MAKING SPIKING NEURONS MORE SUCCINCT WITH MULTI-COMPARTMENT MODELS

NICE workshop, 18.3.2021

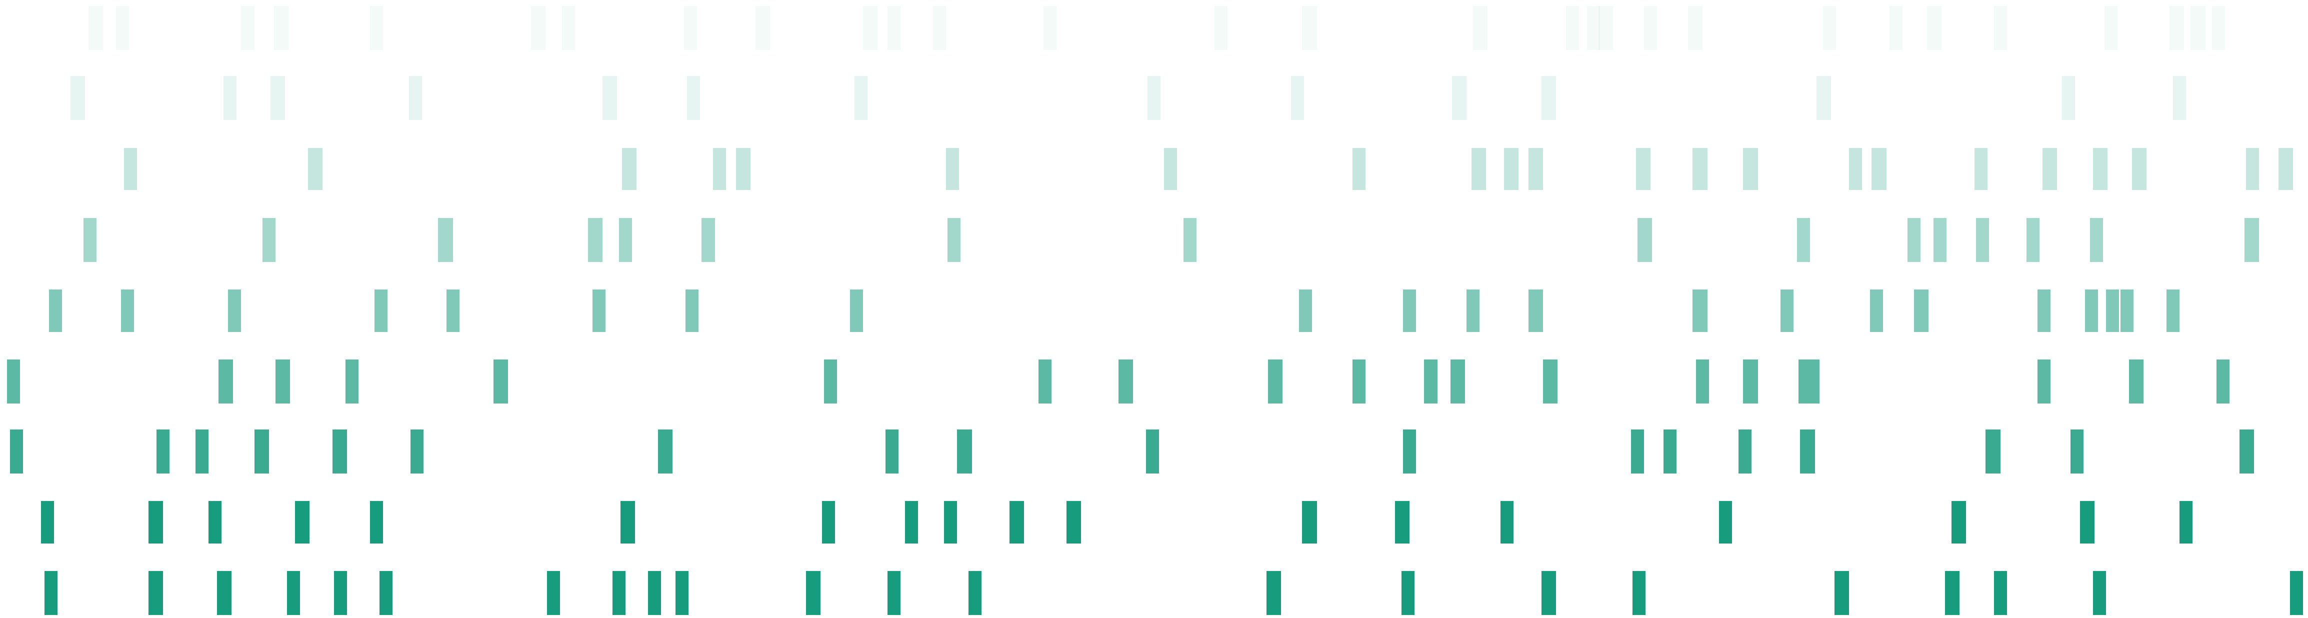

public

UNIVERSITÄT OSNABRÜCK
Fraunhofer IIS

Introductions

- Osnabrück University *(UOS)*
  - Chair for Neuroinformatics *(NI)*
    - Prof. Gordon Pipa
    - Pascal Nieters
    - Johannes Leugering

- Fraunhofer Society *(FhG)*
  - Institute for Integrated Circuits *(IIS)*
    - Communication Systems Division *(KS)*
      - Broadband and Broadcast Department (BB)
        - Embedded AI Group *(eAI)*
          - Johannes Leugering

Insights from Neurobiology

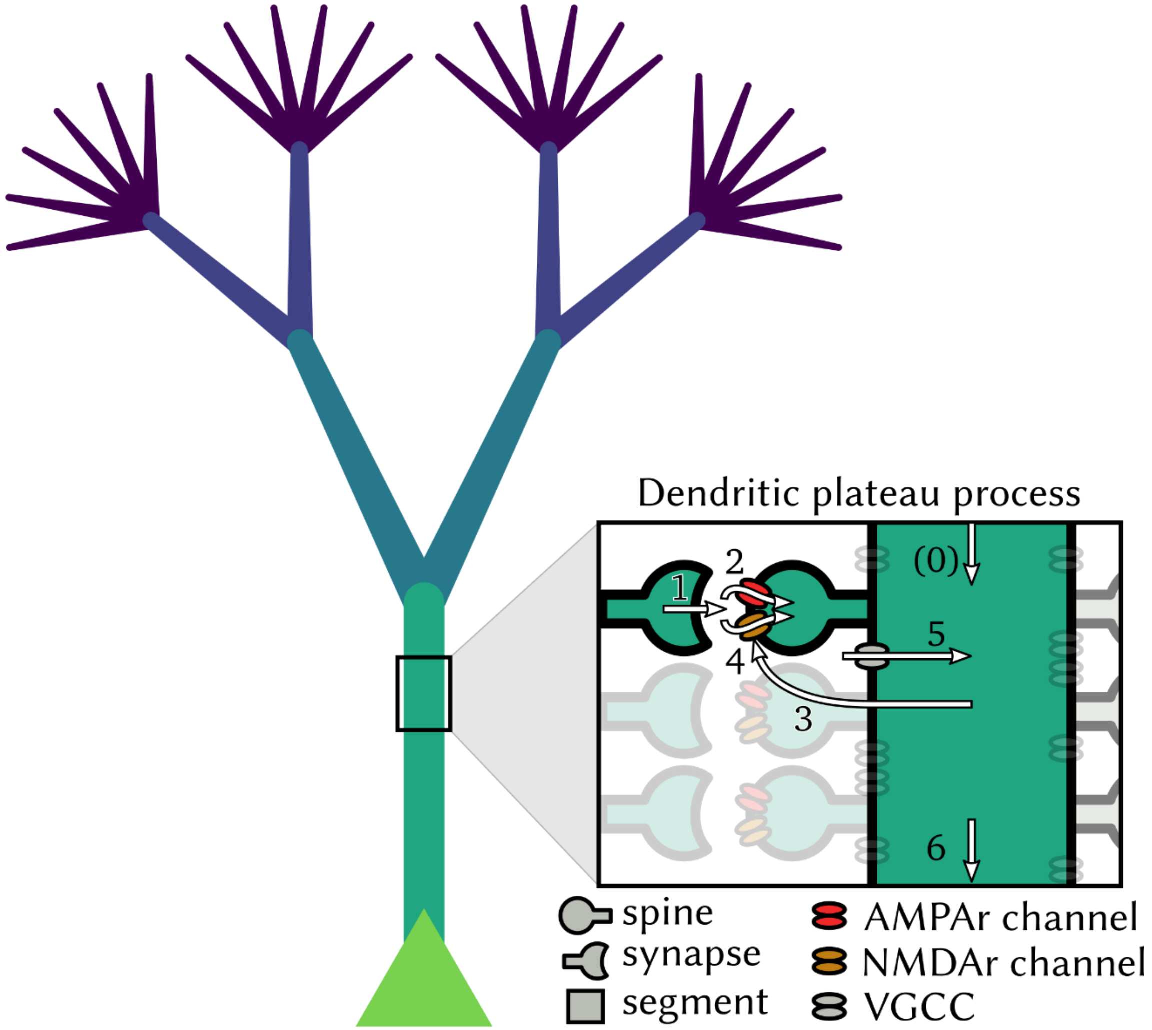

- Active dendritic processes (**plateau potentials**)

Insights from Neurobiology

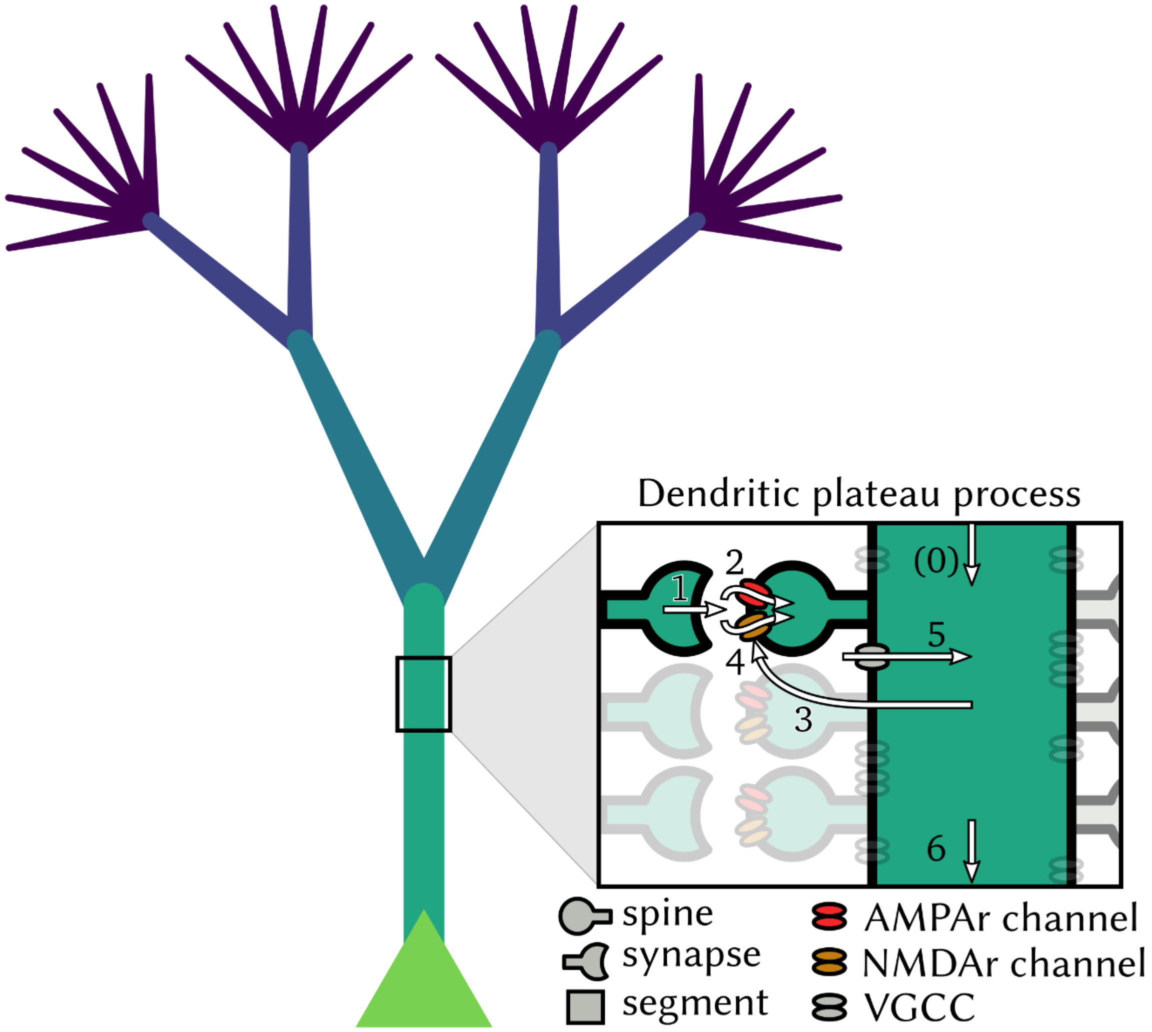

- Active dendritic processes (**plateau potentials**)
  - localized depolarization

Insights from Neurobiology

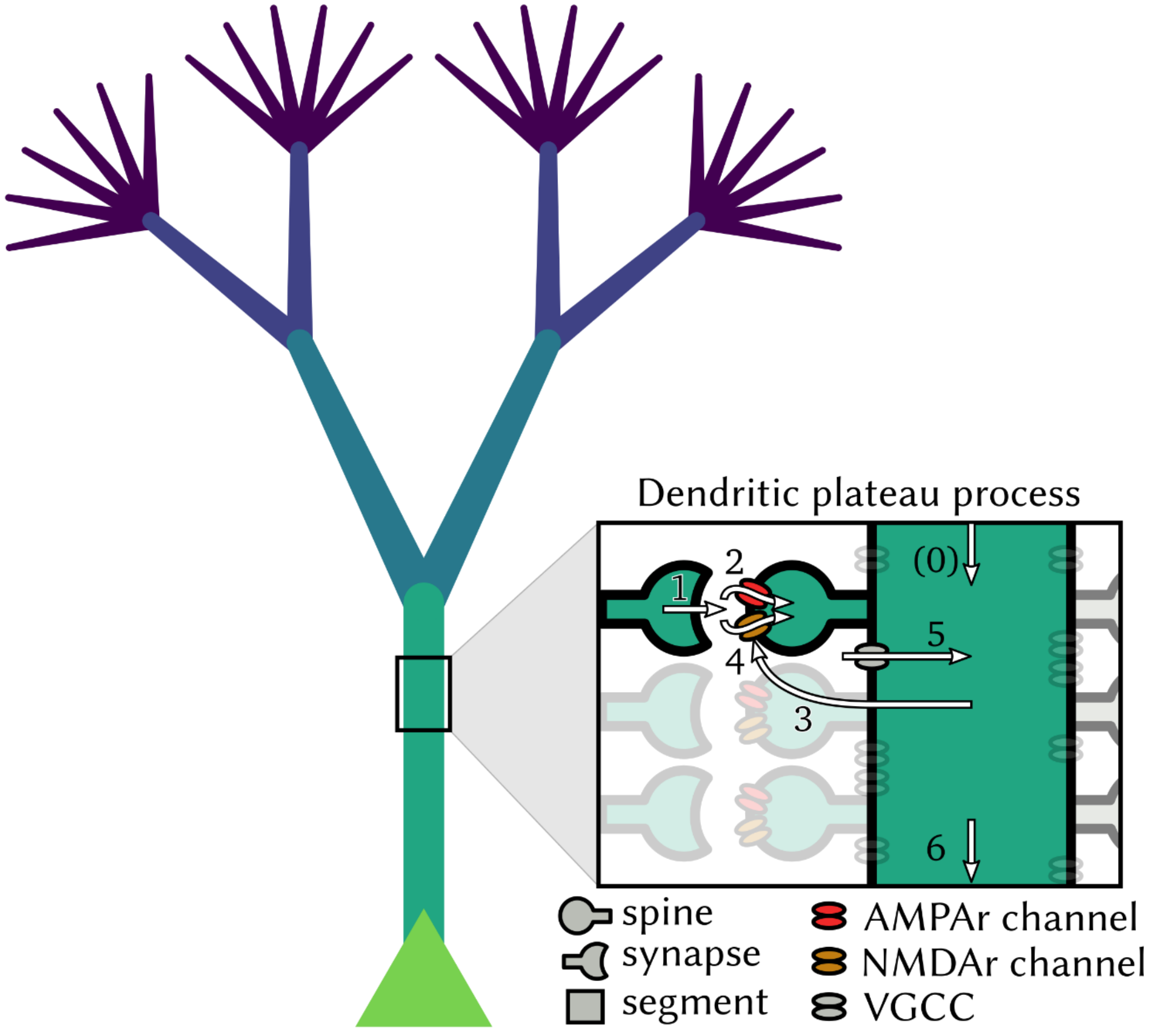

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)

