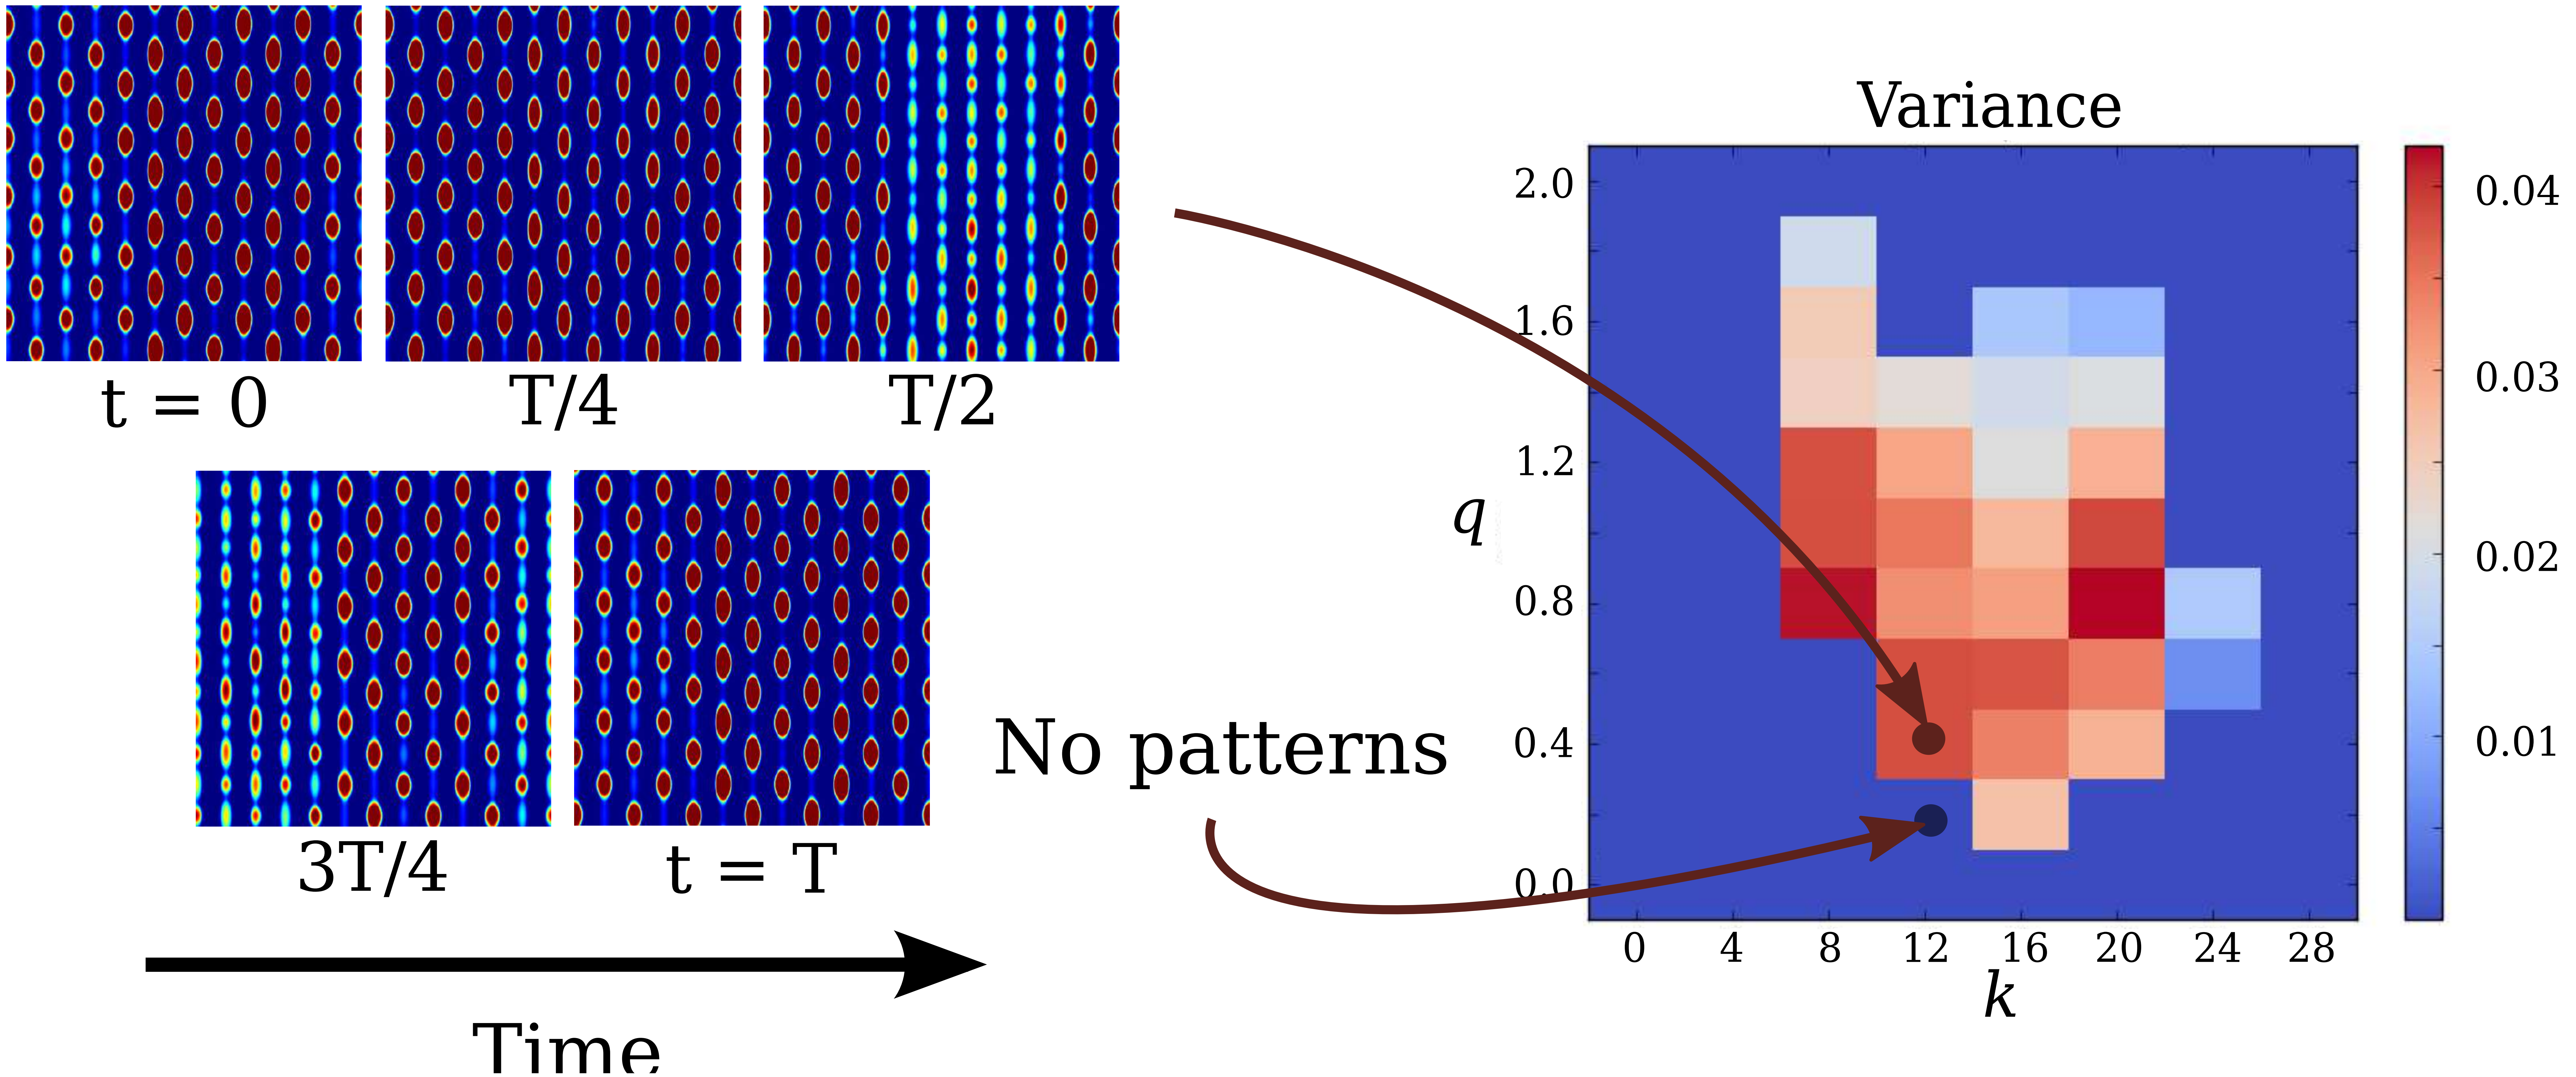

Inverting lower boundary to find sensitivity of network to different wavenumbers results in similar resonance as in experiments

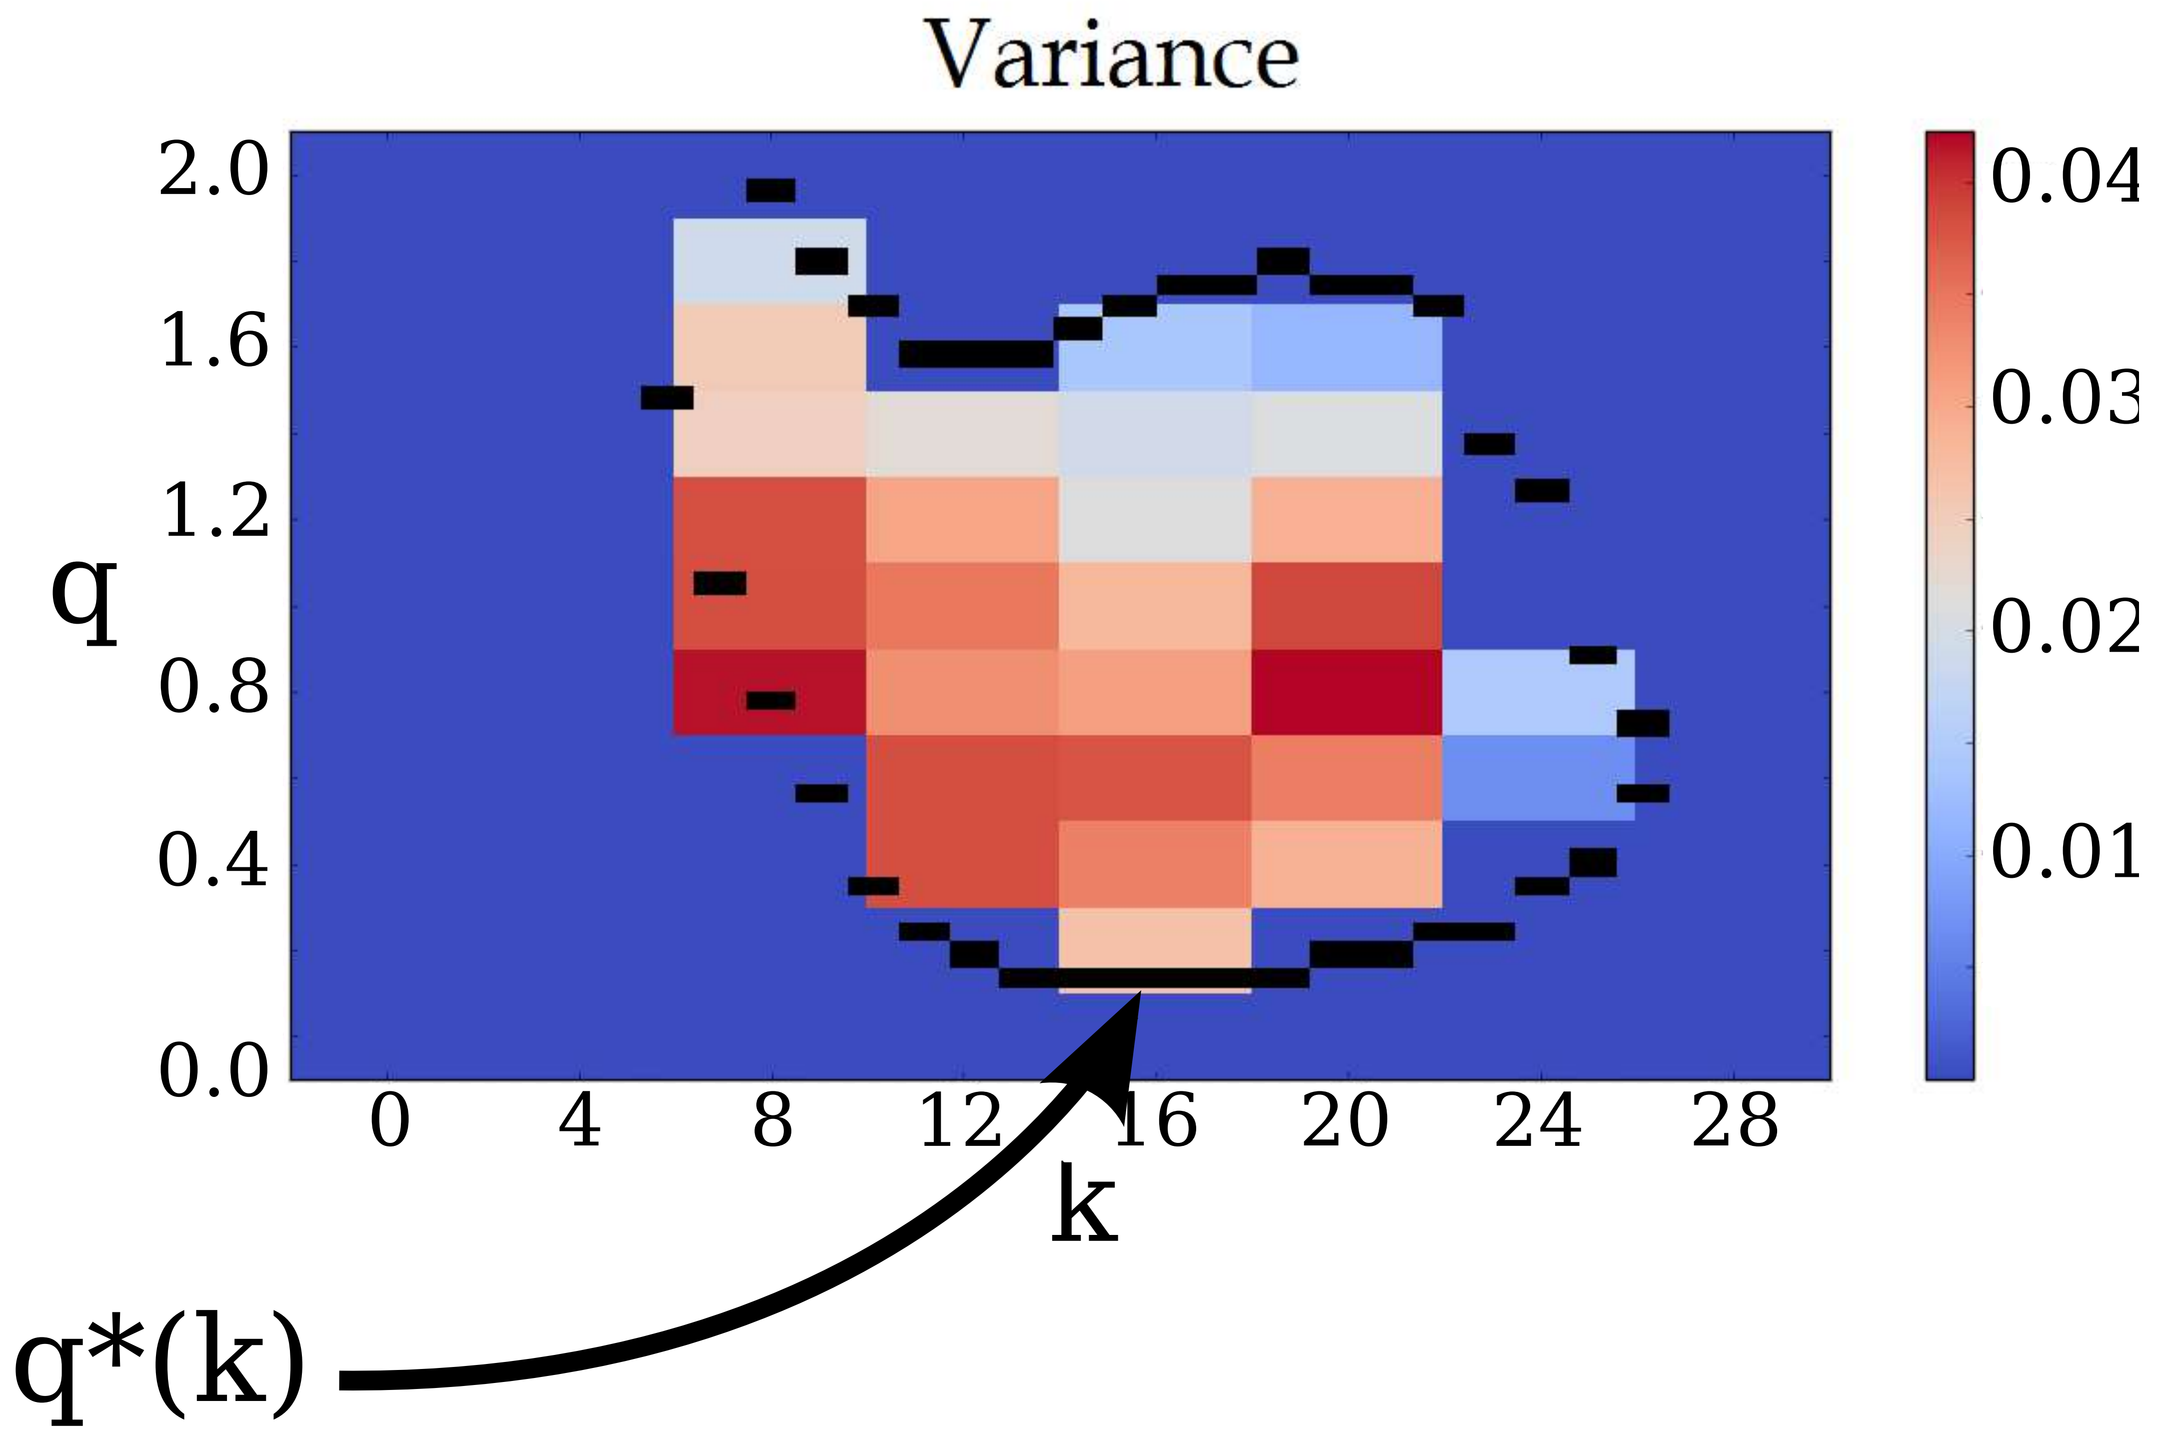

Inverting lower boundary to find sensitivity of network to different wavenumbers results in similar resonance as in experiments

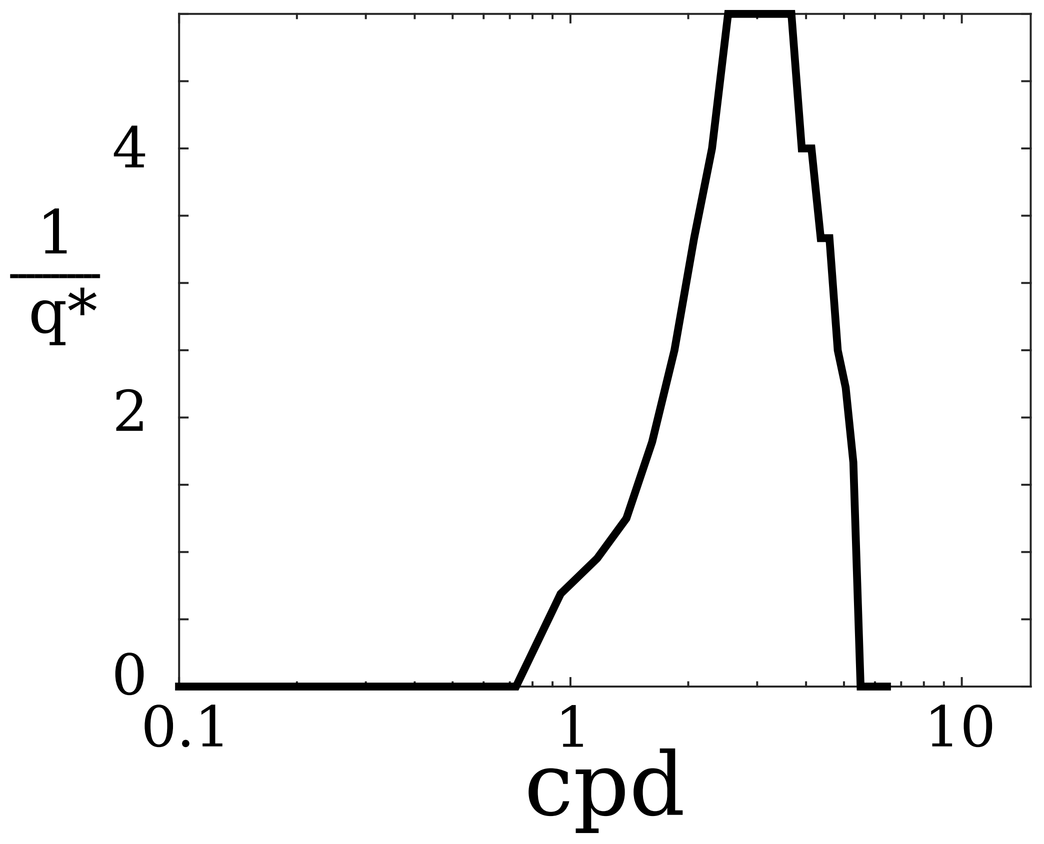

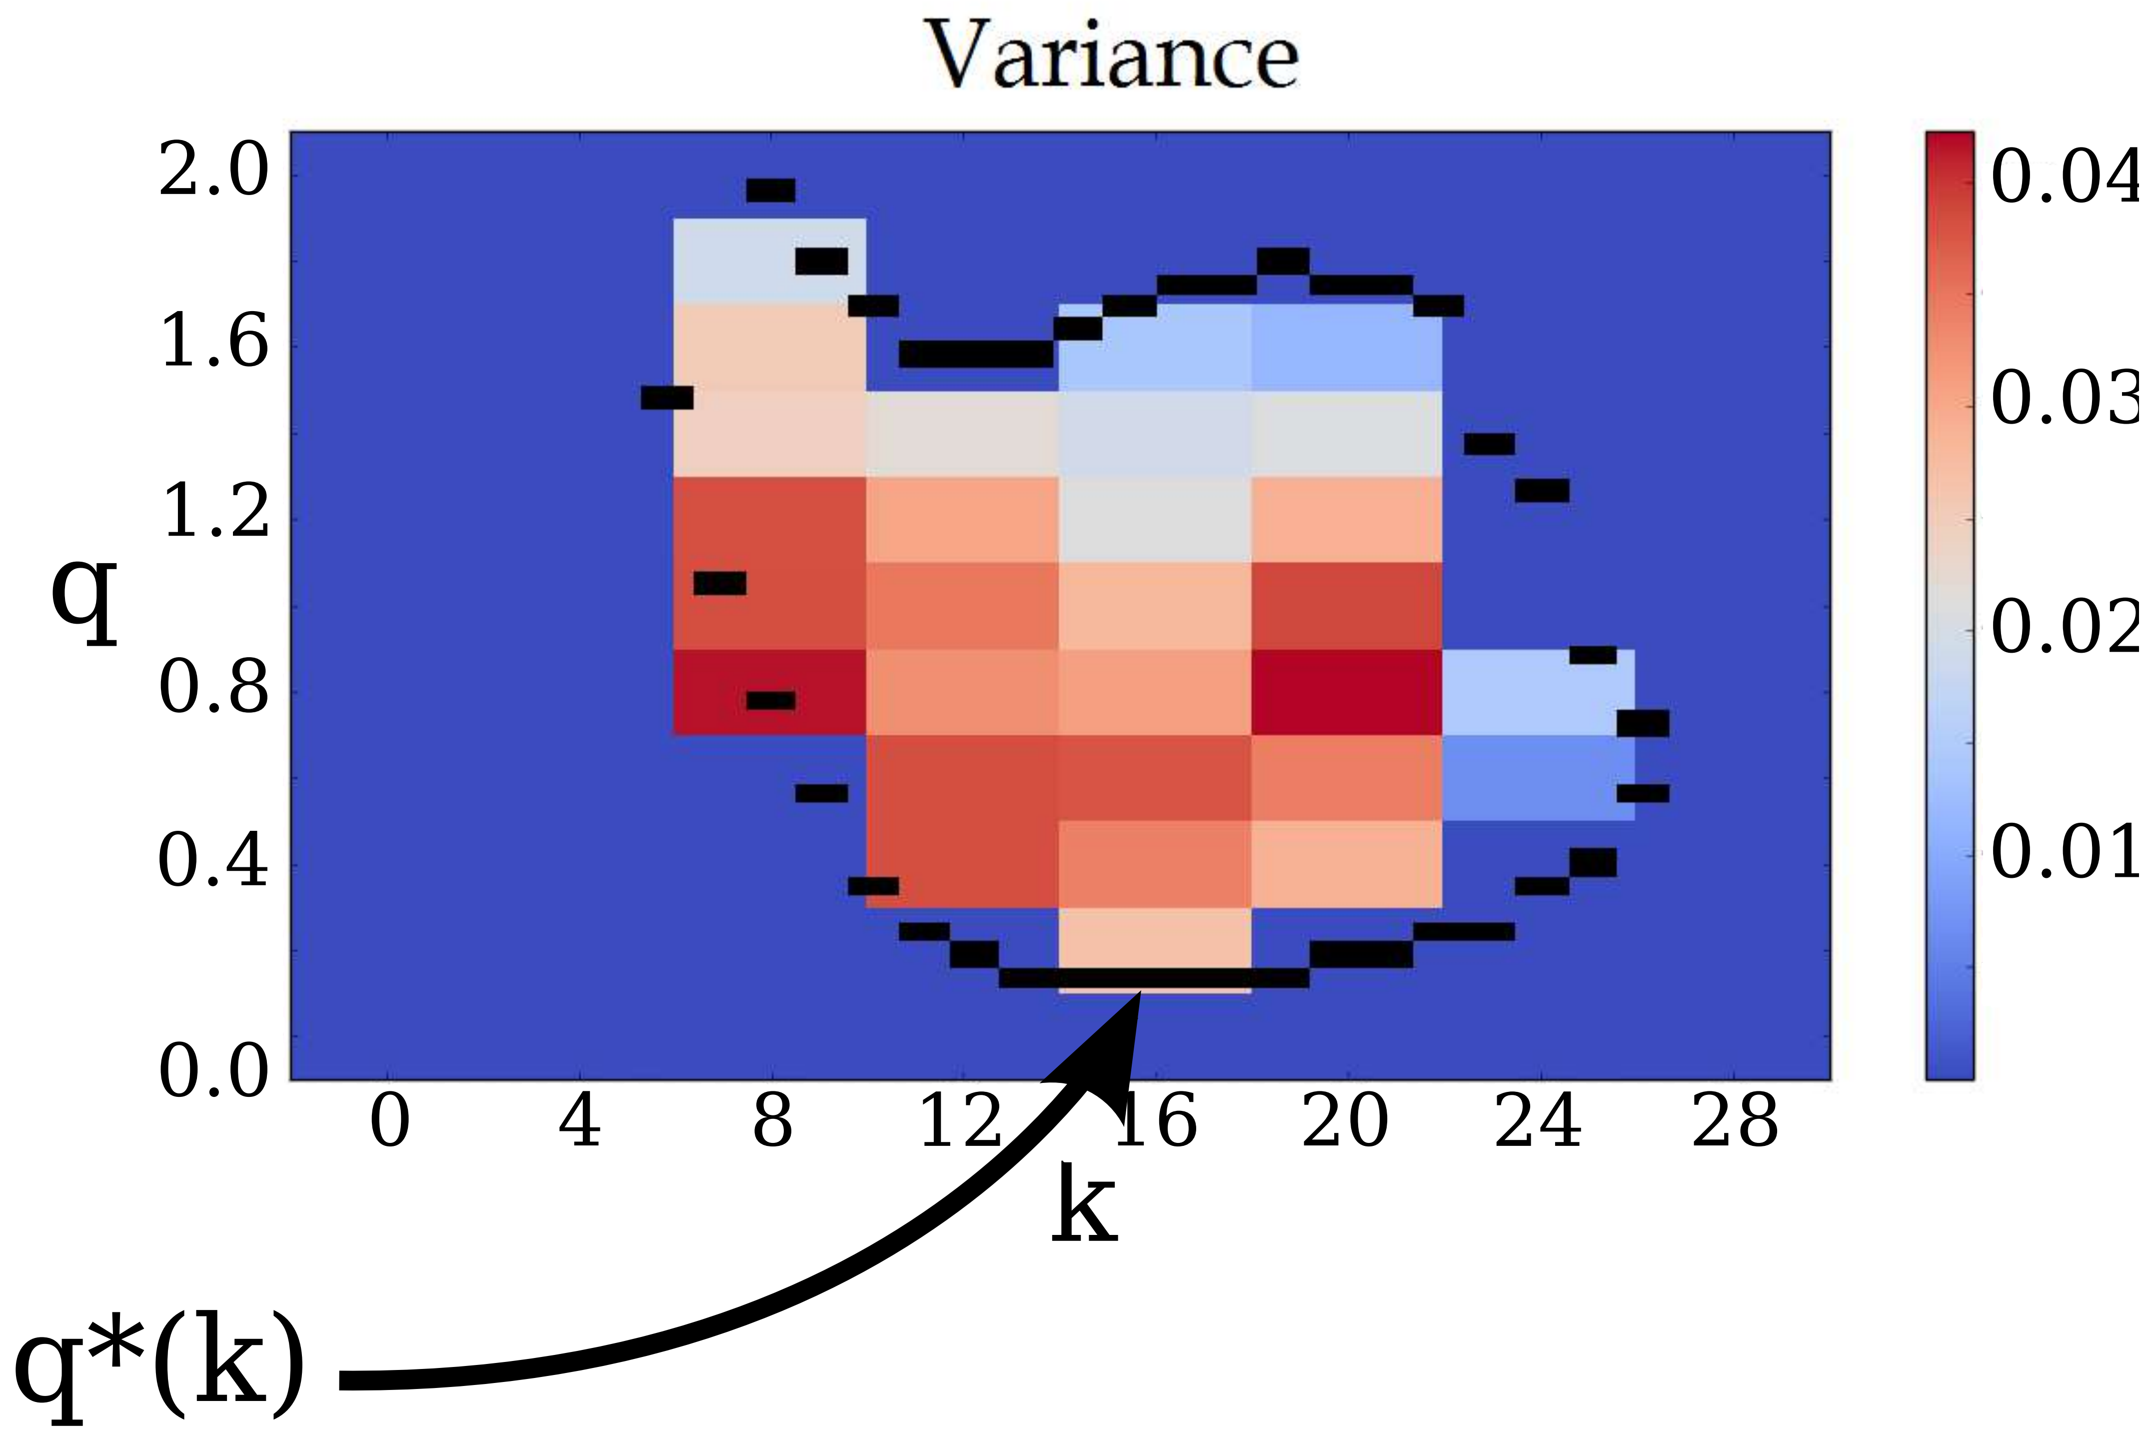

Inverting lower boundary to find sensitivity of network to different wavenumbers results in similar resonance as in experiments

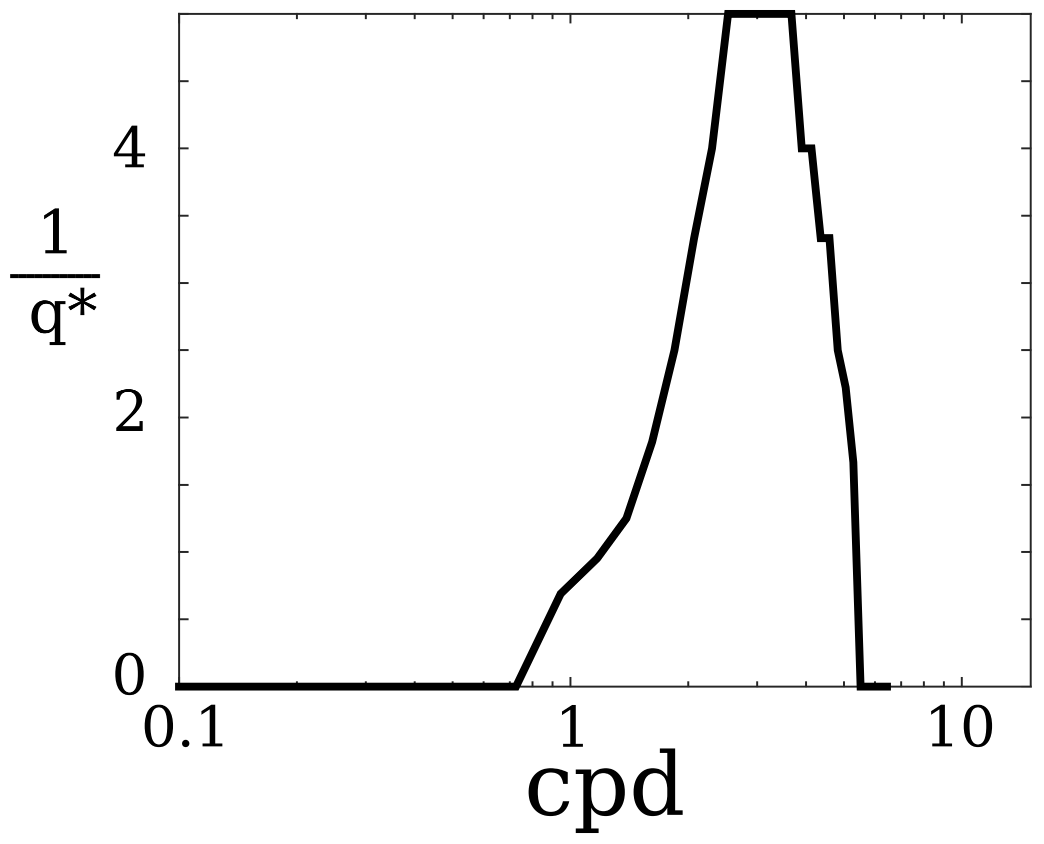

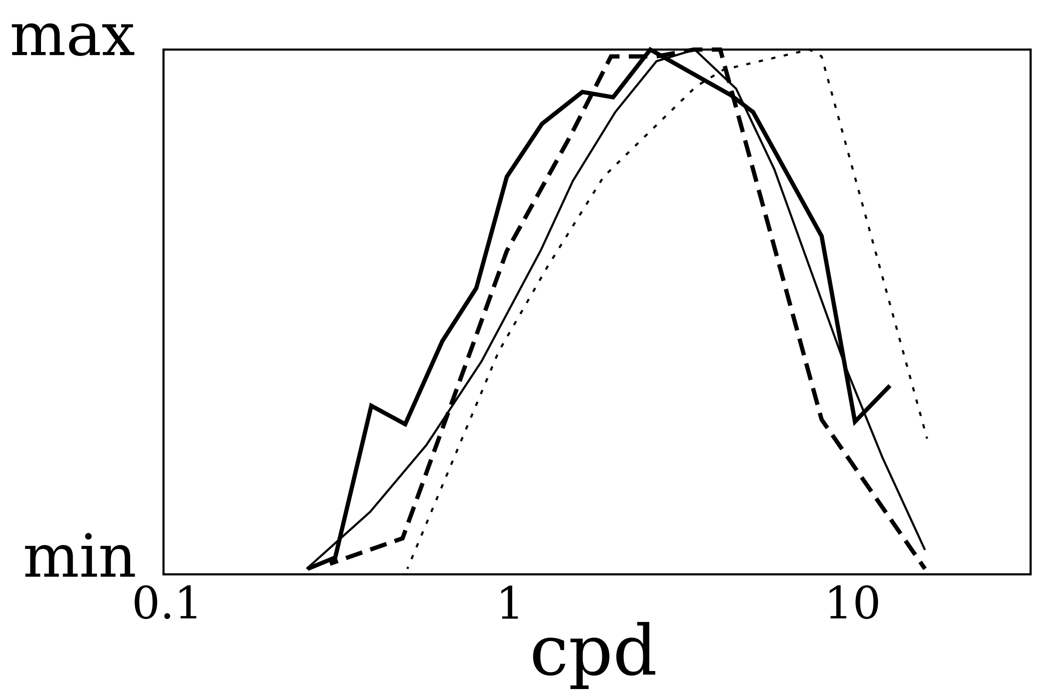

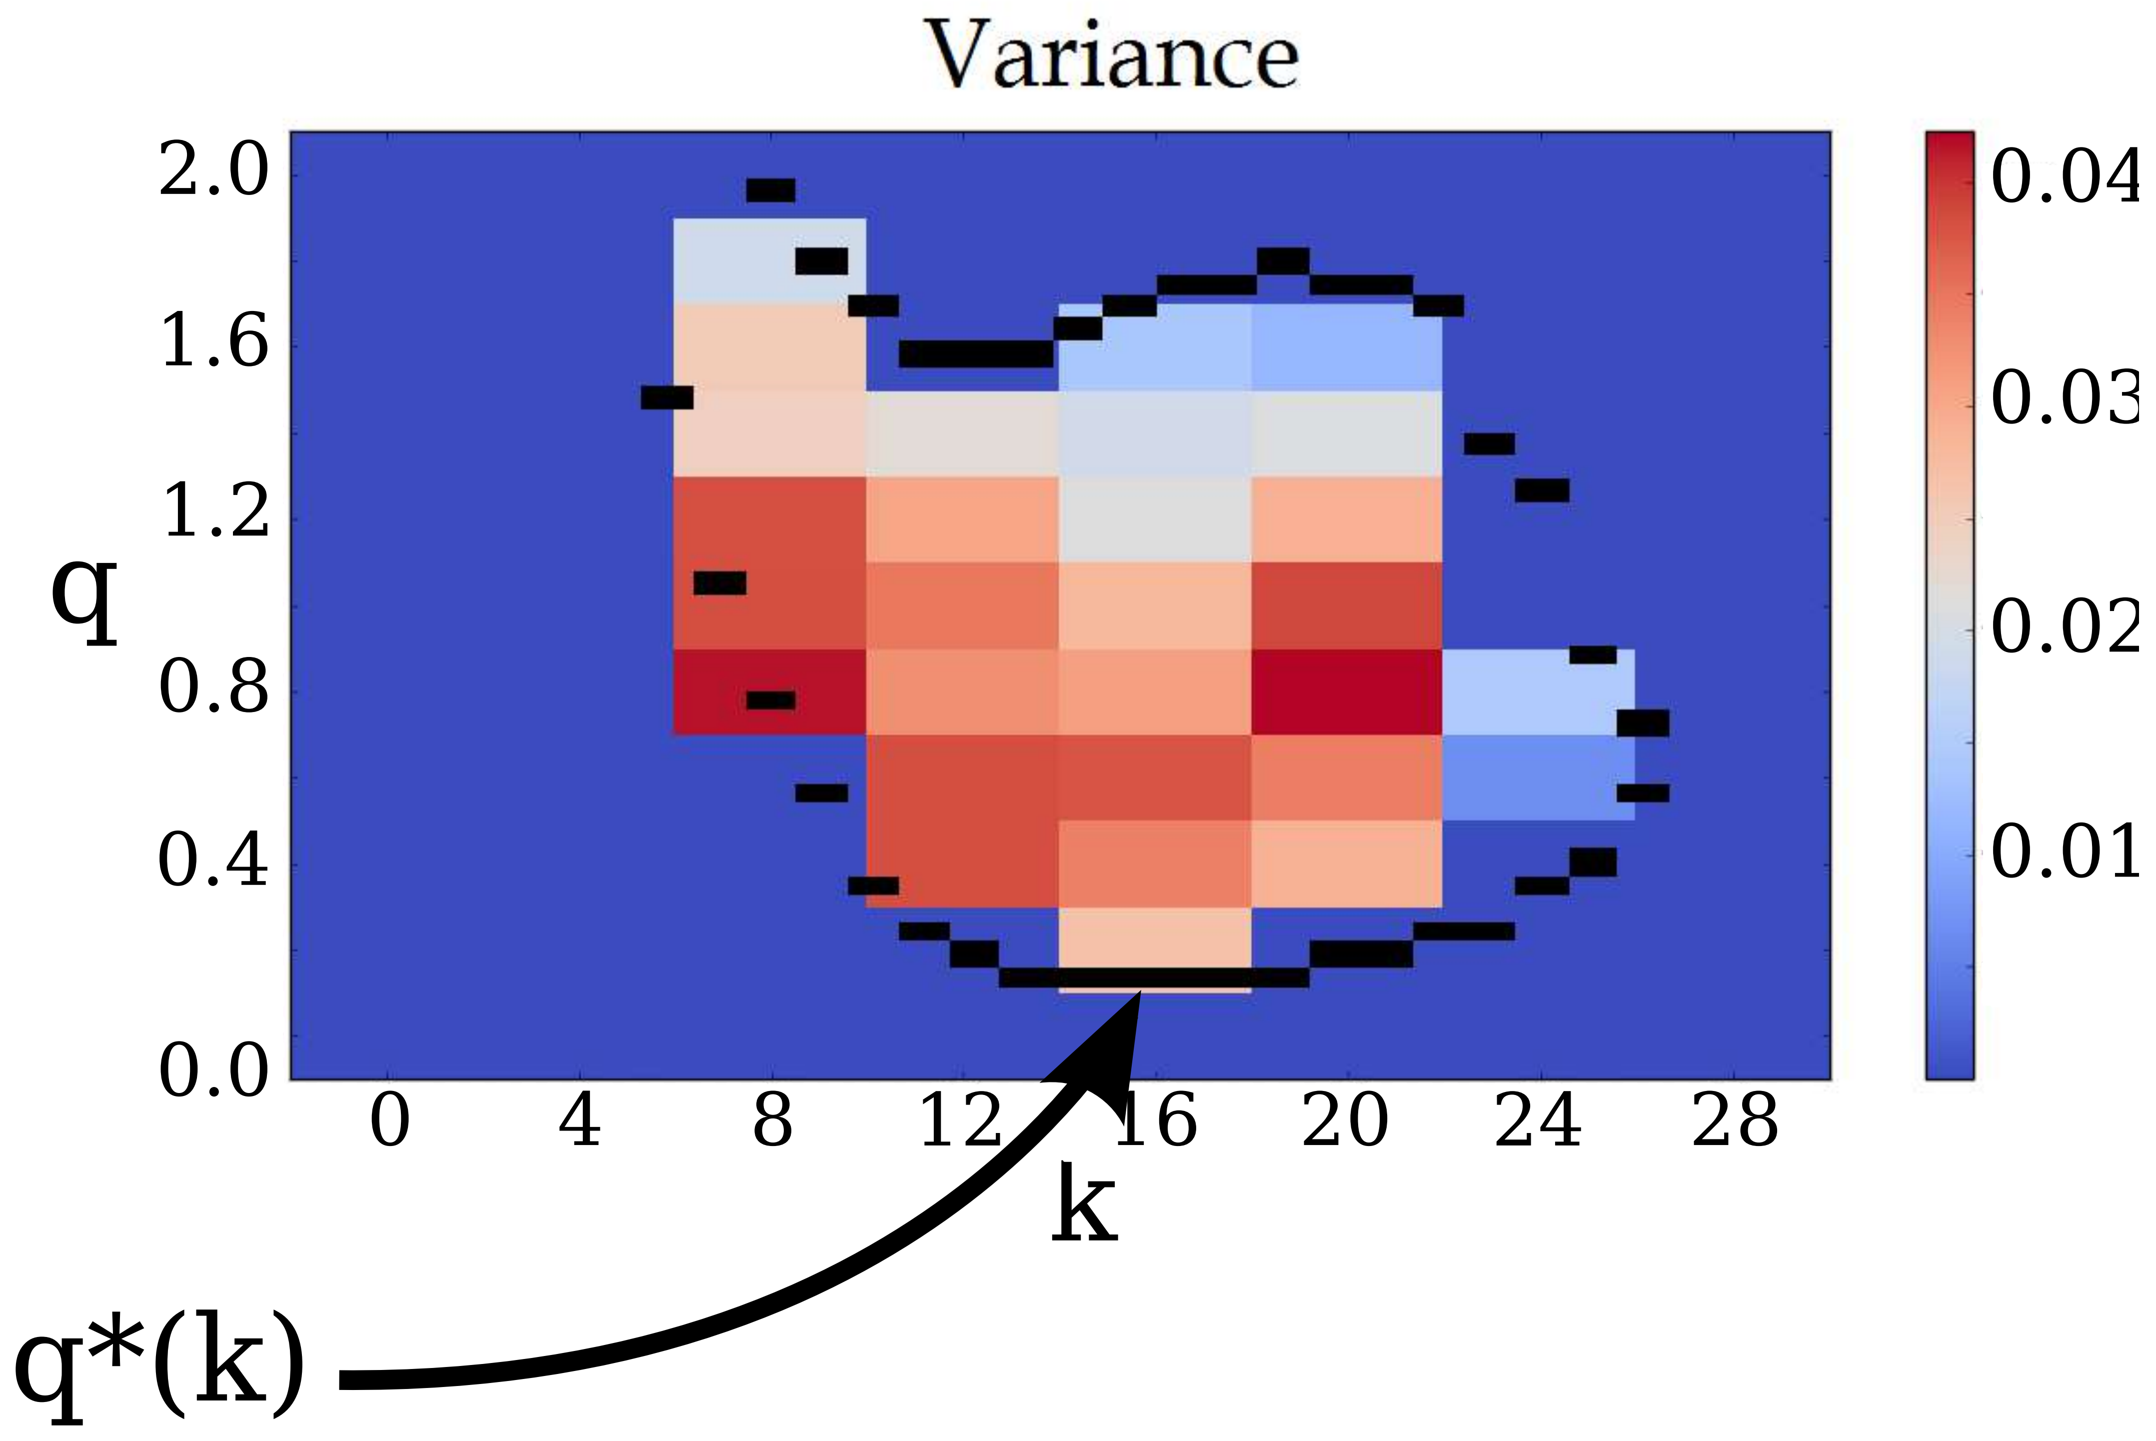

Neural field model in 1 spatial dimension easier to analyze and produces a large subset of spatiotemporal dymnamics