Insights from Neurobiology

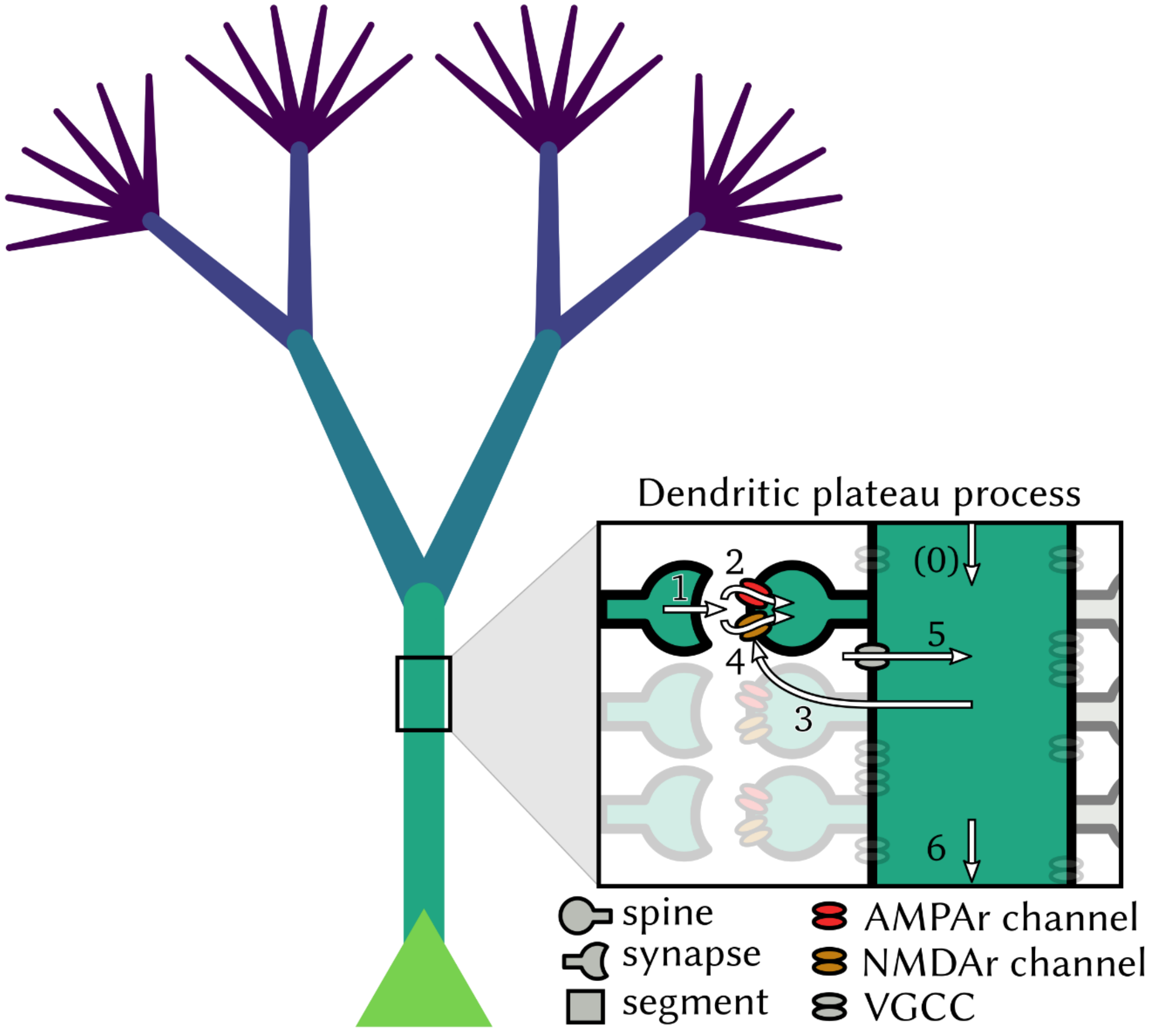

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state")

Insights from Neurobiology

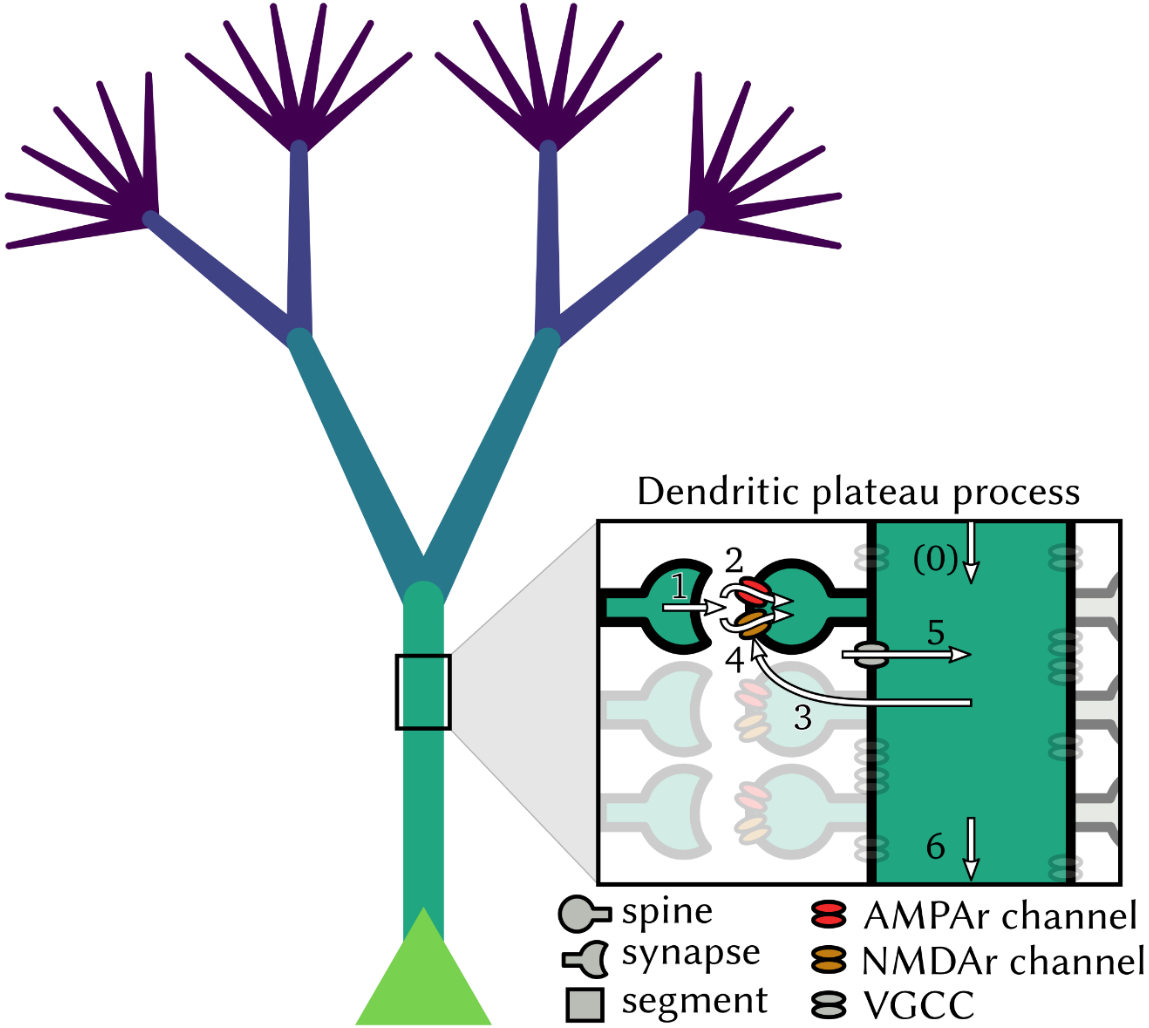

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state")
  - are ubiquitious & critical for learning

Insights from Neurobiology

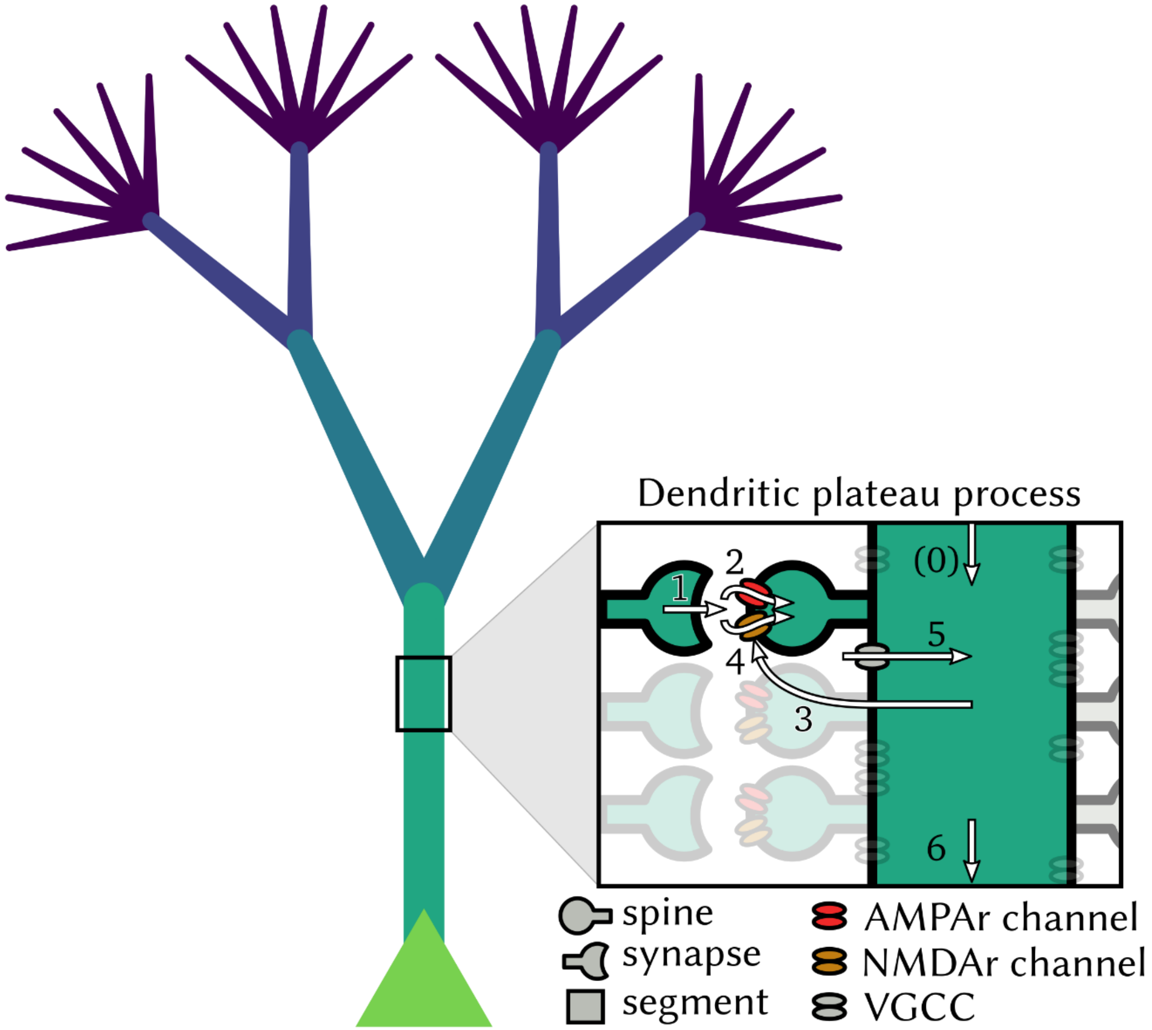

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)

Insights from Neurobiology

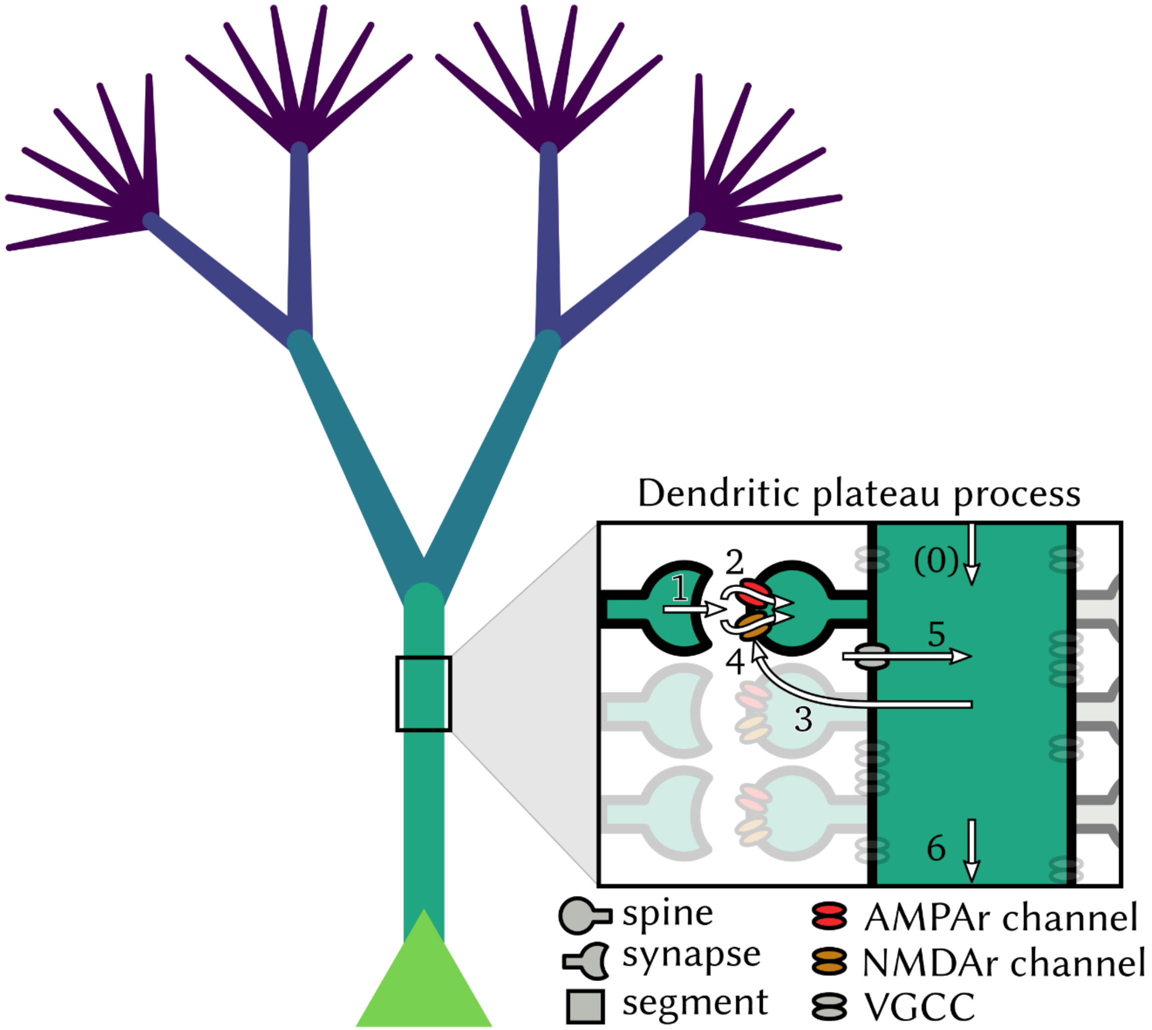

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state")
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**