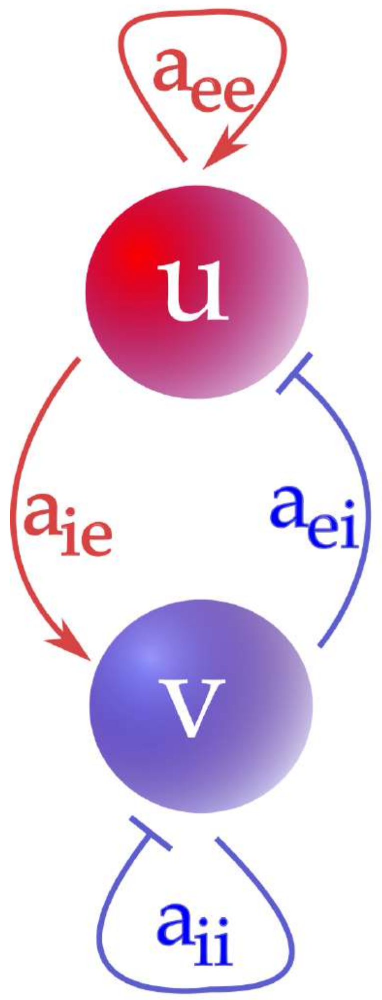

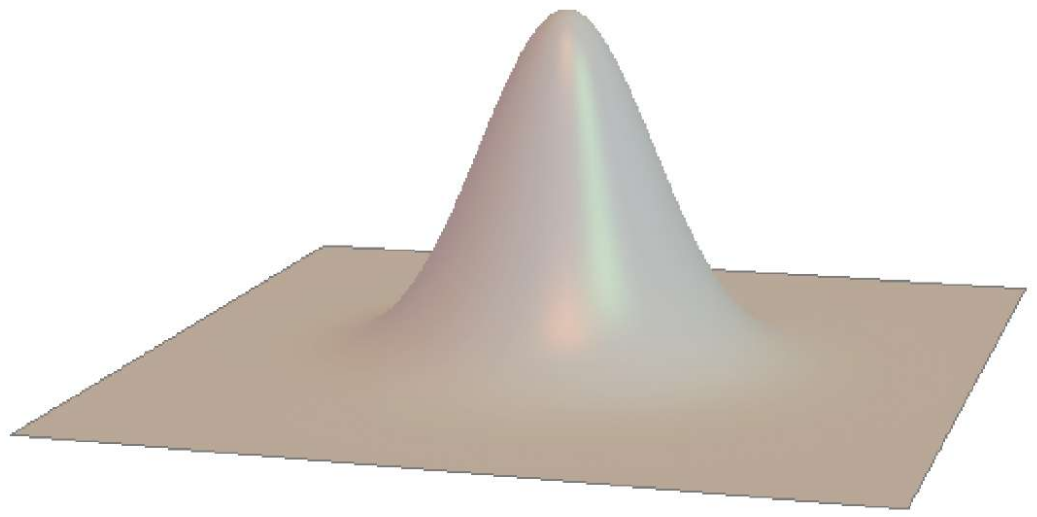

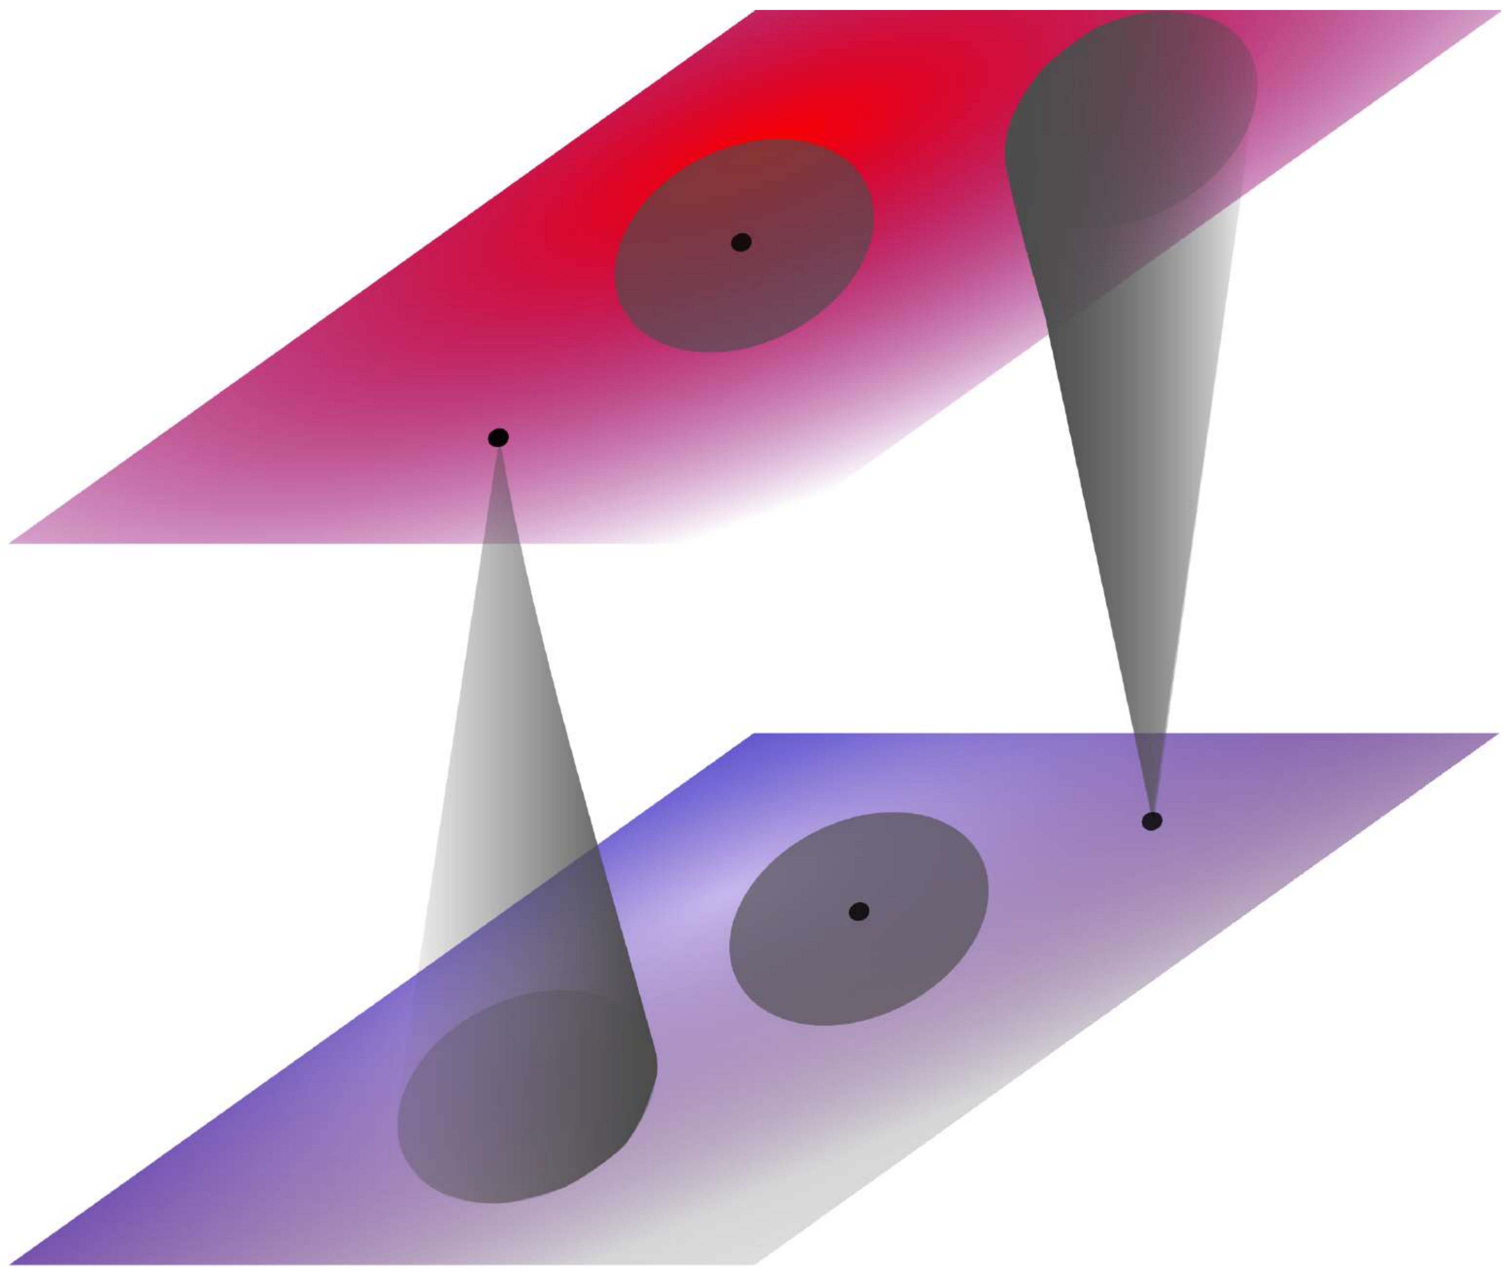

Neural field model in 1 spatial dimension easier to analyze and produces a large subset of spatiotemporal dymnamics

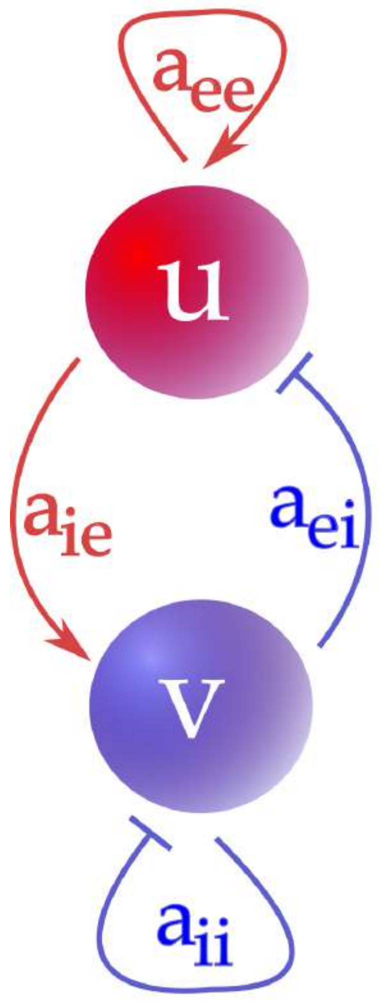

Neural field model in 1 spatial dimension easier to analyze and produces a large subset of spatiotemporal dymnamics

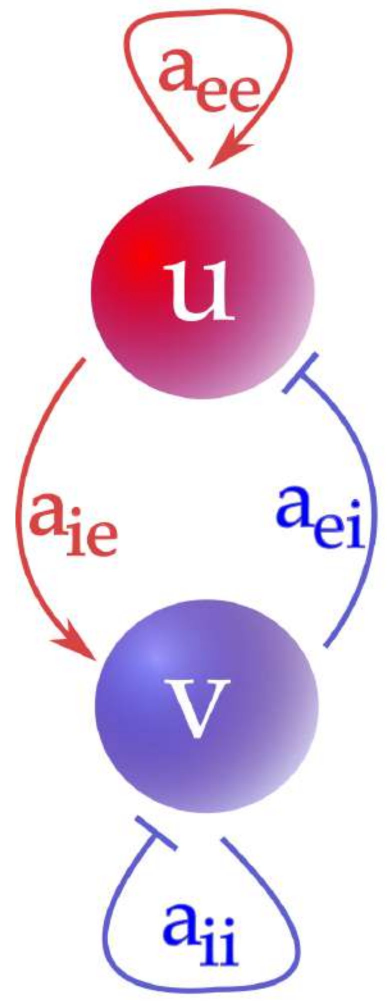

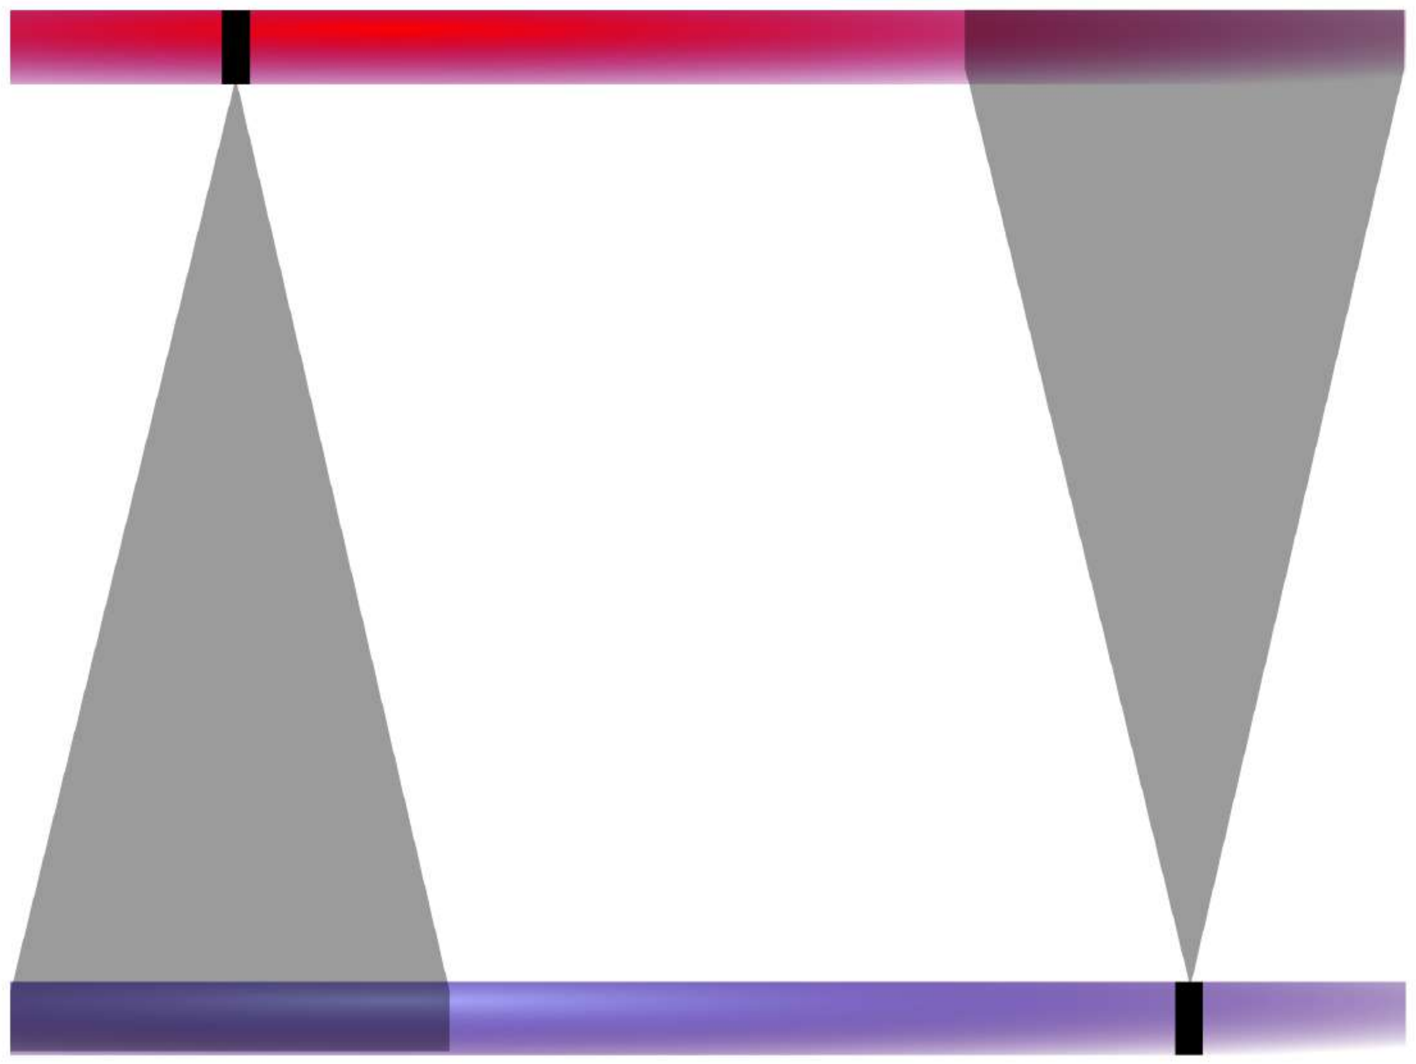

Neural field model in 1 spatial dimension easier to analyze and produces a large subset of spatiotemporal dymnamics

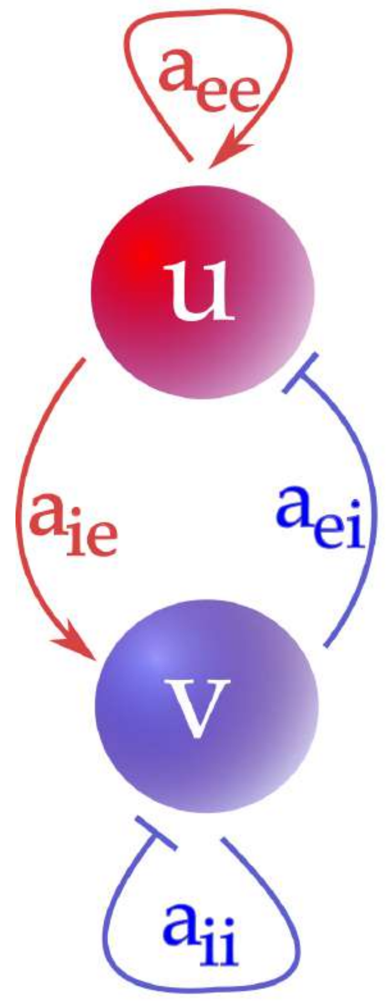

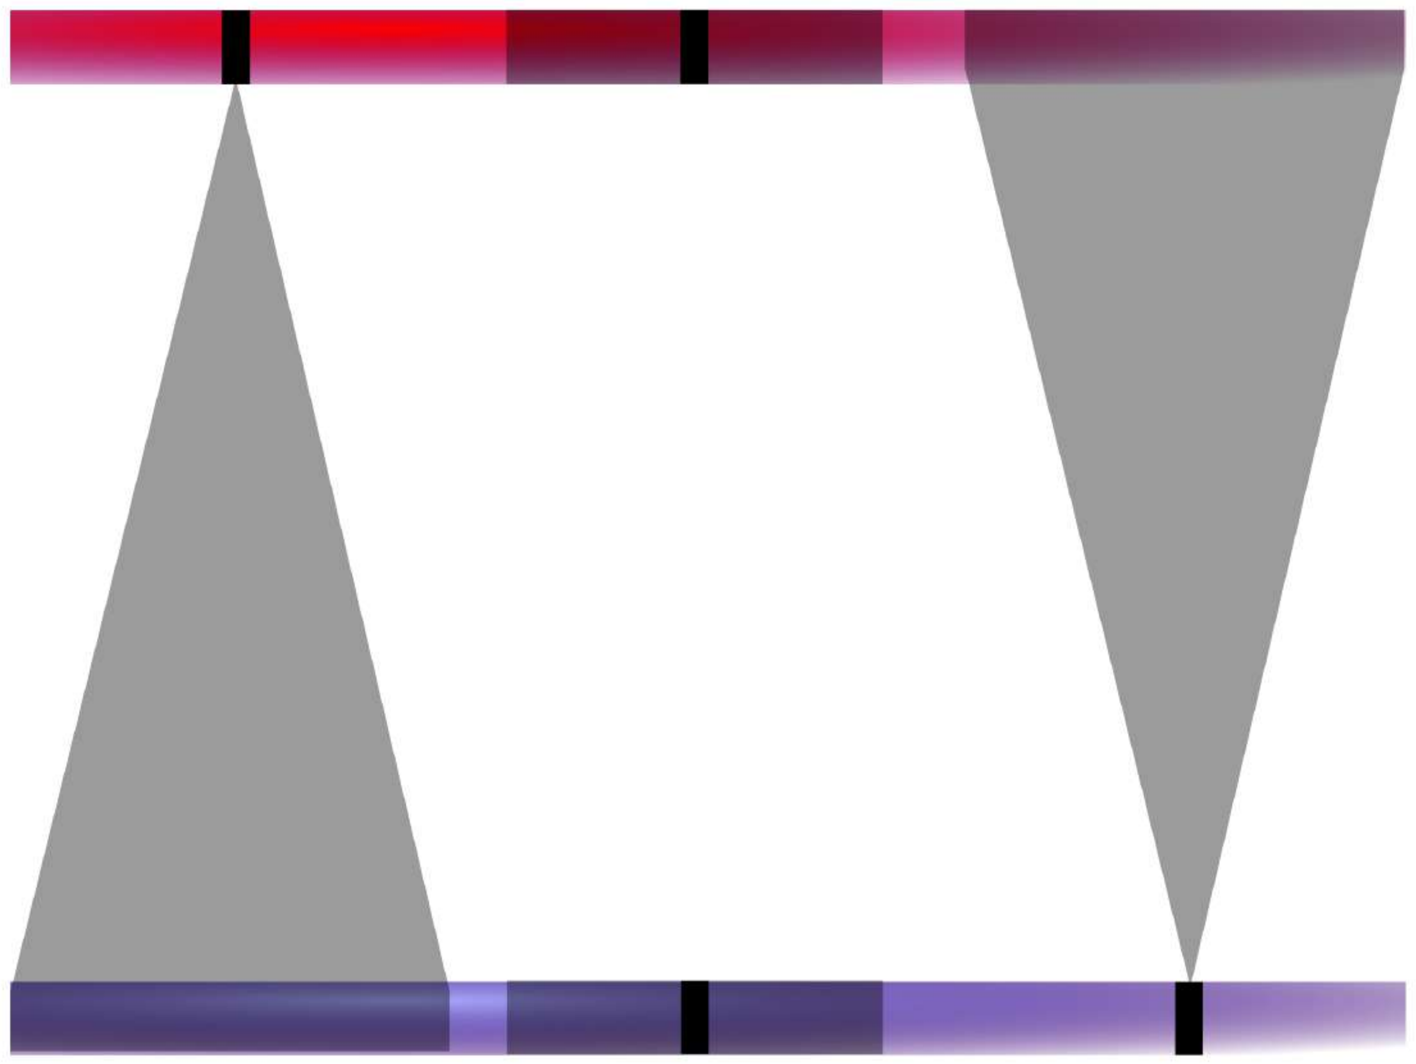

Neural field model in 1 spatial dimension easier to analyze and produces a large subset of spatiotemporal dymnamics

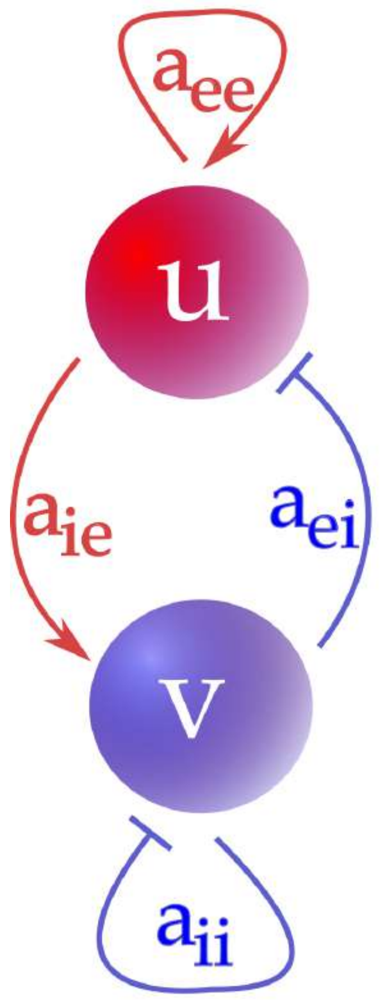

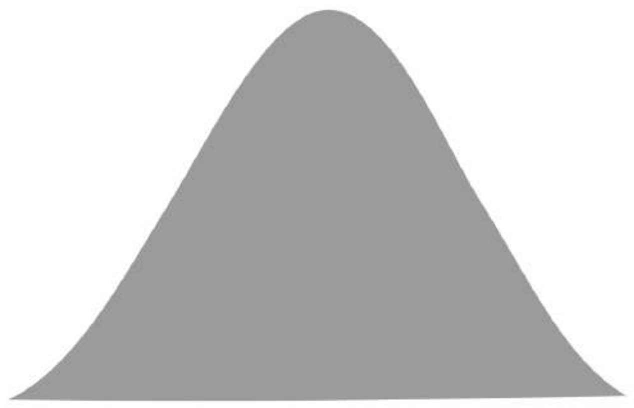

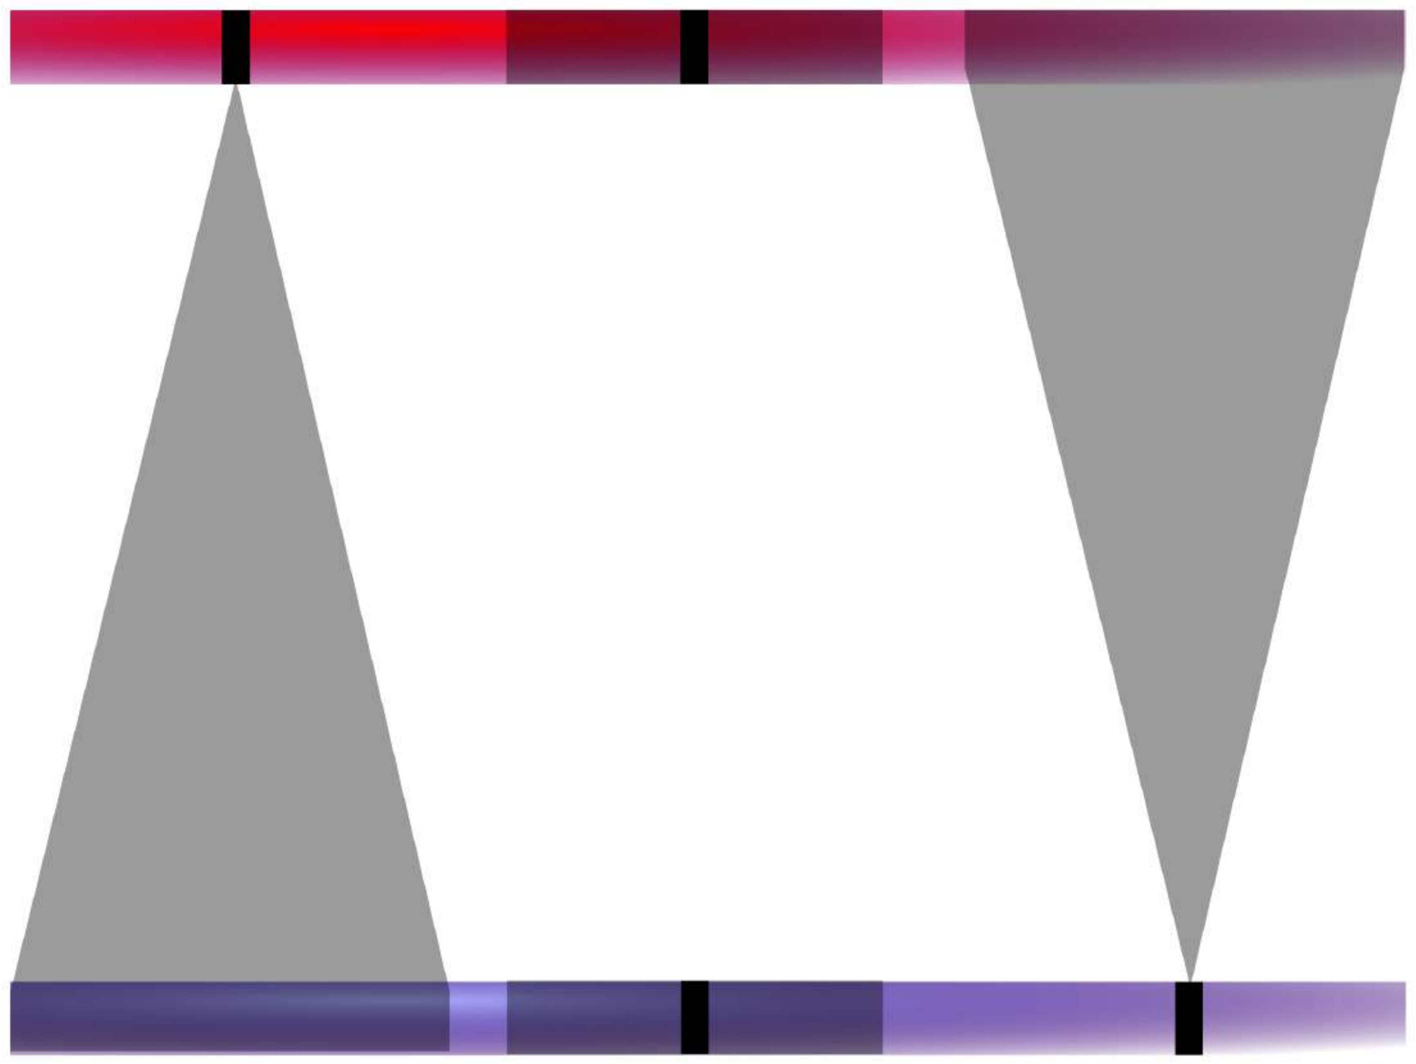

In 1D, we obtain similar resonance / sensitivity as with 2-D model

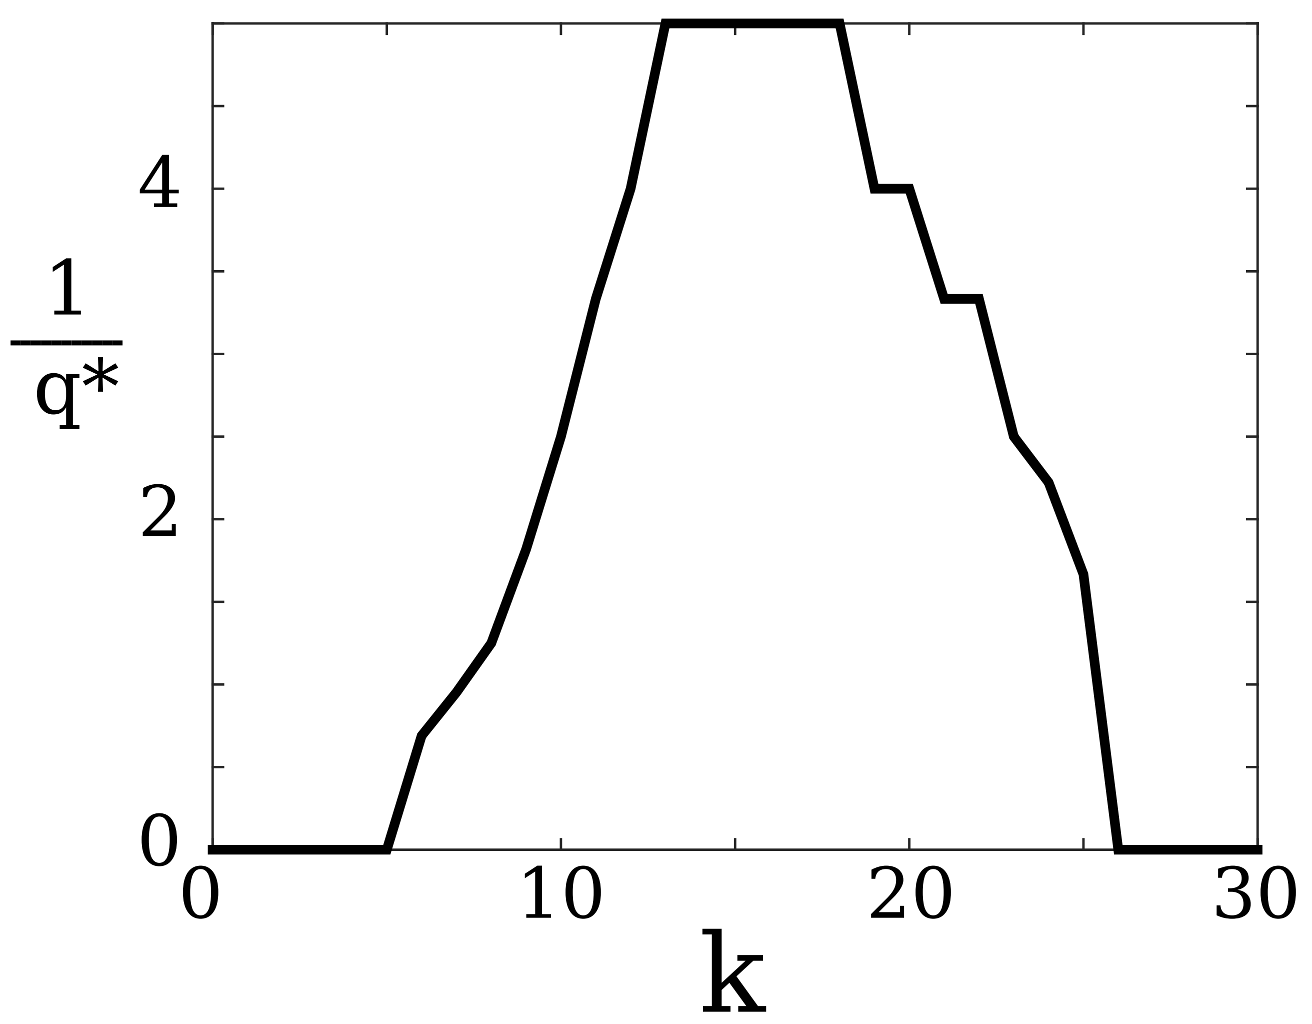

In 1D, we obtain similar resonance / sensitivity as with 2-D model

In 2D

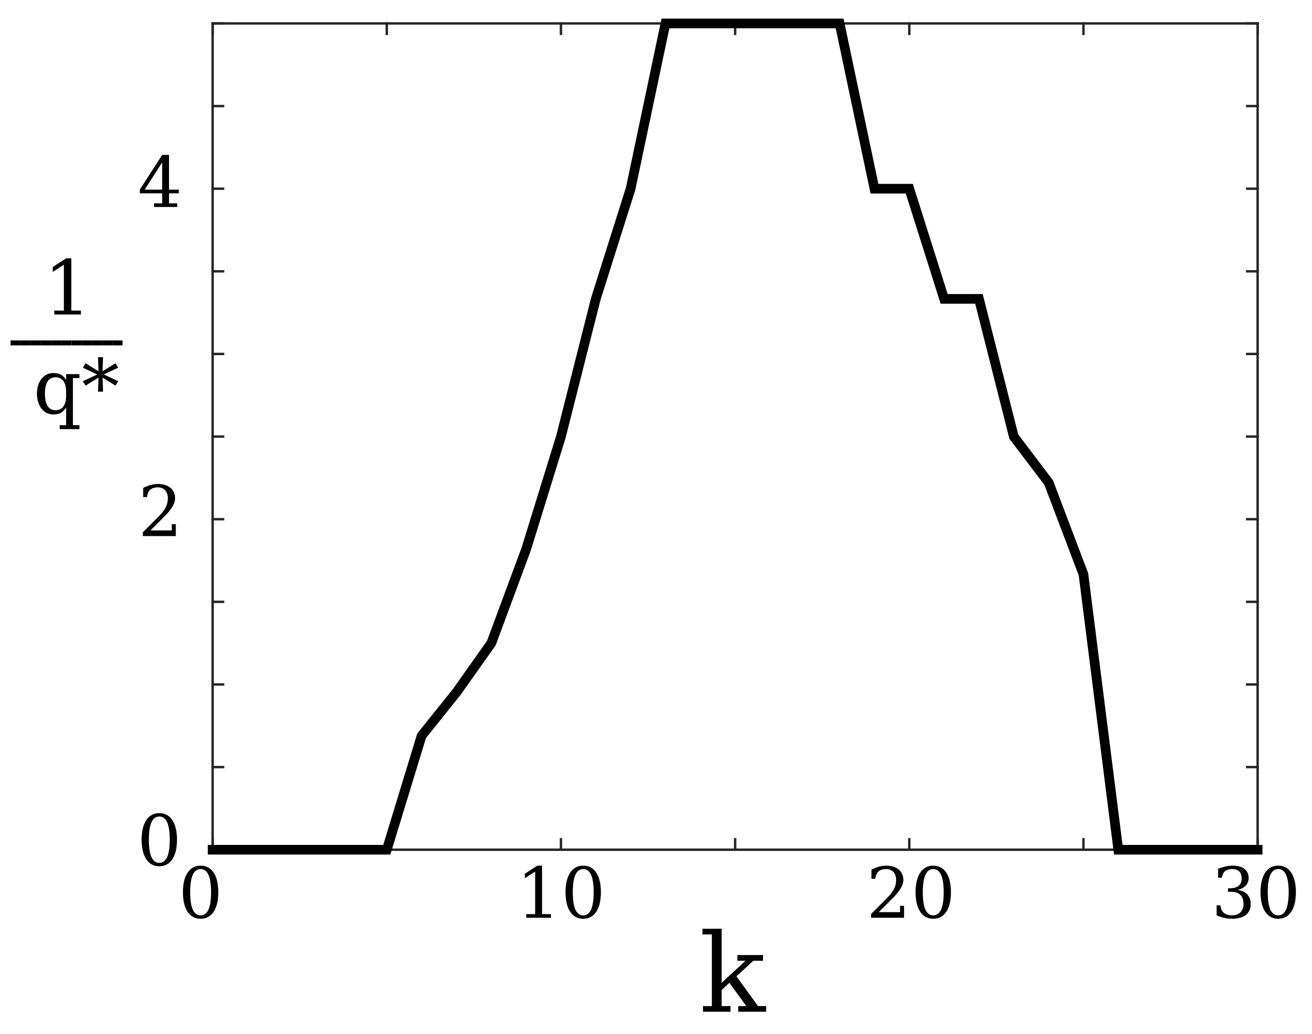

In 1D

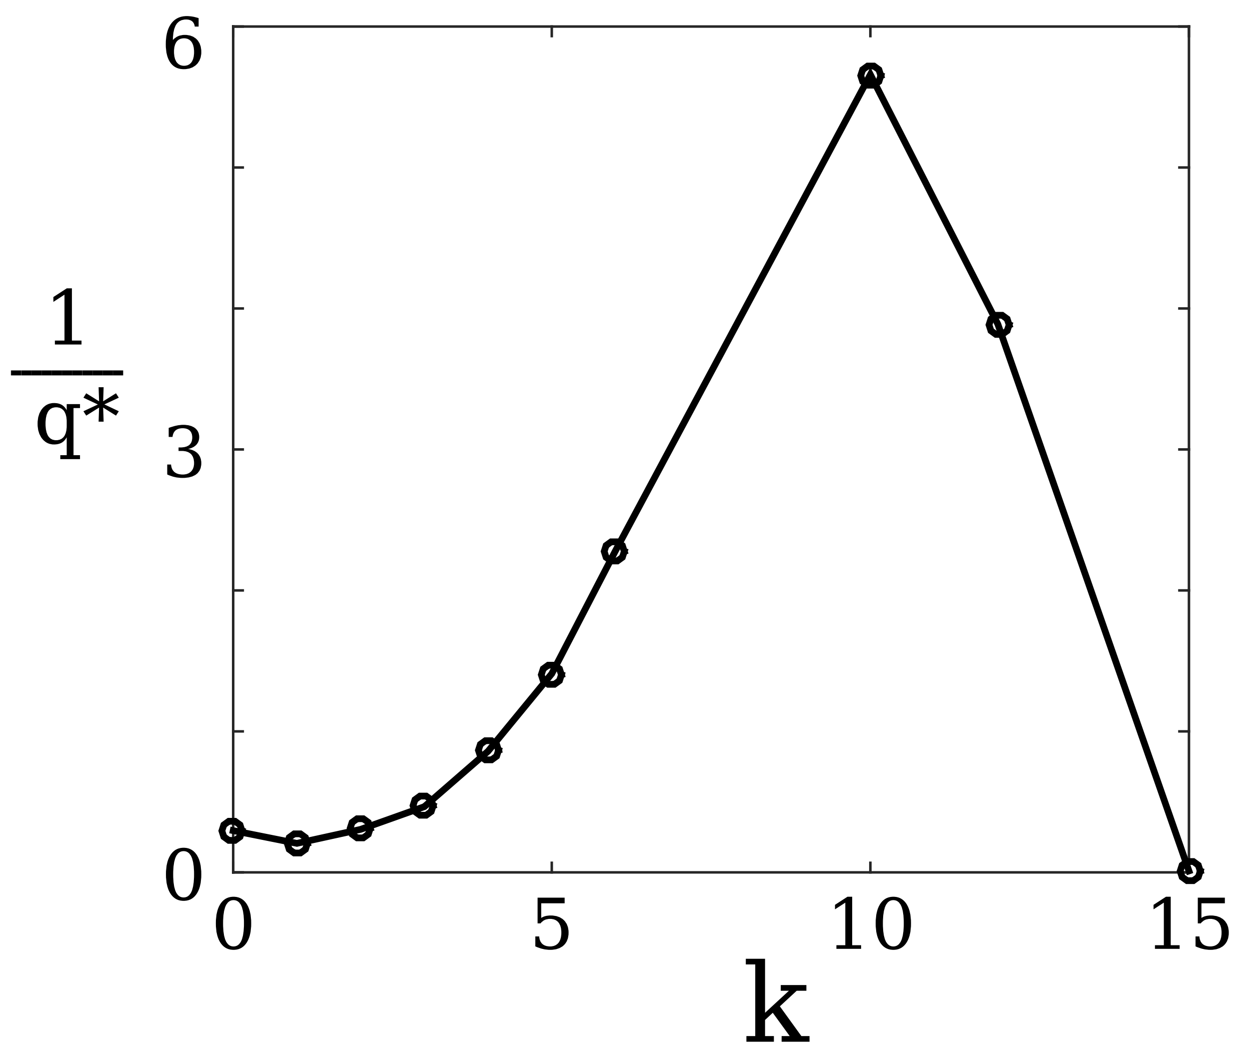

We generally obtain standing-wave patterns in 1D that look similar to those obtained in 2D