Insights from Neurobiology

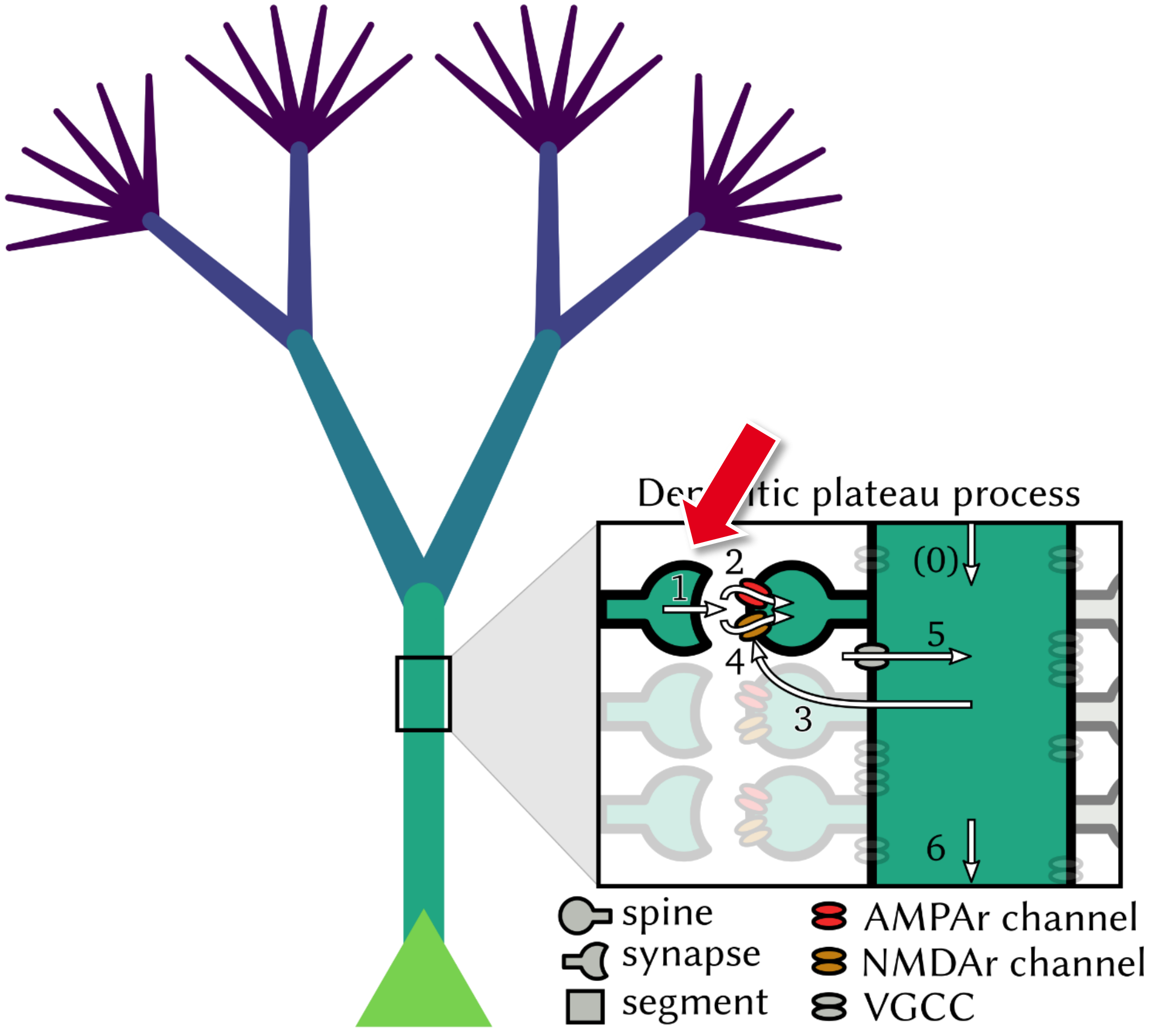

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter

Insights from Neurobiology

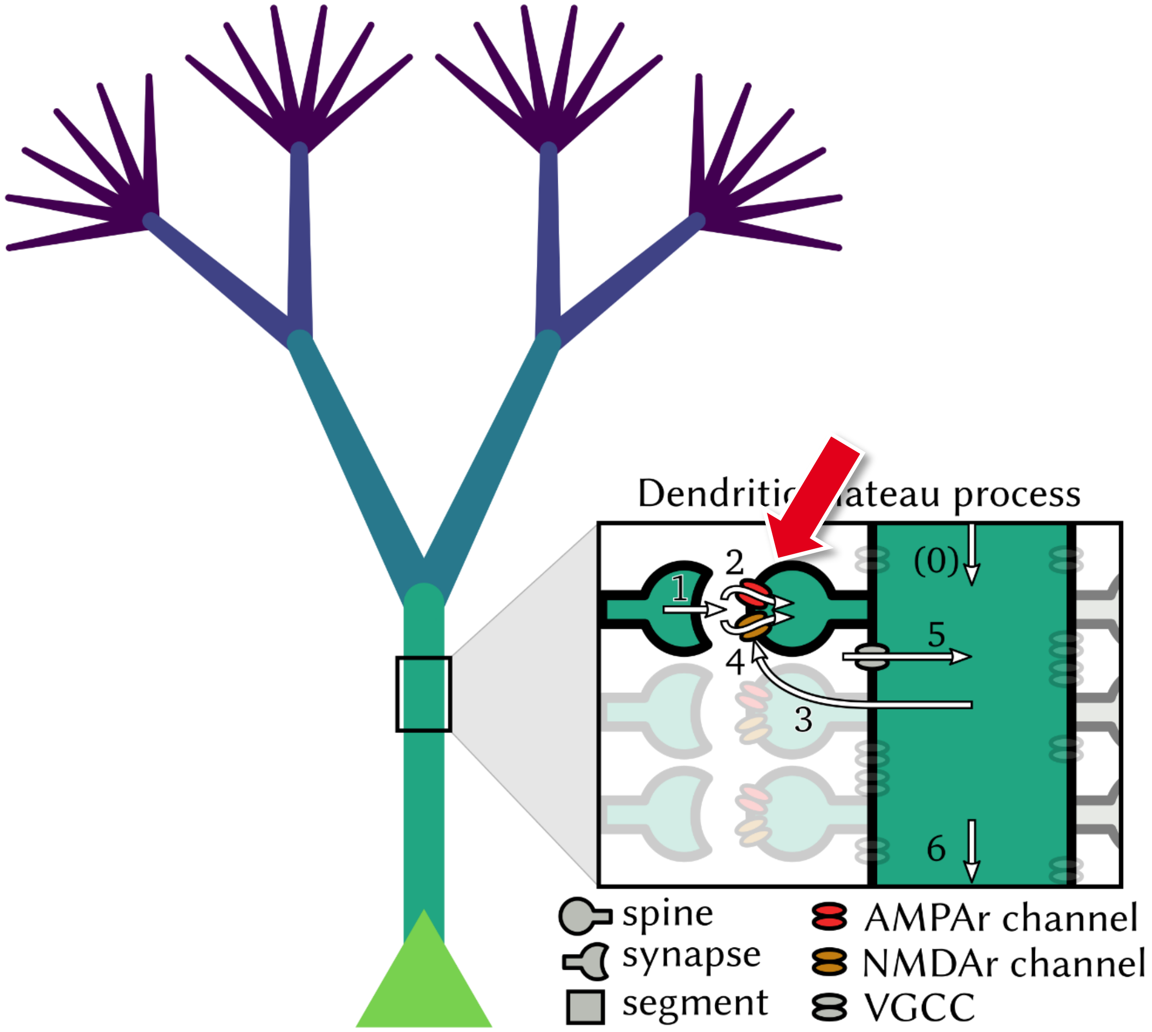

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter
  2. depolarize dendrite (AMPAr)

public

Insights from Neurobiology

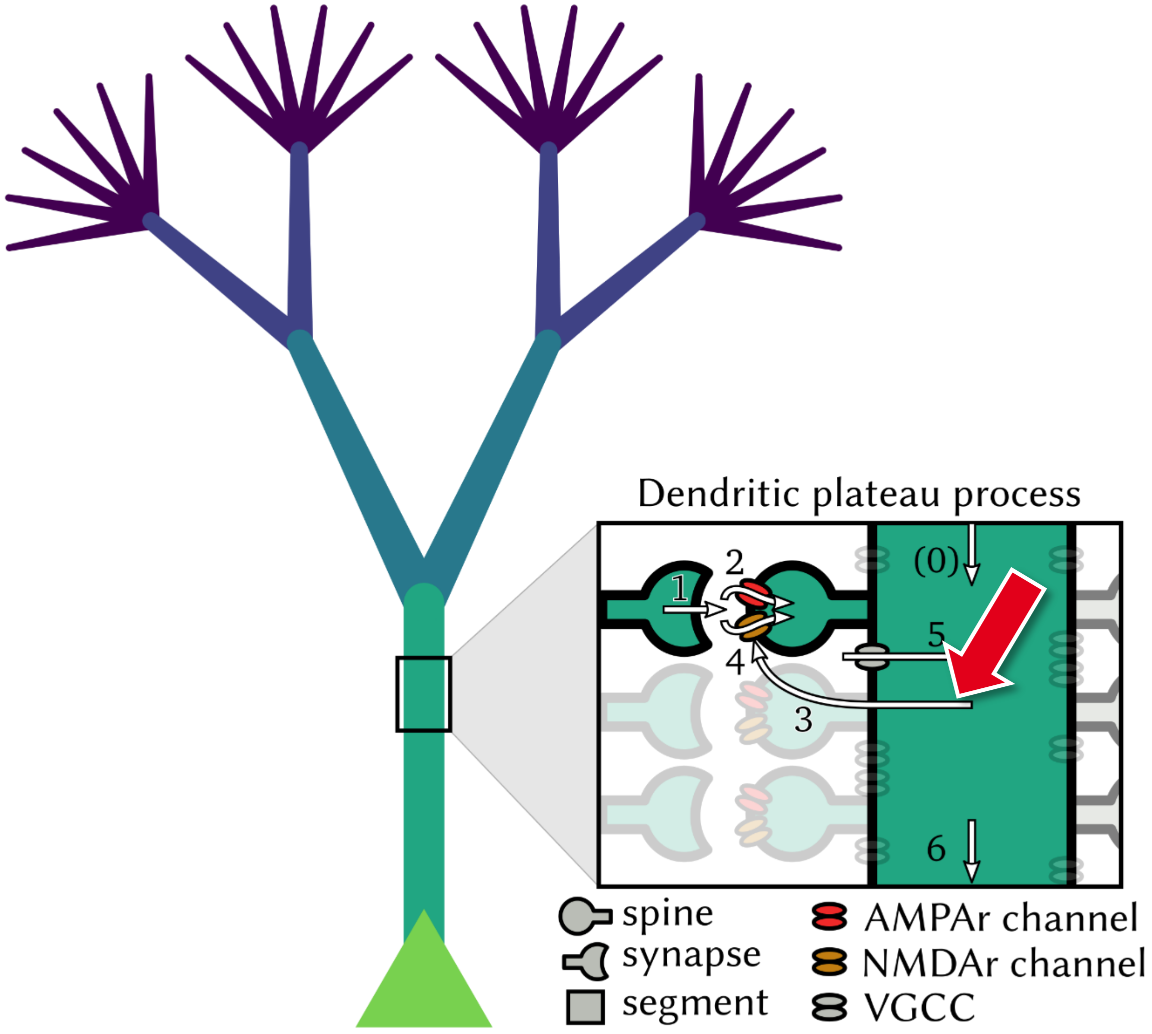

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter
  2. depolarize dendrite (AMPAr)
  3. localized coincidence detection

Insights from Neurobiology

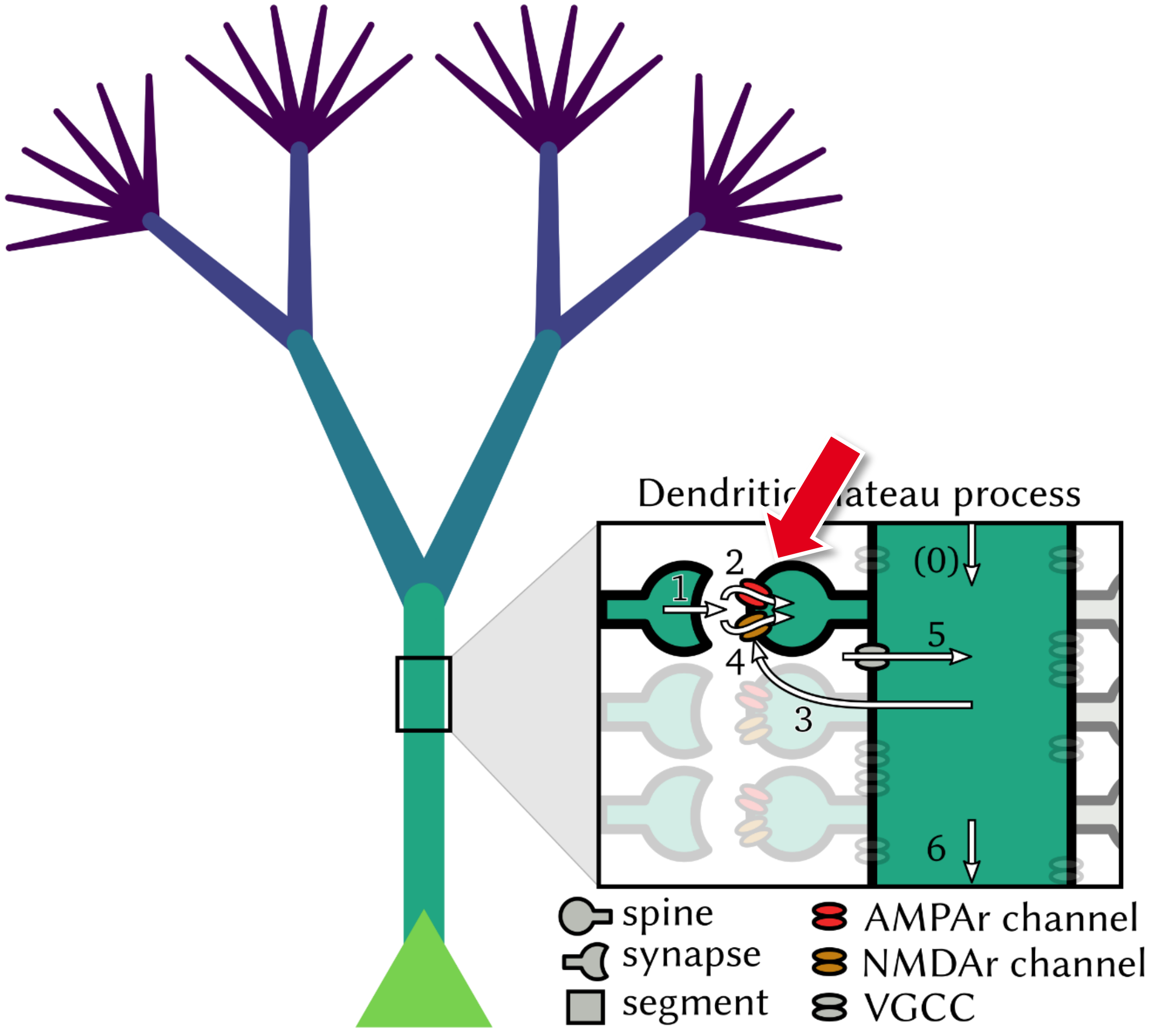

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state")
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter
  2. depolarize dendrite (AMPAr)
  3. localized coincidence detection
  4. triggers plateau (NMDAr)