Stimulus−free

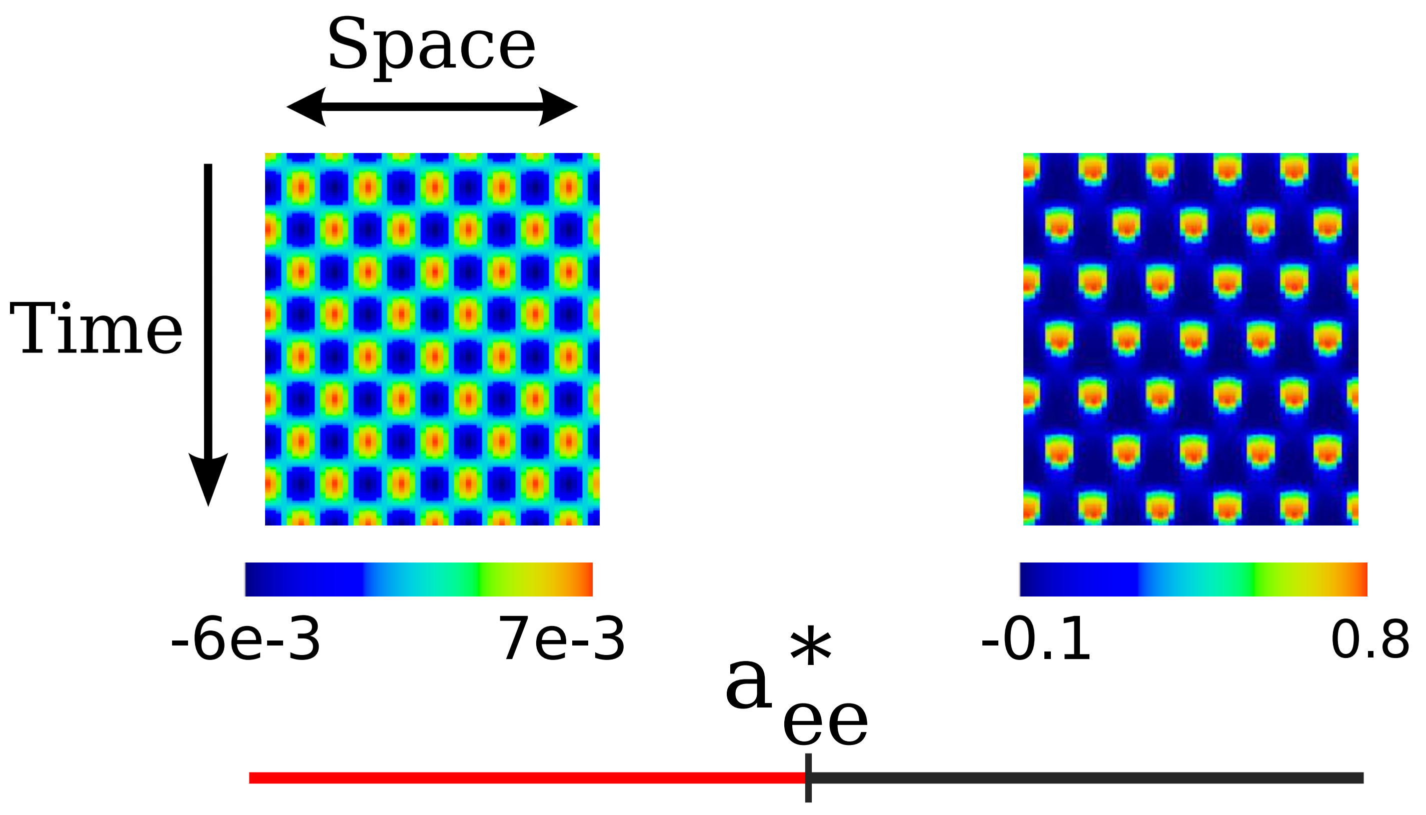

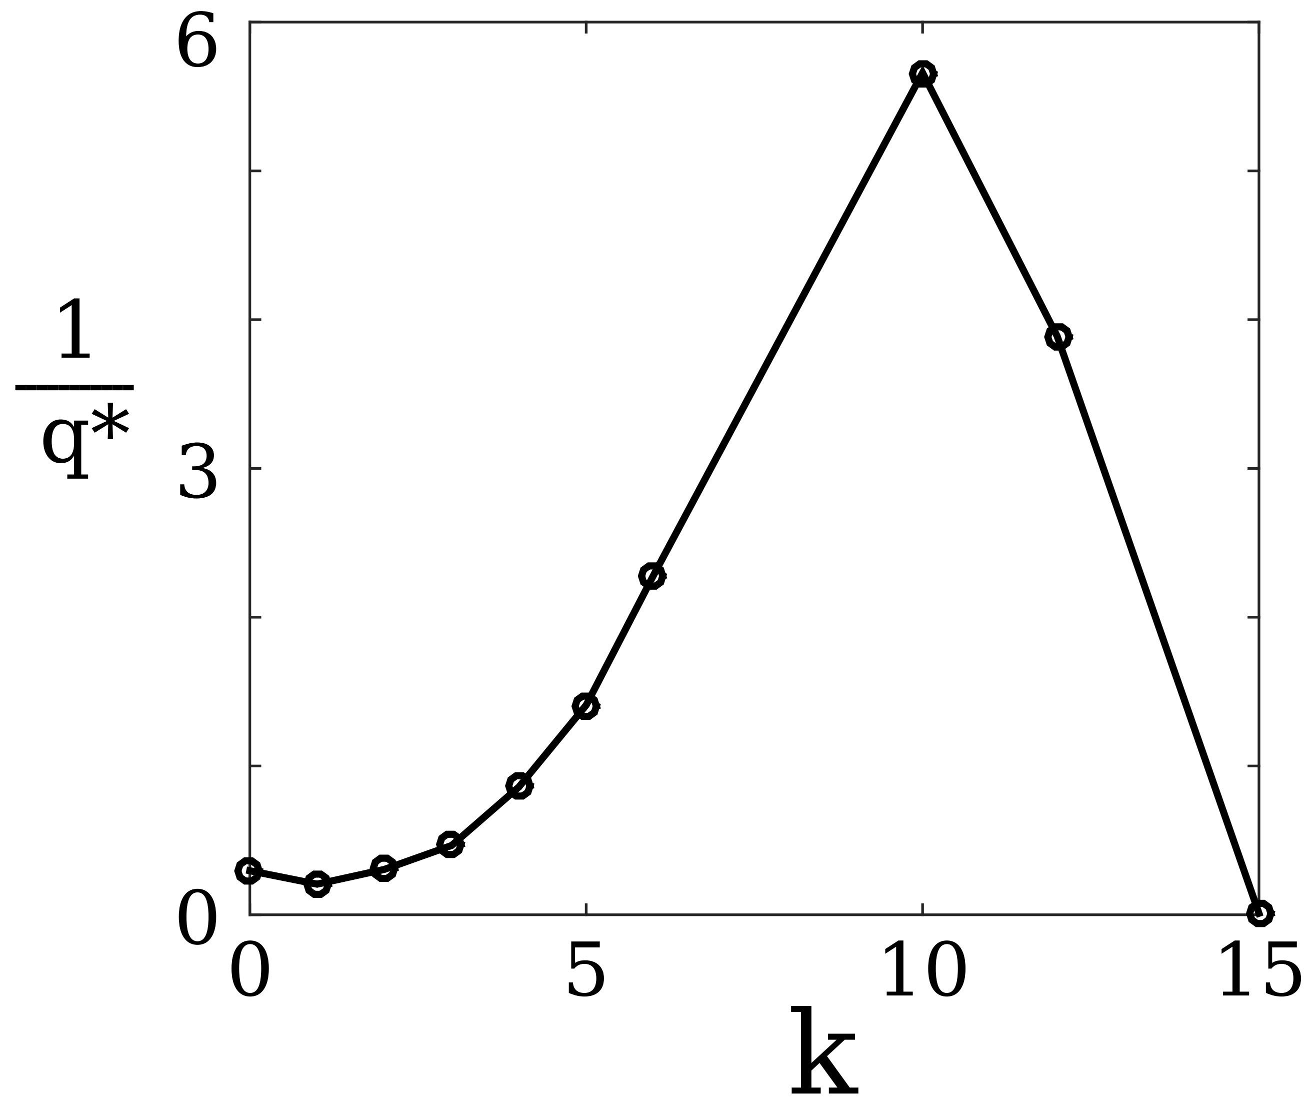

We generally obtain standing-wave patterns in 1D that look similar to those obtained in 2D

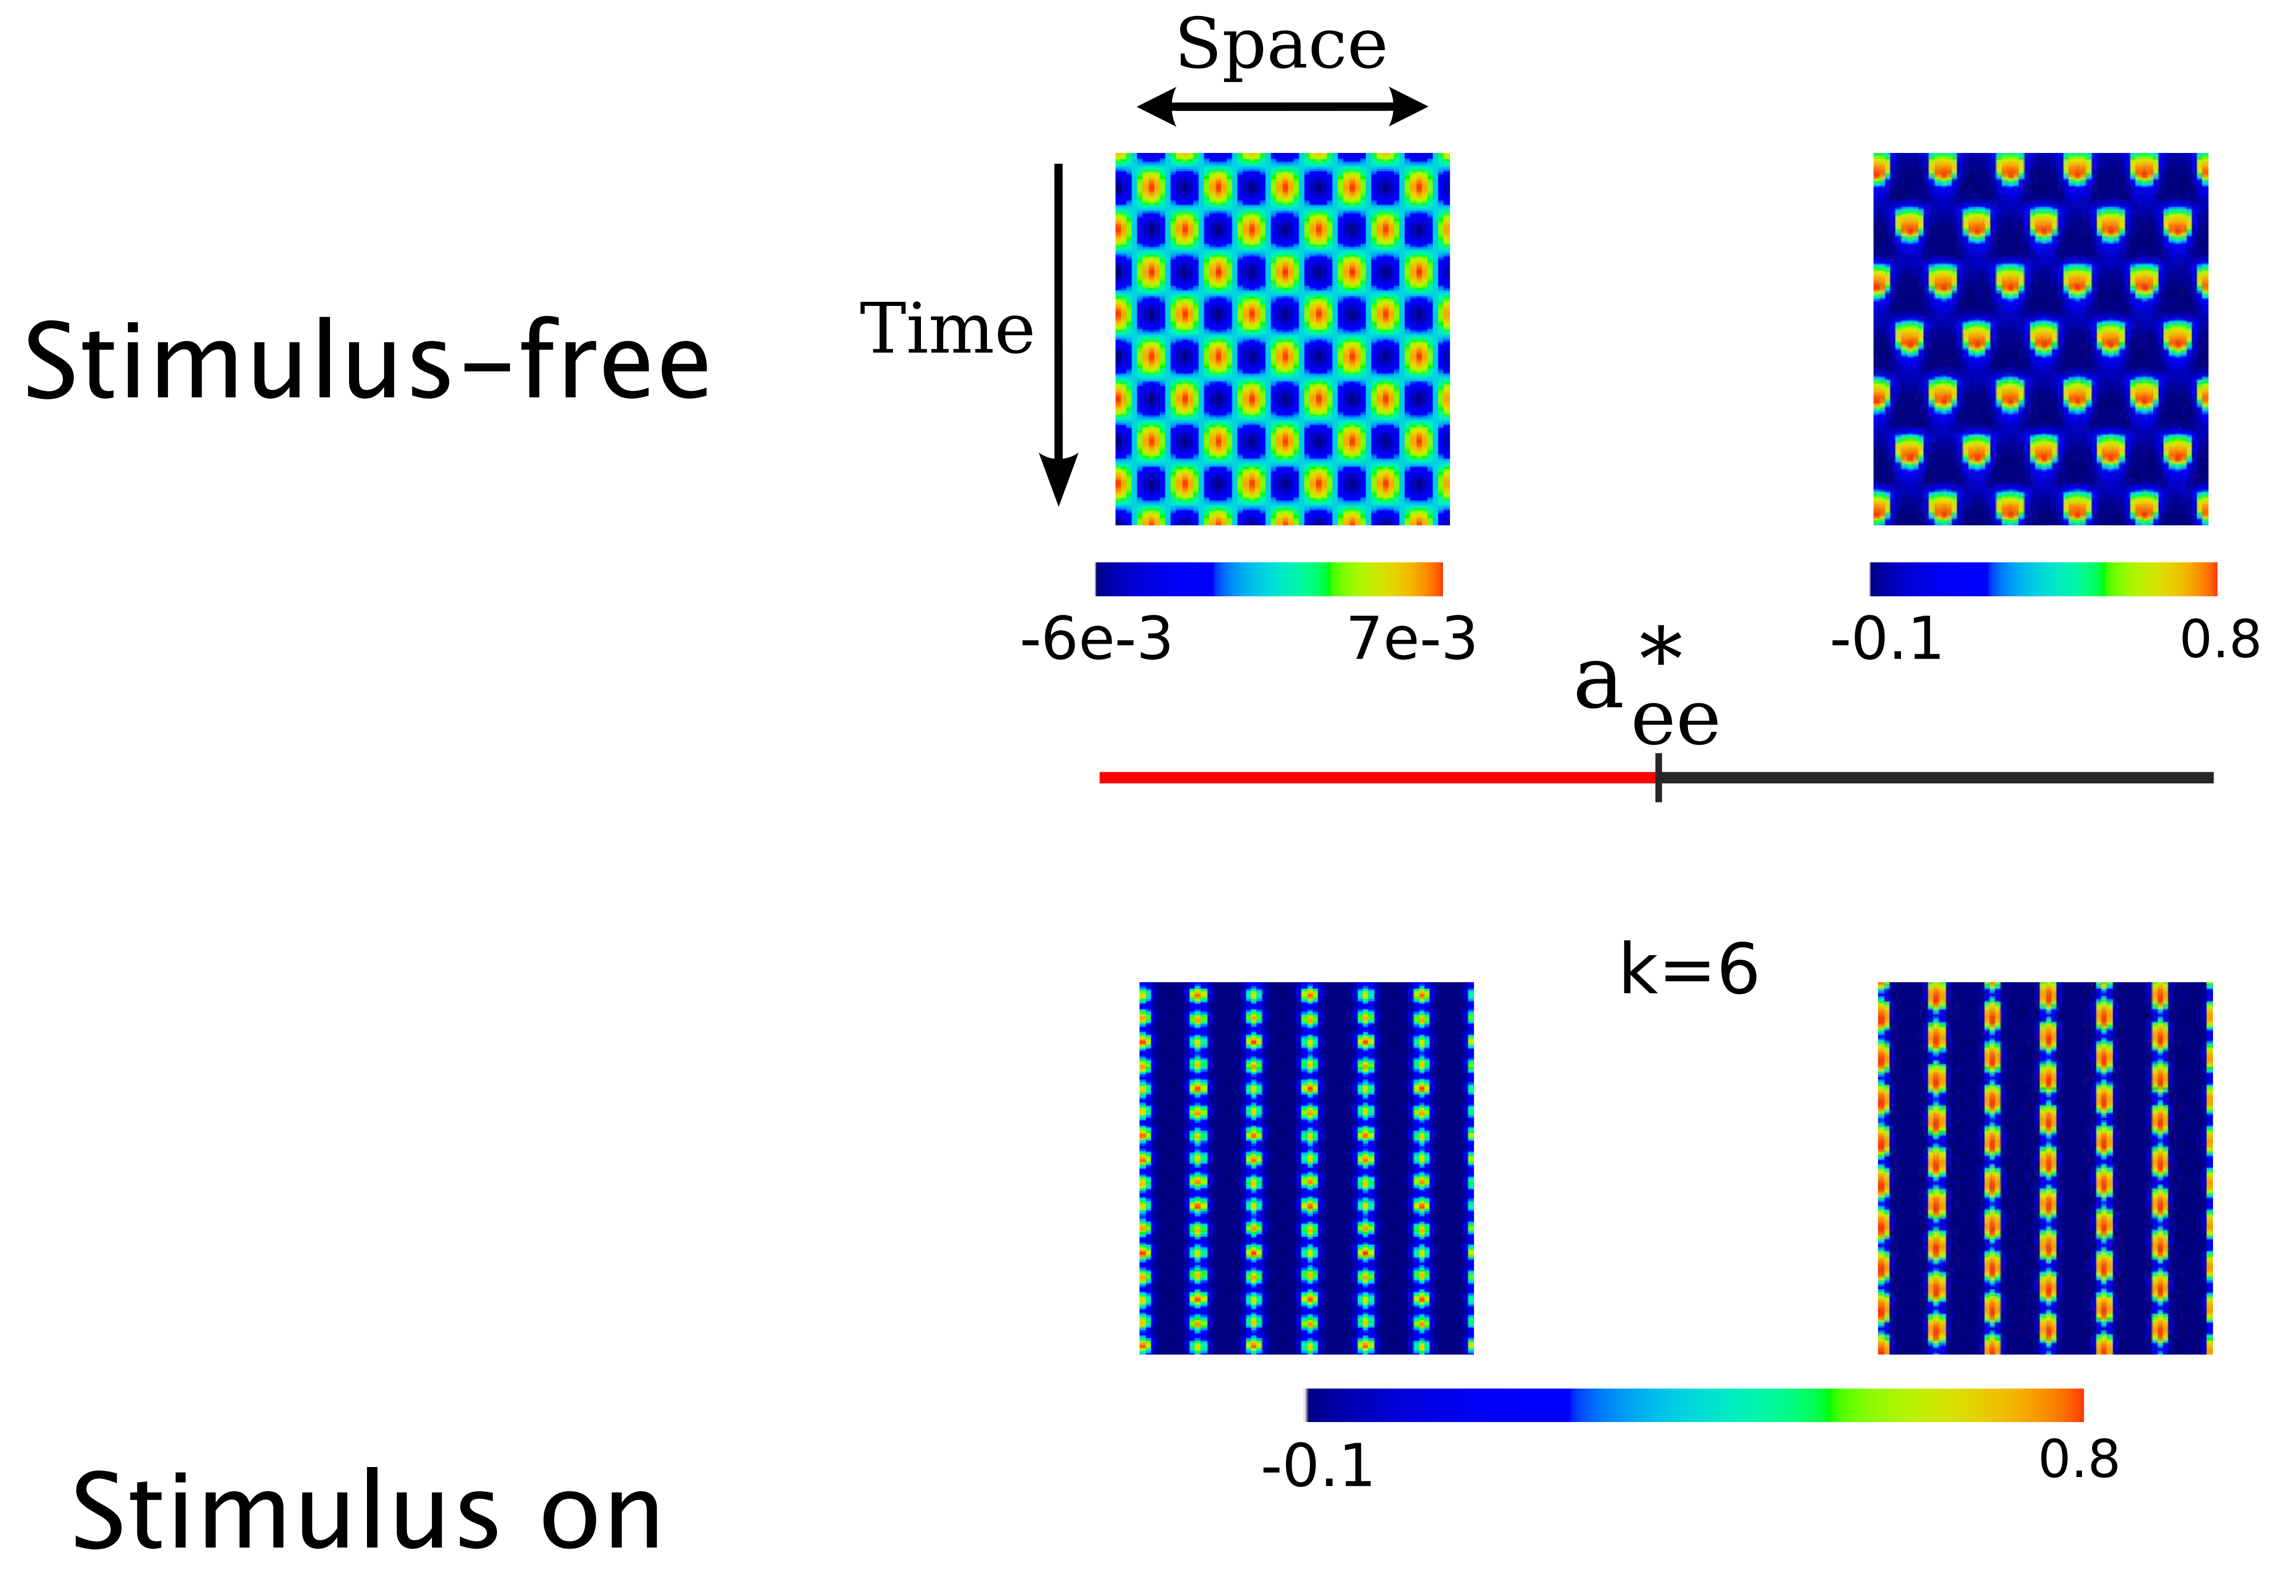

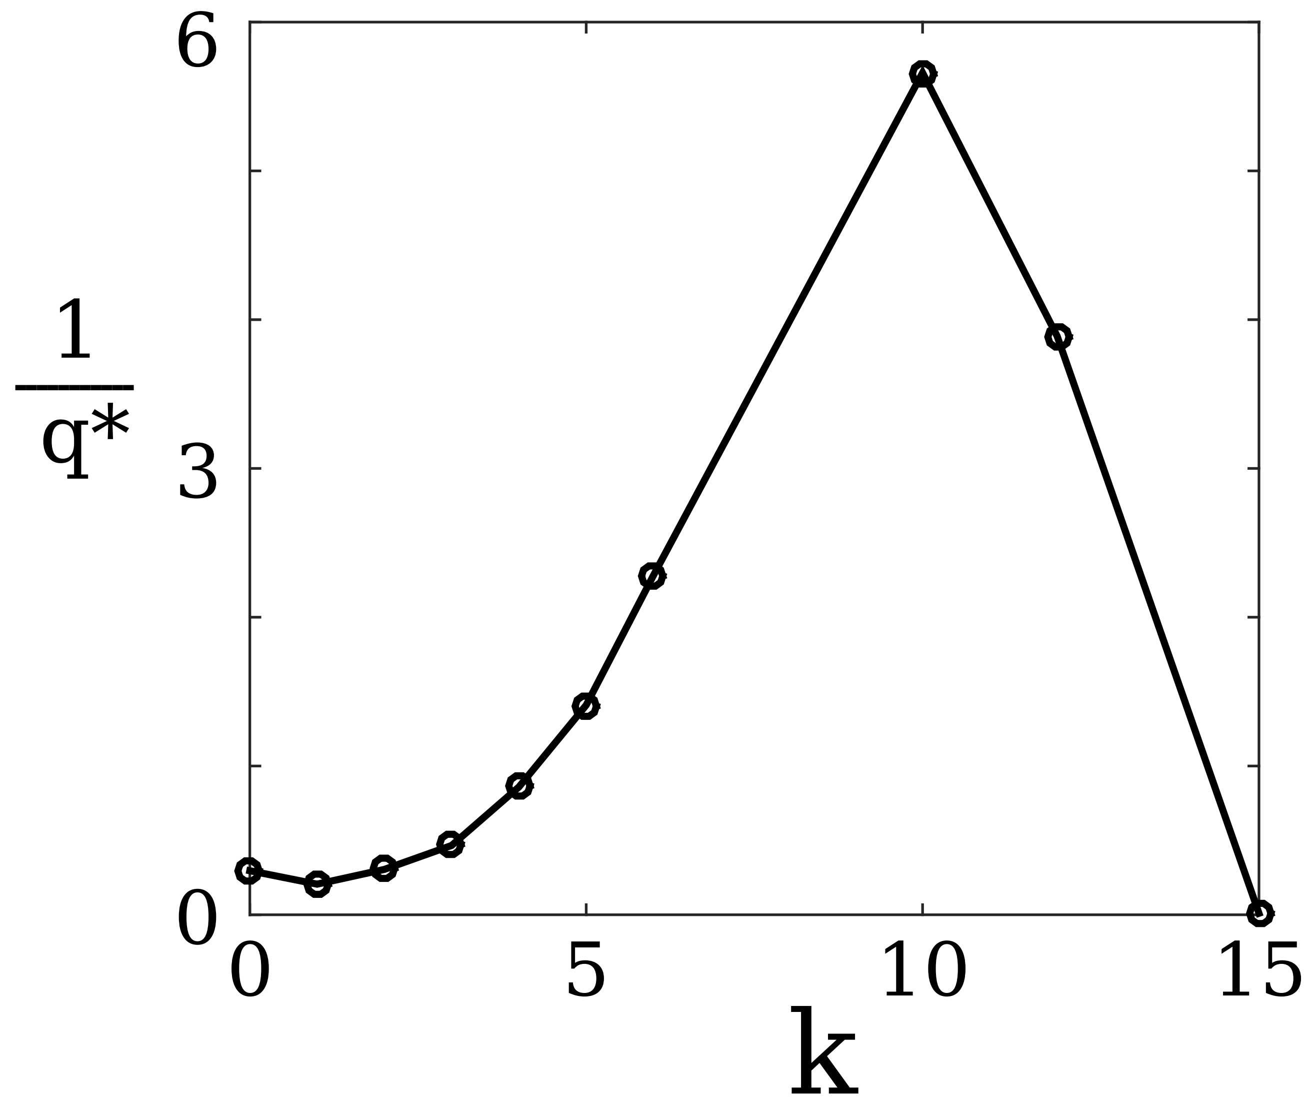

We generally obtain standing-wave patterns in 1D that look similar to those obtained in 2D

Stimulus–free

Stimulus on

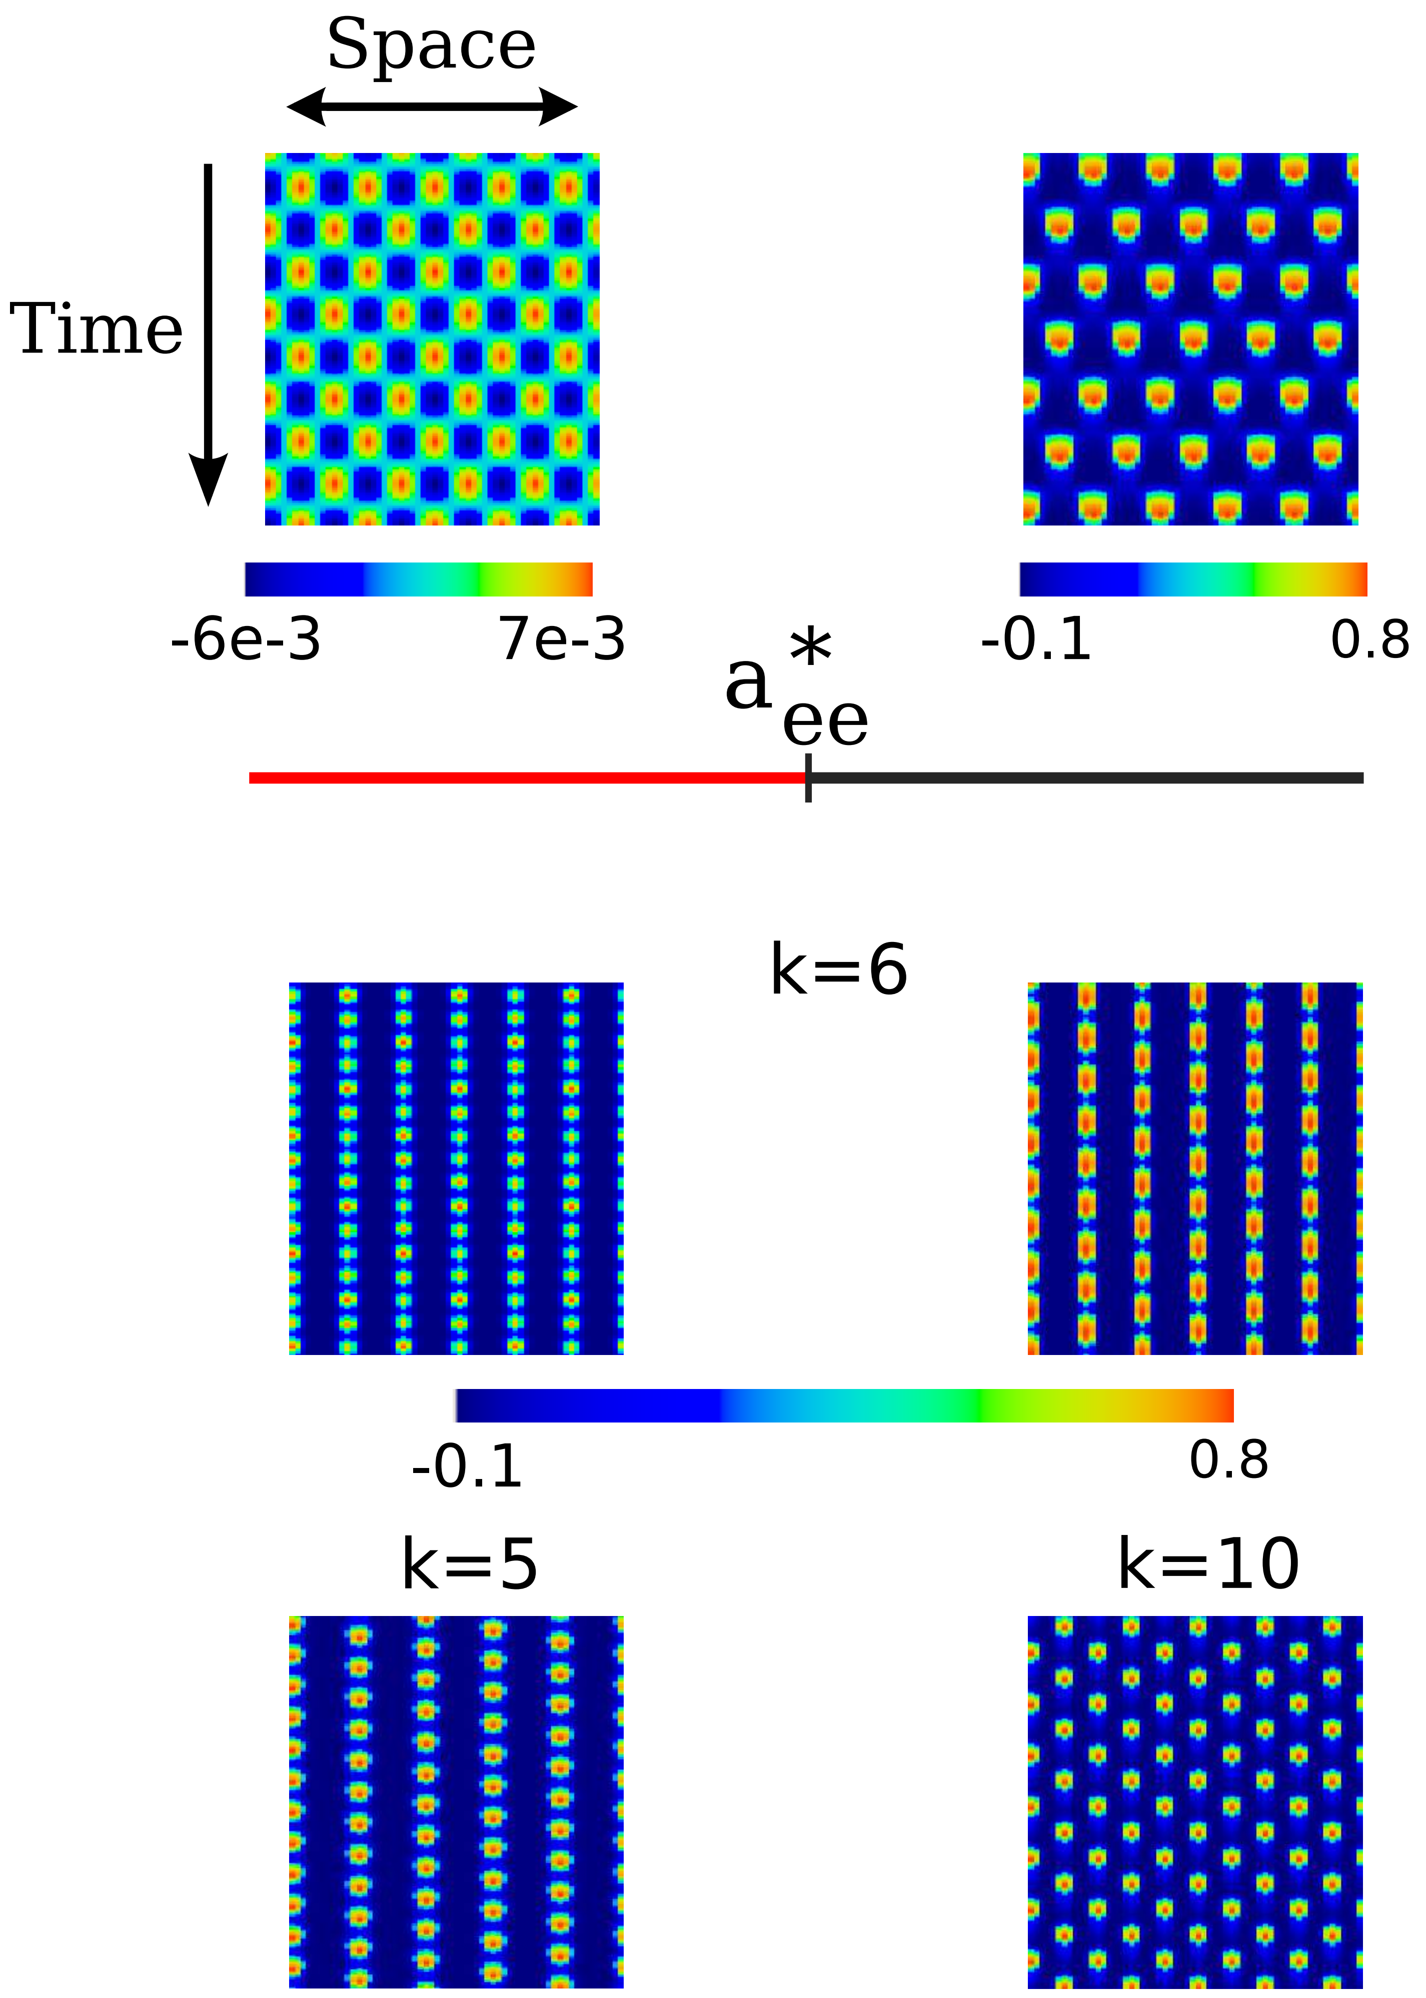

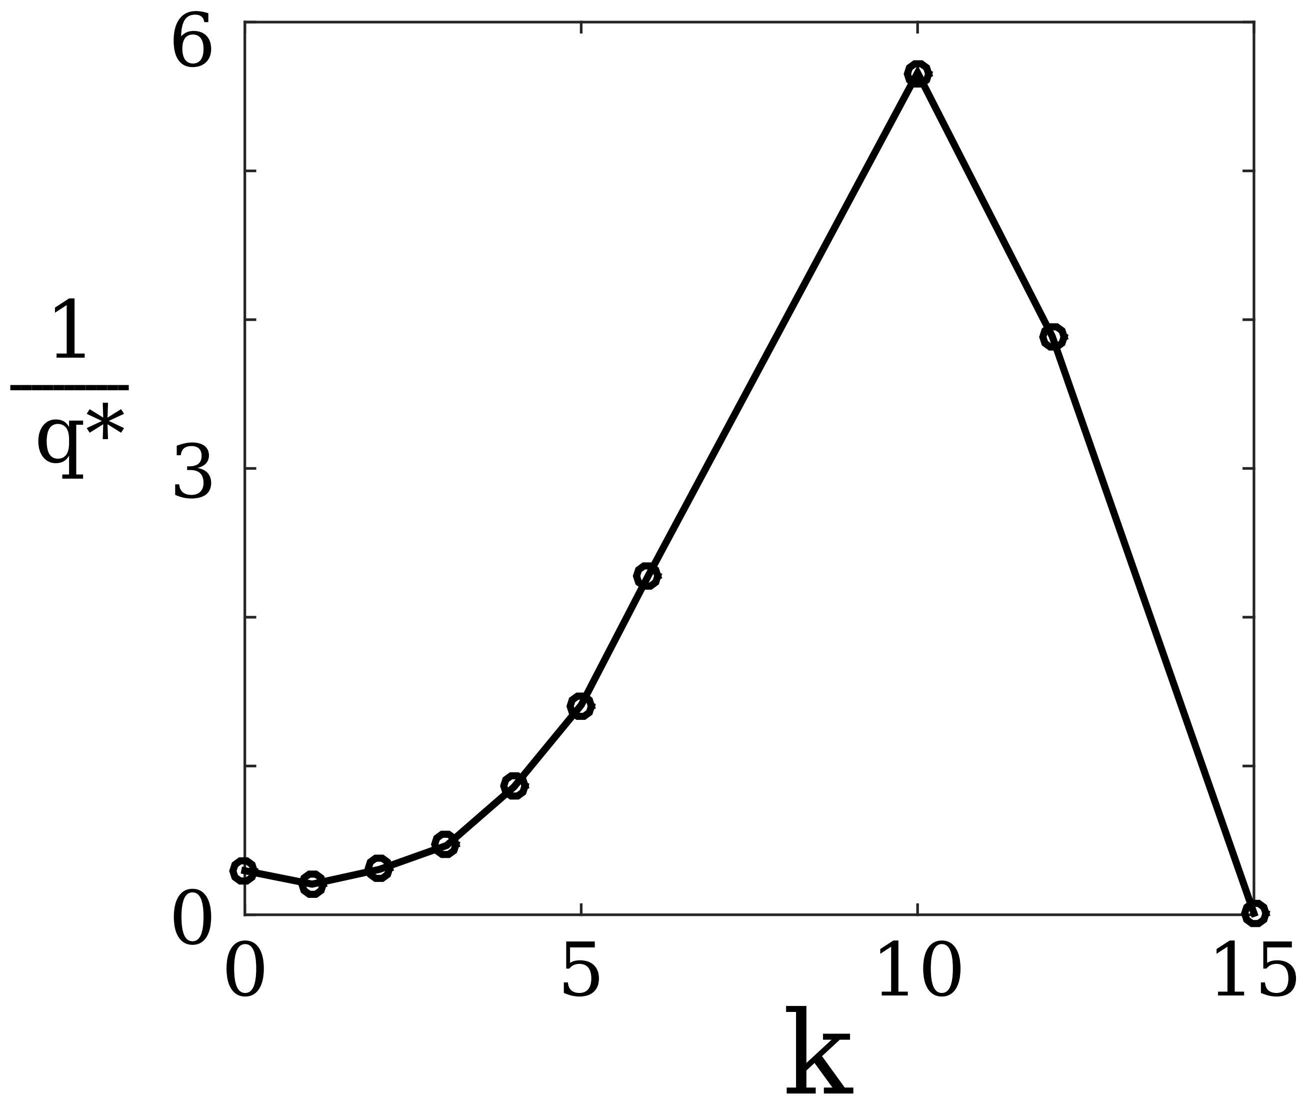

Natural spatial frequency = 5, resonant spatial frequency = 10 due to alternating pattern

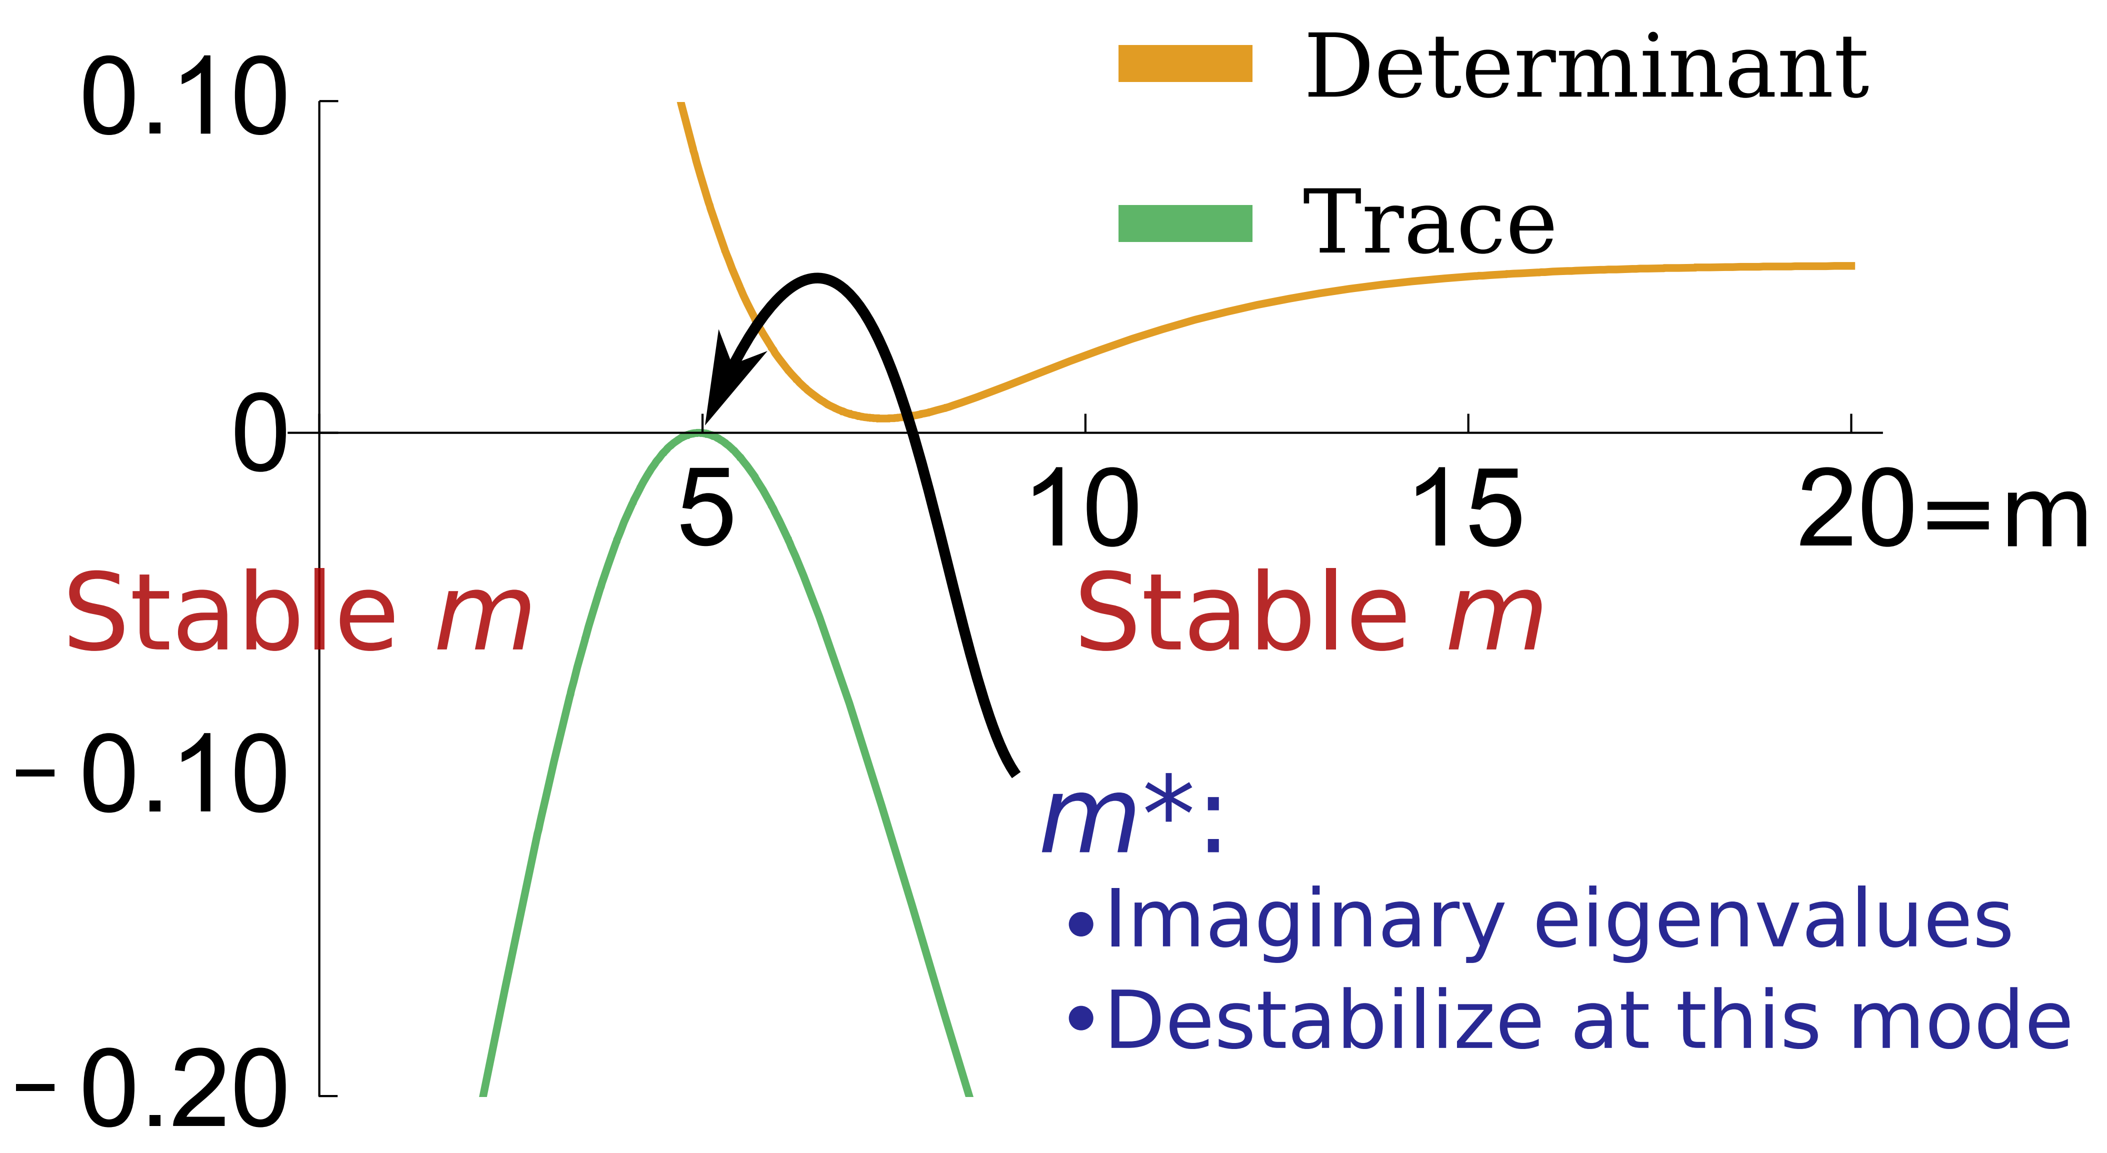

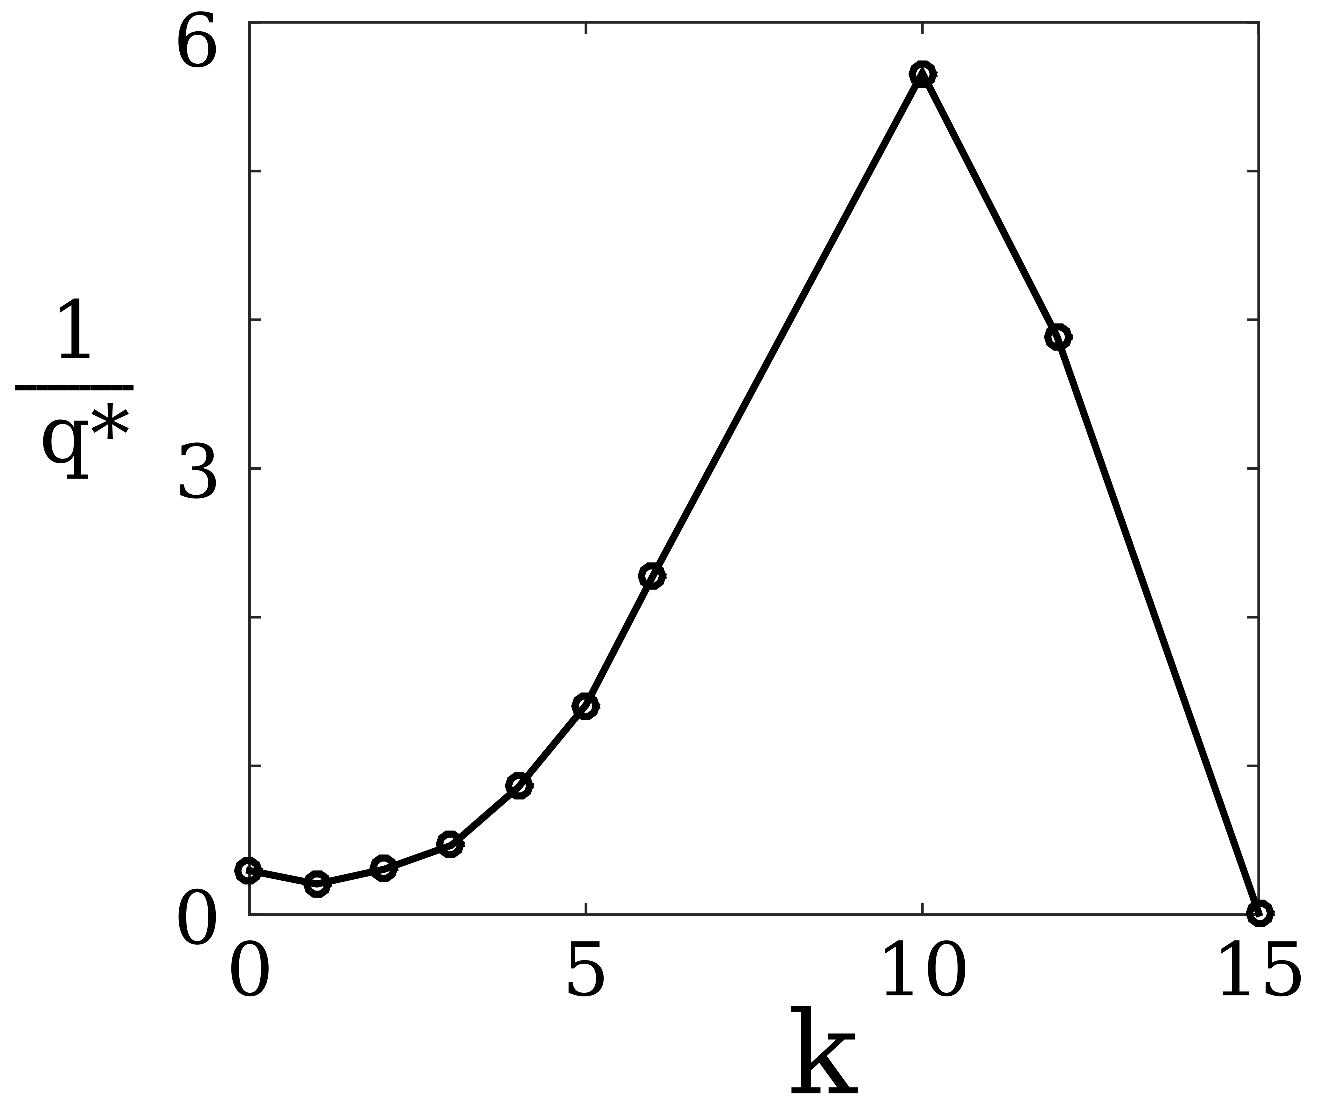

k=5

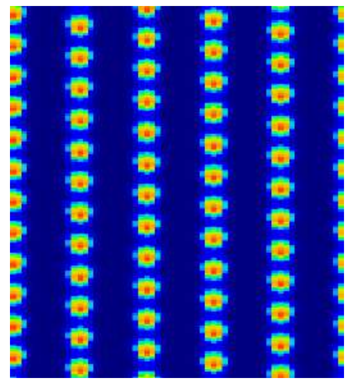

k=10

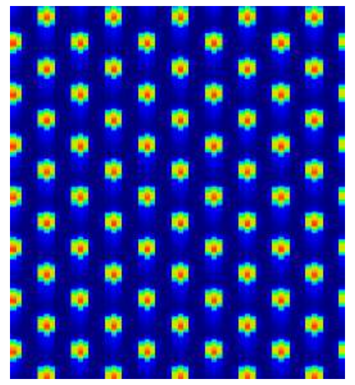

In 1D, we can explore further by producing 2-parameter bifurcation diagrams near the dynamic instability

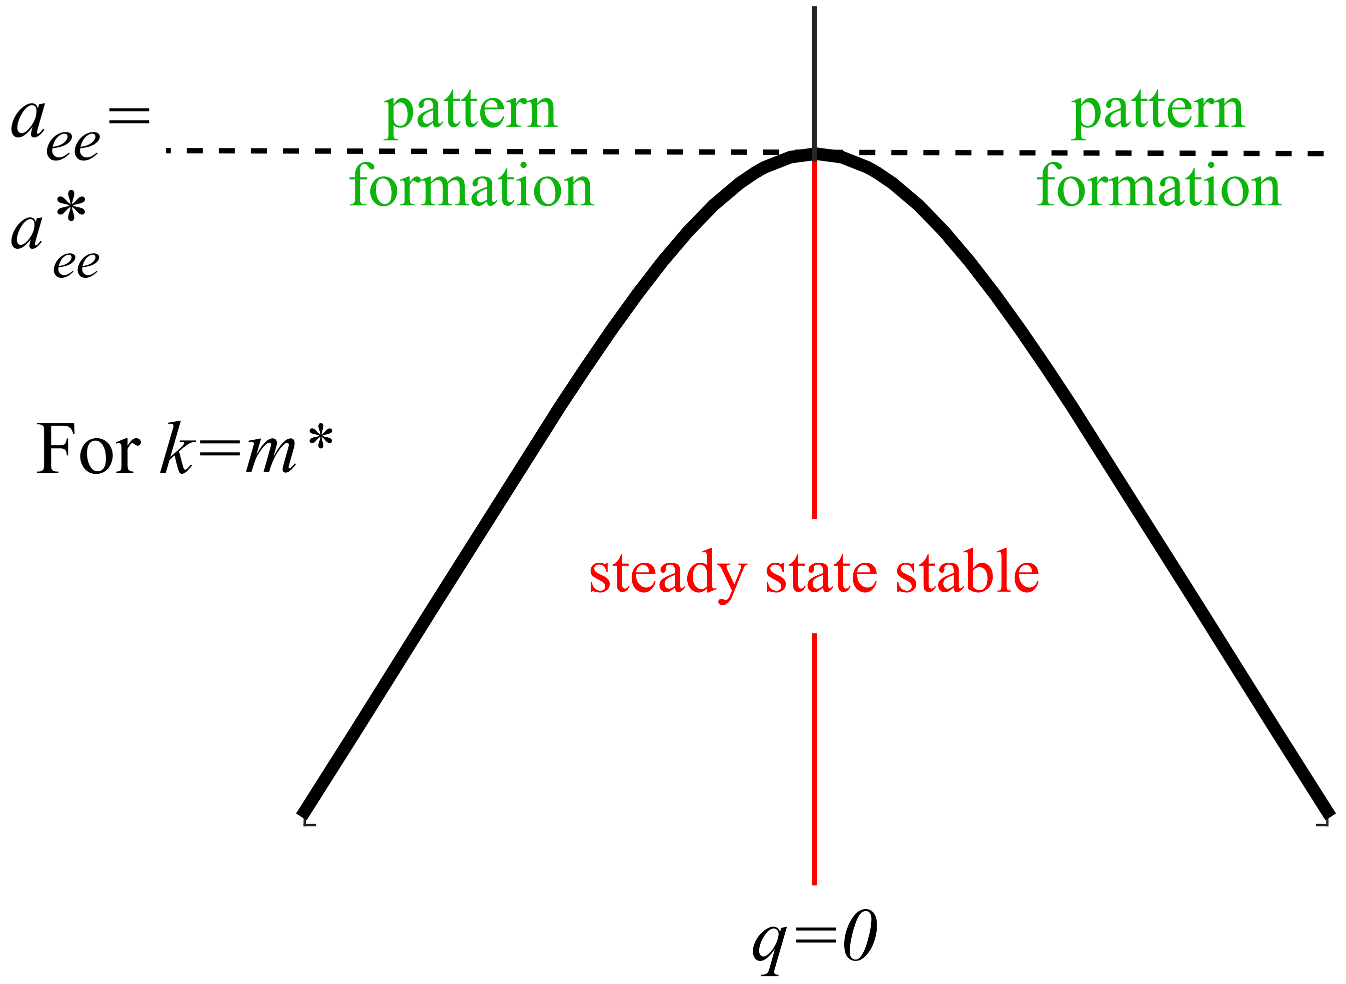

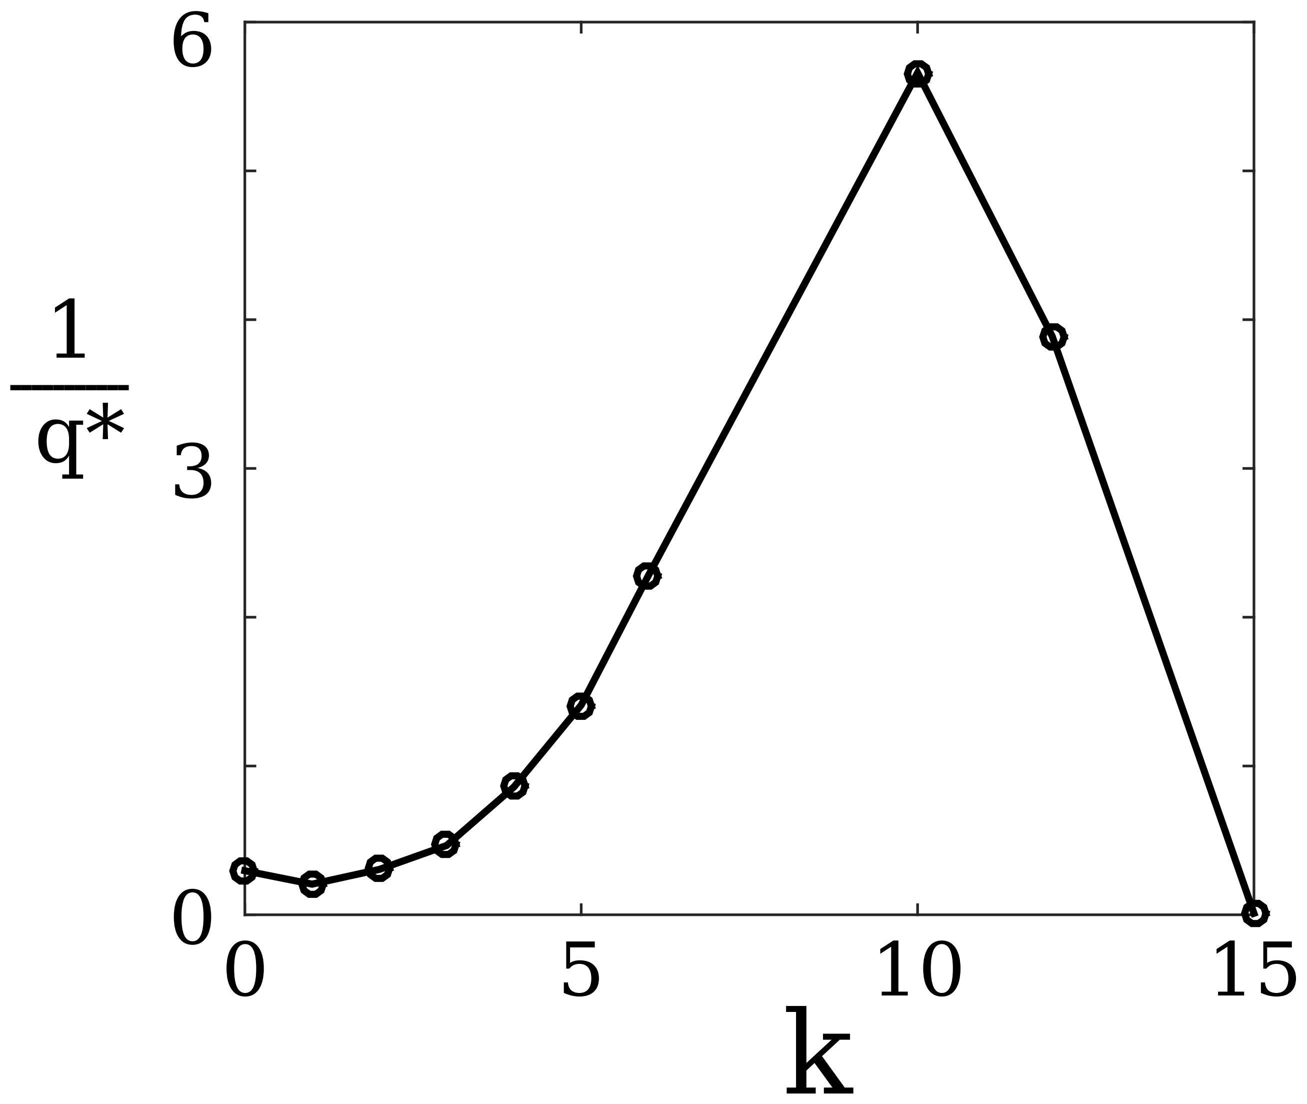

k=5

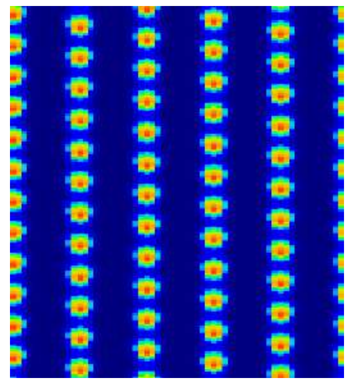

k=10

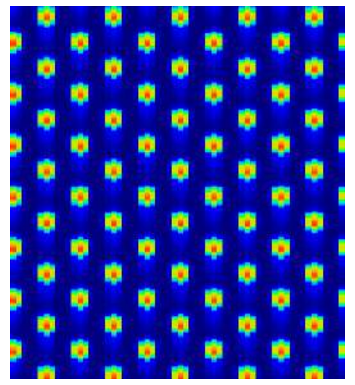

In 1D, we can explore further by producing 2-parameter bifurcation diagrams near the dynamic instability

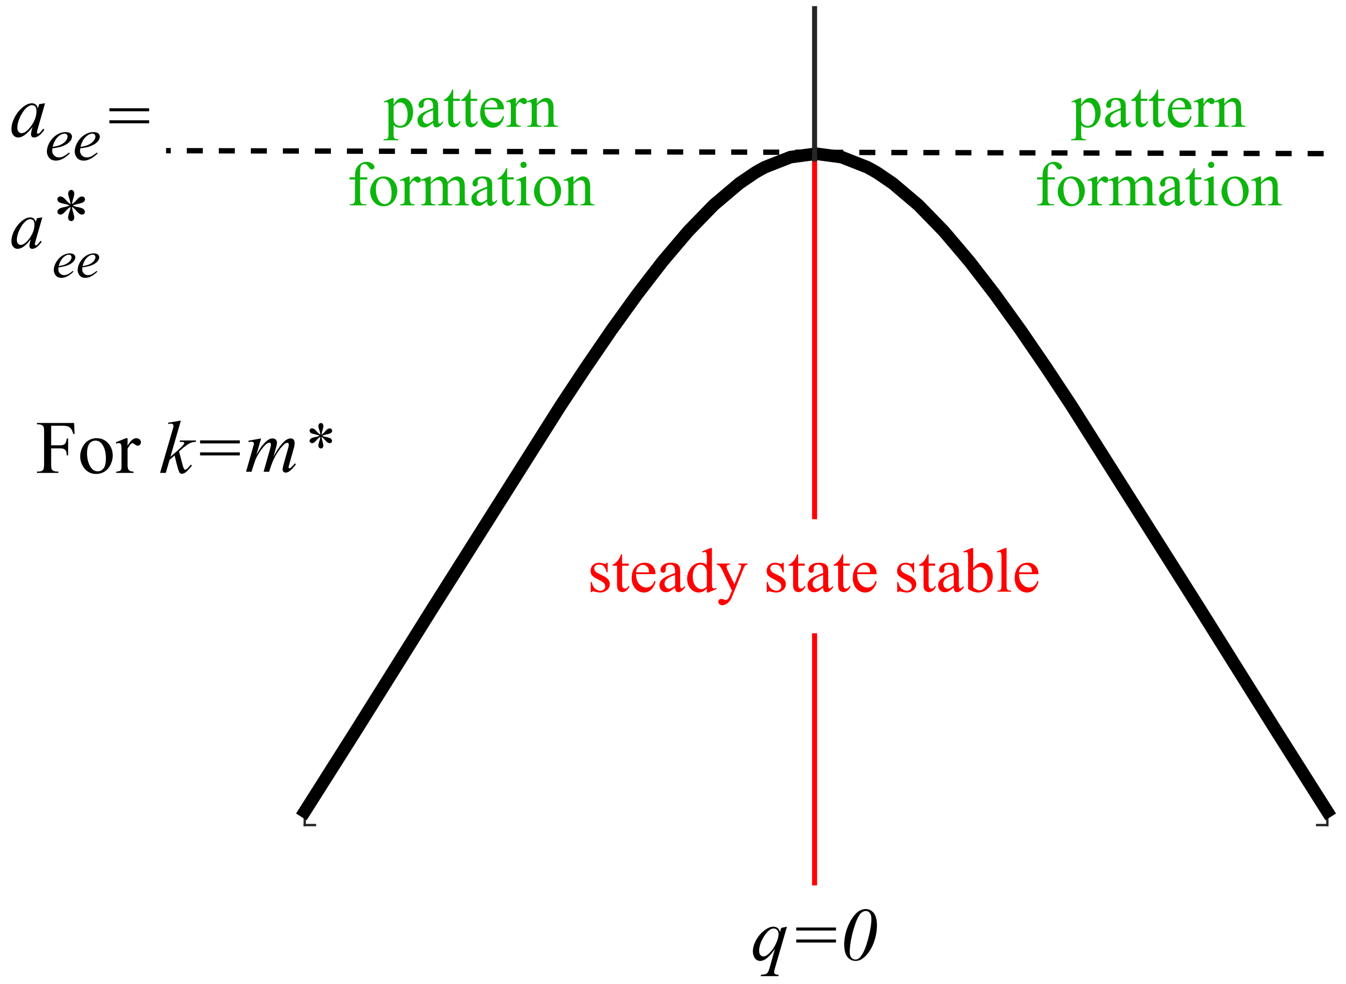

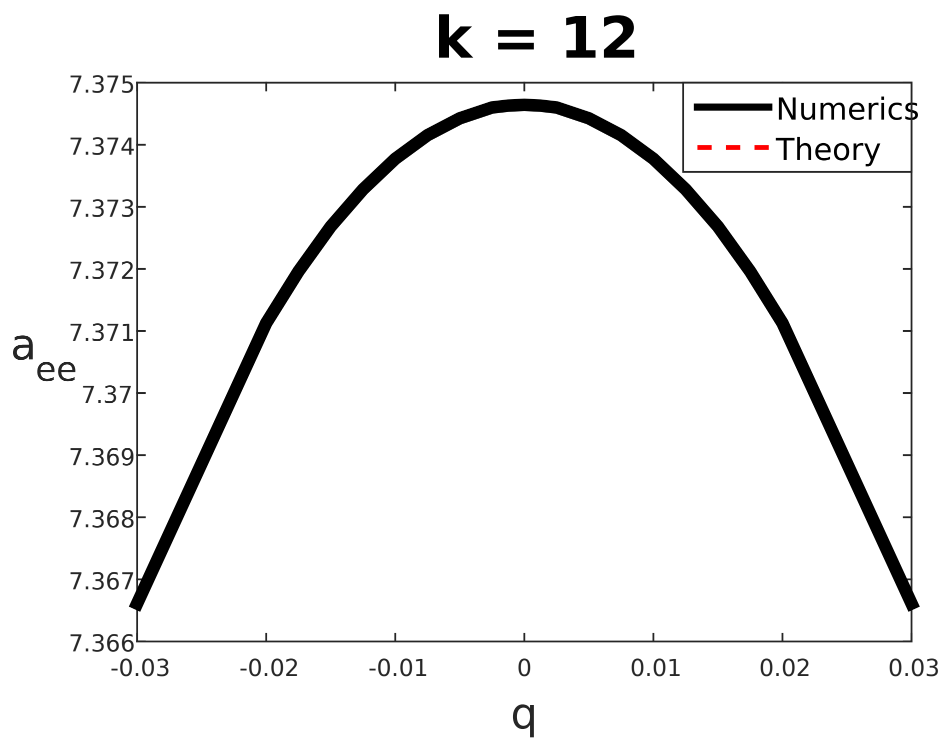

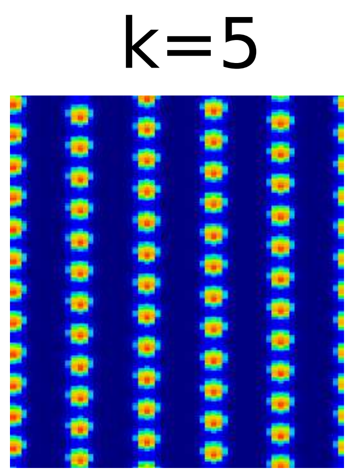

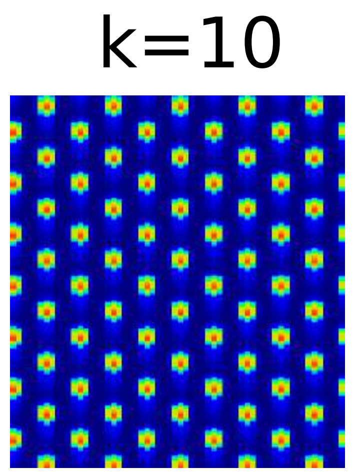

In 1D, we can explore further by producing 2-parameter bifurcation diagrams near the dynamic instability

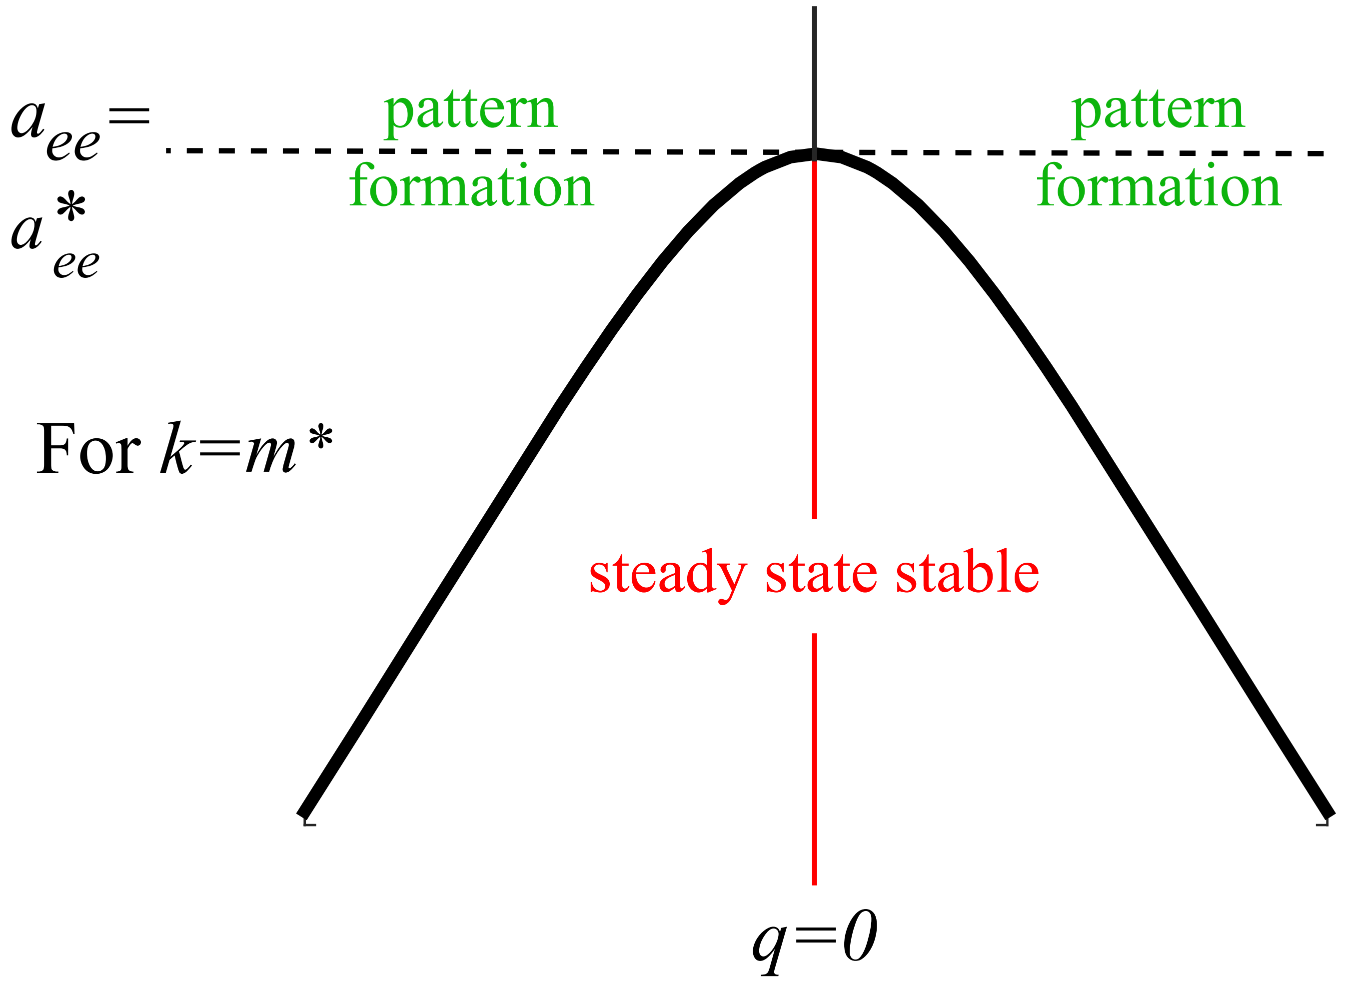

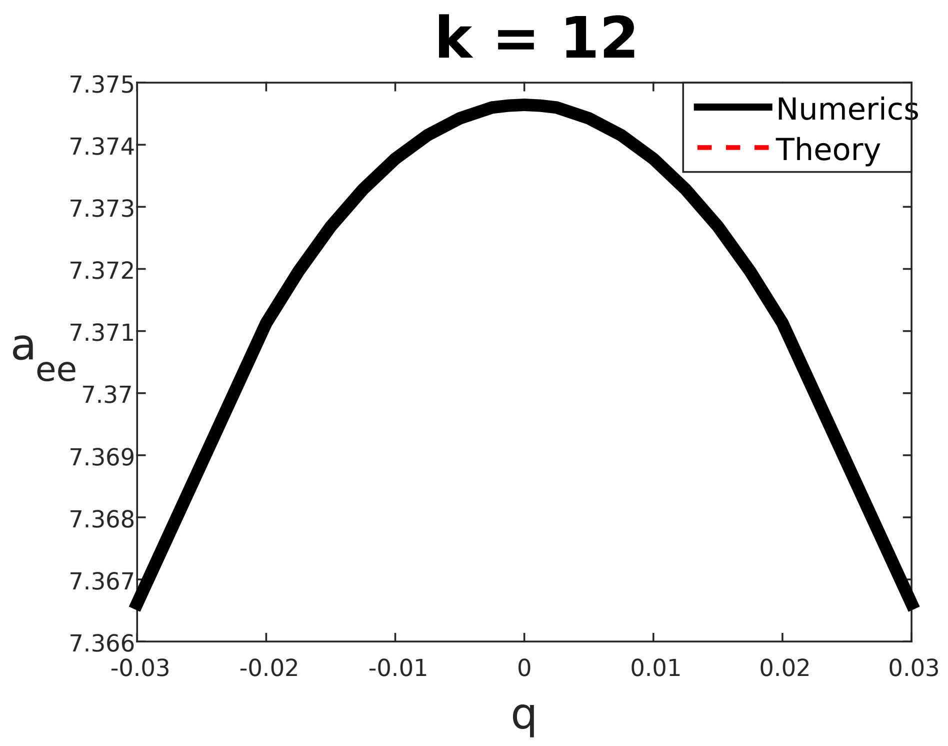

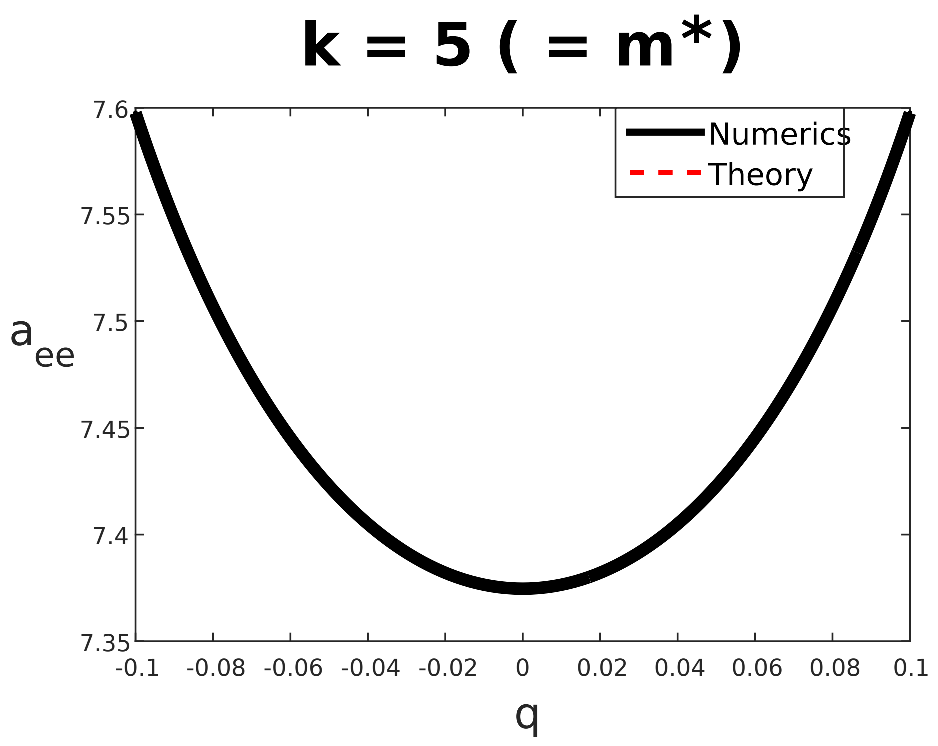

k=5

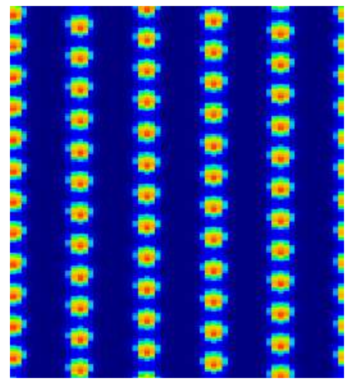

k=10

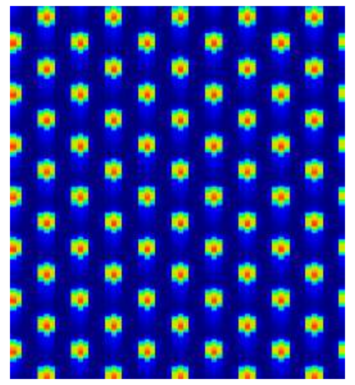

Curves are quadratic for k≠10, linear for k=10!