Insights from Neurobiology

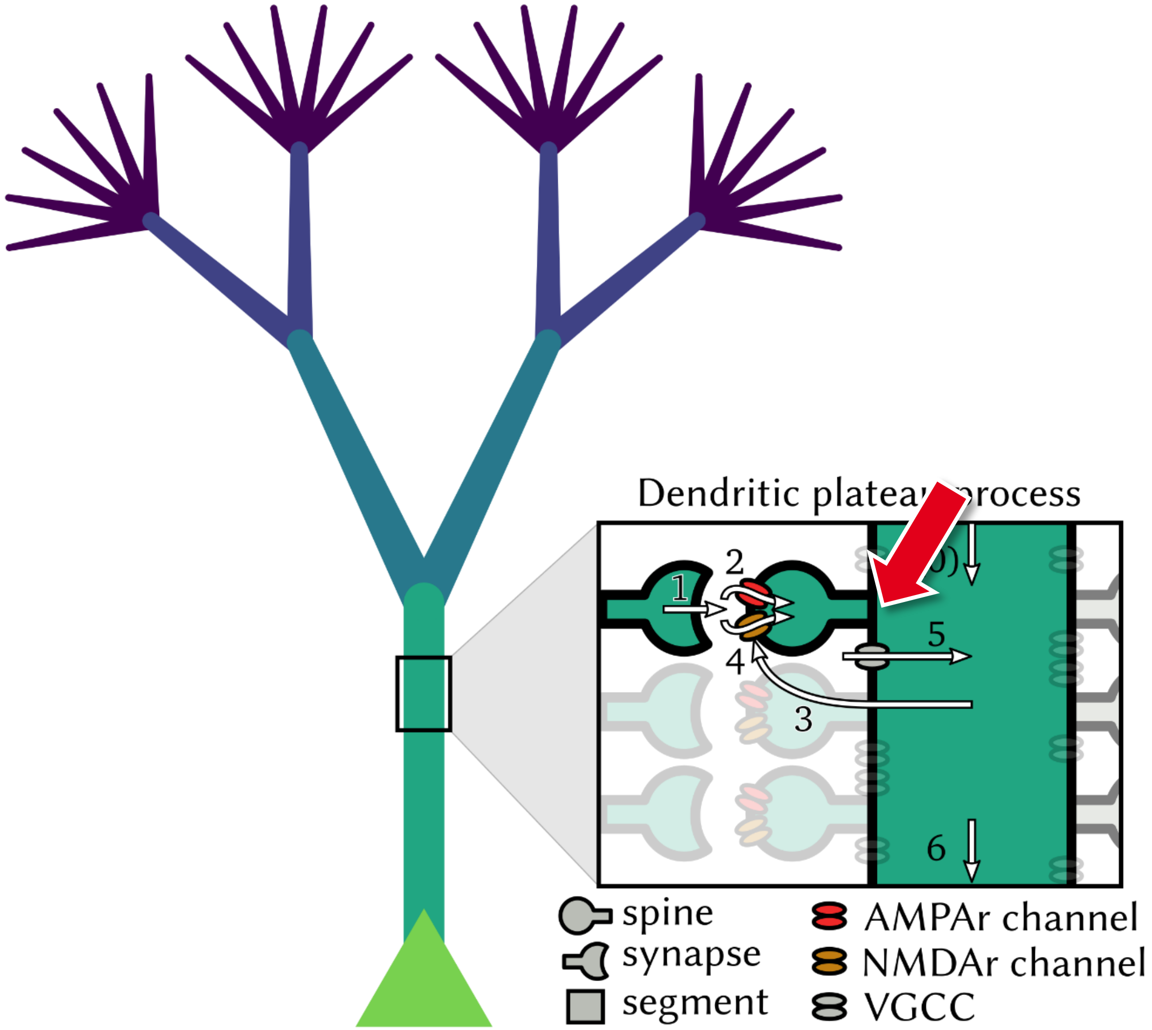

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter
  2. depolarize dendrite (AMPAr)
  3. localized coincidence detection
  4. triggers plateau (NMDAr)
  5. self-sustained activation (VGCC)

Insights from Neurobiology

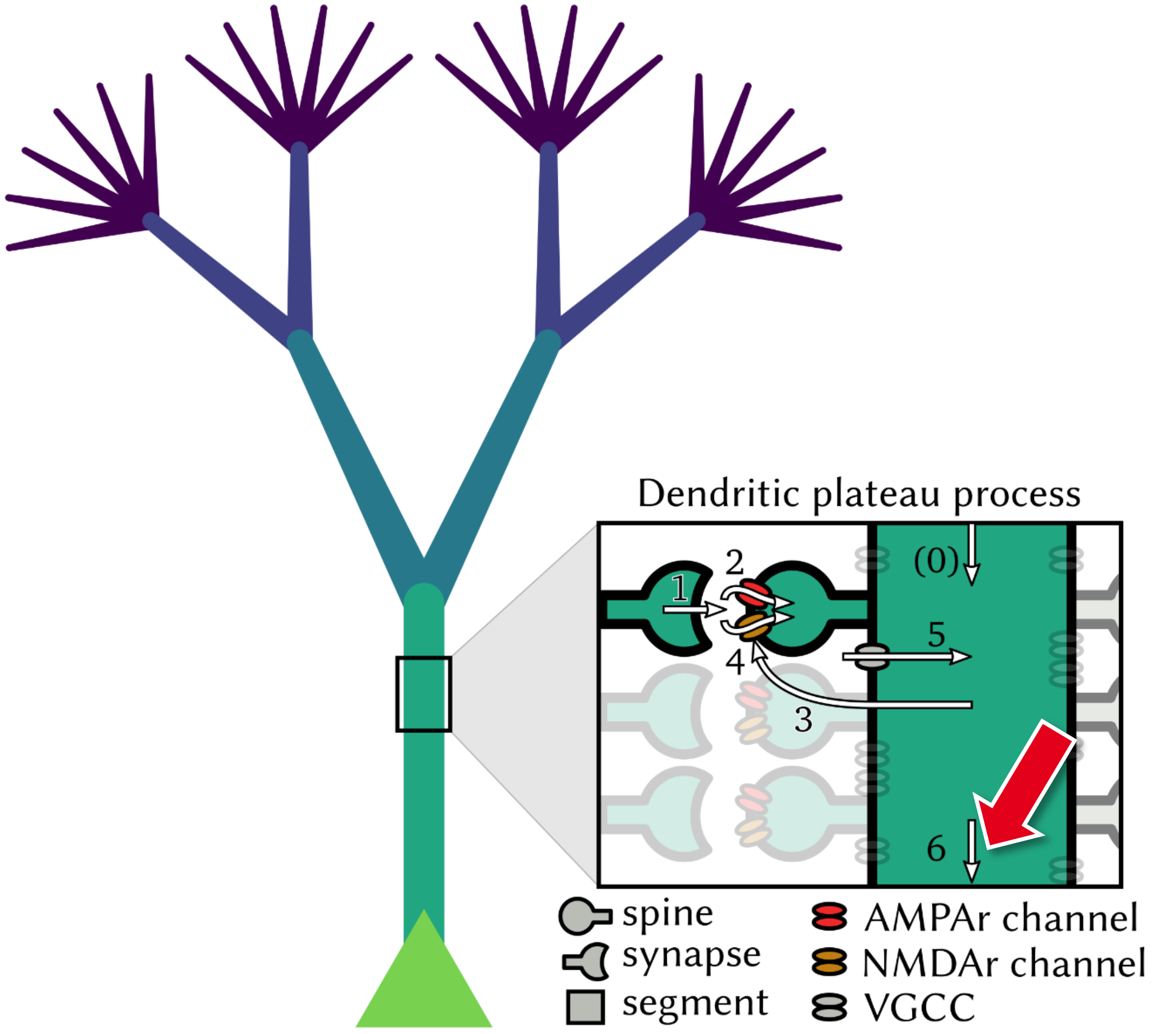

- Active dendritic processes (**plateau potentials**)
  - localized depolarization
  - last long (>100ms)
  - increase probability to fire („UP-state“)
  - are ubiquitious & critical for learning
  - strangely absent from most neuron models (notable exceptions: Hawkins, Senn, Heusser)
- **Biological mechanism (simplified):**
  1. spikes release neurotransmitter
  2. depolarize dendrite (AMPAr)
  3. localized coincidence detection
  4. triggers plateau (NMDAr)
  5. self-sustained activation (VGCC)
  6. Diffusion of depolarization along dendrite

A simplified model

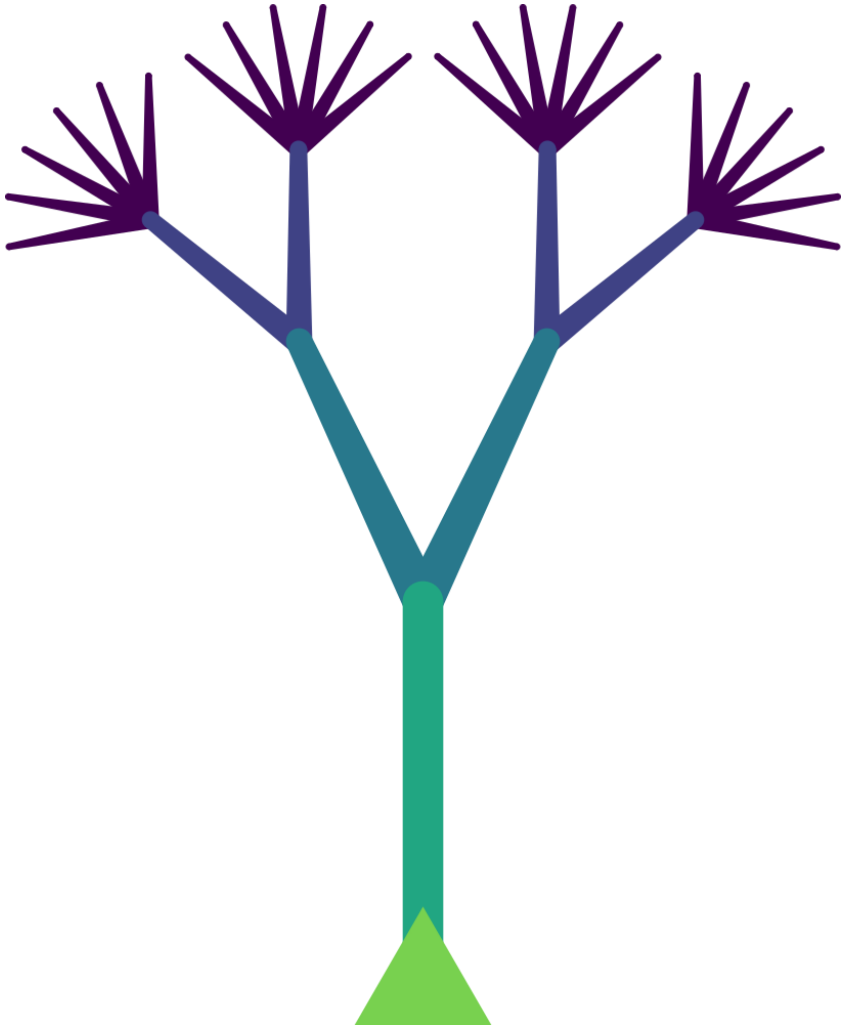

- A **neuron** is a tree of dendrite **segments**

A simplified model

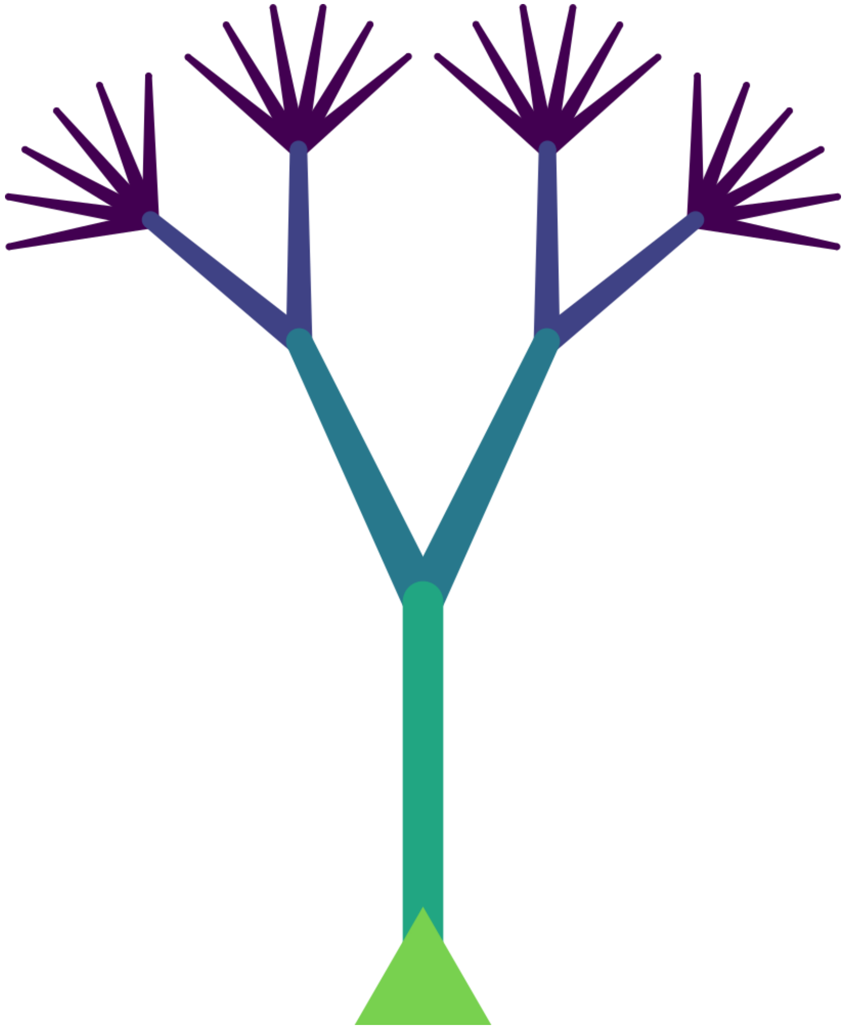

- A **neuron** is a tree of dendrite **segments**
- Segments can fire local **plateau potentials** that
  - are **binary**
  - last for **long** period (~100ms)
  - can be extended (~retriggerable monoflop)
  - depolarize („enable“) downstream neighbour

A simplified model

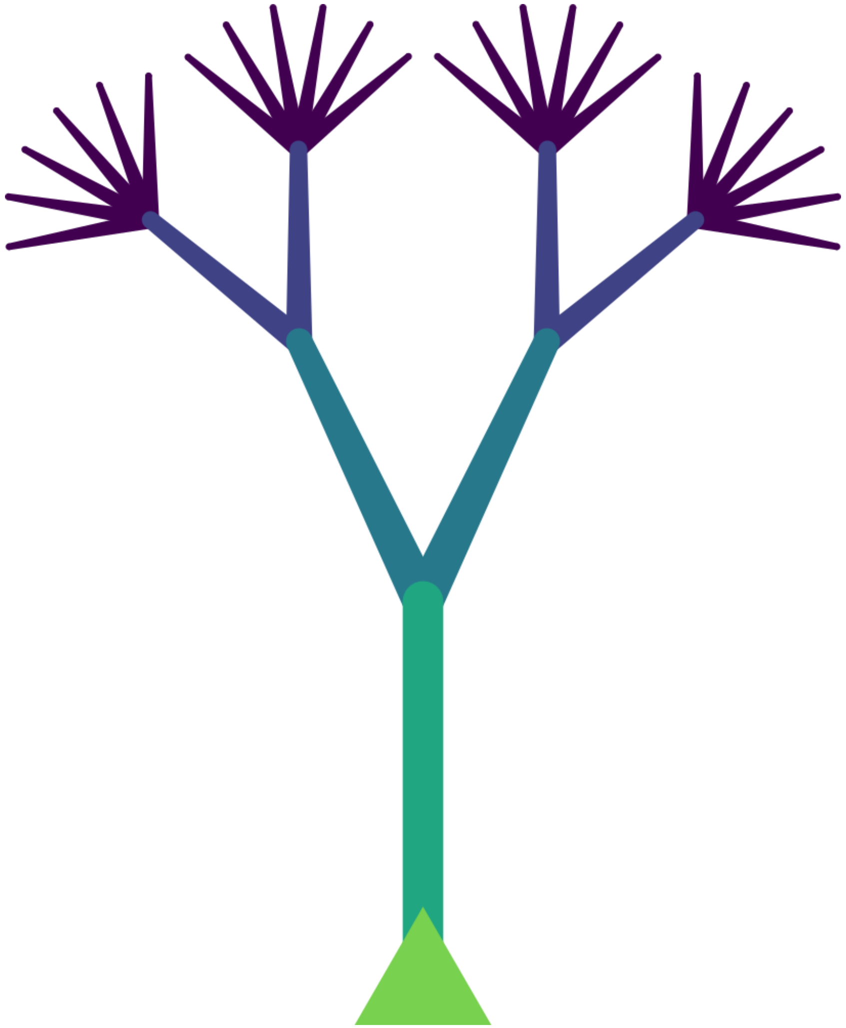

- A **neuron** is a tree of dendrite **segments**
- Segments can fire local **plateau potentials** that
  - are **binary**
  - last for **long** period (~100ms)
  - can be extended (~retriggerable monoflop)
  - depolarize („enable“) downstream neighbour
- Each segment (and the soma) fires if and only if:
  - receives **volley** of **coincident** (~5ms) spikes
  - and **depolarized** by upstream neigbour(s) (or if it is apical dendrite)