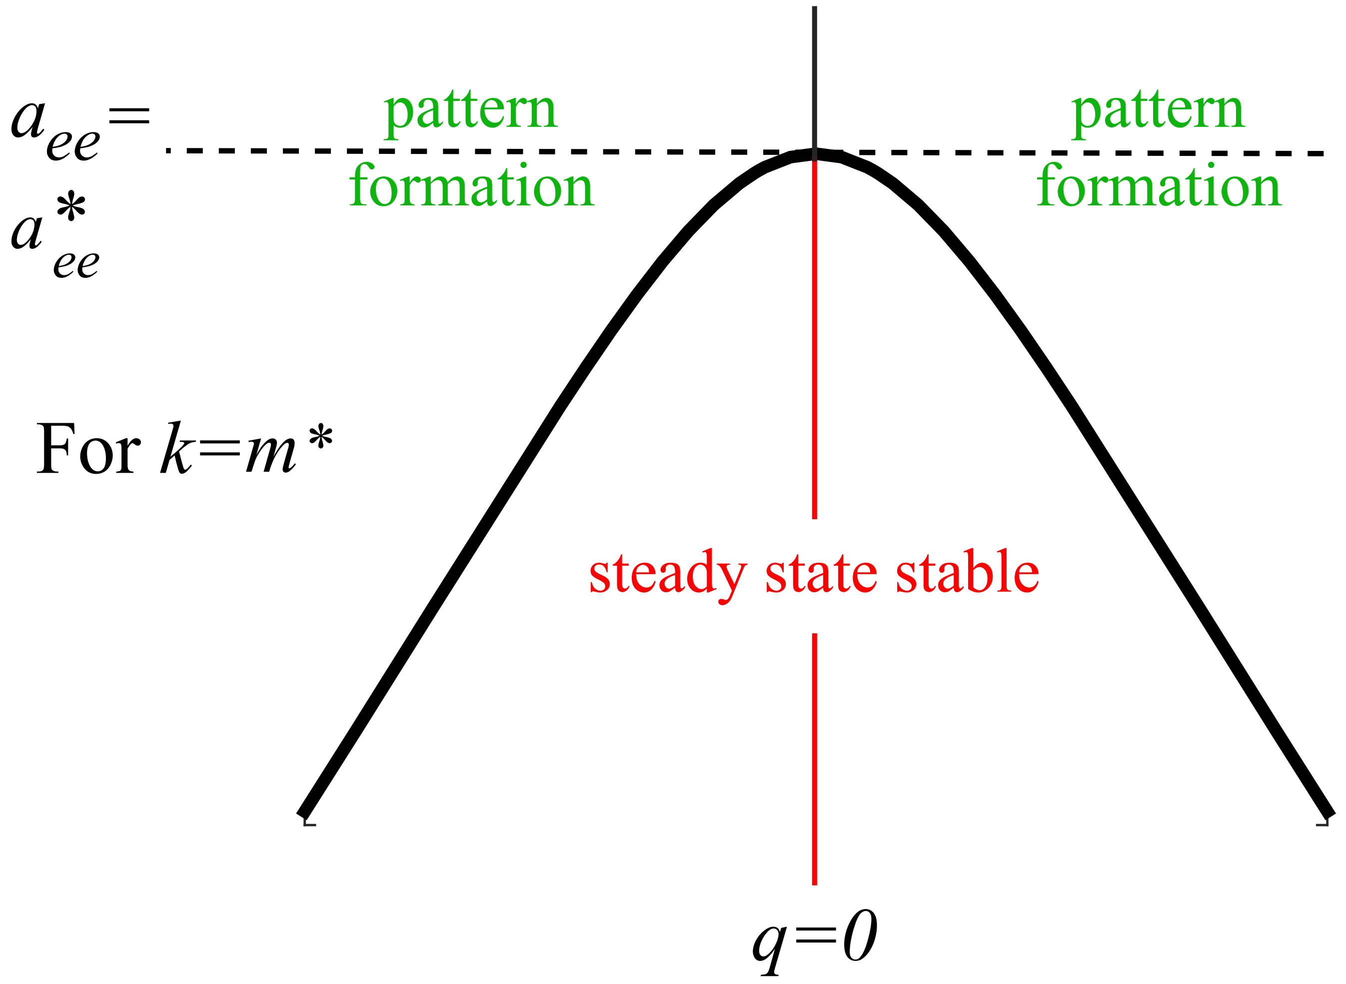

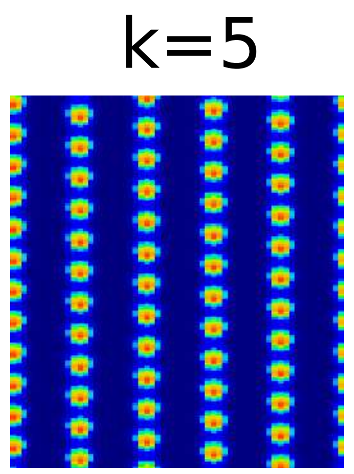

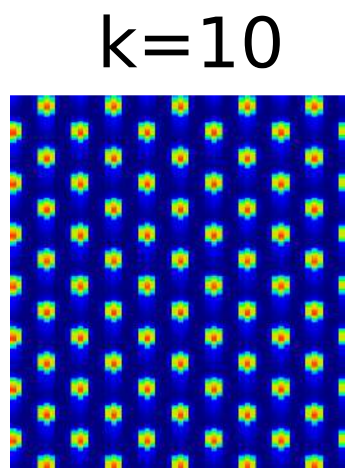

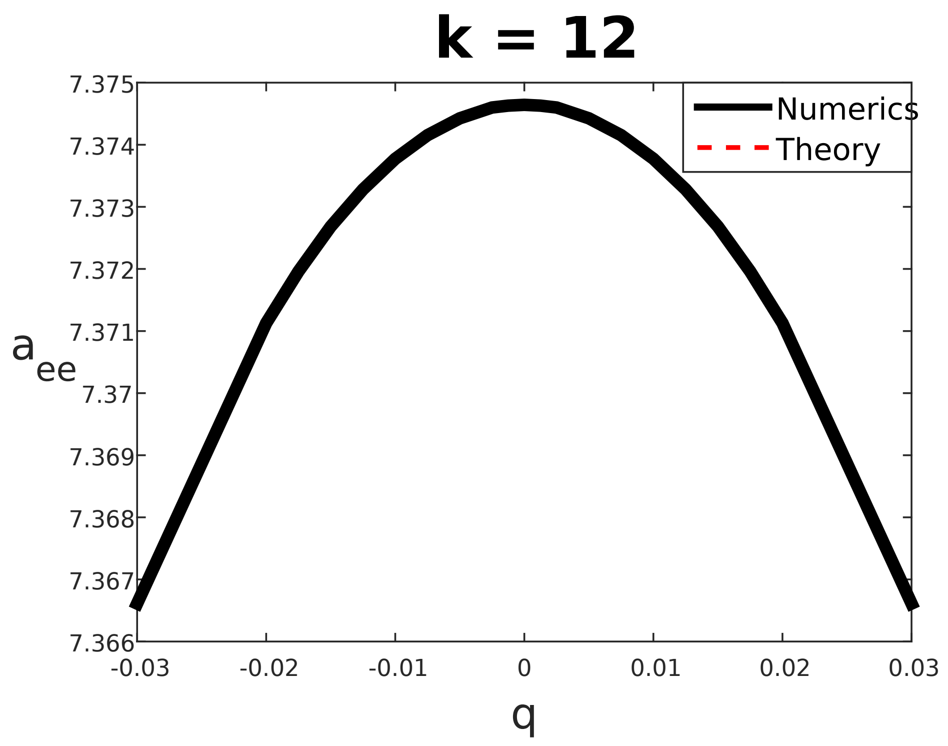

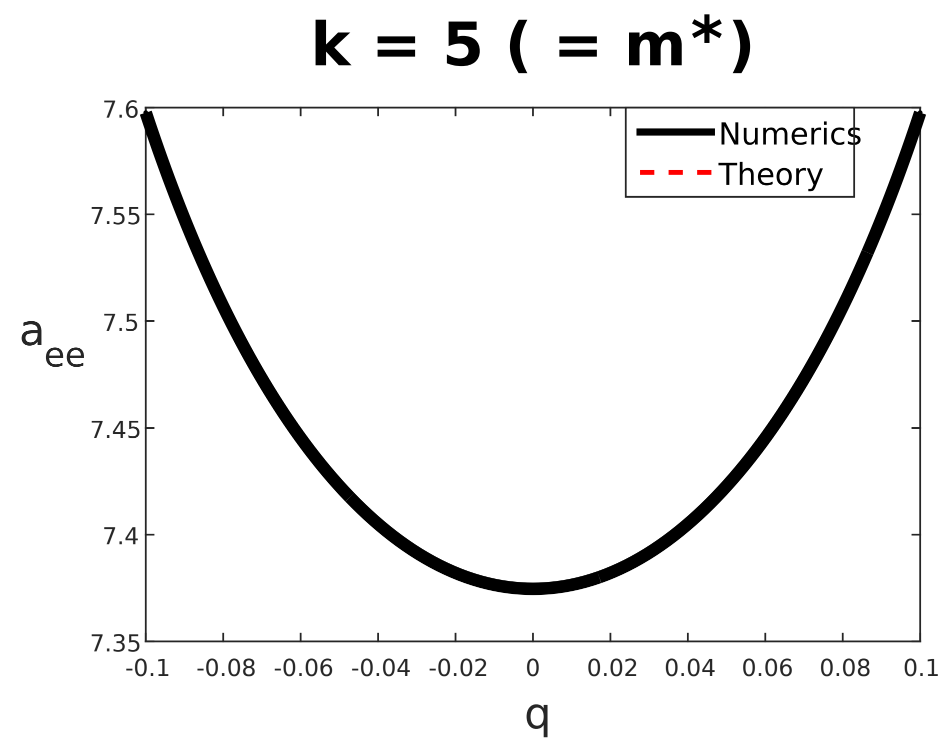

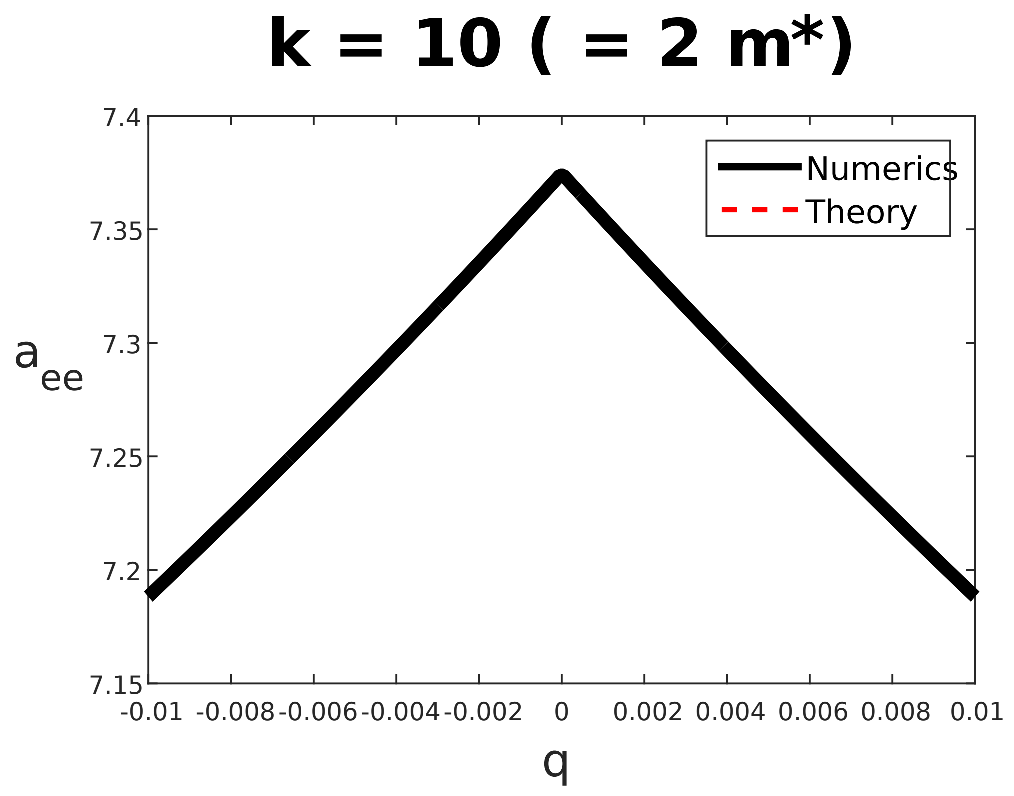

Linear stability analysis for small q!

Theoretical curves very closely match numerically-computed curves near instability

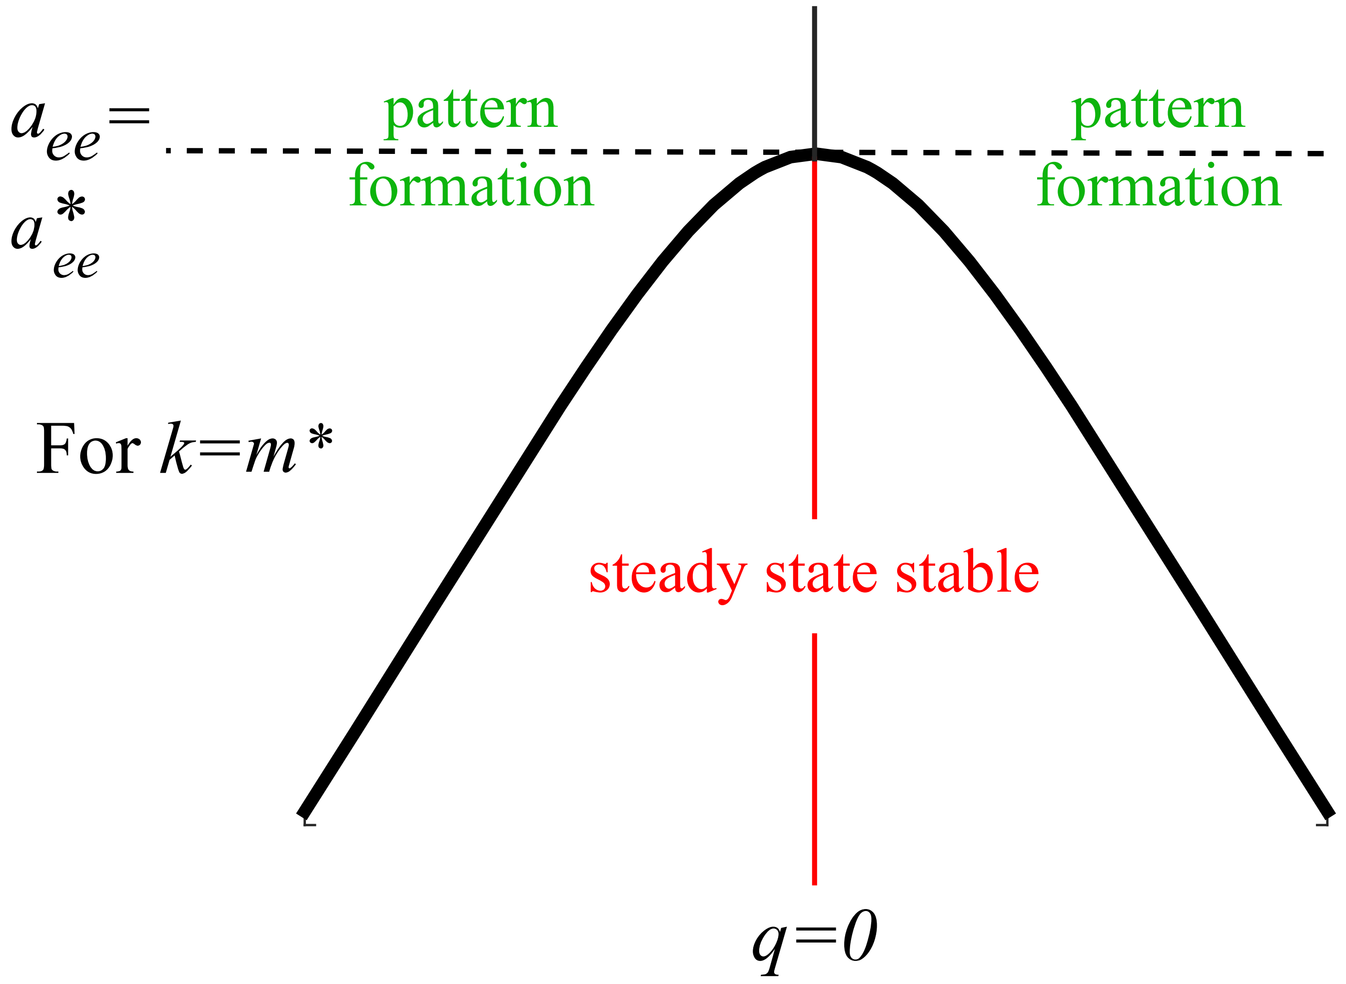

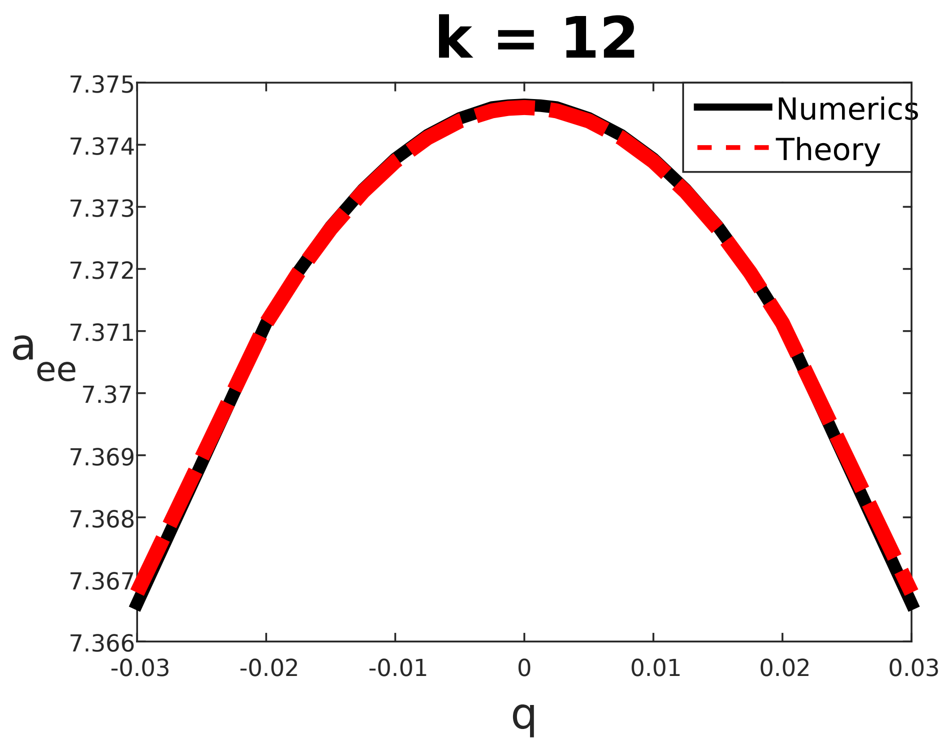

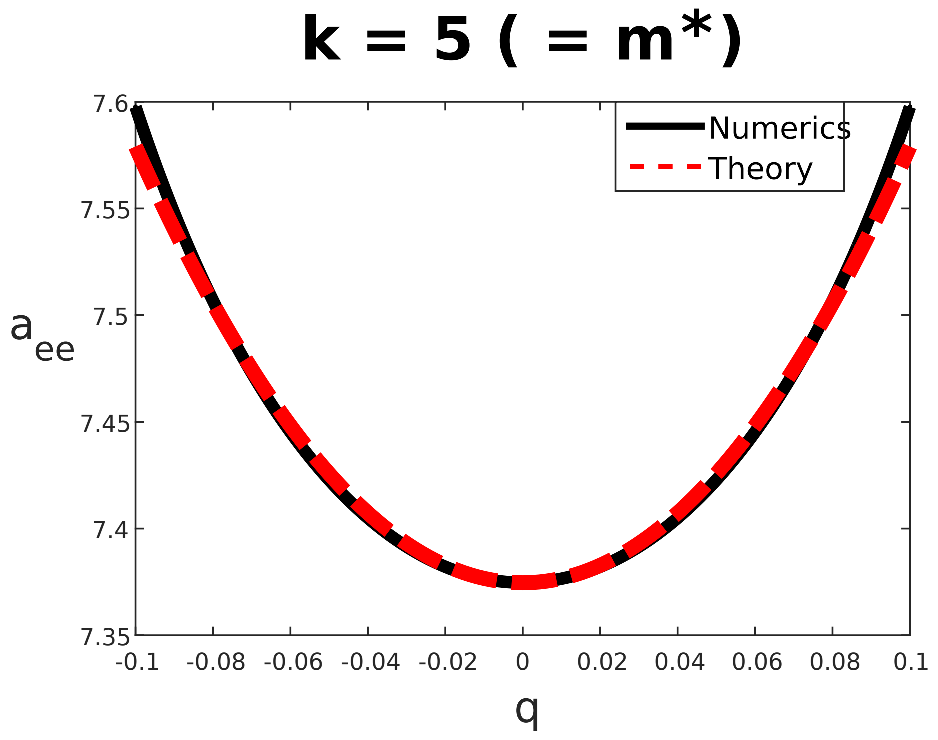

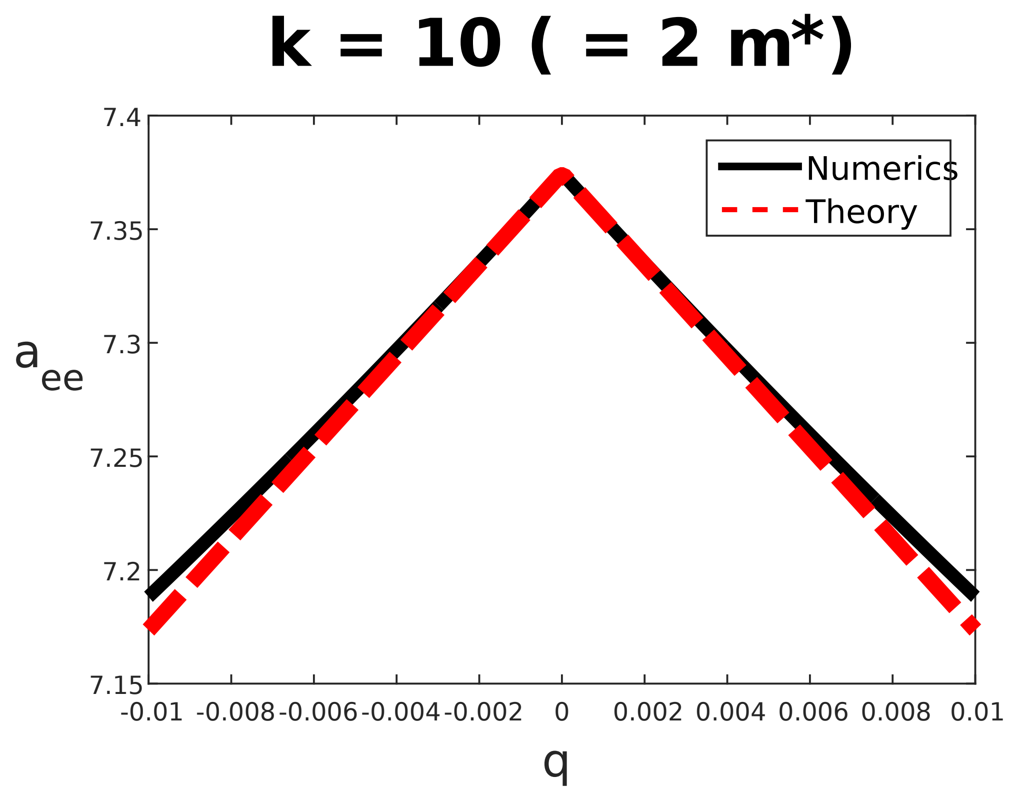

k=5

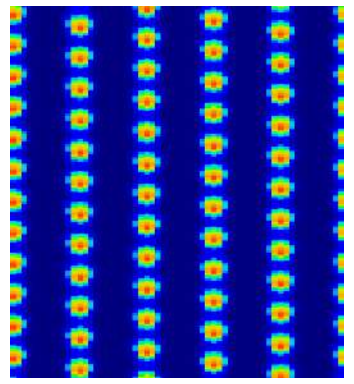

k=10

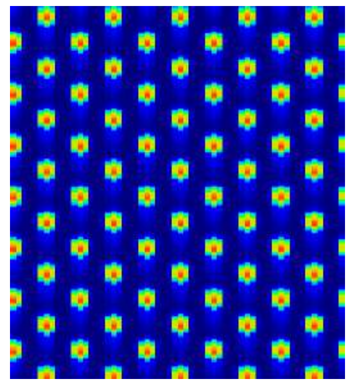

Since k=10 curve is linear, fits within quadratic curves. Hence, more sensitive near onset to k=10

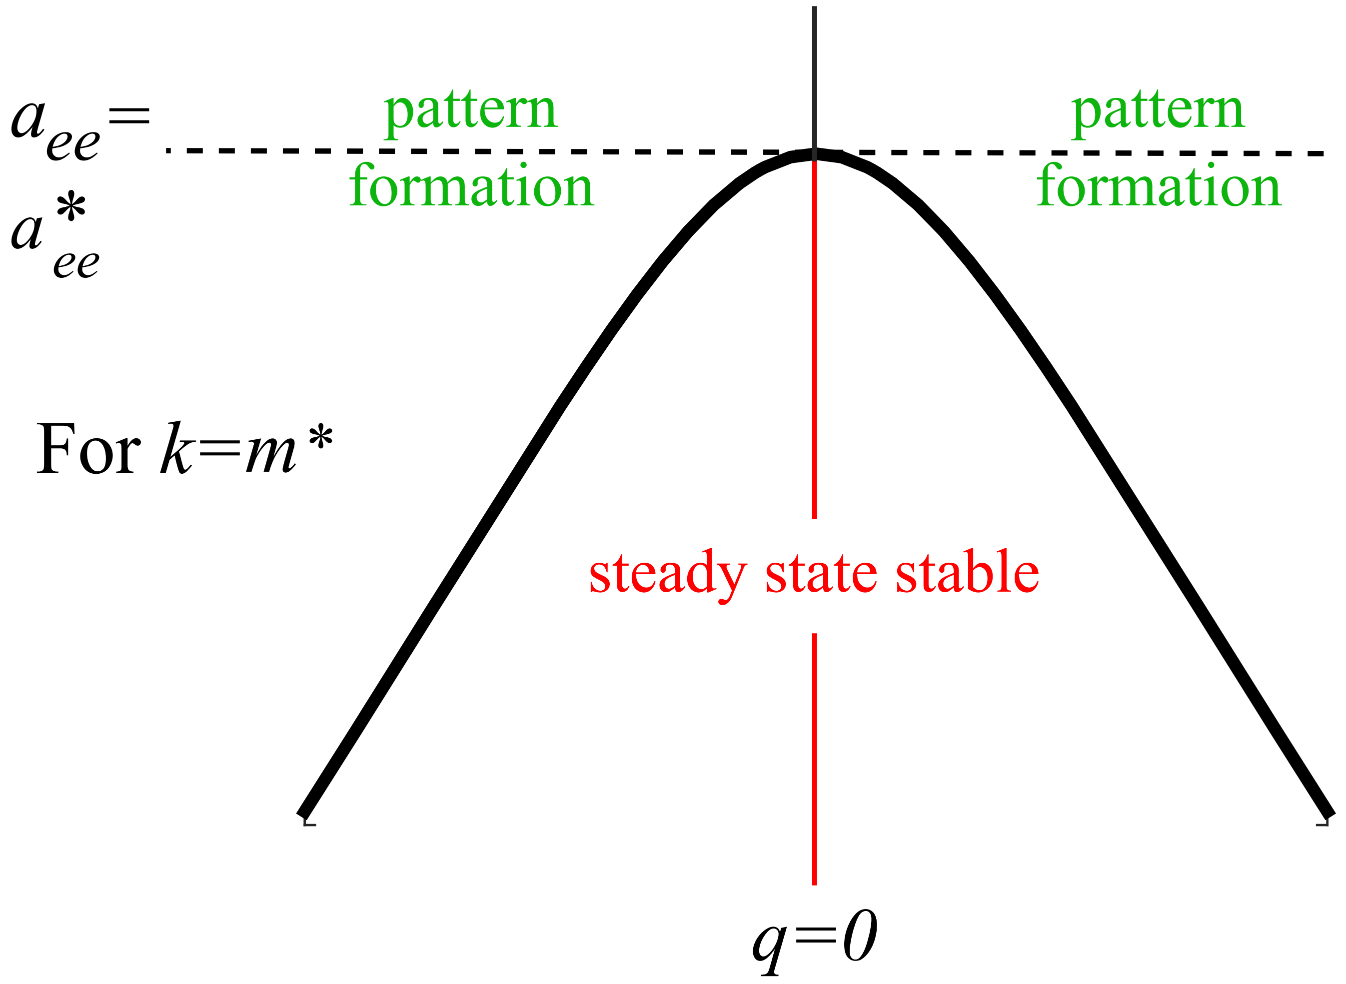

k=5

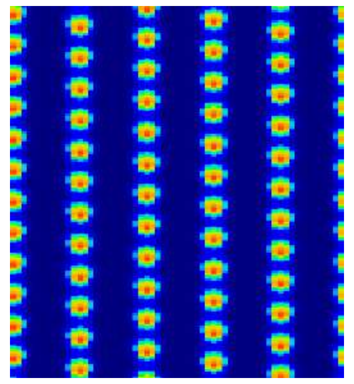

k=10

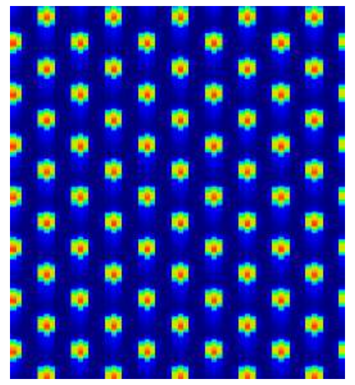

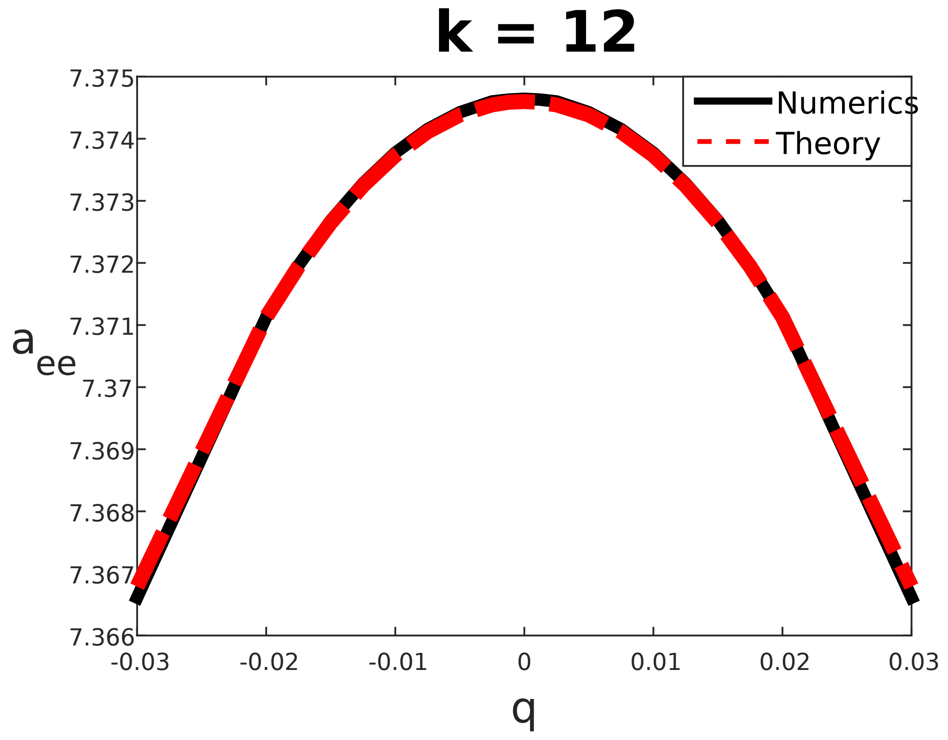

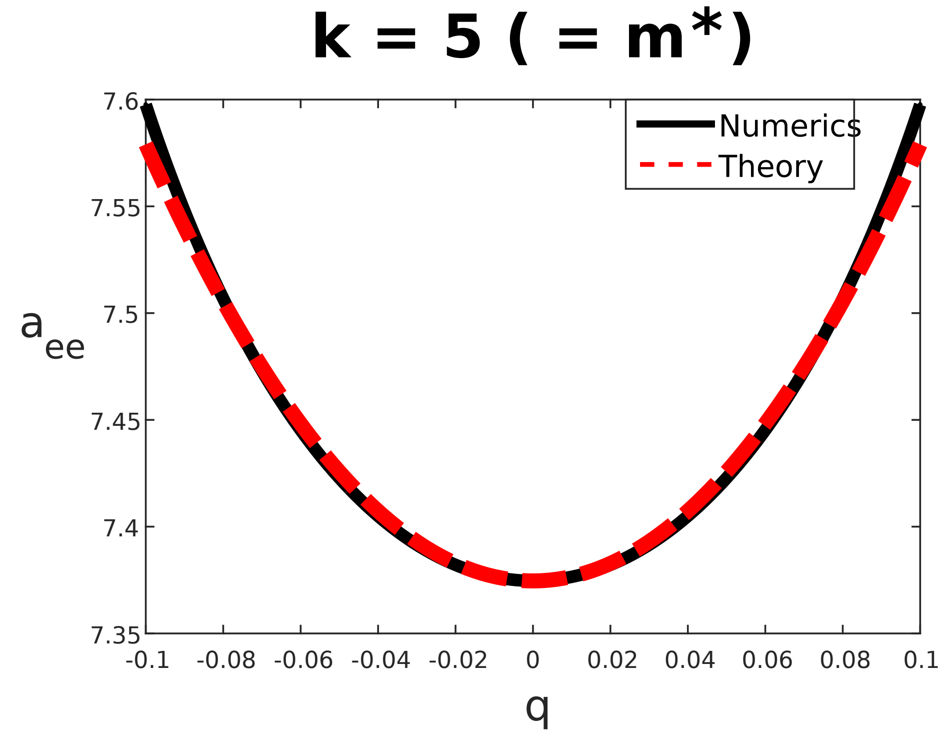

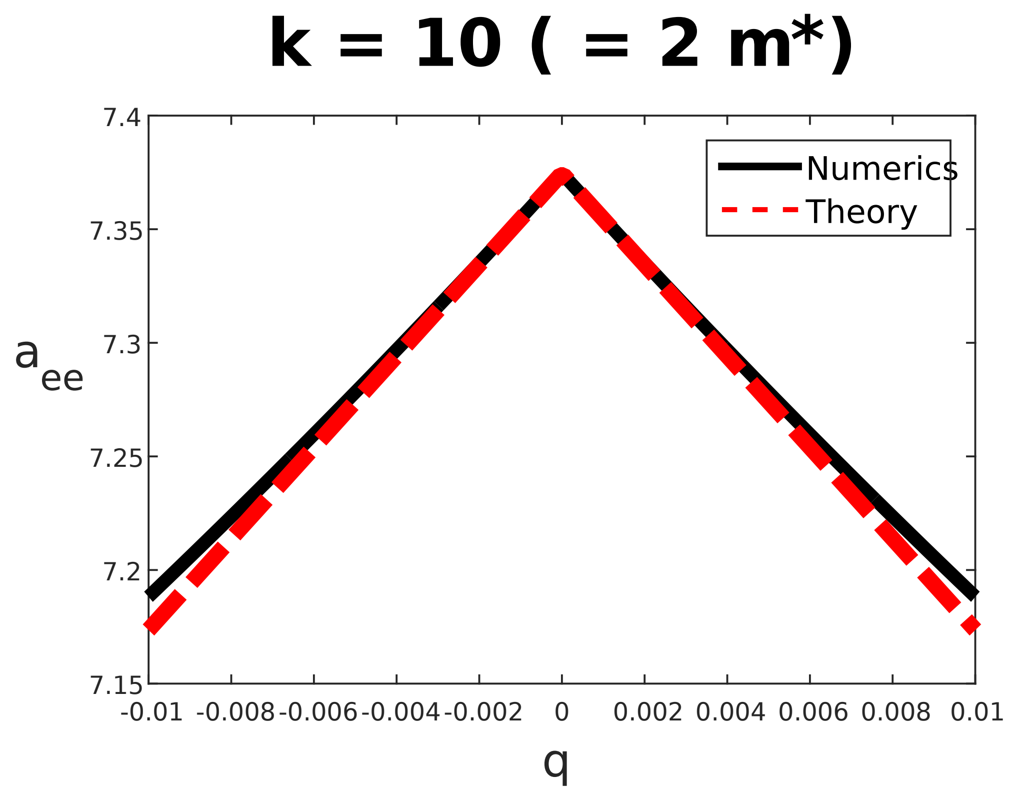

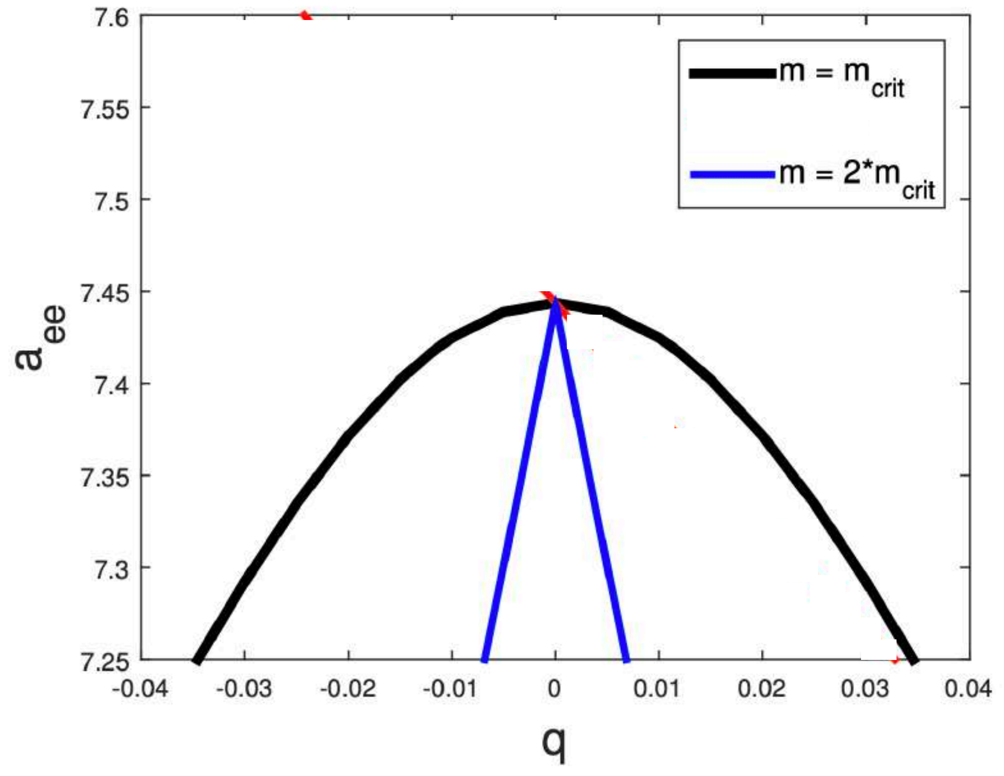

Pattern formation summary

- Spatial resonances: oscillatory neural responses to static images with dominant frequencies in a narrow band
- Spatially extended neural field model captures resonance when placed near Turing-Hopf bifurcation
- 2-D and 1-D networks show similar behaviors
  - Both: resonances near those found psychophysically
- Mathematically show that network more sensitive to stimuli with twice the natural frequency

Does the visual system implement a deep network?

One class of deep network models that the brain is hypothesized to implement are *predictive hierarchical models*

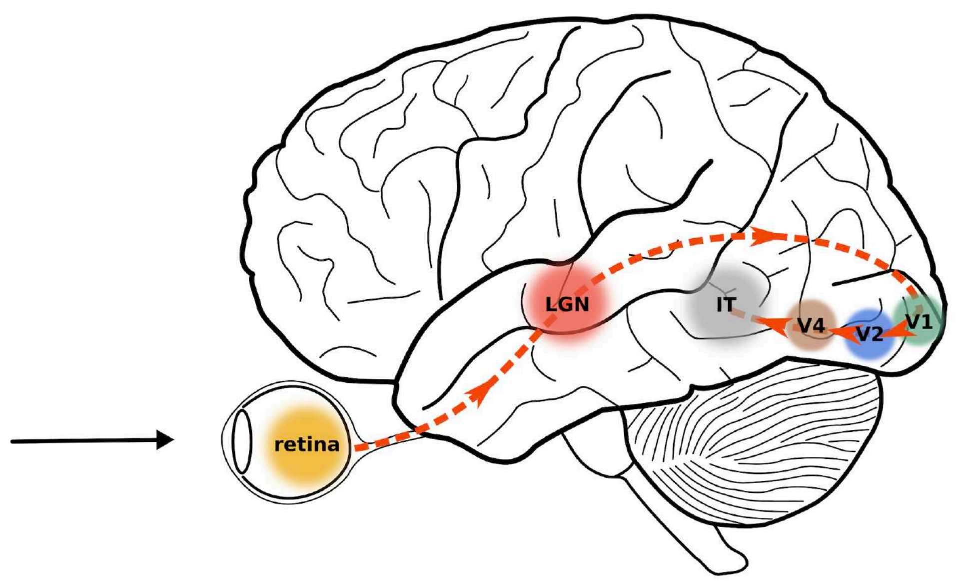

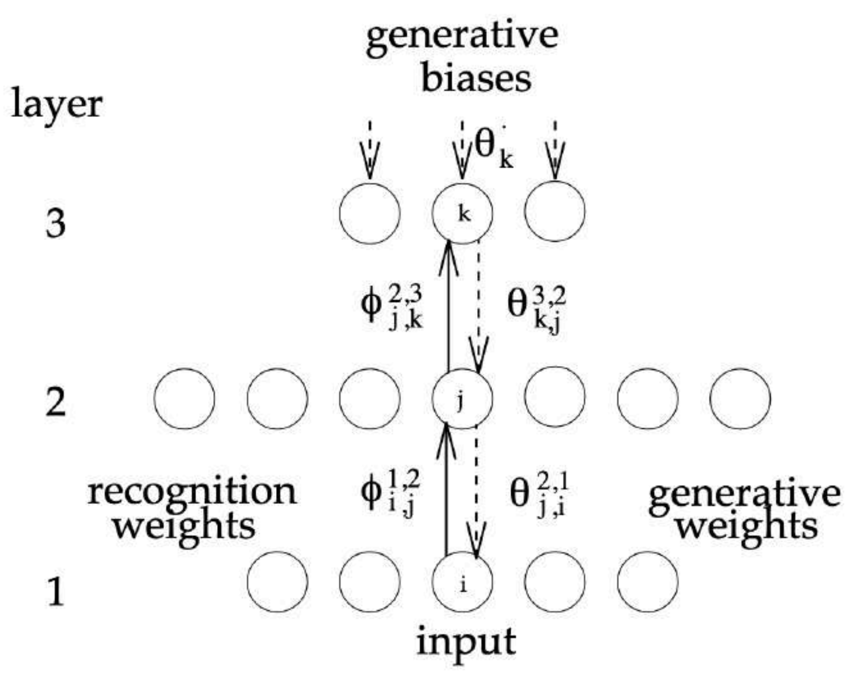

Helmholtz machine schematic[1]

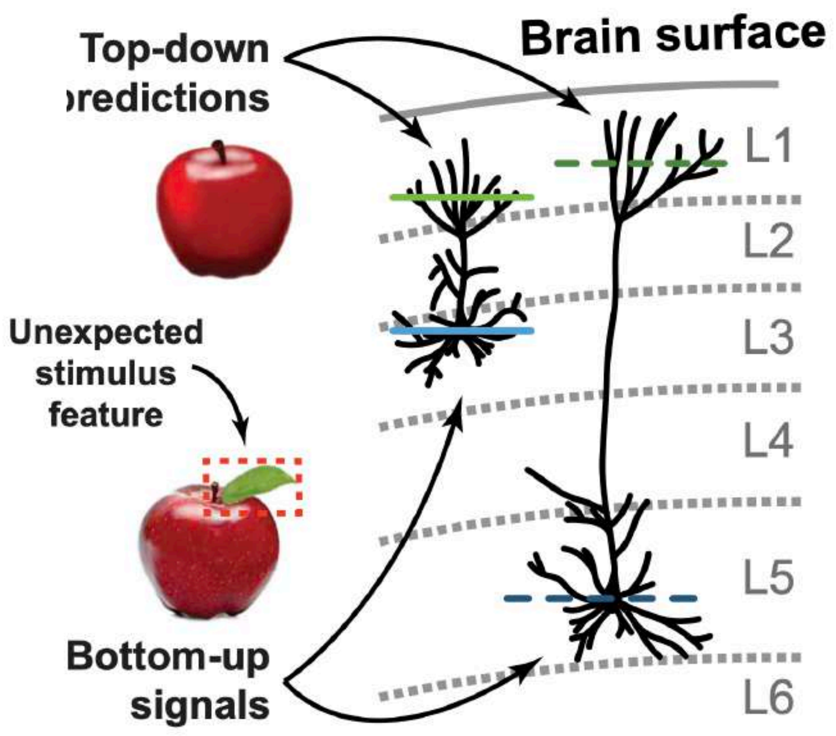

Preprint:

THE PREPRINT SERVER FOR BIOLOGY

doi: https://doi.org/10.1101/2021.01.15.426915

[1]- Dayan … Zemel, *Neur. Comp.*, 1995

[2] - Rao & Ballard, *Nat. Neur.*, 1999

Does the visual system implement a predictive hierarchical model of the world?

- Predictive hierarchical models (e.g., Helmholtz machines[1], Rao & Ballard[2]) comprise a broad class of models of how the visual system is hypothesized to function. Briefly:
    - Higher brain areas make predictions about incoming stimuli based on prior experience (possibly evolutionary)
    - These predictions are compared to the incoming stimuli
    - The predictions, comprising the internal model of the world, are updated based on these comparisons
        - i.e., differences between predictions and stimuli drive learning

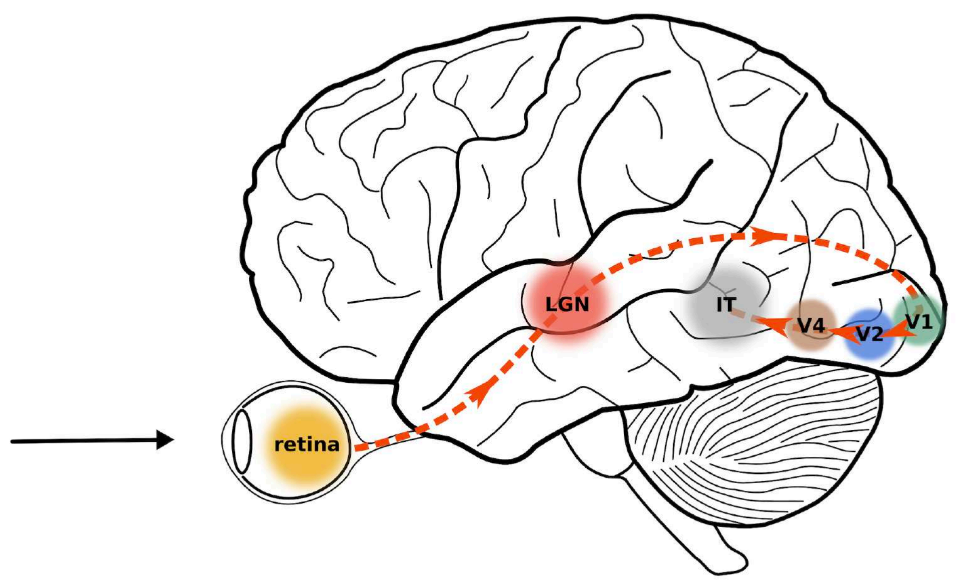

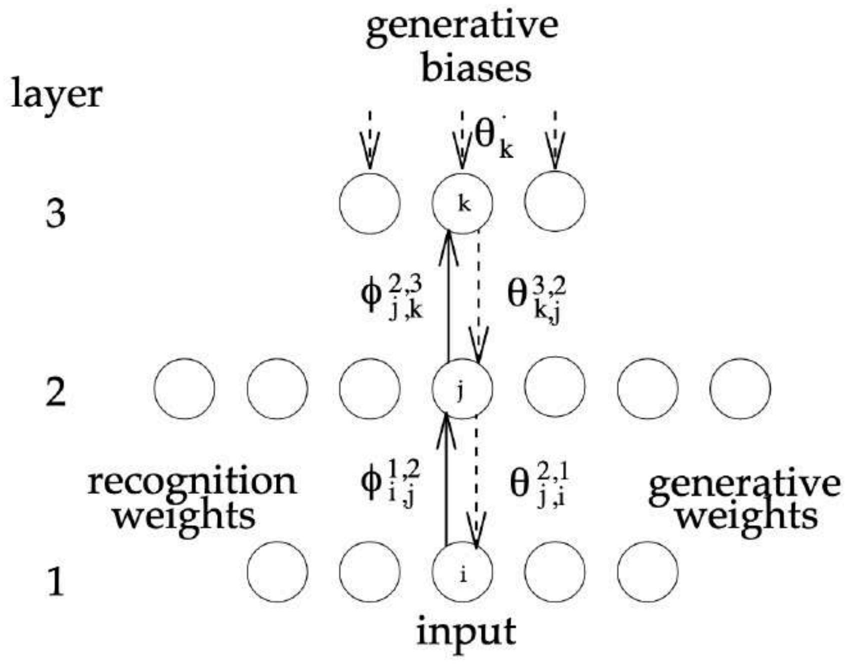

Helmholtz machine schematic[1]

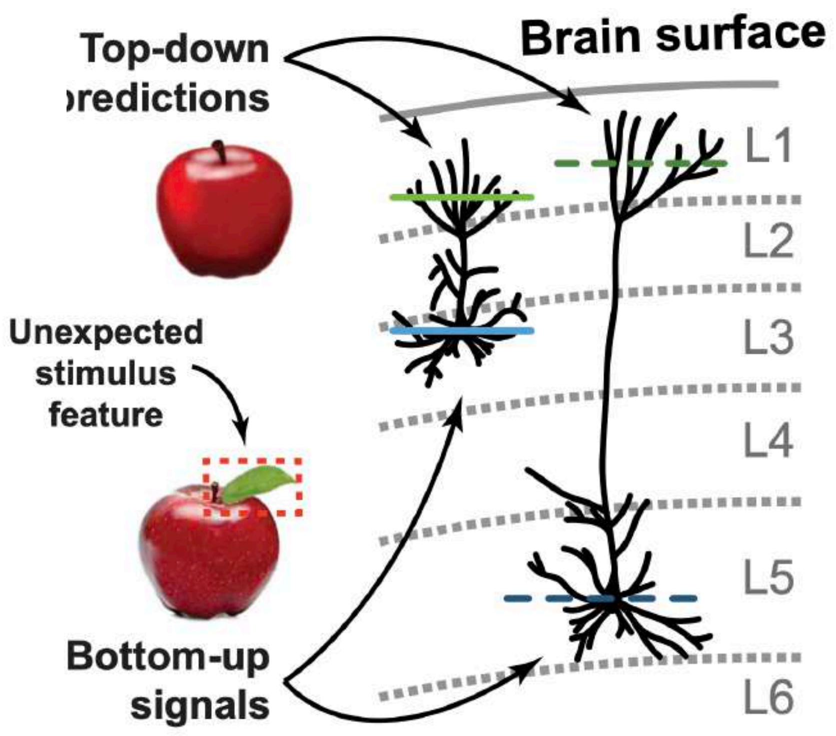

[1] - Dayan … Zemel, *Neur. Comp.*, 1995

[2] - Rao & Ballard, *Nat. Neur.*, 1999

Does the visual system implement a predictive hierarchical model of the world?

Logical consequents of such predictive hierarchical models:

1. There should be distinct responses to expected and unexpected stimuli
2. These responses should change with experience
3. Top-down and bottom-up responses should evolve differently due to hierarchical structure
4. **Unexpected responses should predict how they evolve in time in indiv. neurons**

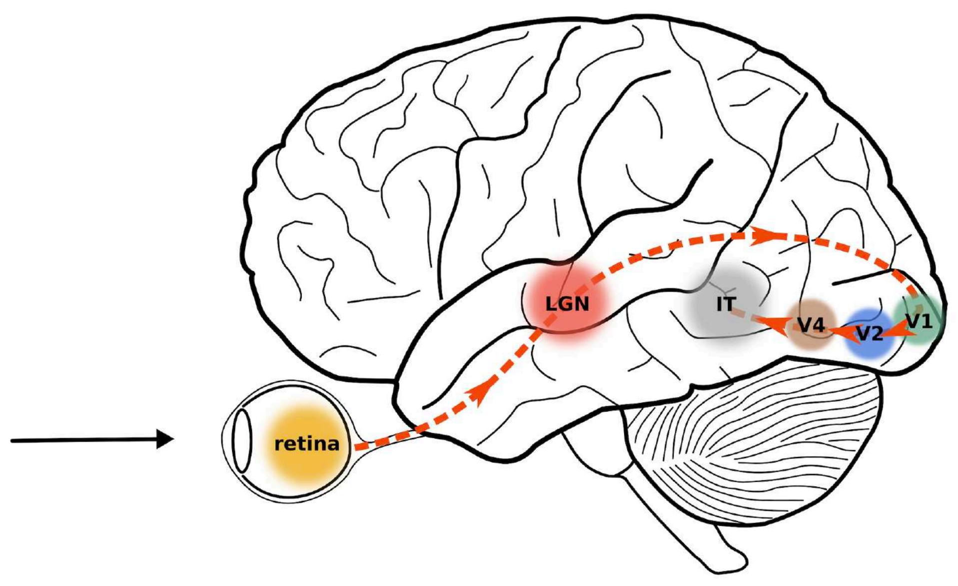

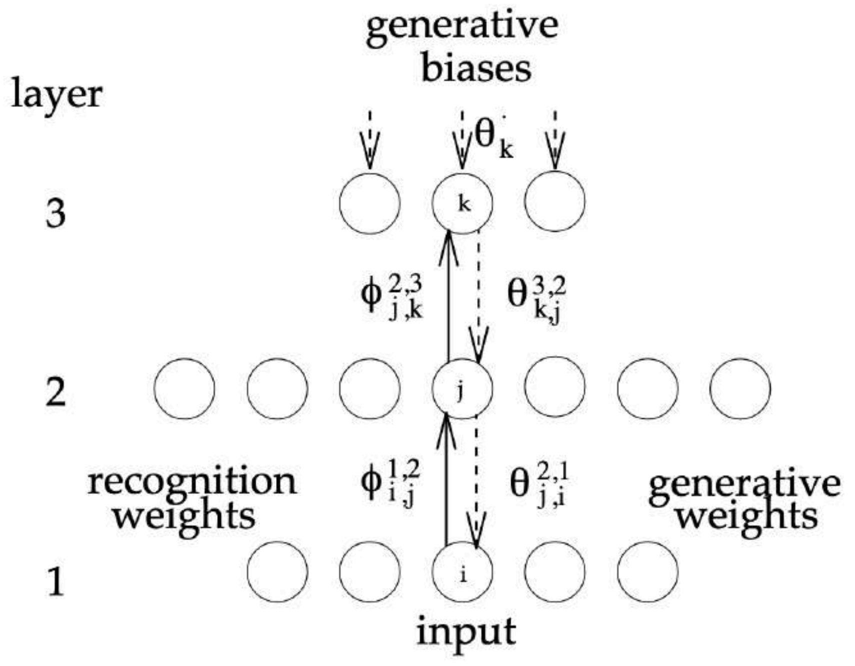

Helmholtz machine schematic[1]

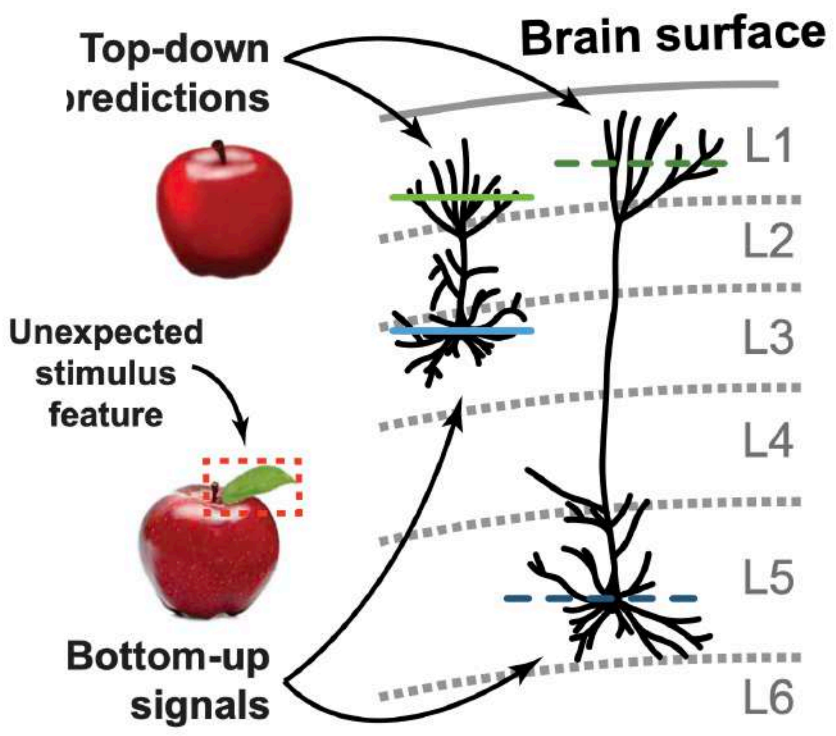

[1]- Dayan … Zemel, *Neur. Comp.*, 1995

Seed mouse with expectations and observe responses to expectation violations over multiple days

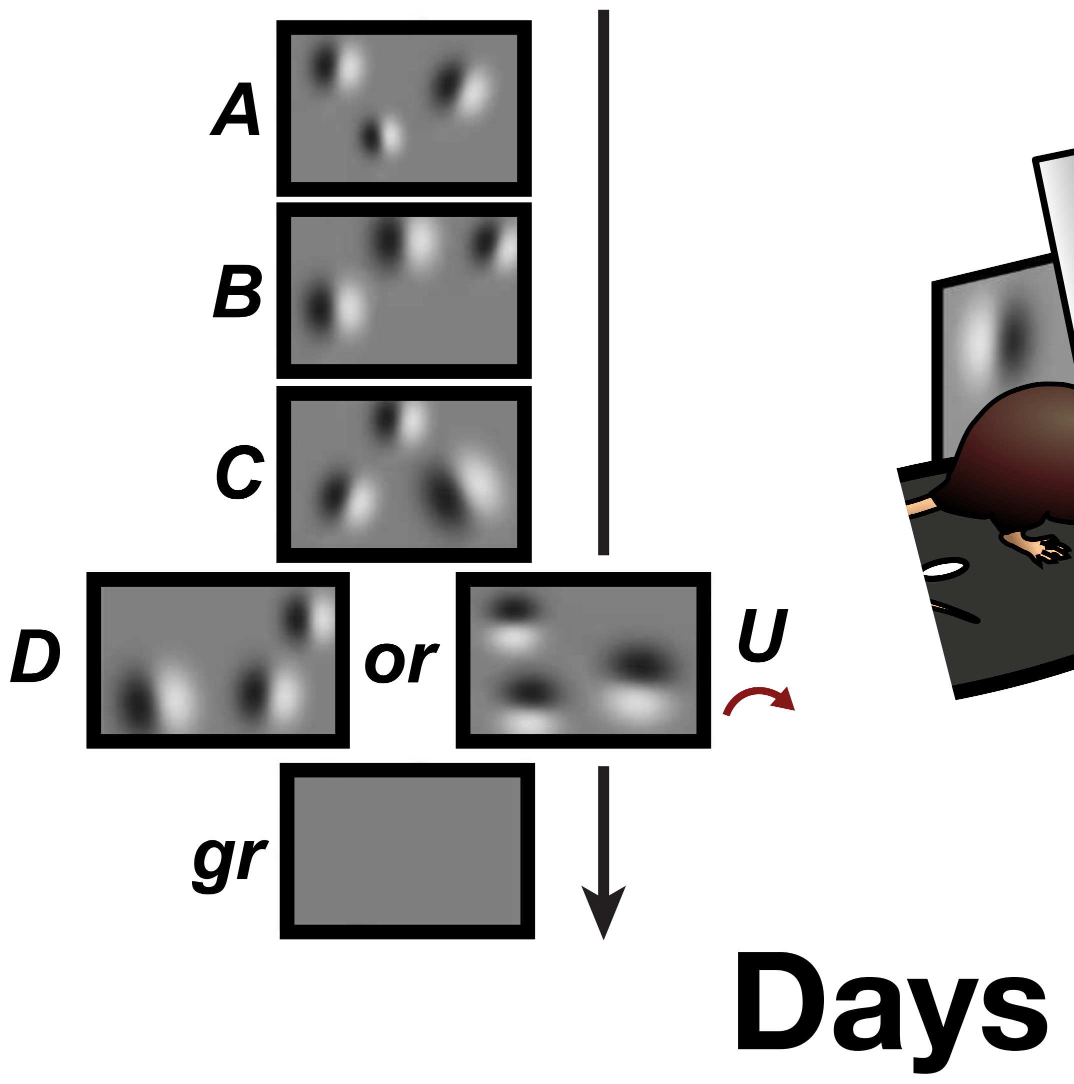

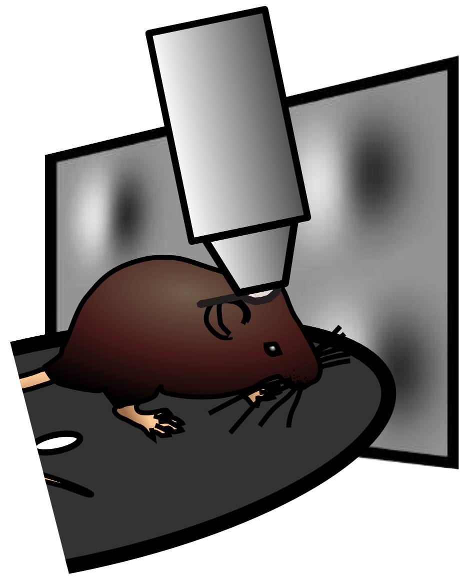

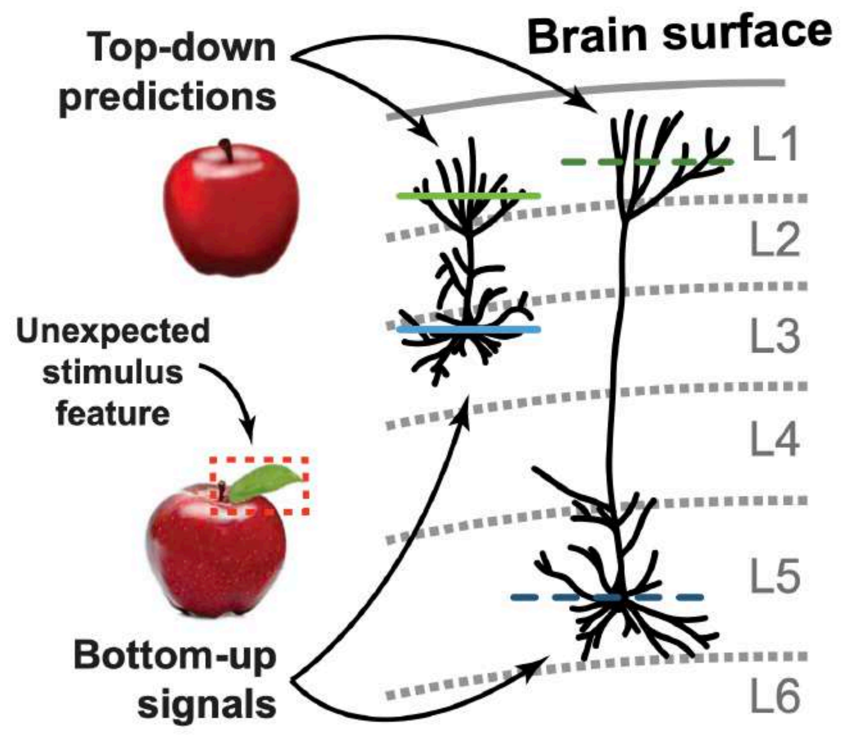

Experimental timeline (in days)

H1 H2 H3 H4 H5	H6 H7 H8 H9 H10 H11	1 2 3
Habituation to recording rig without stimulus	Habituation with expected sequences (no U frames)	Optical imaging with full stimulus (incl. U frames)

ALLEN INSTITUTE for BRAIN SCIENCE

- Habituate mice to *A-B-C-D* Gabor-patch sequences
- Then substitute ~8% of *D* frames with unexpected *U* Gabor-patch frames (diff. positions and orientations than *D*)
- Image 2-photon calcium activity over 3 recording days
- Segment and match ROIs to follow activity over different days
  - Distal apical dendrites (top-down signals) and somata (bottom-up signals)
  - Error signals? Matching signals?

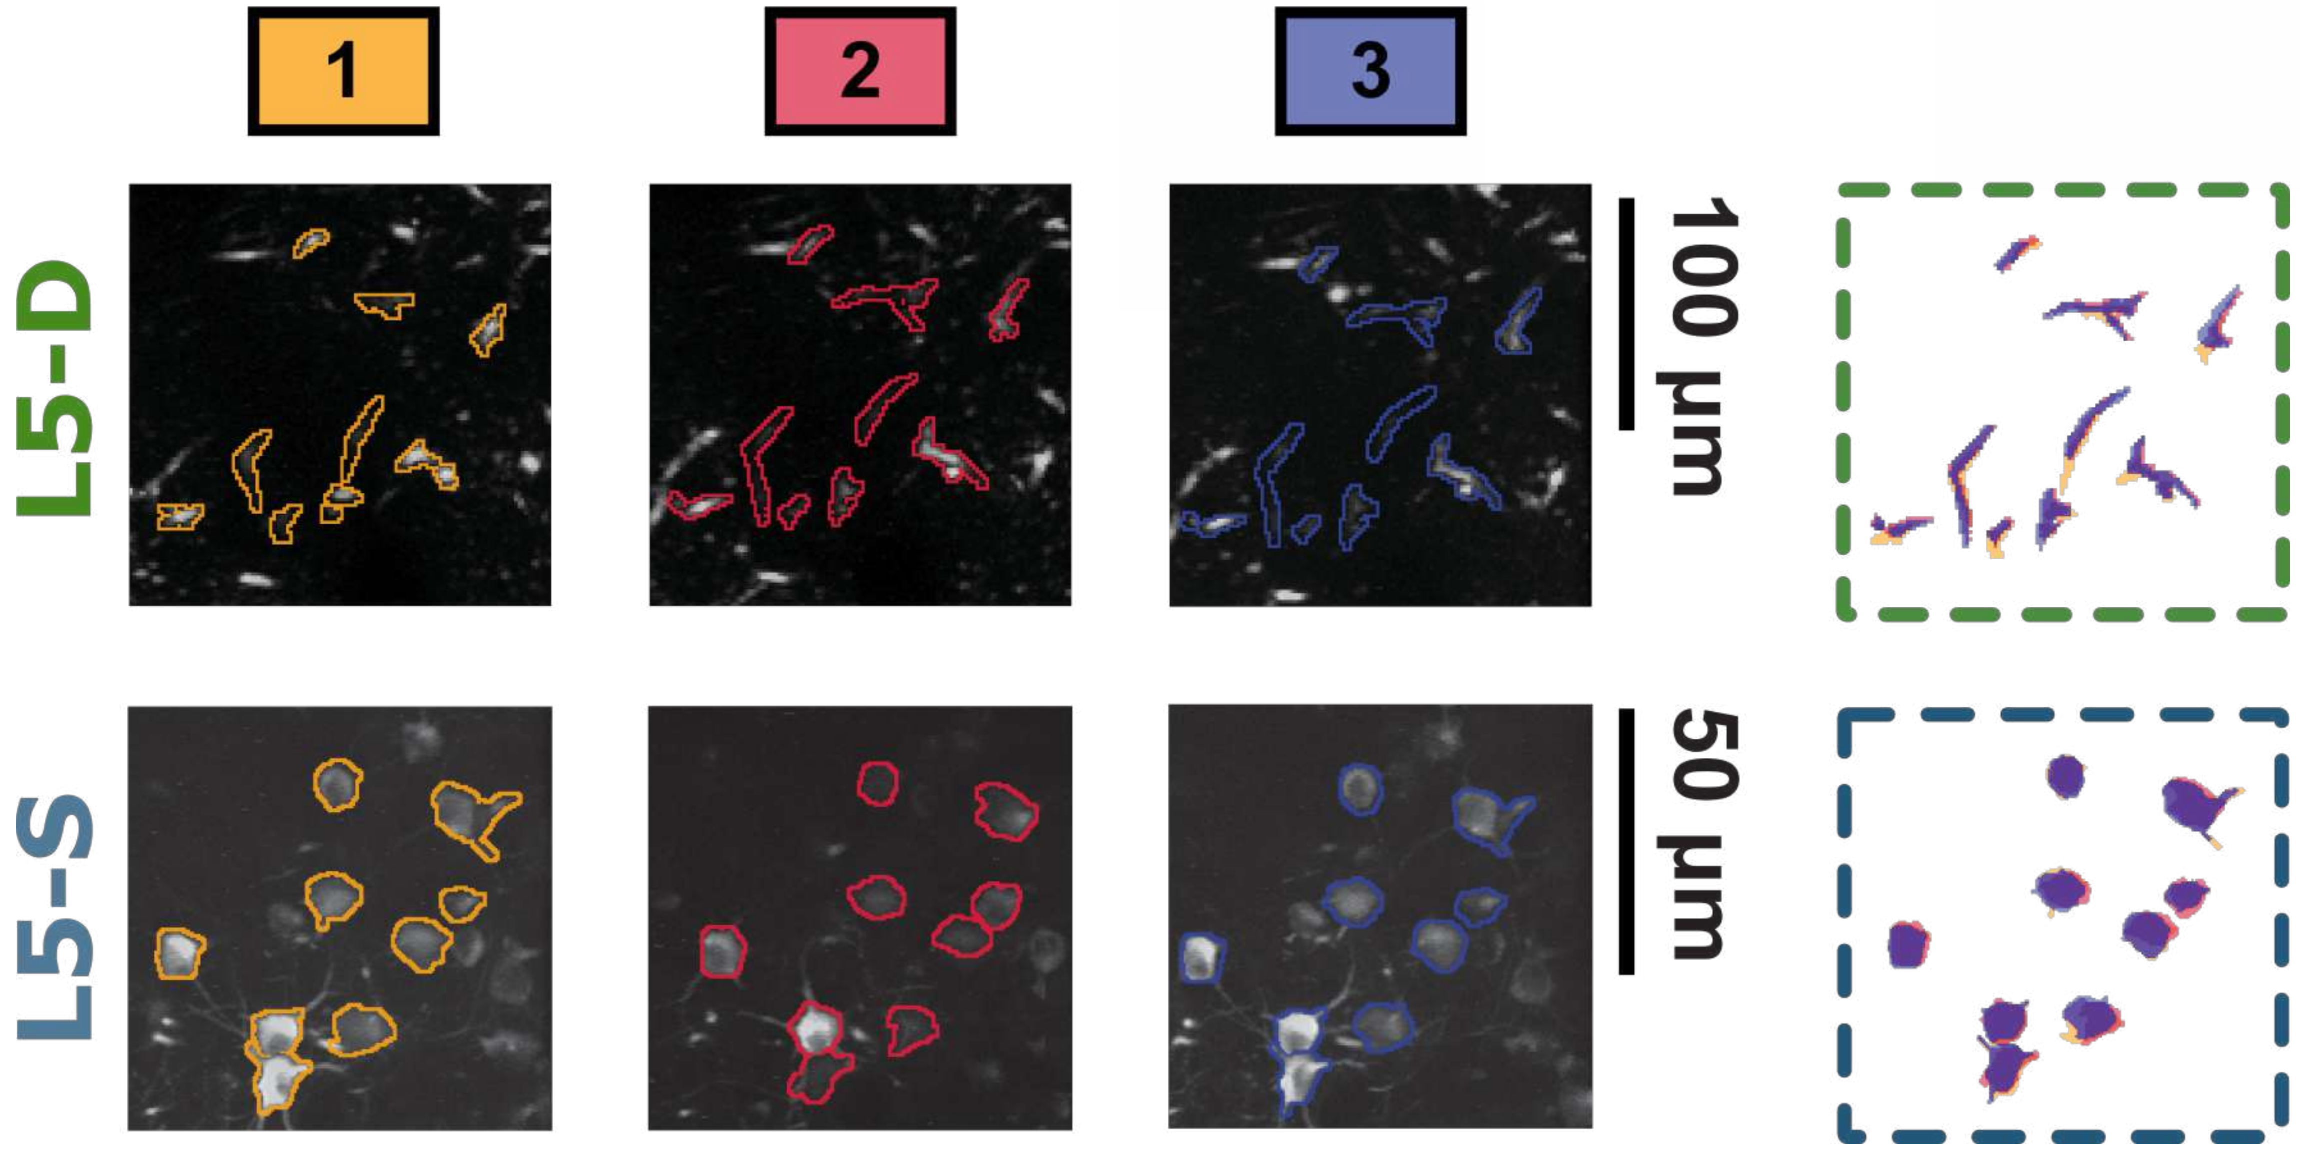

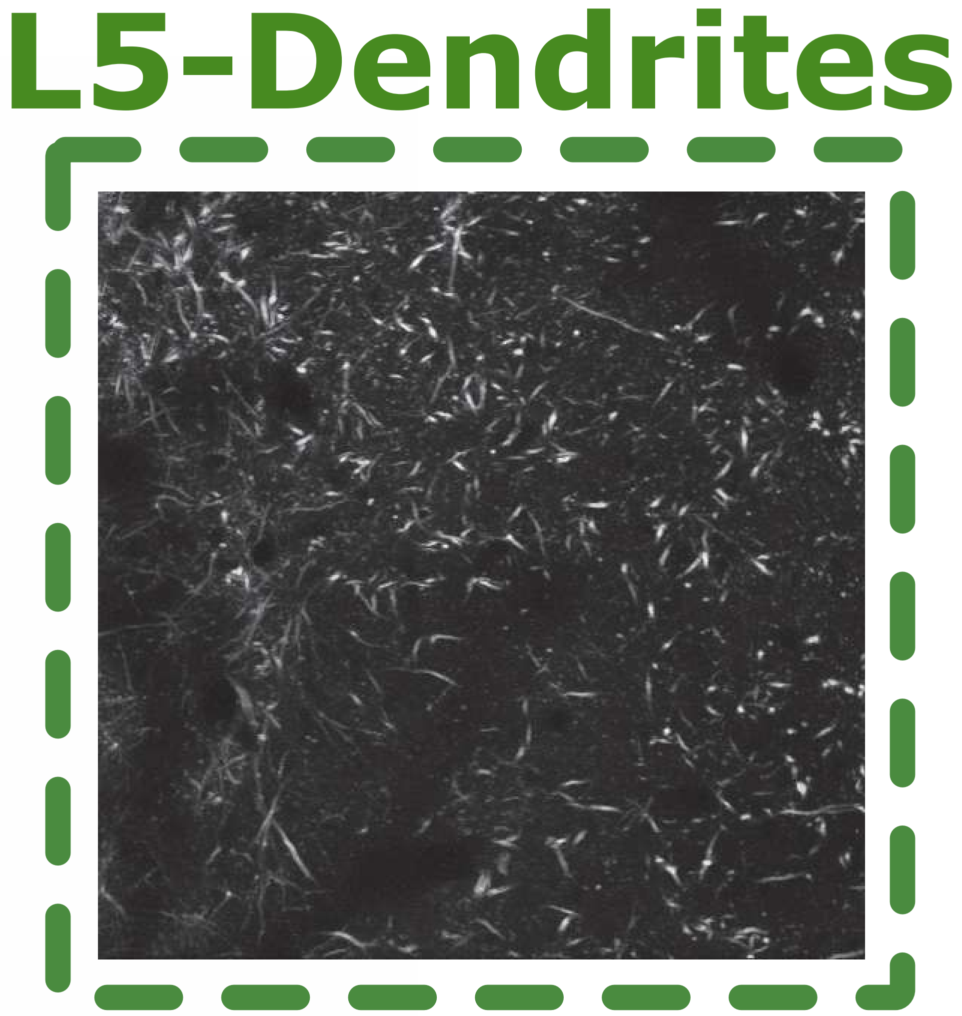

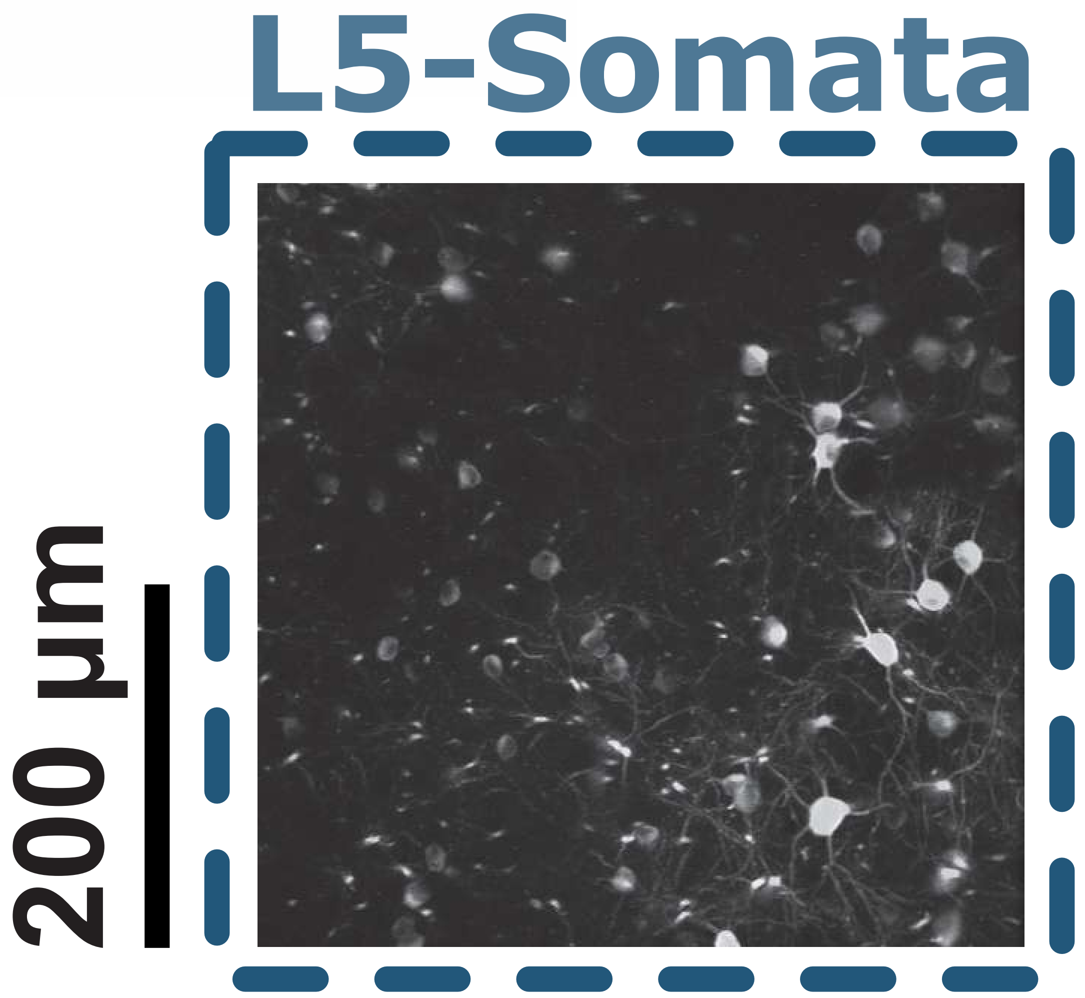

(1) Are there distinct responses to expected and unexpected stimul?

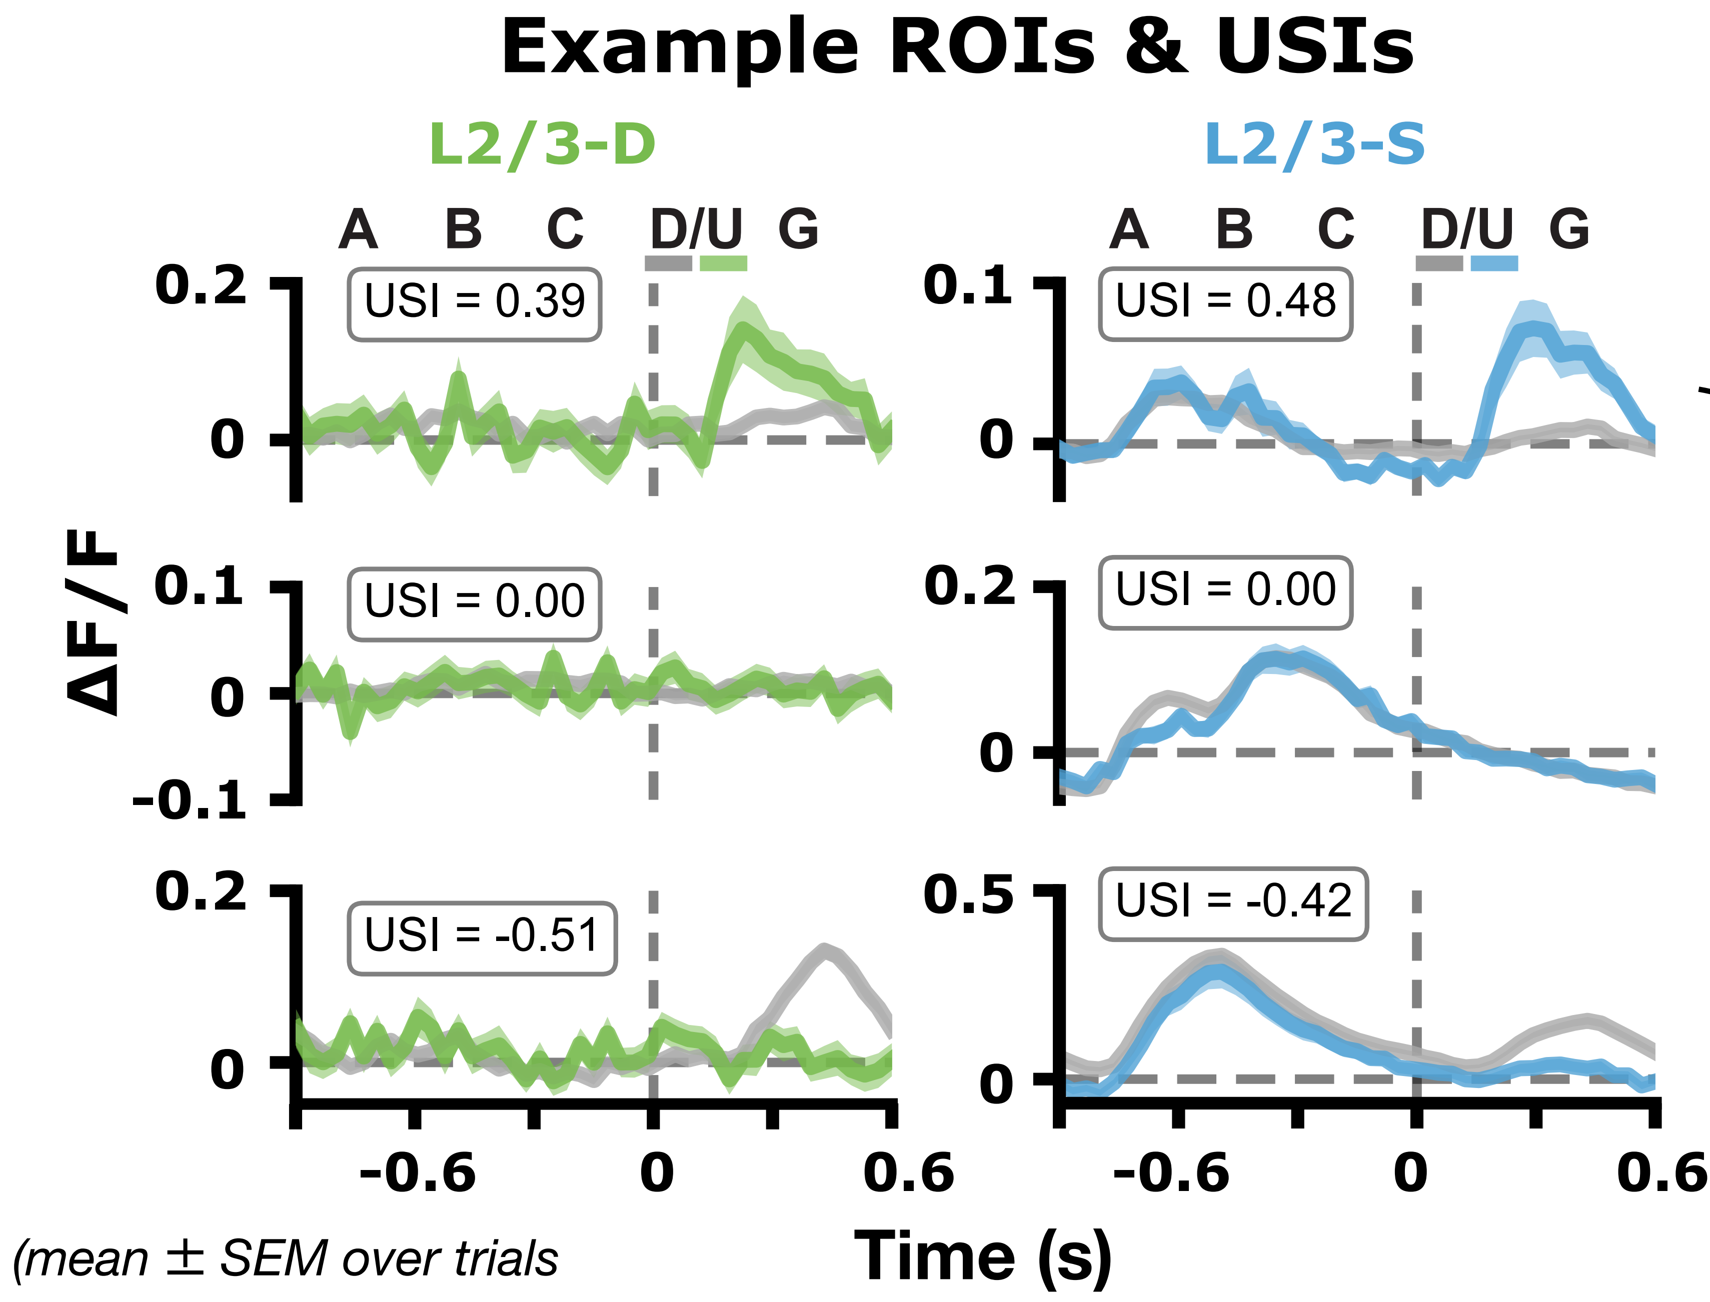

compare ea. ROI's USI to its null distro (shuffle D/U labels and recalculate USI 1×10^5 *times)*