A simplified model

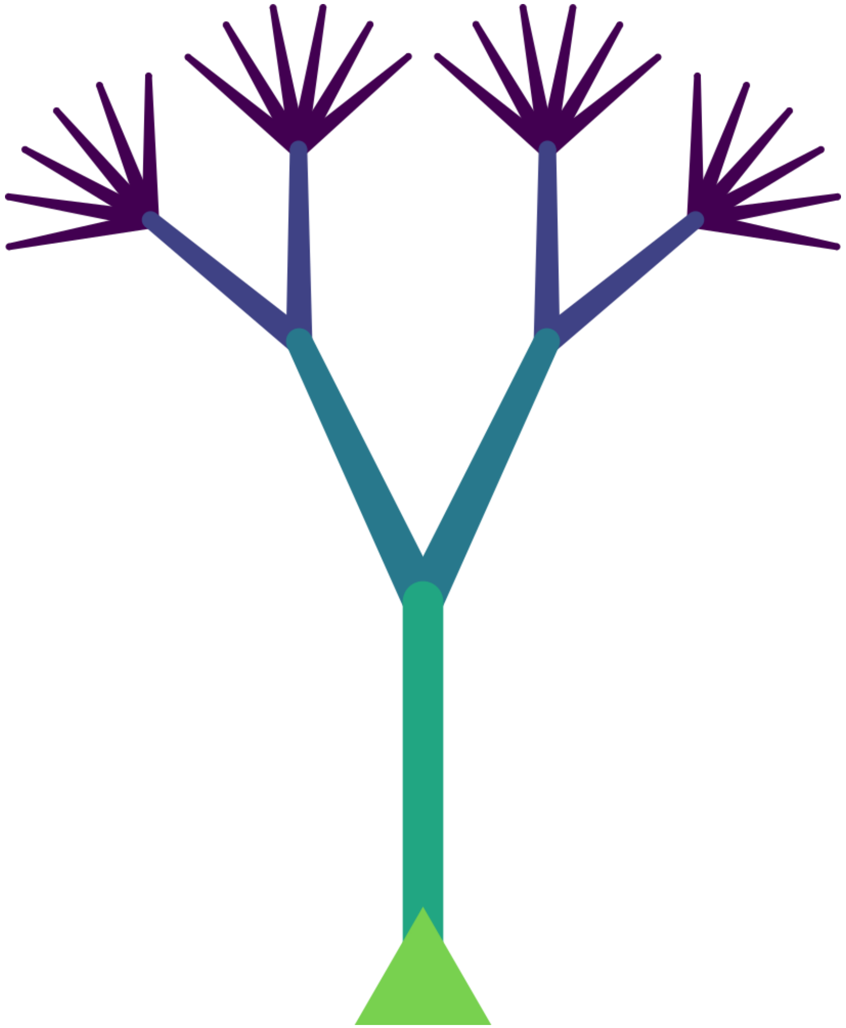

- A **neuron** is a tree of dendrite **segments**
- Segments can fire local **plateau potentials** that
  - are **binary**
  - last for **long** period (~100ms)
  - can be extended (~retriggerable monoflop)
  - depolarize („enable“) downstream neighbour
- Each segment (and the soma) fires if and only if:
  - receives **volley** of **coincident** (~5ms) spikes
  - and **depolarized** by upstream neigbour(s) (or if it is apical dendrite)
- Synapses have a **weight** & **transmission prob.**

Dendritic plateaus provide memory traces

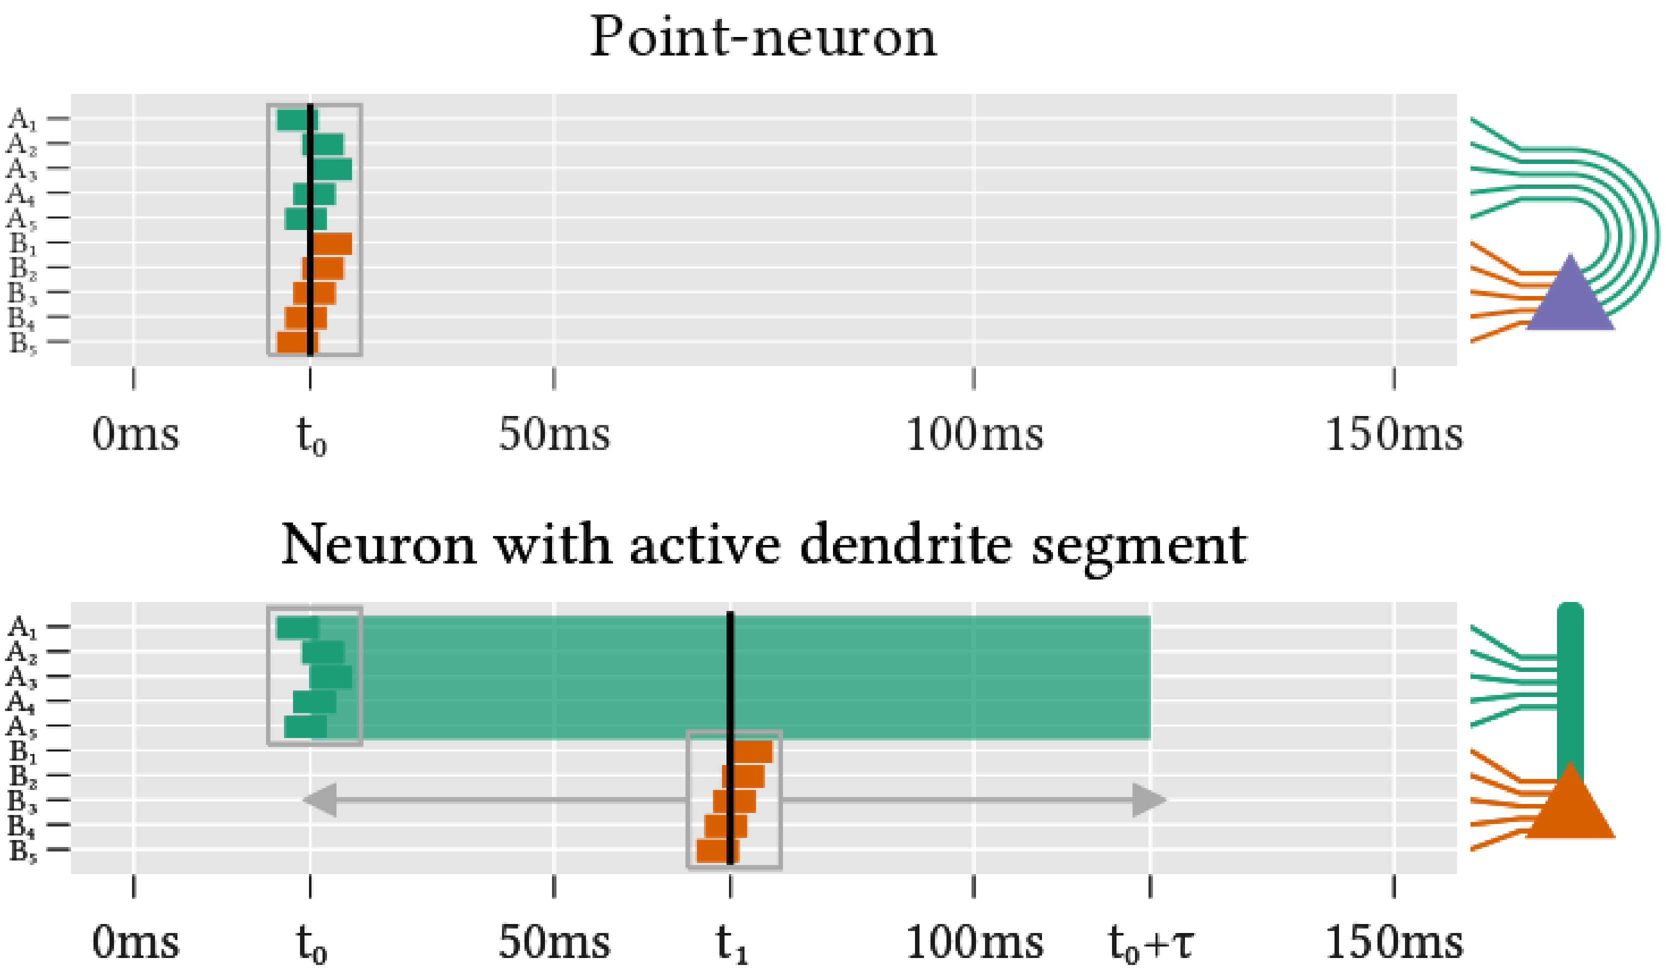

Dendritic plateaus provide memory traces

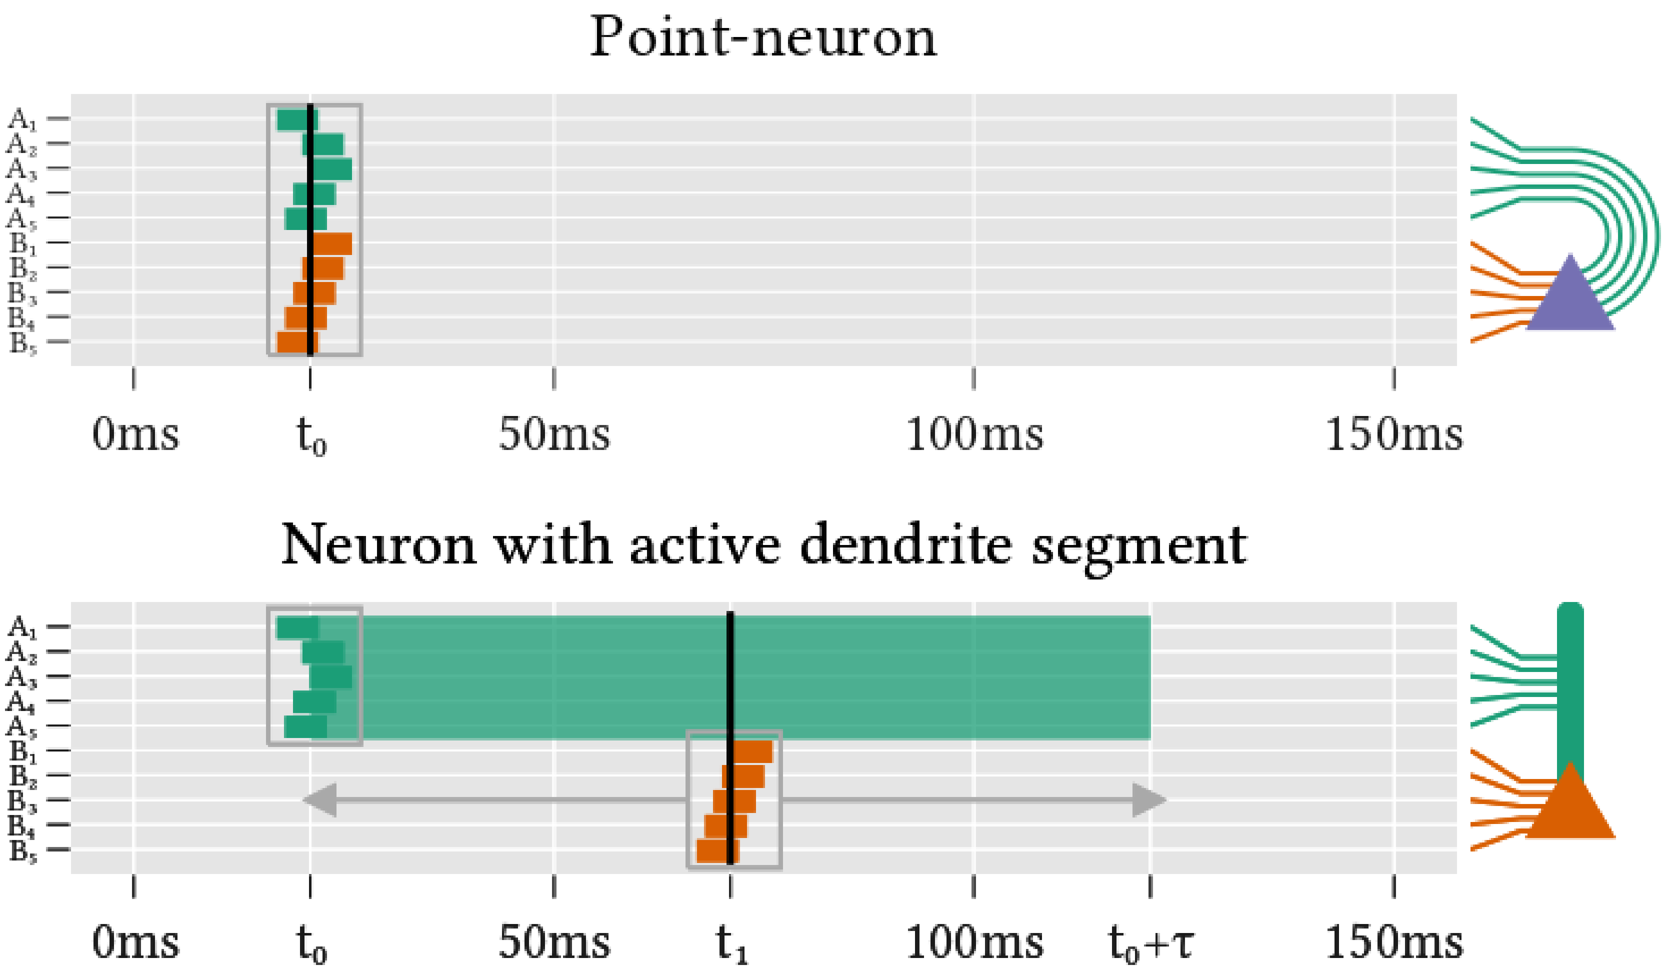

→sensitive to **temporal order**!