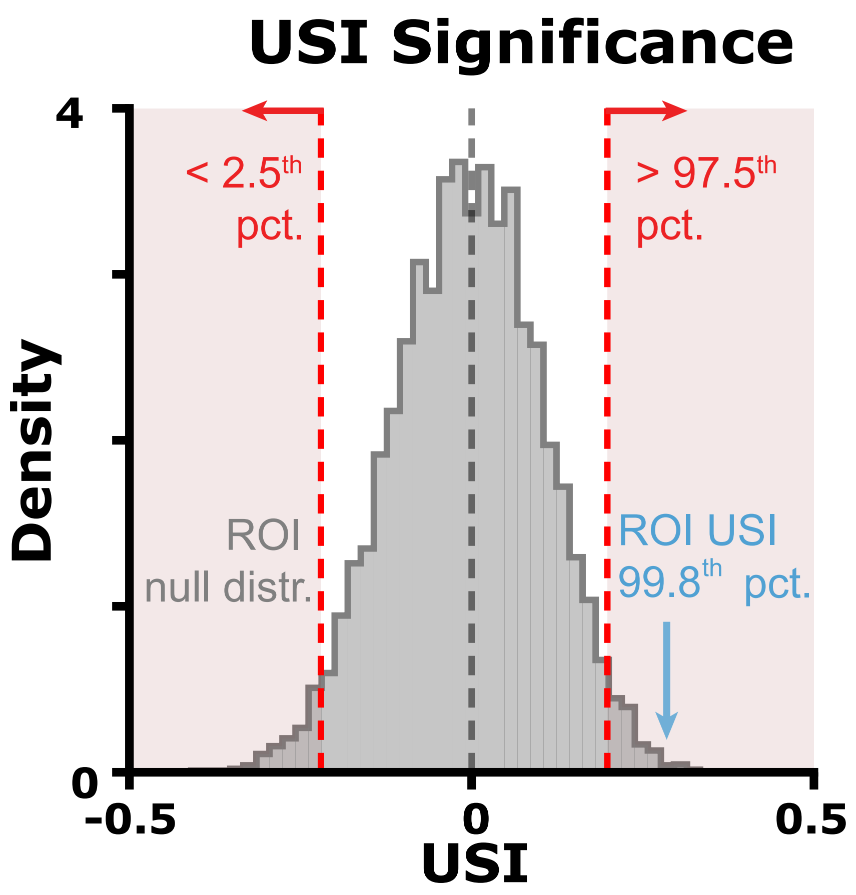

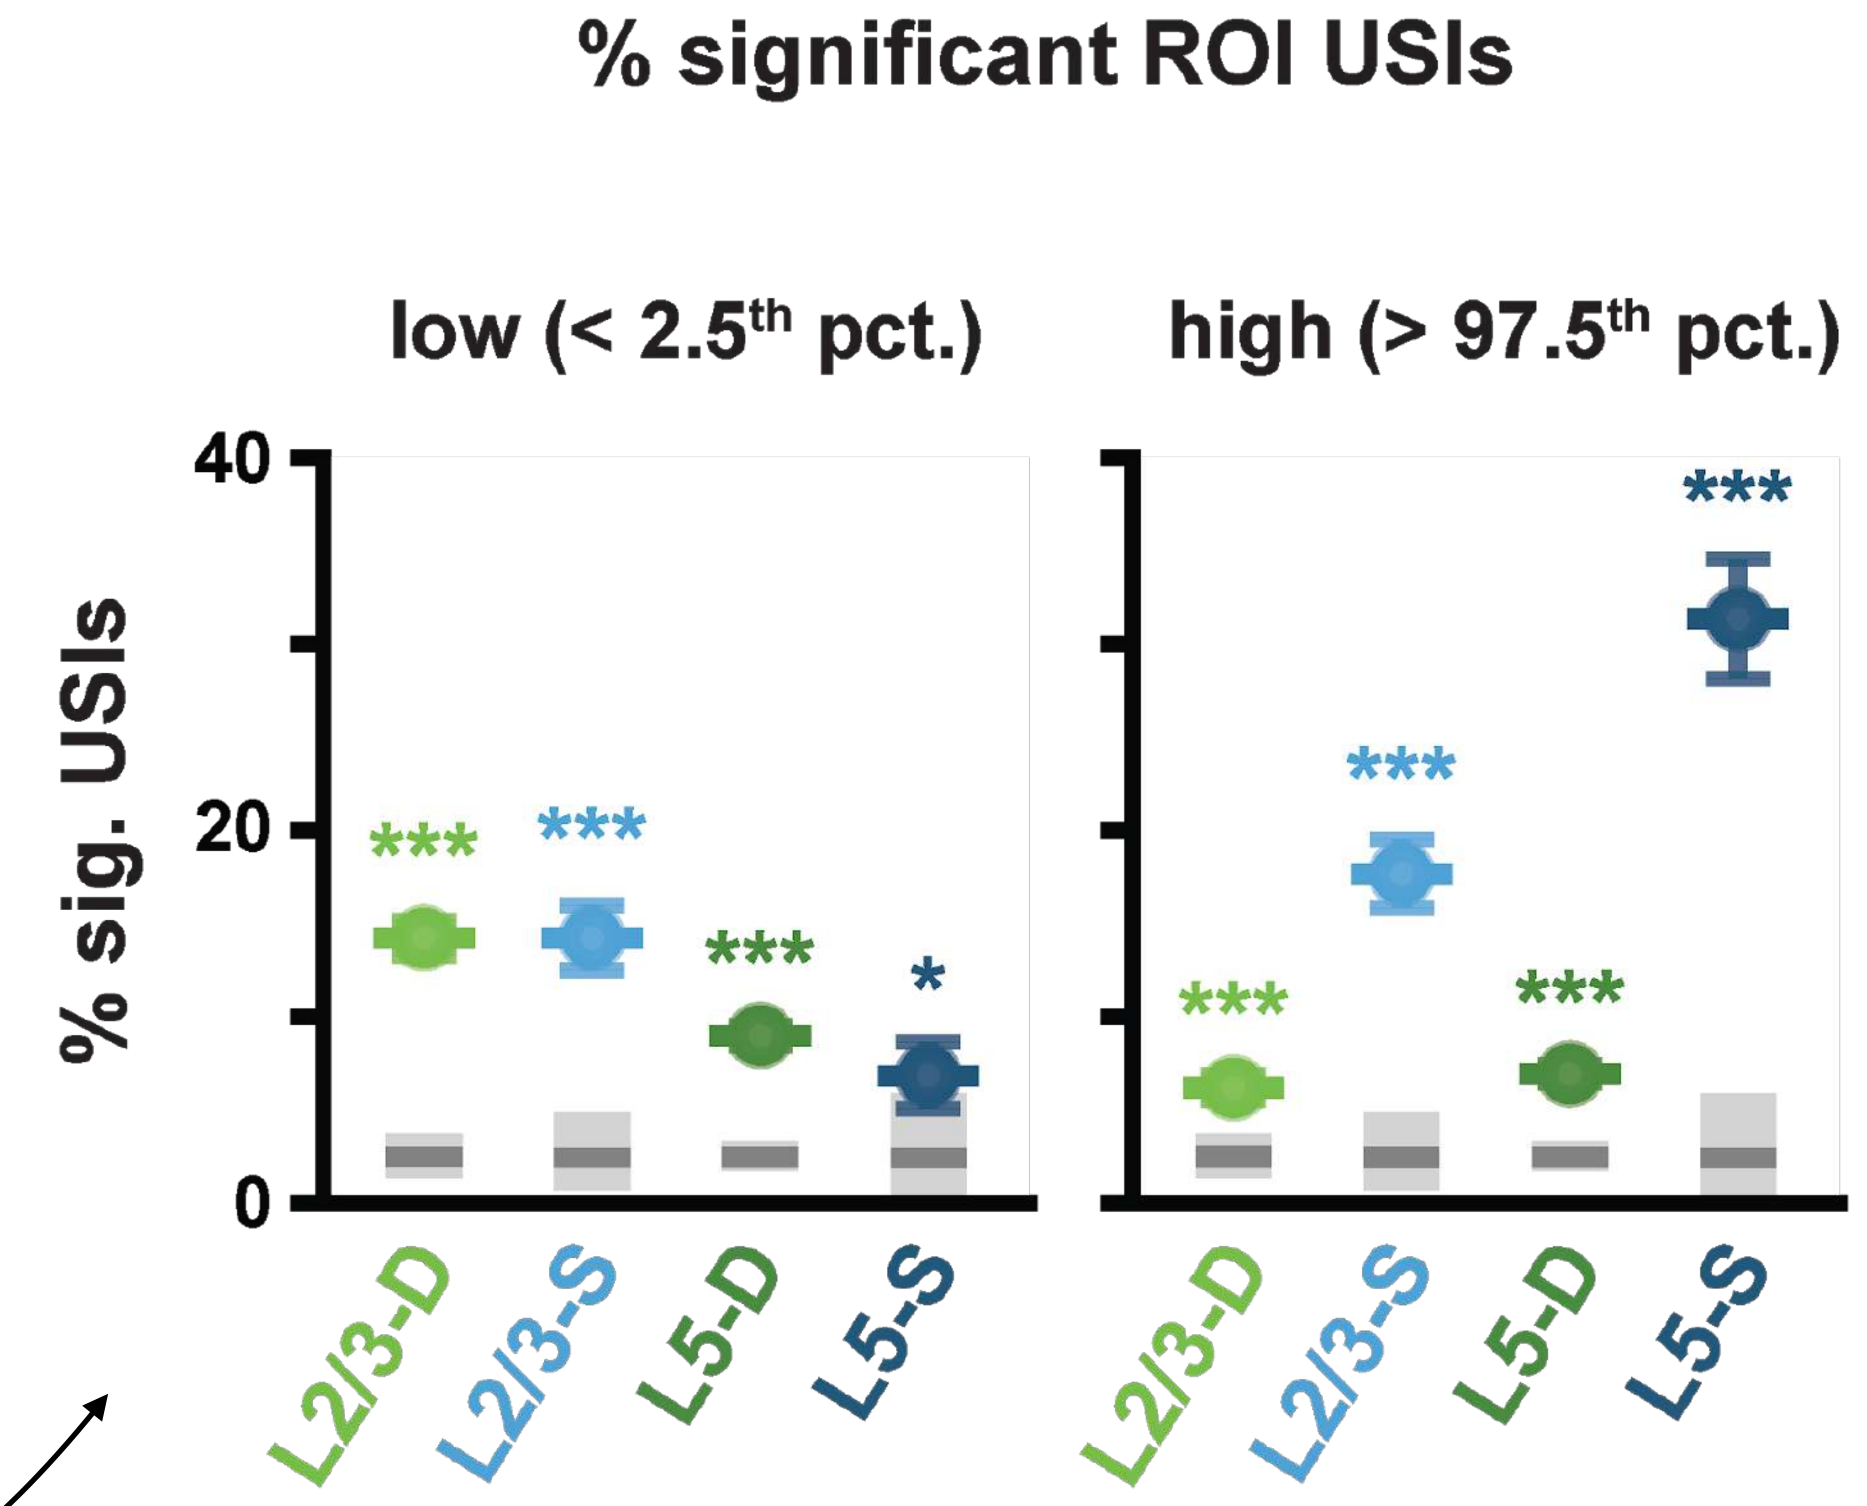

*U*nexpected event *S*electivity *I*ndex
(measure of sensitivity to *U* frames)

$$\mathrm{USI} = \frac{\mu_{\mathrm{UG}} - \mu_{\mathrm{DG}}}{\sqrt{\frac{1}{2}\left(\sigma^2_{\mathrm{UG}} + \sigma^2_{\mathrm{DG}}\right)}}$$

Take-home: Many more USIs are sensitive (±) to unexpected vs. expected events than predicted by chance, demonstrating distinct responses

(2) Do they change with experience?
(3) Do bottom-up and top-down representations evolve differently?

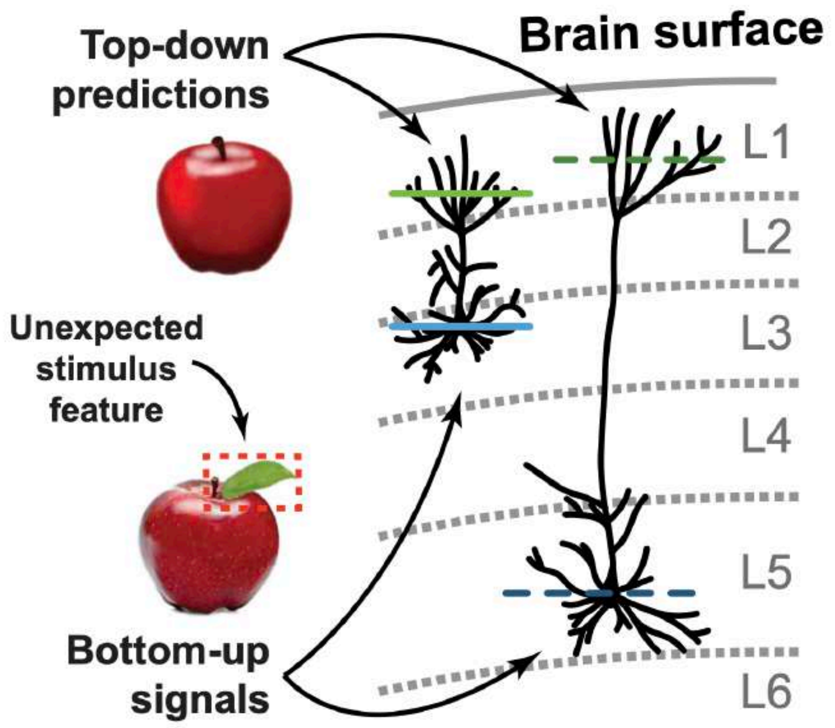

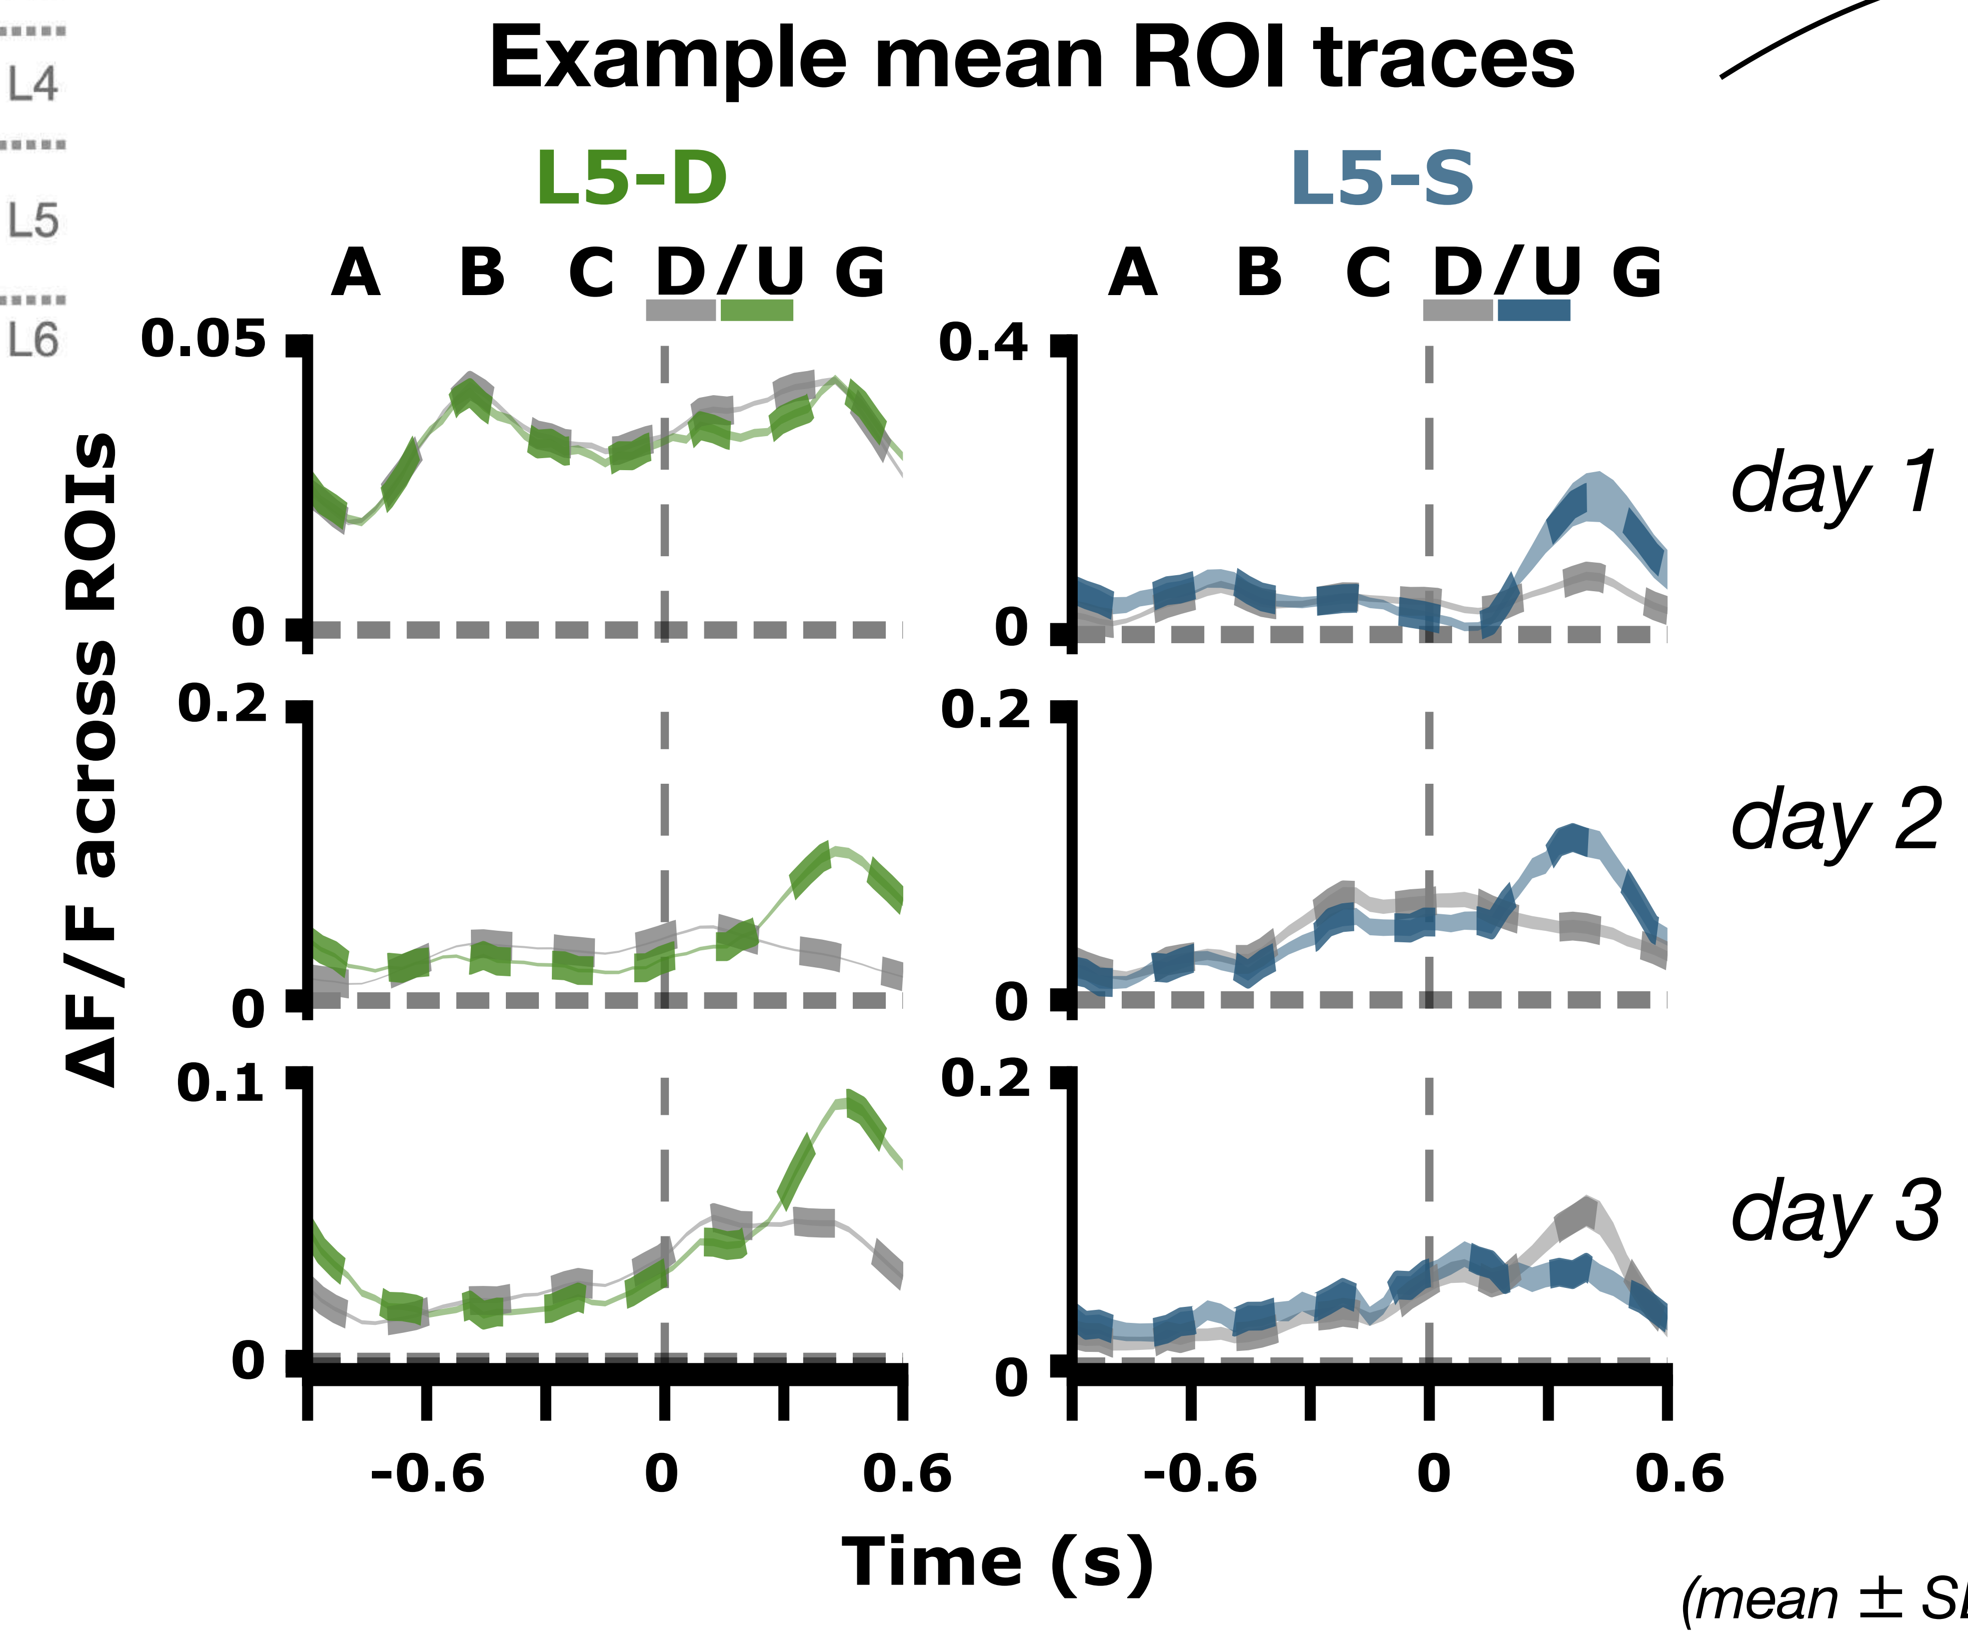

(mean ± SEM over ROIs)

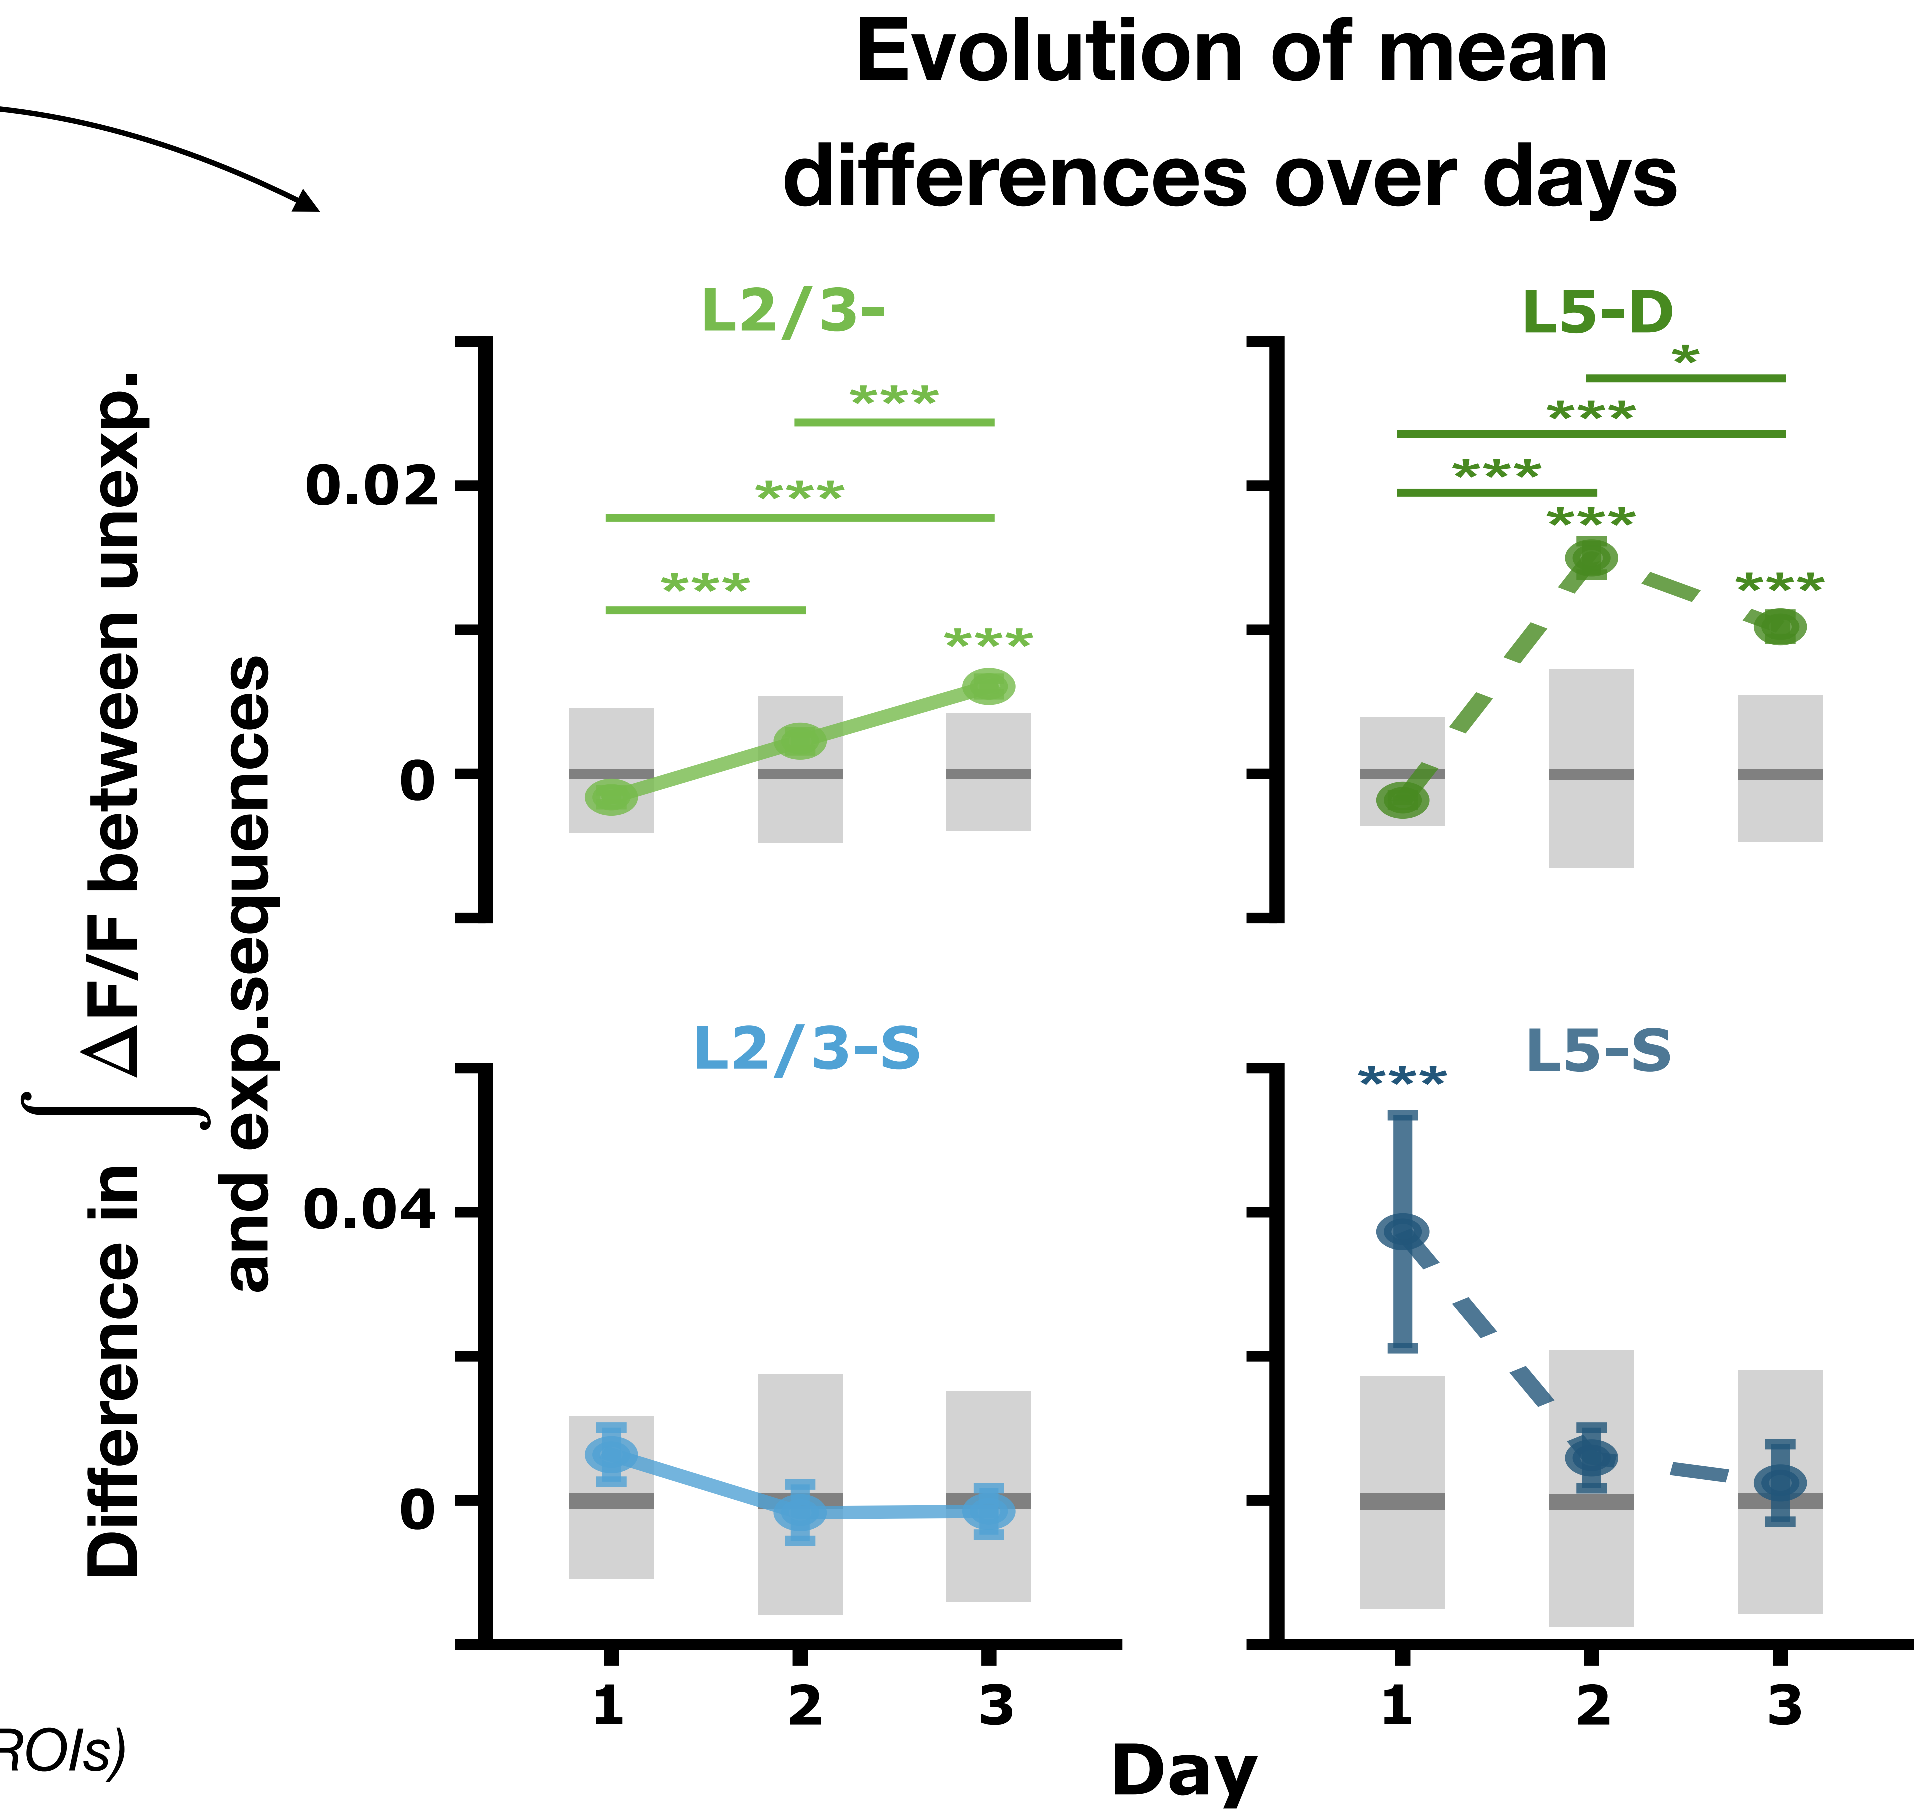

Take-home: Dendritic (top-down) responses increase with experience over days, while somatic (bottom-up) responses decrease

(4) Do the unexpected responses predict how they evolve in time?

Individual ROI USI evolution

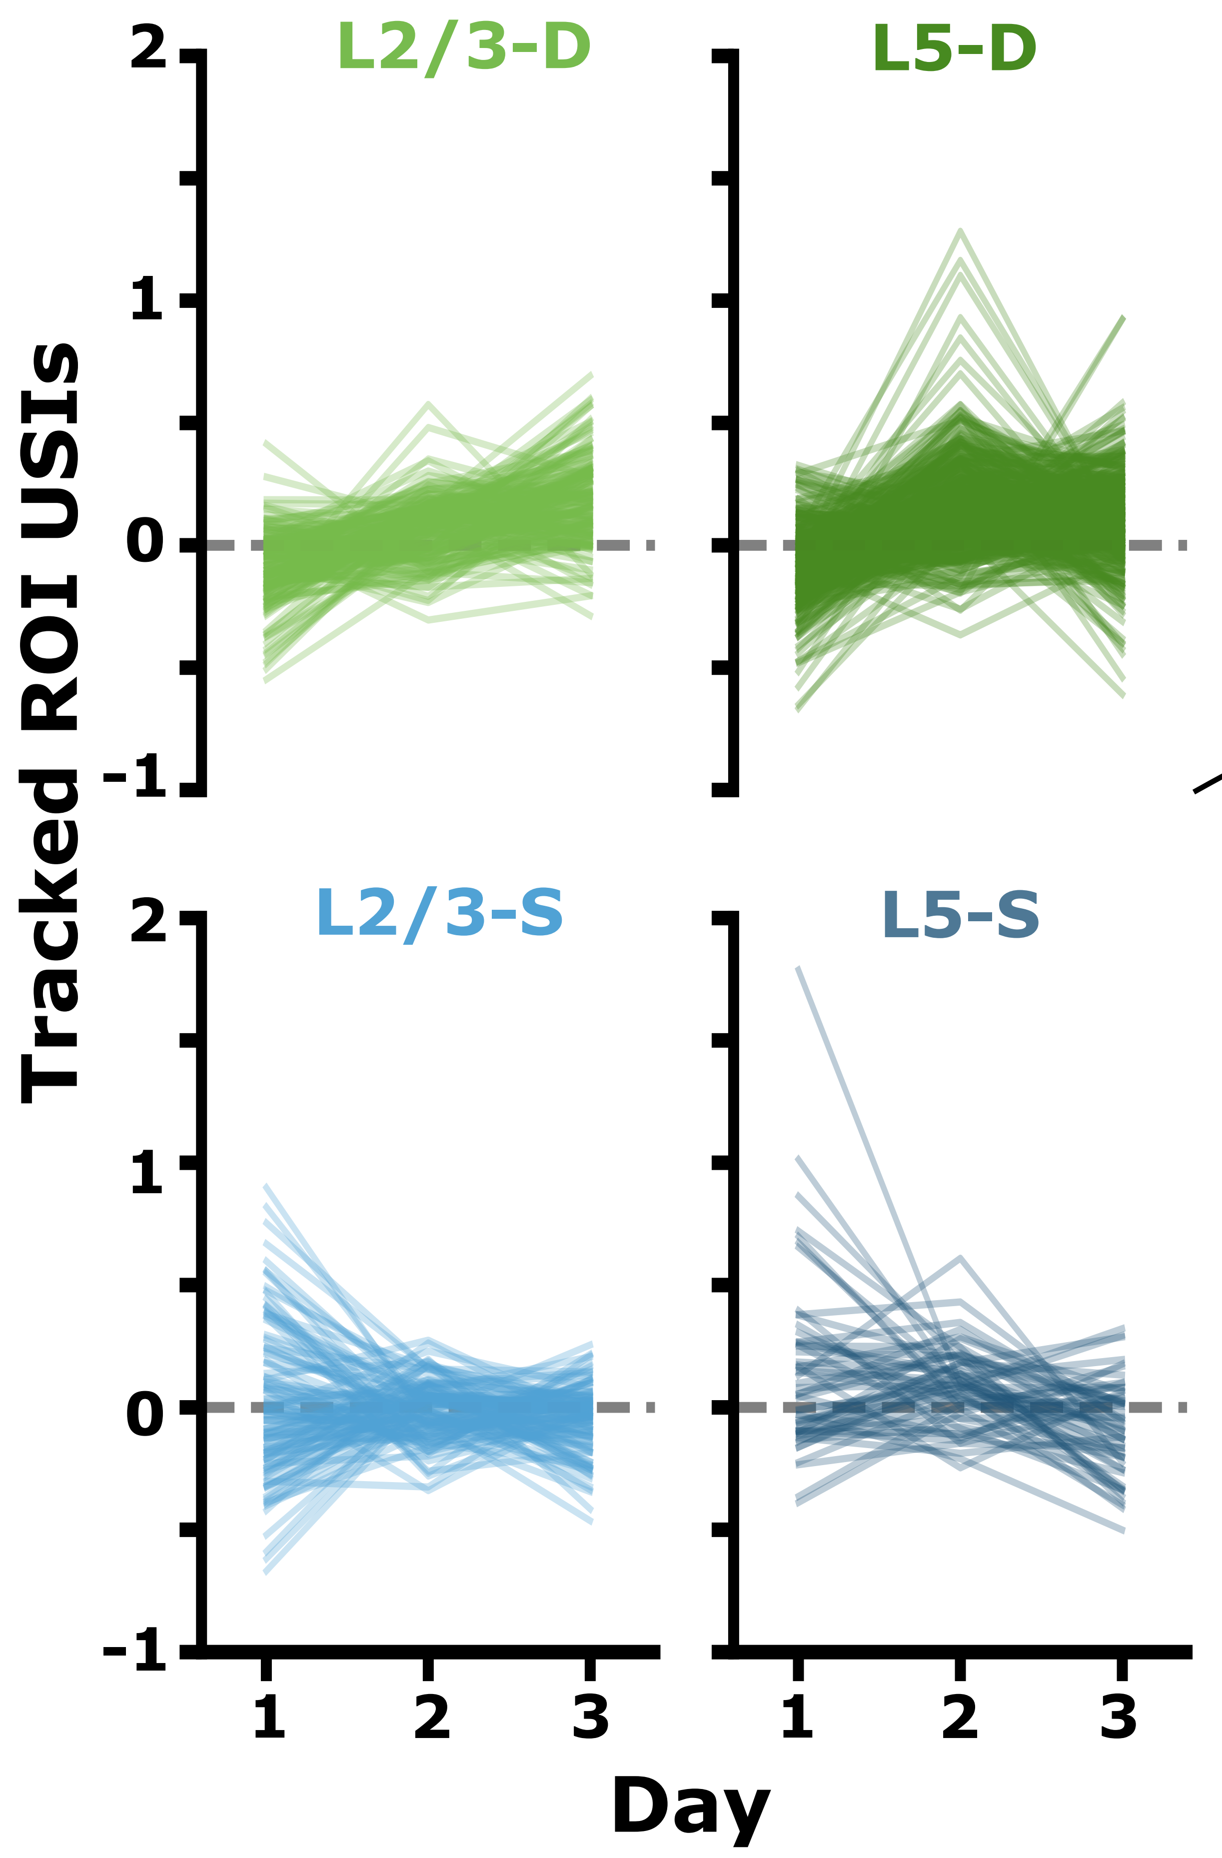

(each line corresponds to 1 tracked ROI)

Examine correlation between USIs on 1 day and the change in value on the next day