→**invariant** to small-scale timing jitter

→combines **fast ESPS timescale** with **slow plateau timescale**

Different topologies detect different patterns

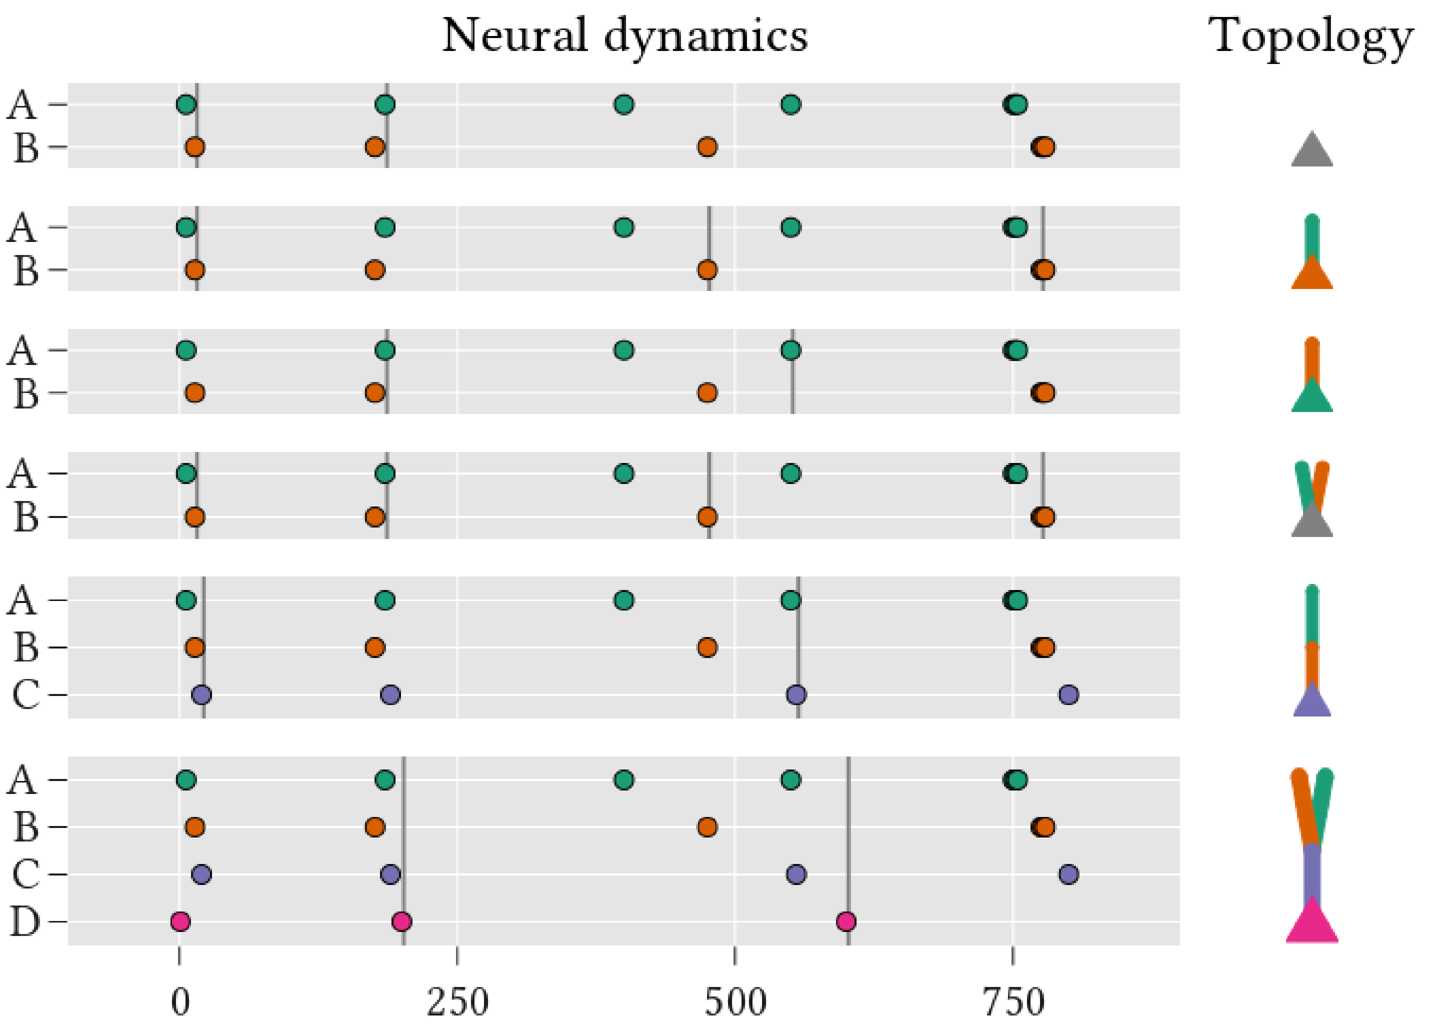

UNIVERSITÄT OSNABRÜCK

Fraunhofer IIS

Different topologies detect different patterns

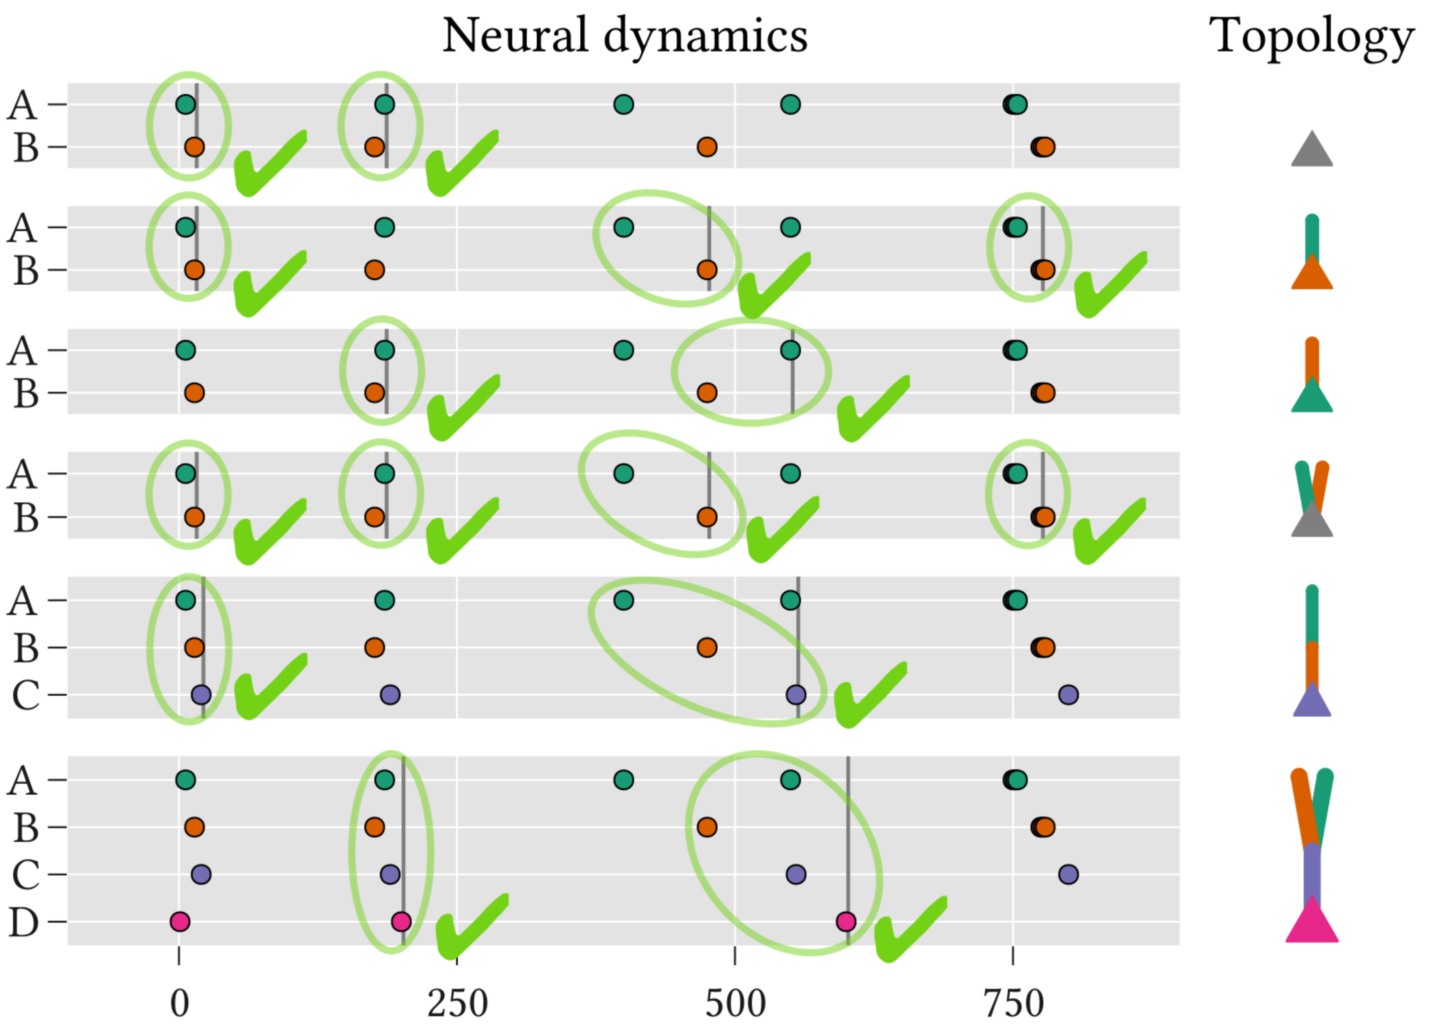

public

Different topologies detect different patterns

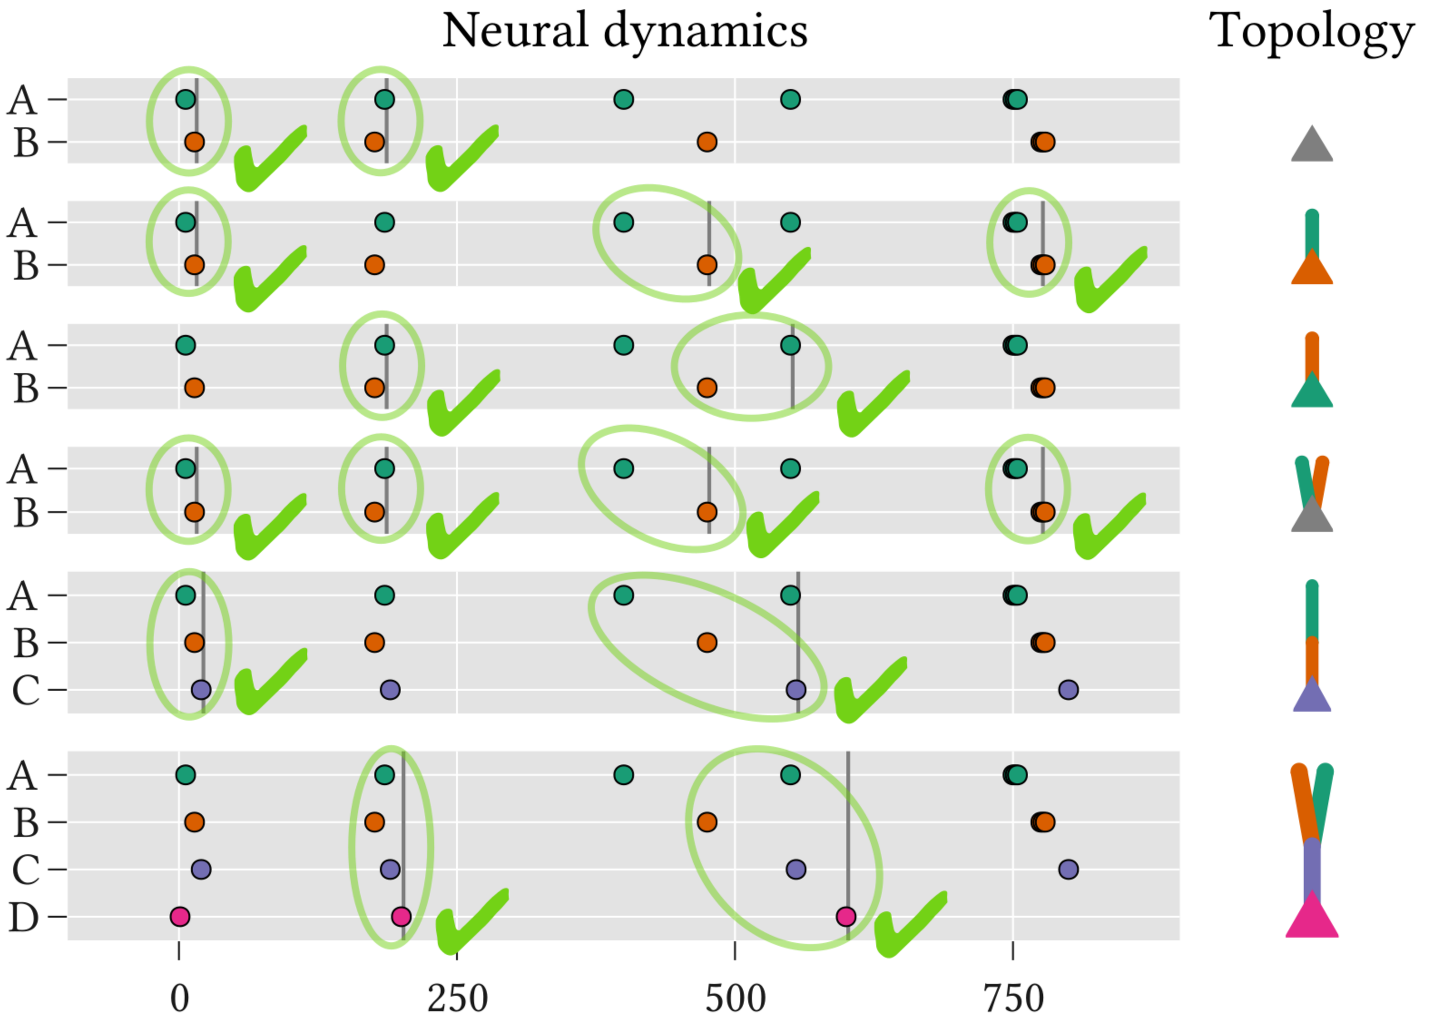

→segments compose like **gates**

→**topology** determines **pattern**

→can provide **longer memory**

Probabilistic synapses allow graded responses

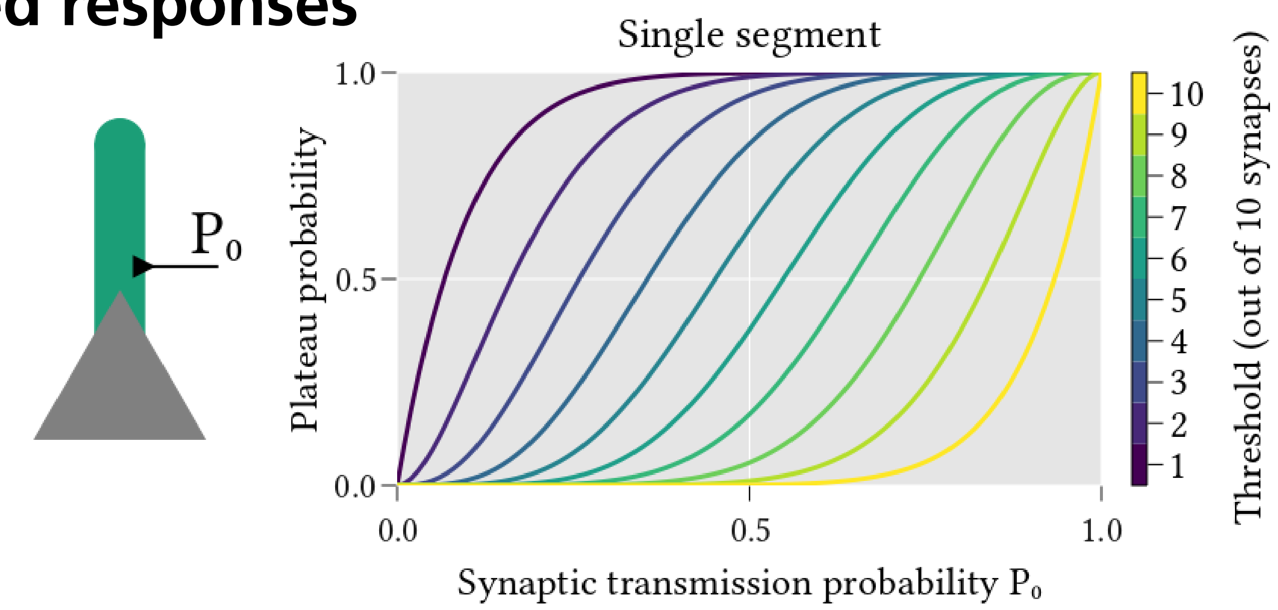

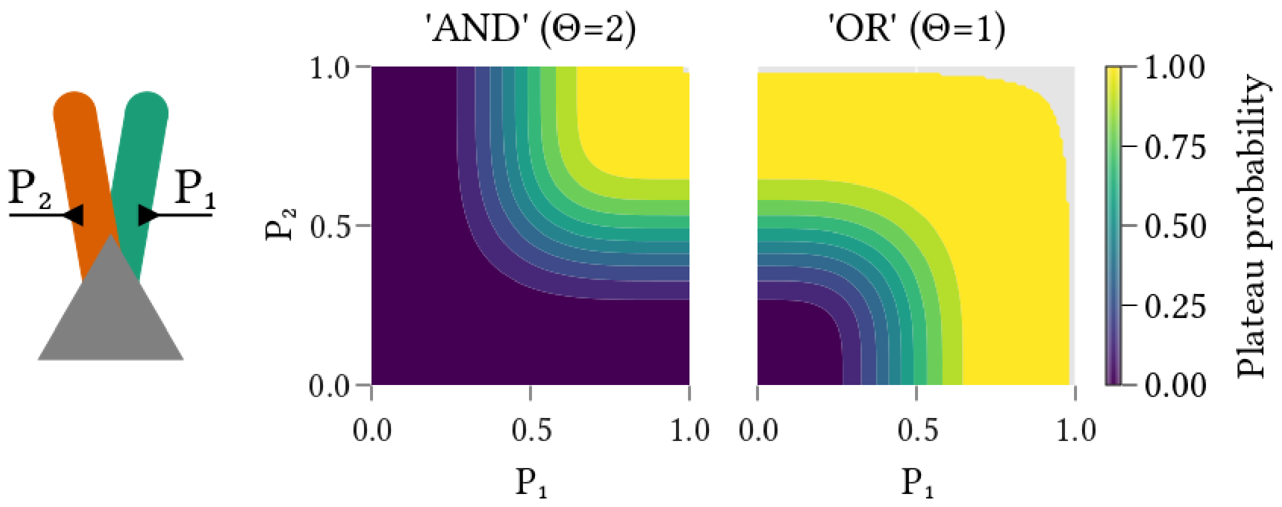

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

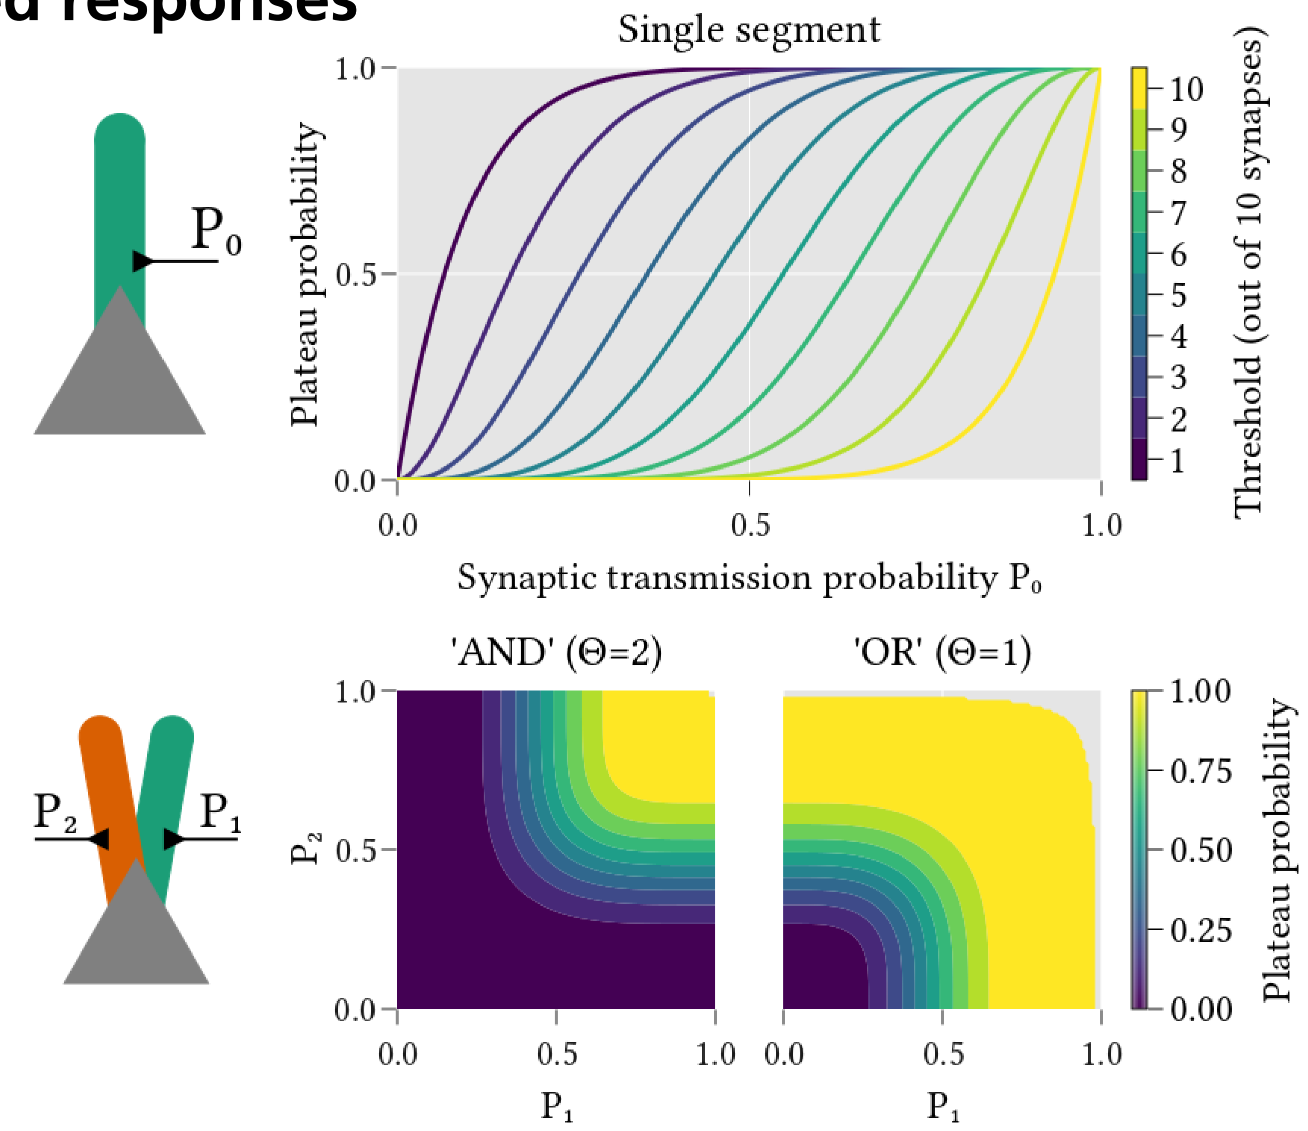

public

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

- → but: **stochastic input** results in **stochastic output**

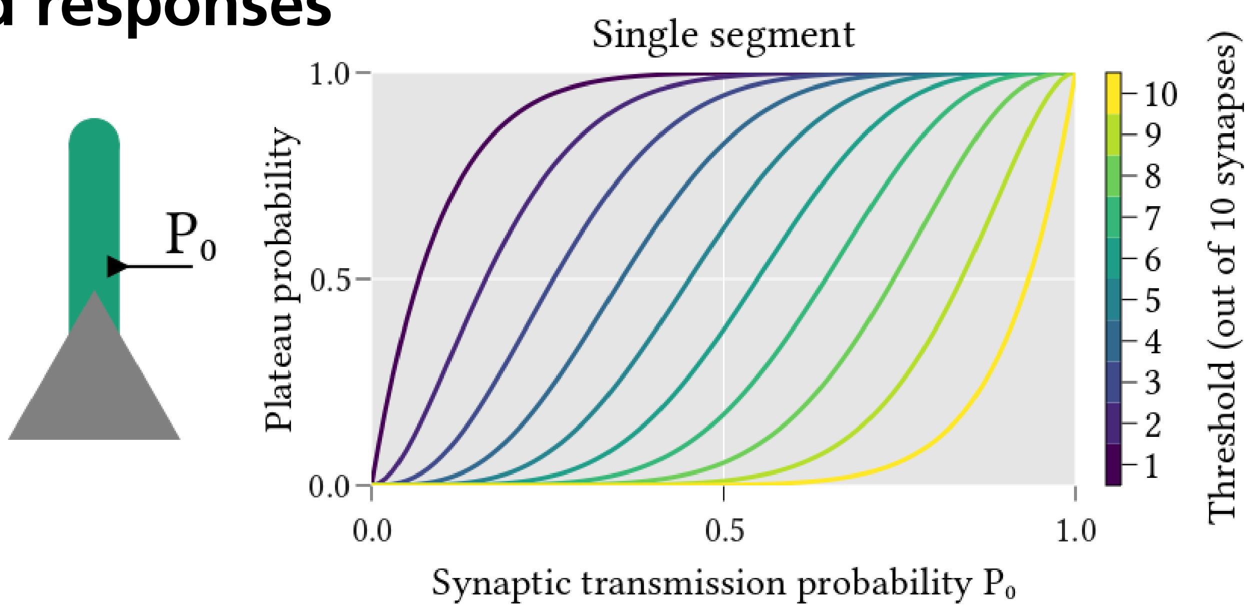

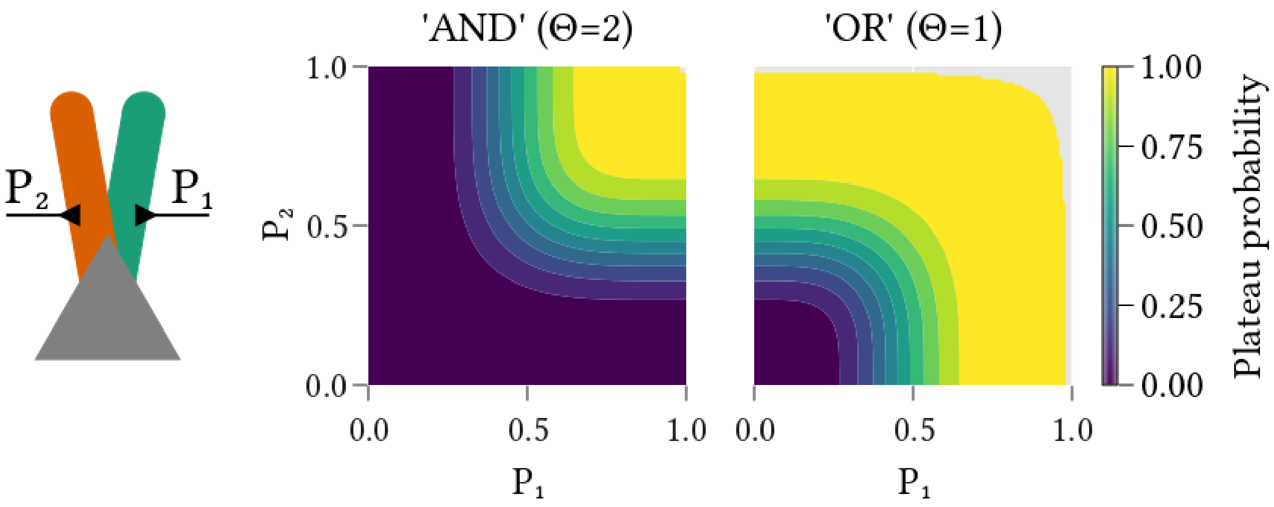

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

- → but: **stochastic input** results in **stochastic output**
- → the **probability** is graded!

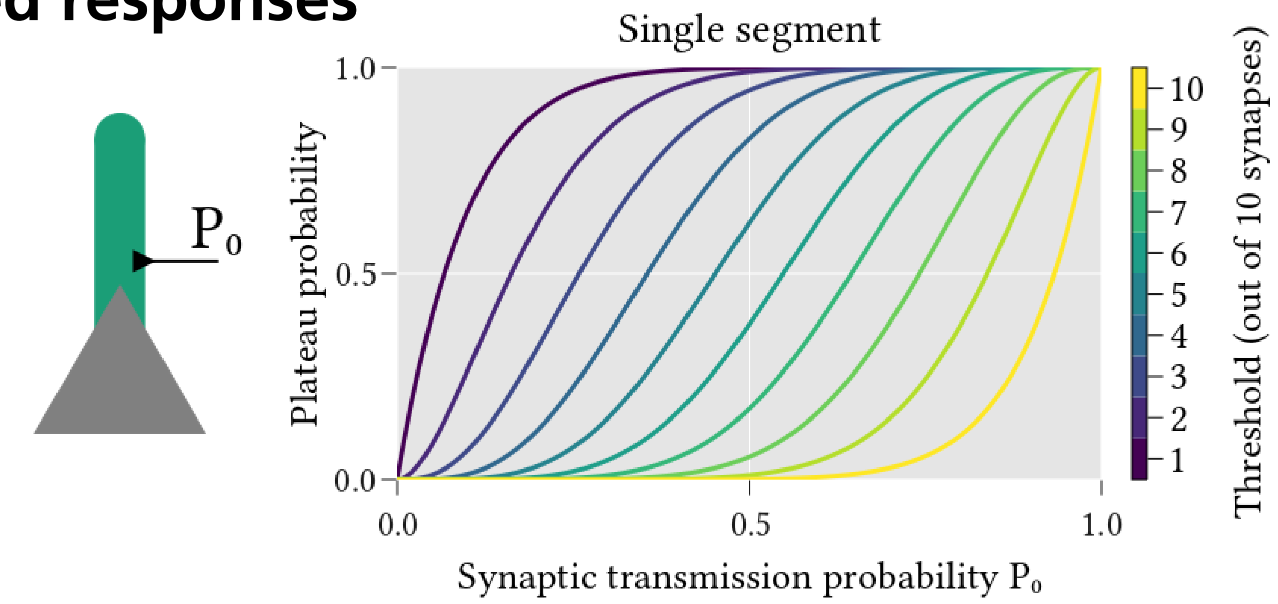

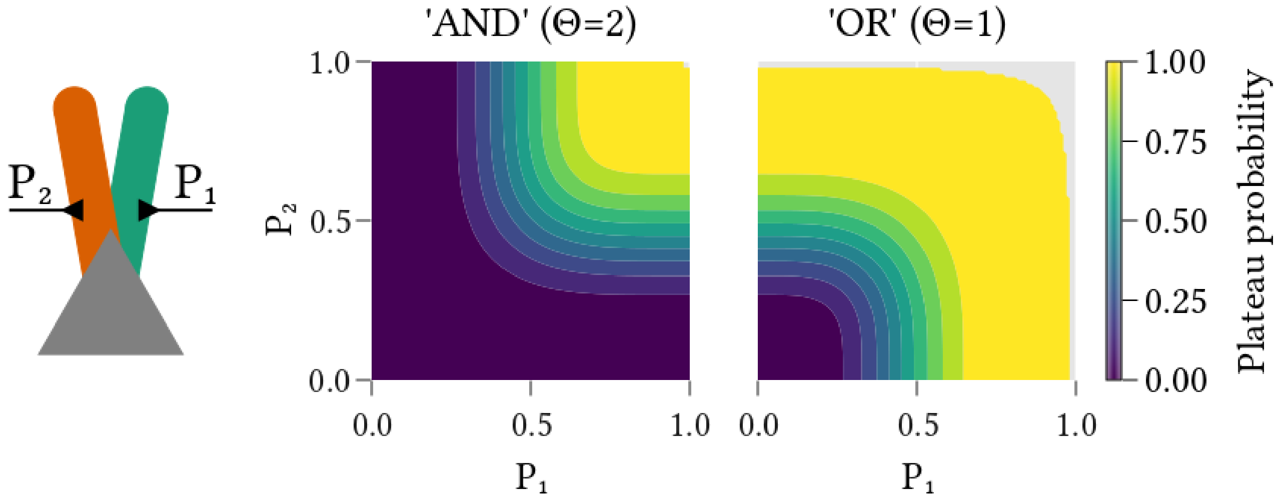

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

- → but: **stochastic input** results in **stochastic output**
- → the **probability** is graded!
- → could be put into **ensembles**