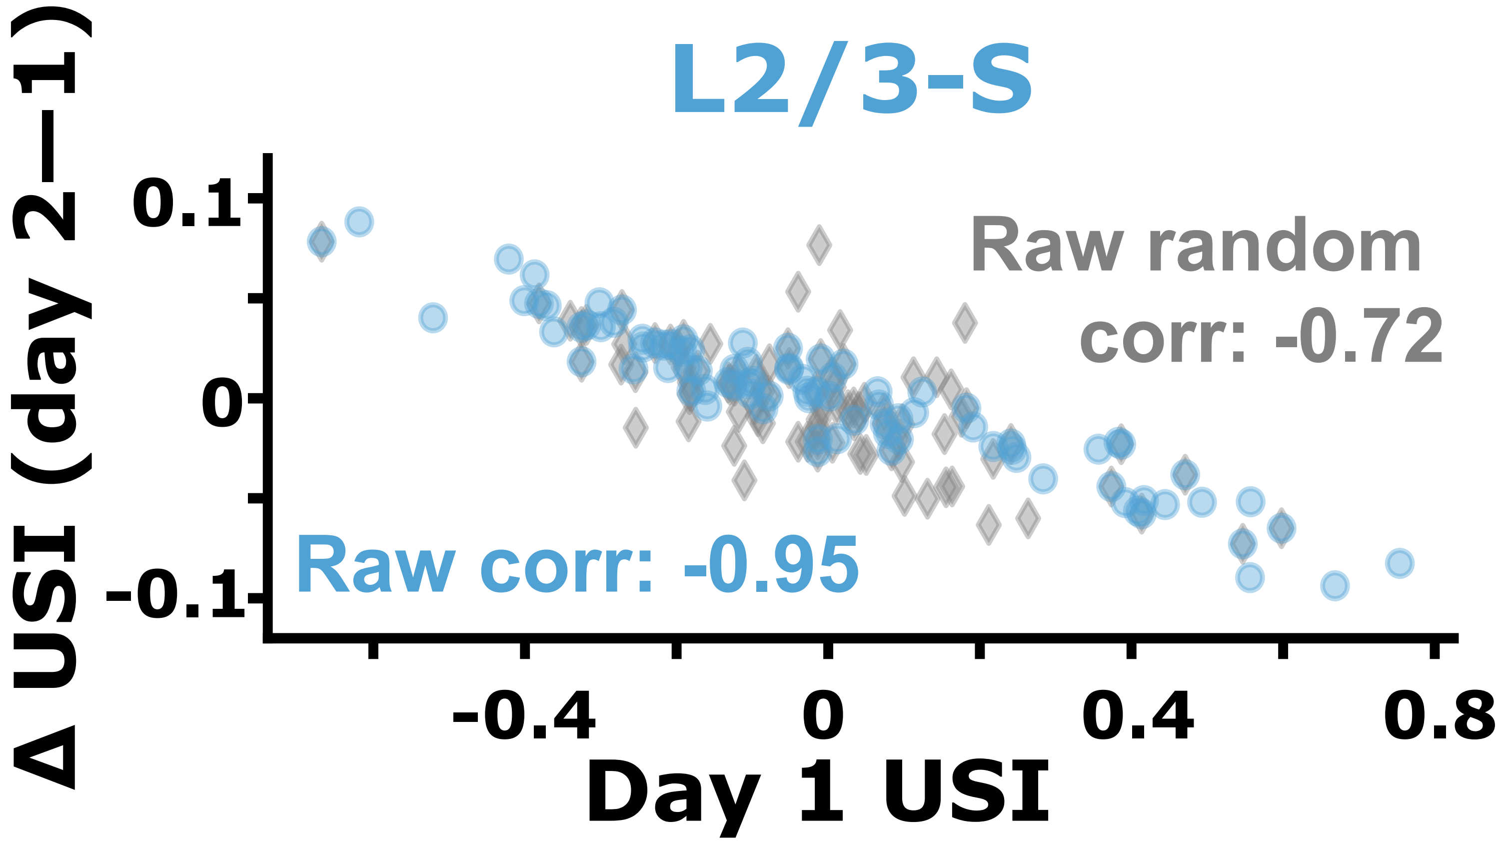

1. *Shuffle day labels* 1×10^5 *times for ROIs and recalculate each correlation, producing correlation distribution*
2. *Calculate normalized resid. corr. as shown*

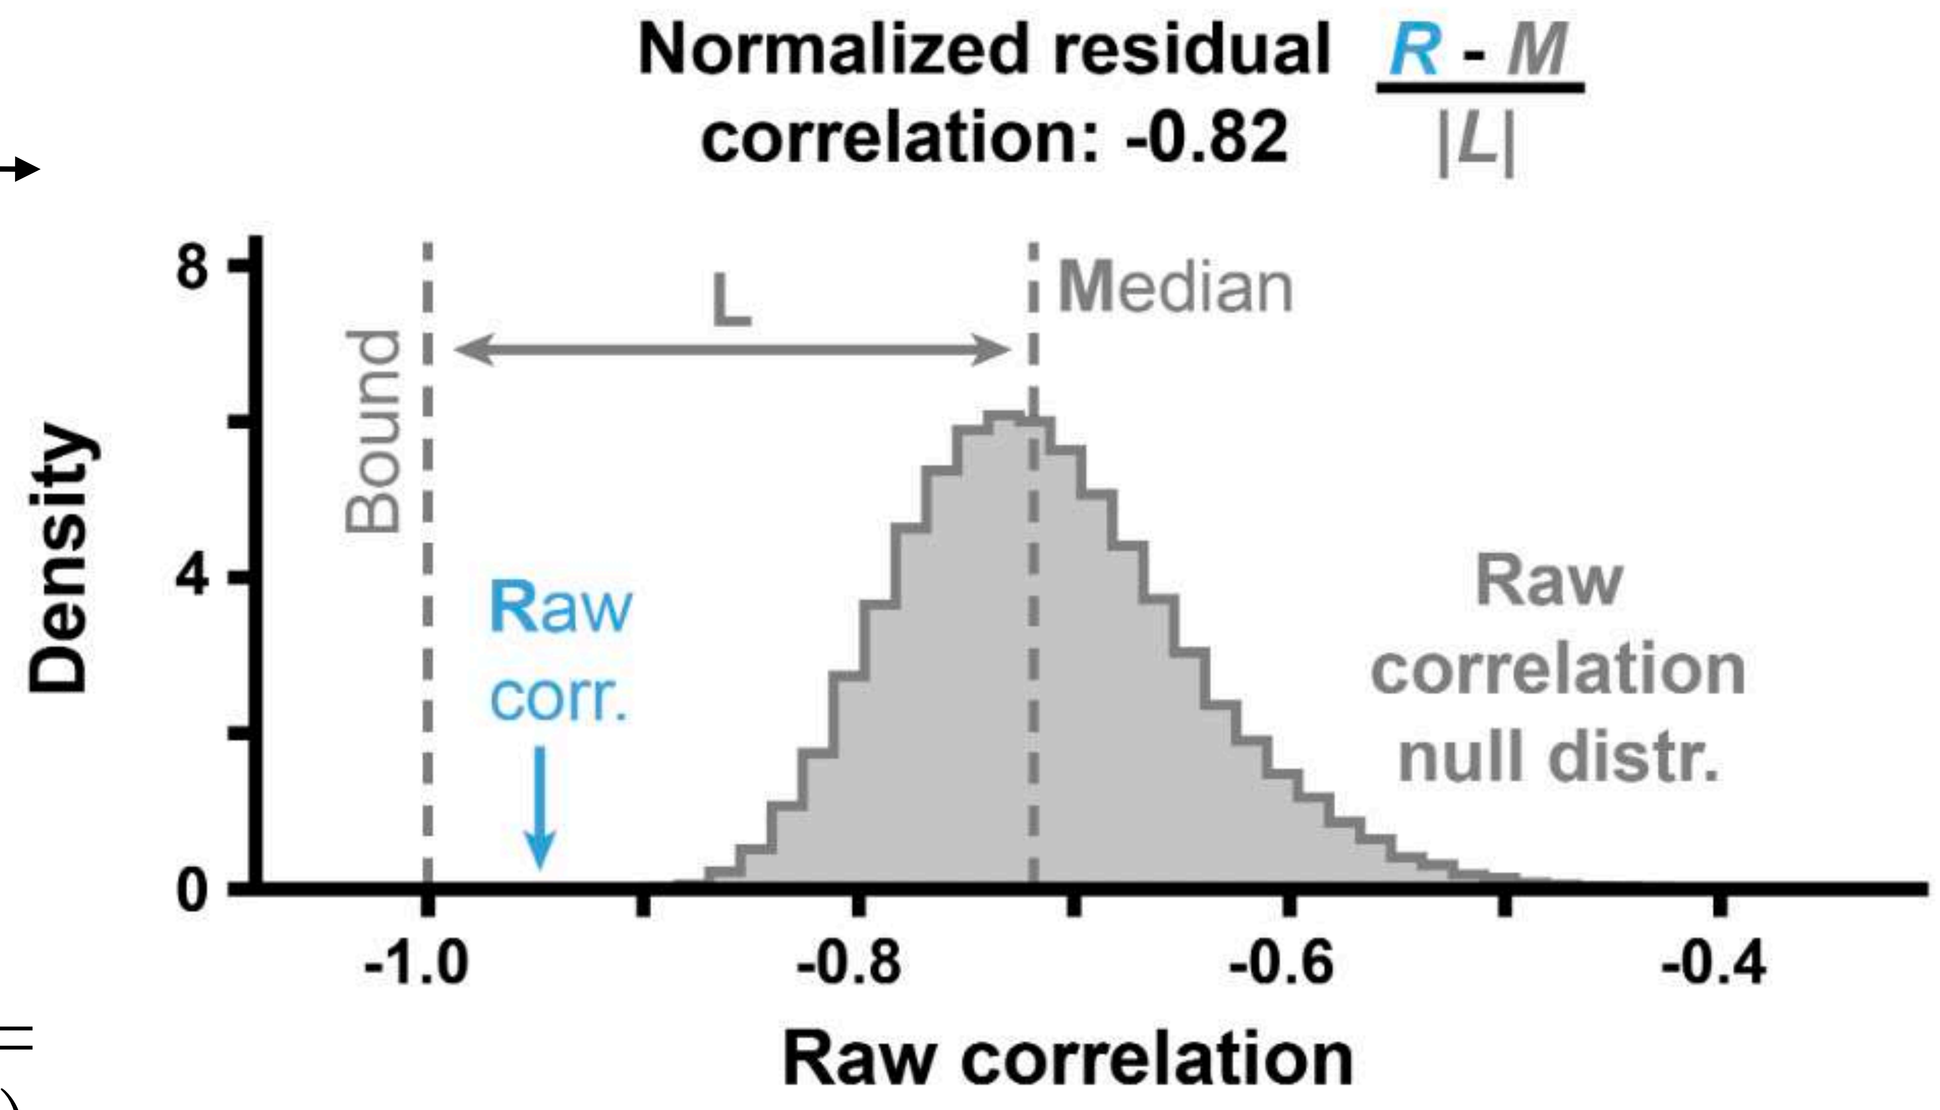

$$\mathrm{USI} = \frac{\mu_{\mathrm{UG}} - \mu_{\mathrm{DG}}}{\sqrt{\frac{1}{2}\left(\sigma^2_{\mathrm{UG}} + \sigma^2_{\mathrm{DG}}\right)}}$$

(4) Do the unexpected responses predict how they evolve in time?

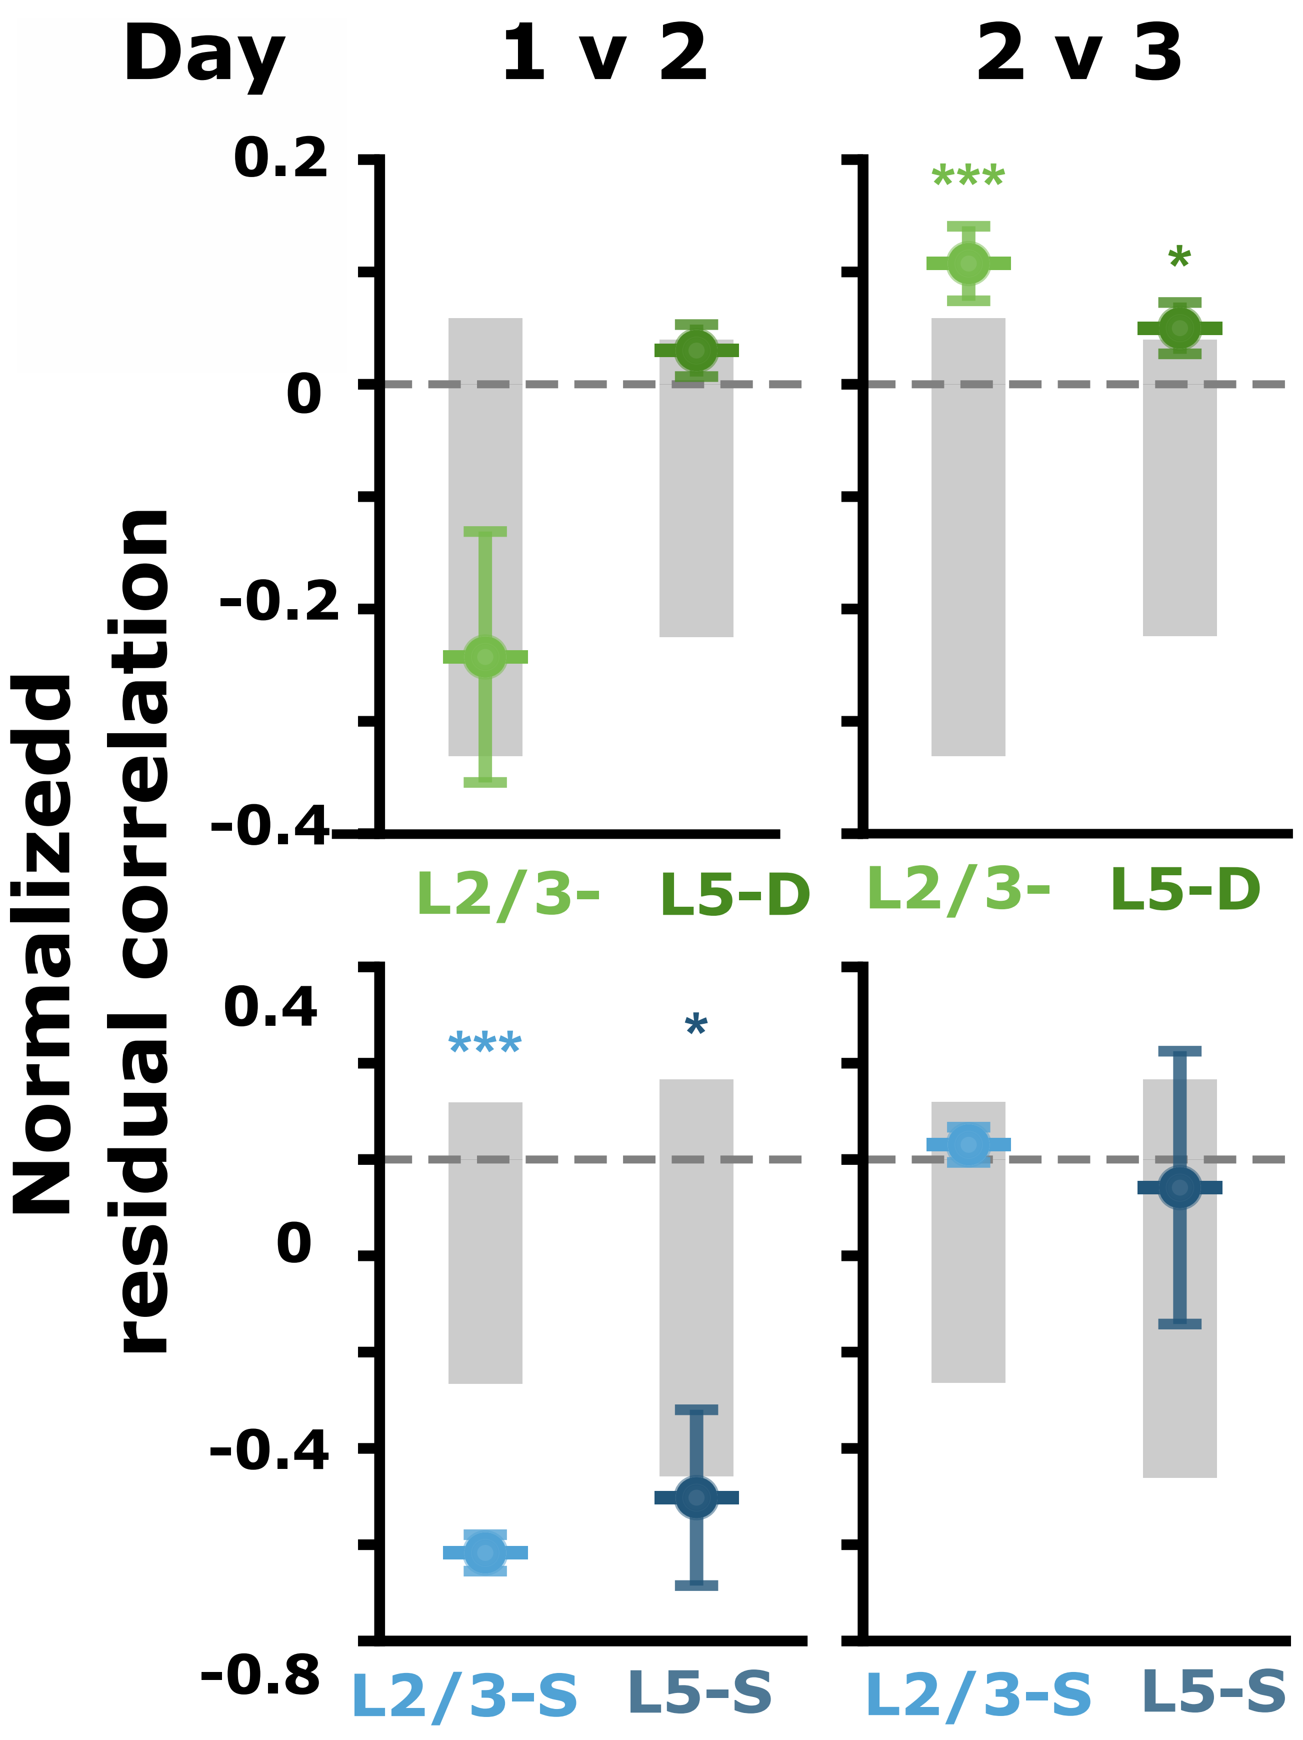

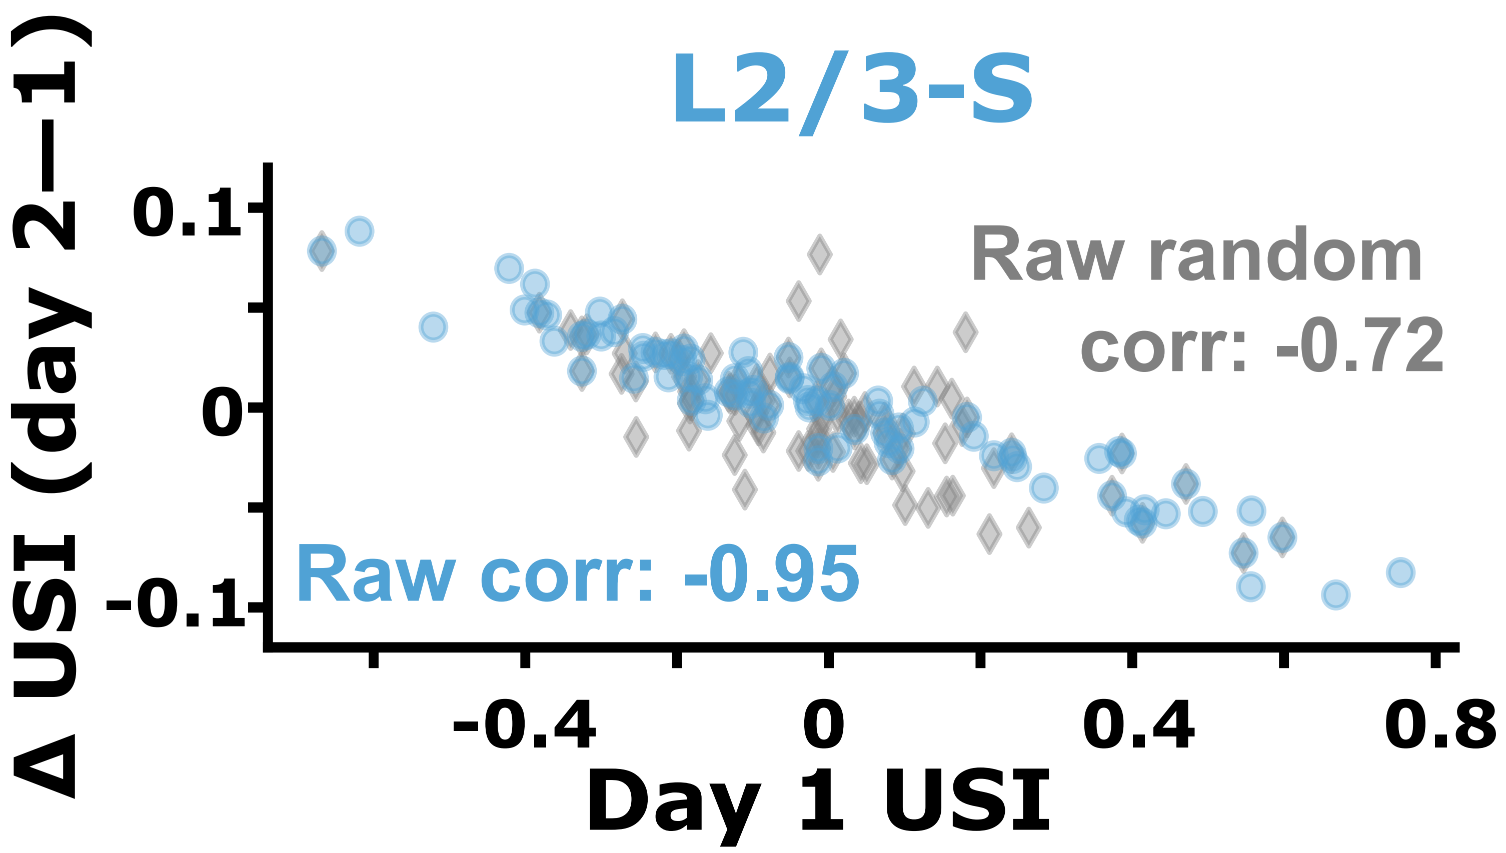

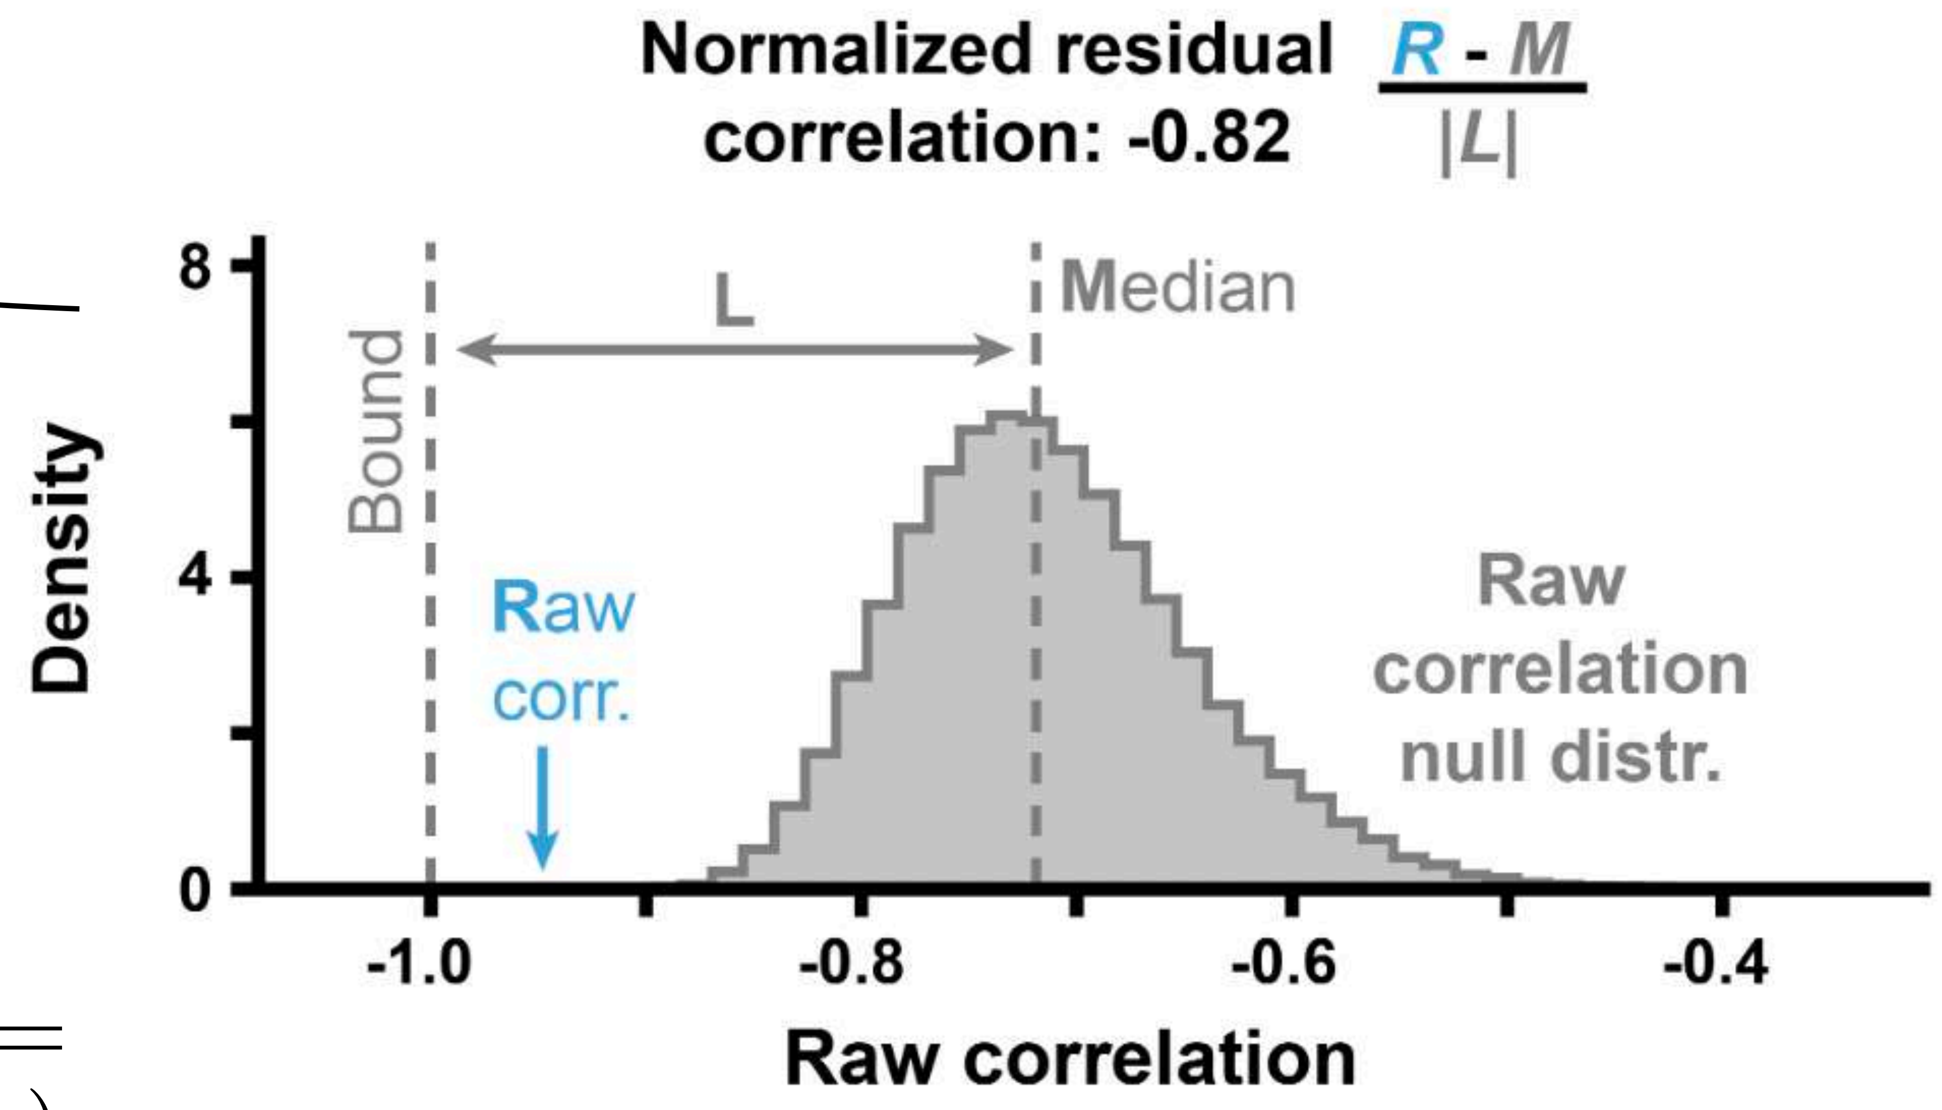

Take-homes:

- Somatic USIs ***decrease as a function of their USIs*** (statistically significantly) from day 1 to 2, consistent with a reduction in error.
- Dendrittic USIs ***increase as a function of their USIs*** (statistically significantly) from day 2 to 3, becoming more sensitive to unexpected stimuli

$$\mathrm{USI} = \frac{\mu_{\mathrm{UG}} - \mu_{\mathrm{DG}}}{\sqrt{\frac{1}{2}\left(\sigma_{\mathrm{UG}}^2 + \sigma_{\mathrm{DG}}^2\right)}}$$

Predictive Hierarchical Network Summary

1. Violations of expectation are observable in neural activity in very early layers (*predictions*)

2. The sensitivity to violations changes over days (*learning of novel stimuli*)

3. Bottom-up and top-down signals evolve differently (*hierarchical learning model*)

4. Furthermore, the sensitivities specifically guide the evolution of responses in individual neurons (*specific differences in responses may drive learning*)

- Overall, visual cortex may instantiate a predictive hierarchical model that is updated as a result of unexpected events

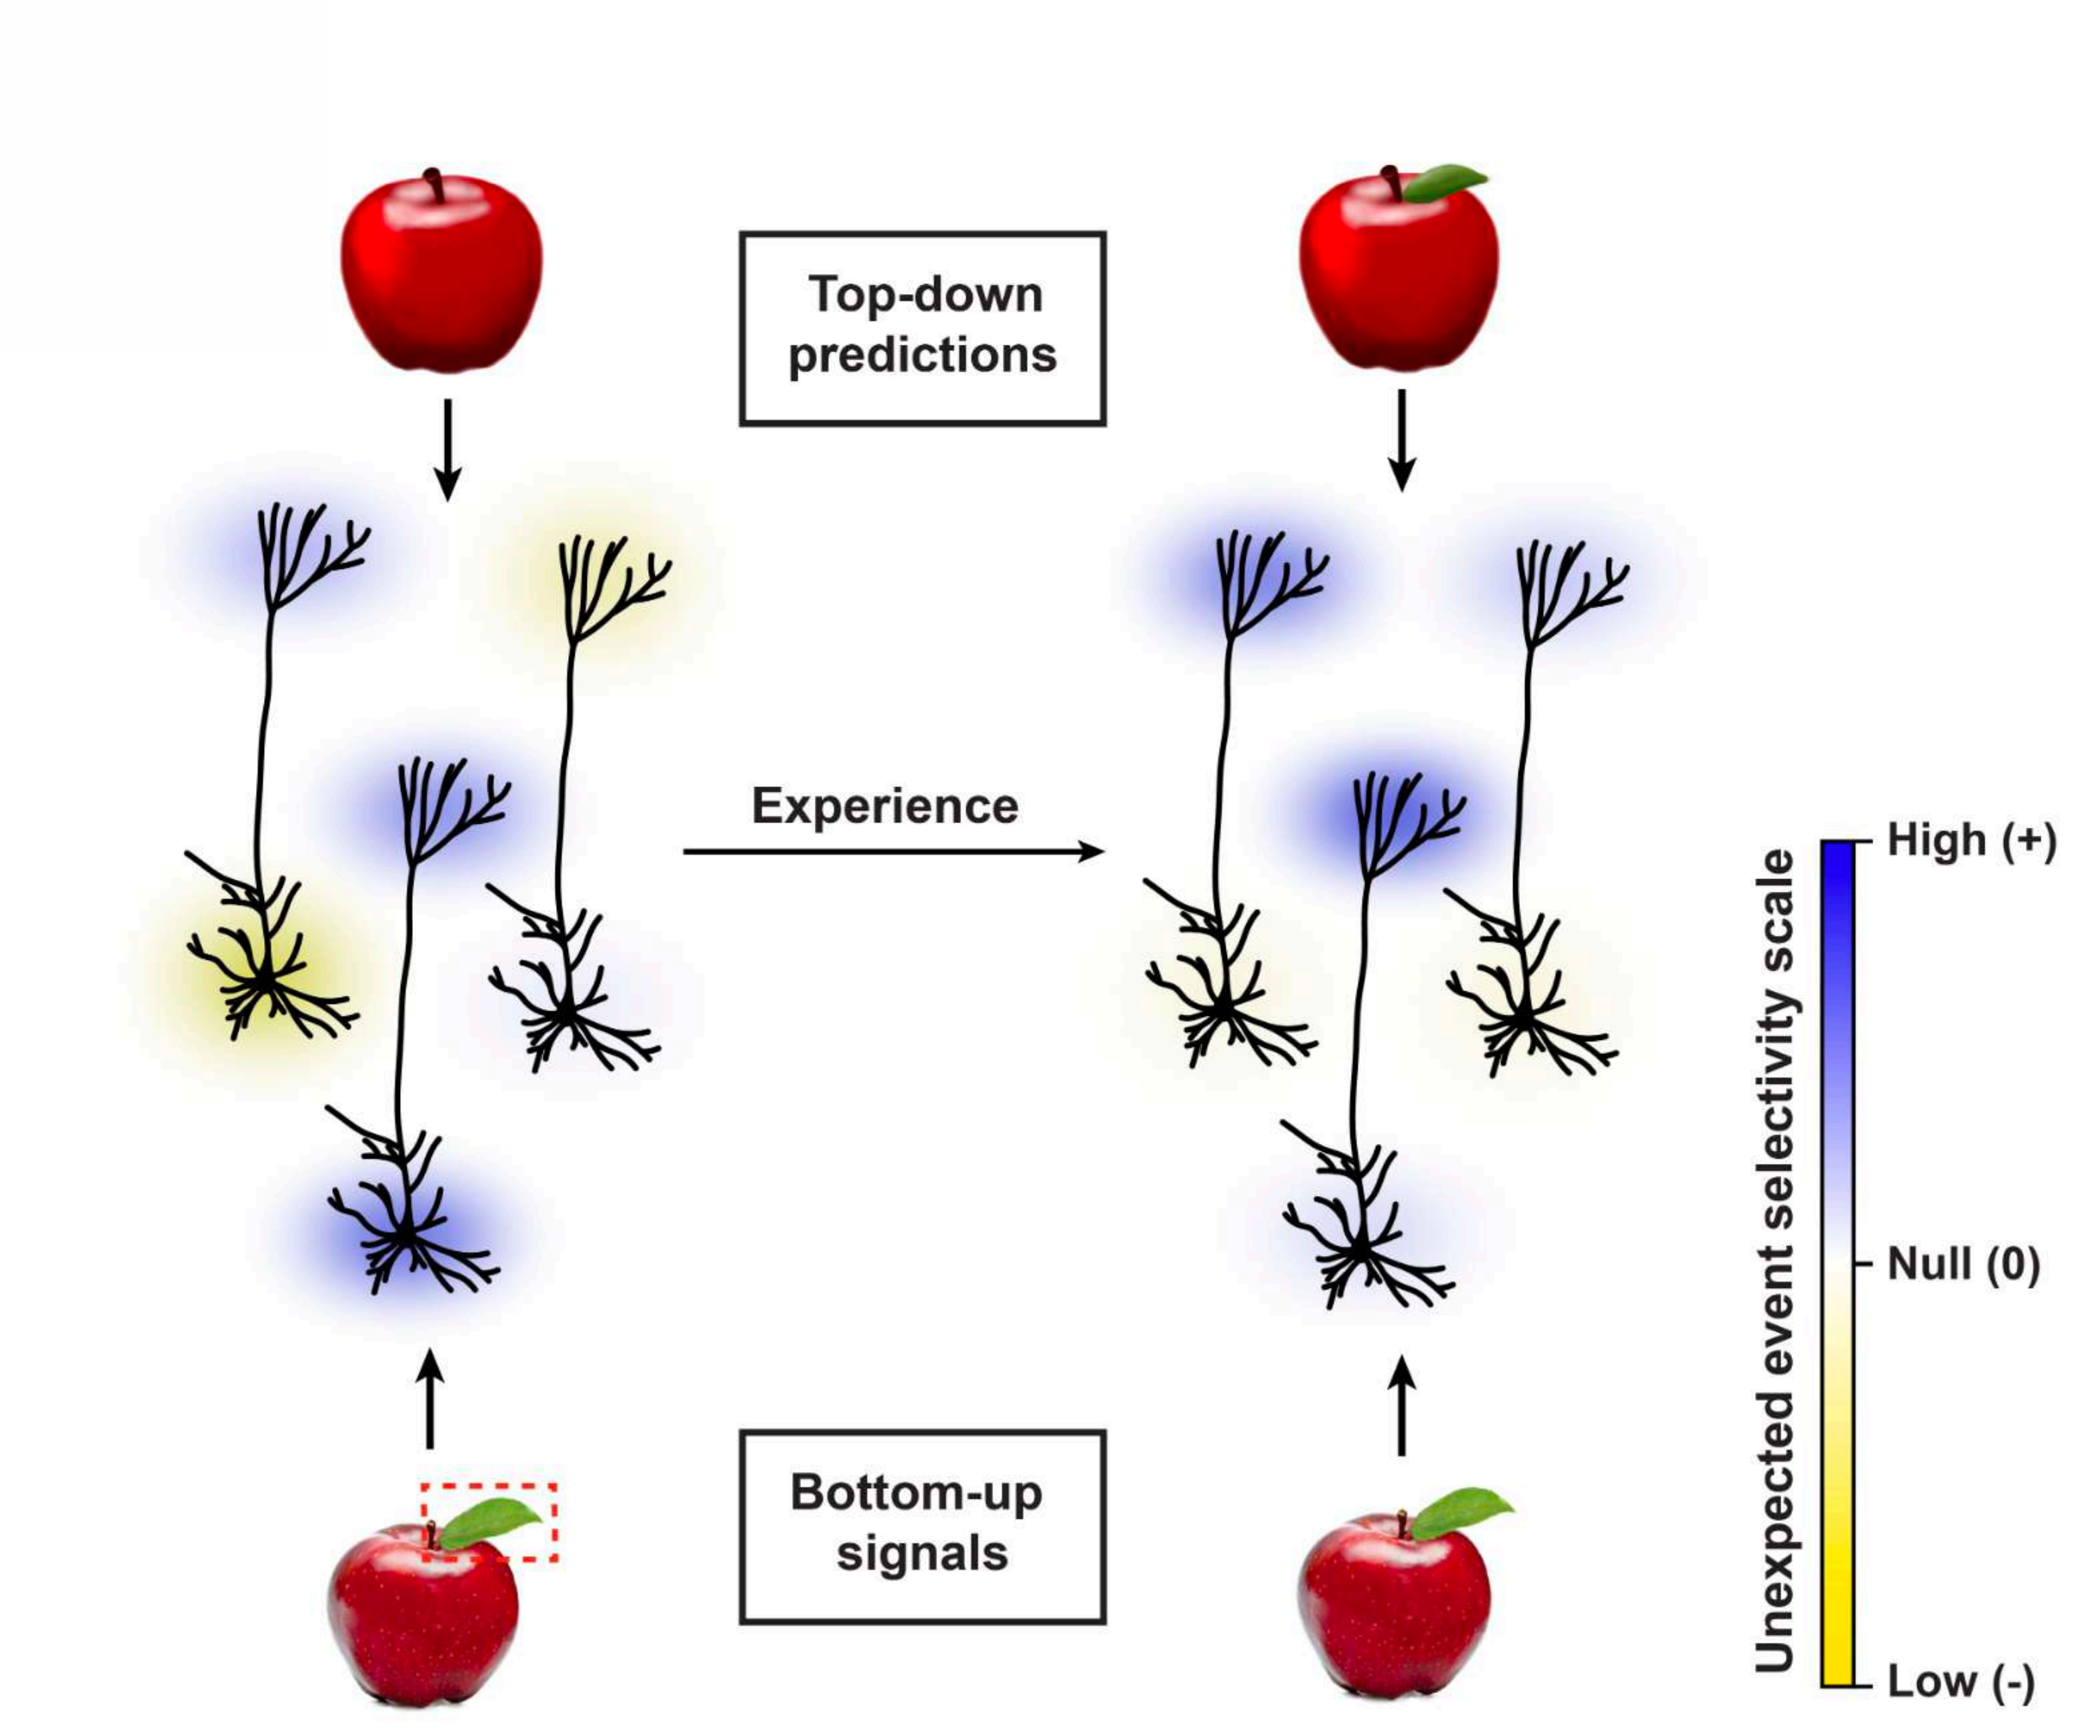

Acknowledgements

Lead labs

Bard Ermentrout† (Univ. Pittsburgh)
Joel Zylberberg† (York University)
Blake Richards† (McGill University)

Collaborators

Colleen Gillon* (University of Toronto)
Yoshua Bengio (MILA)
Timothy Lillicrap (DeepMind)
Timothy M. Henley (York University)

YORK U UNIVERSITÉ UNIVERSITY

University of Pittsburgh

McGill

UNIVERSITY OF TORONTO SCARBOROUGH

Mila

ALLEN INSTITUTE for BRAIN SCIENCE

Collaborators

Jérôme Lecoq*	Kyla Mace
Ruweida Ahmed	Chelsea Nayan
Yazan N. Billeh	Thuyanh V. Nguyen
Shiella Caldejon	Kat North
Peter Groblewski	Jed Perkins
India Kato	Sam Seid
Eric Lee	Matthew Valley
Jennifer Luviano	Ali Williford

National Science Foundation

CIFAR CANADIAN INSTITUTE FOR ADVANCED RESEARCH

NSERC CRSNG

DeepMind

VECTOR INSTITUTE