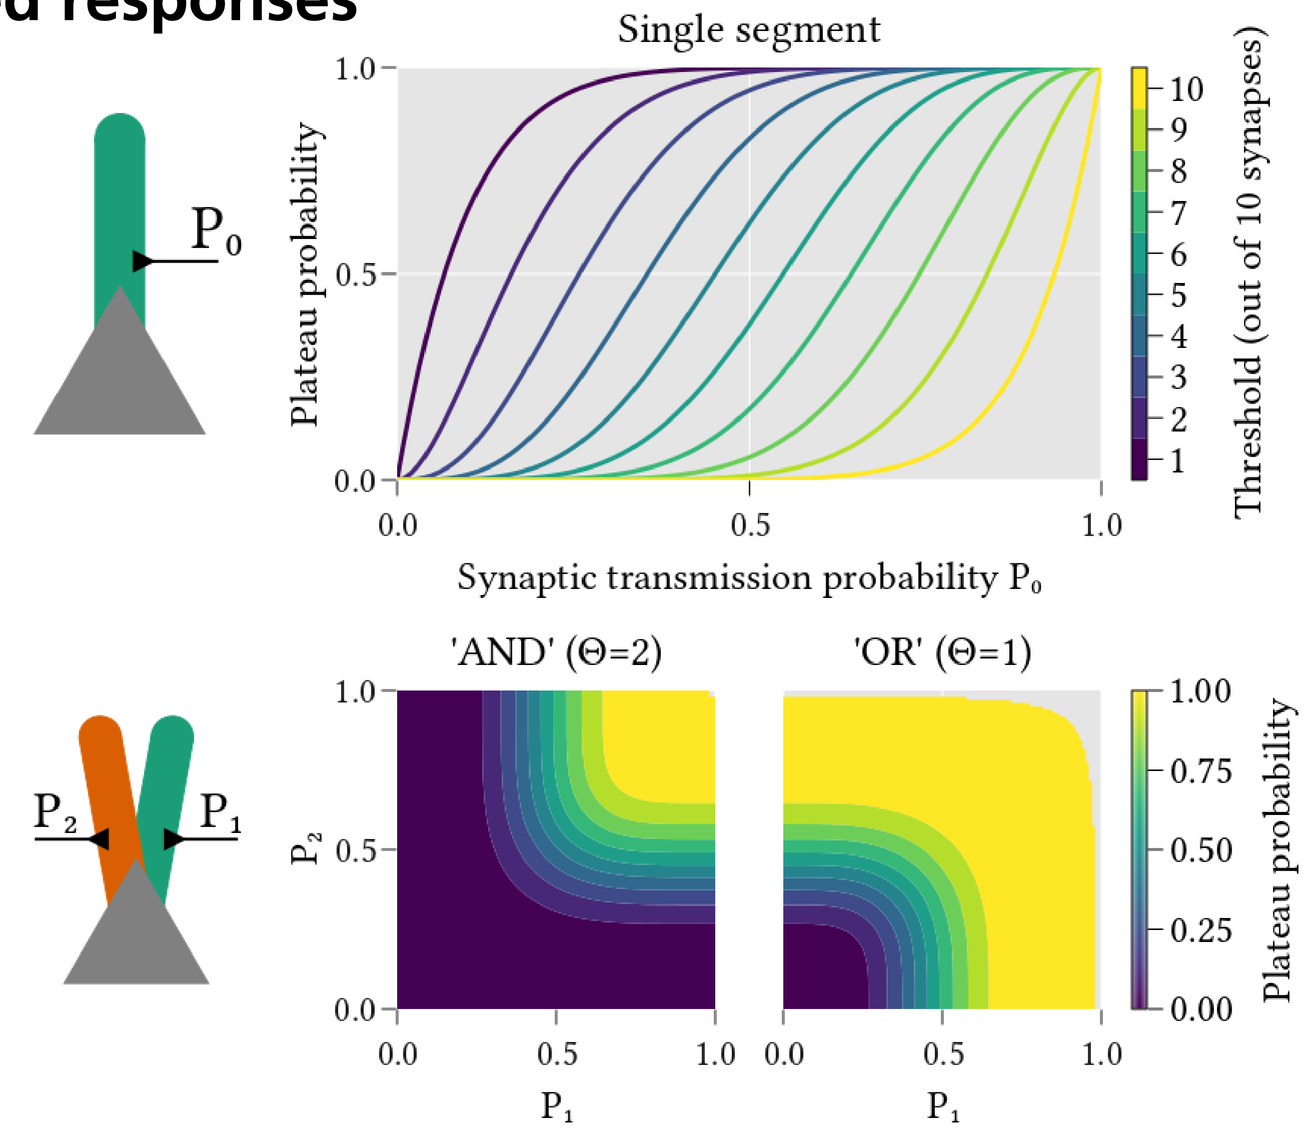

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

- → but: **stochastic input** results in **stochastic output**
- → the **probability** is graded!
- → could be put into **ensembles**

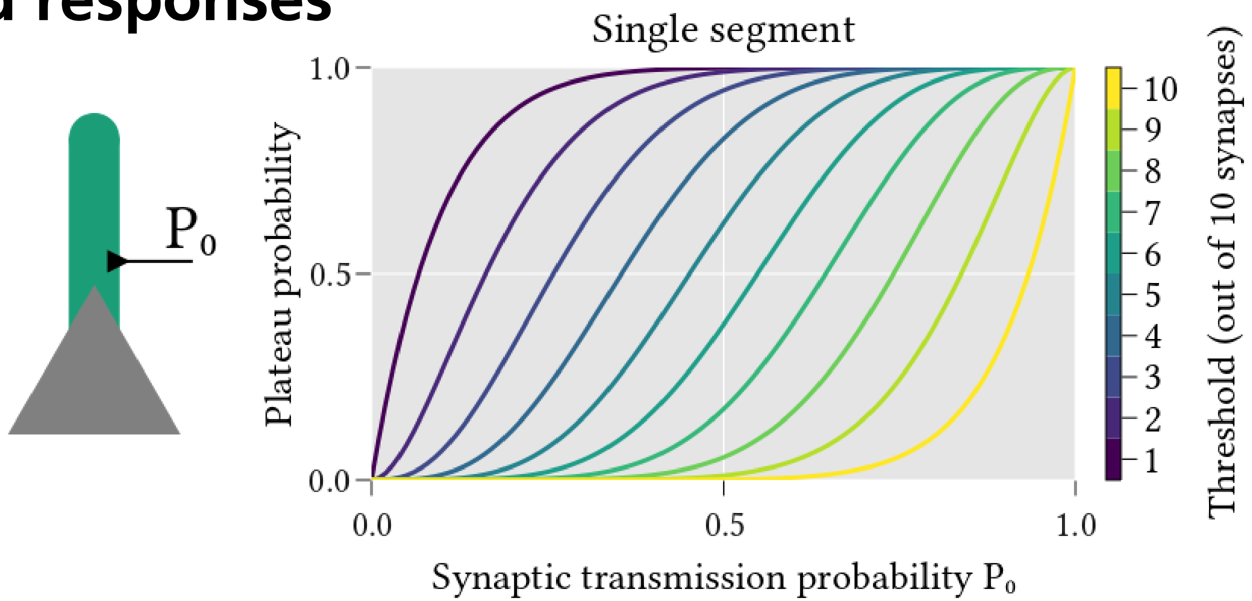

→firing **probability** ~ certainty of stimulus

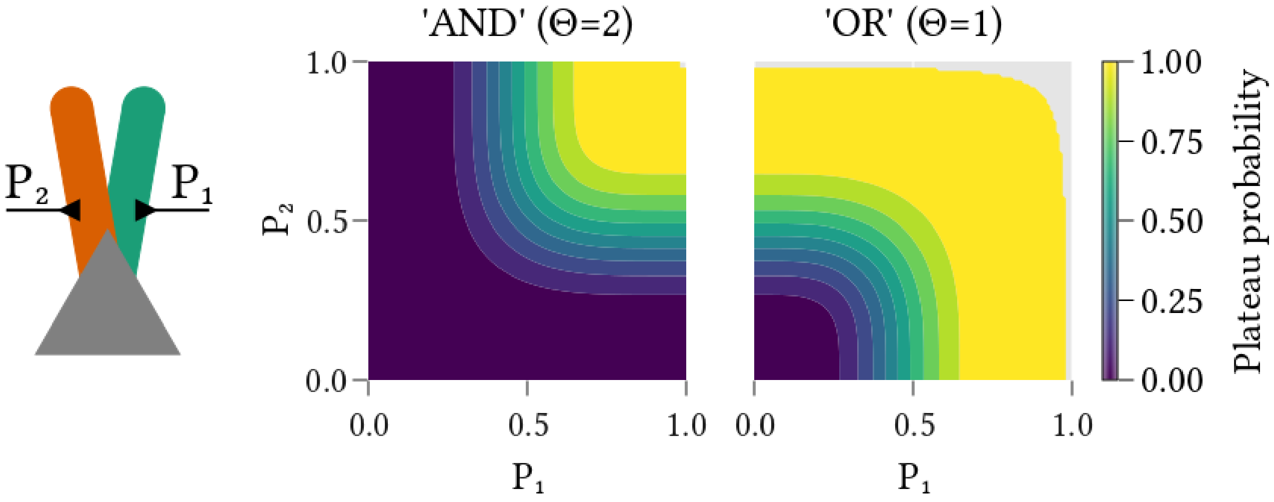

public

UNIVERSITÄT OSNABRÜCK

Fraunhofer IIS

Probabilistic synapses allow graded responses

→each individual neuron is **deterministic**

- → but: **stochastic input** results in **stochastic output**
- → the **probability** is graded!
- → could be put into **ensembles**

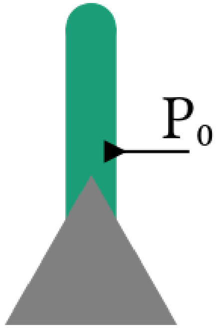

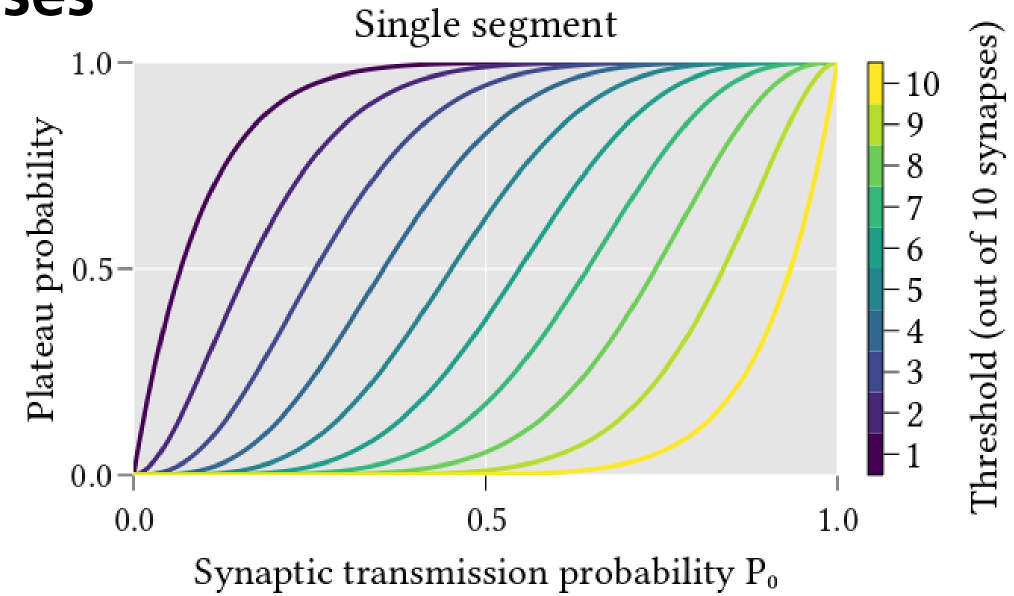

→firing **probability** ~ certainty of stimulus

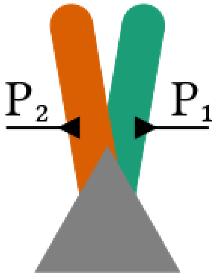

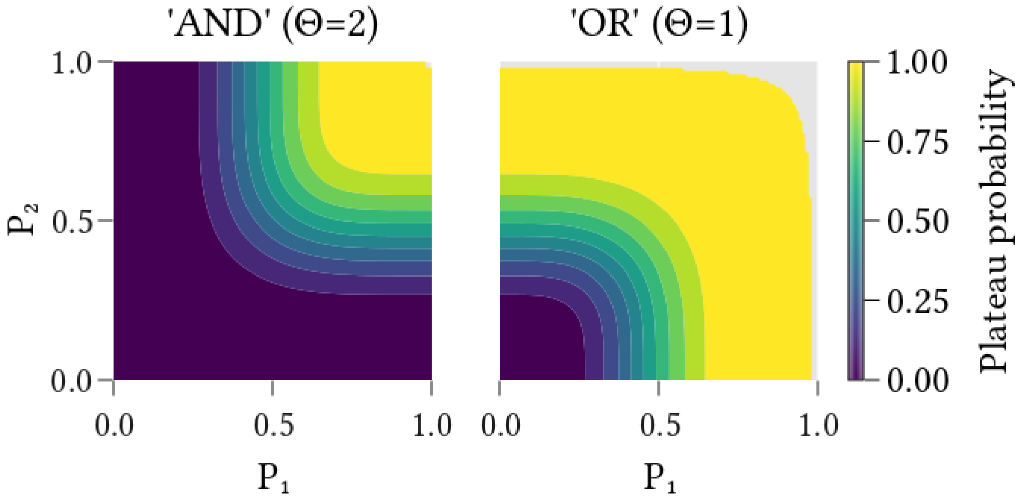

→the neuron becomes a **probabilistic, event-based pattern detector**

Making sense of it all: A toy example

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

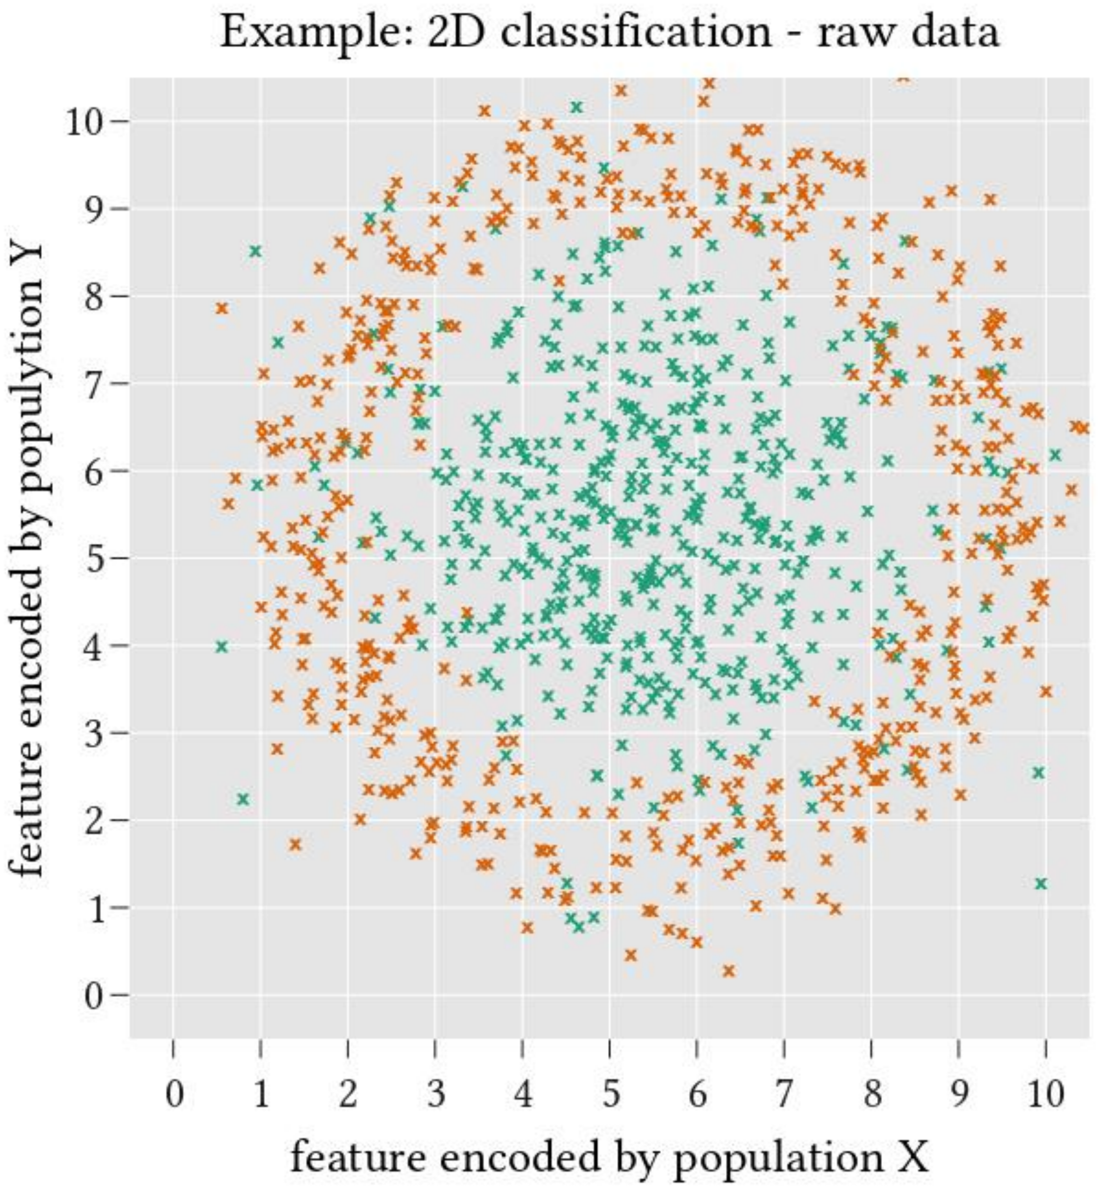

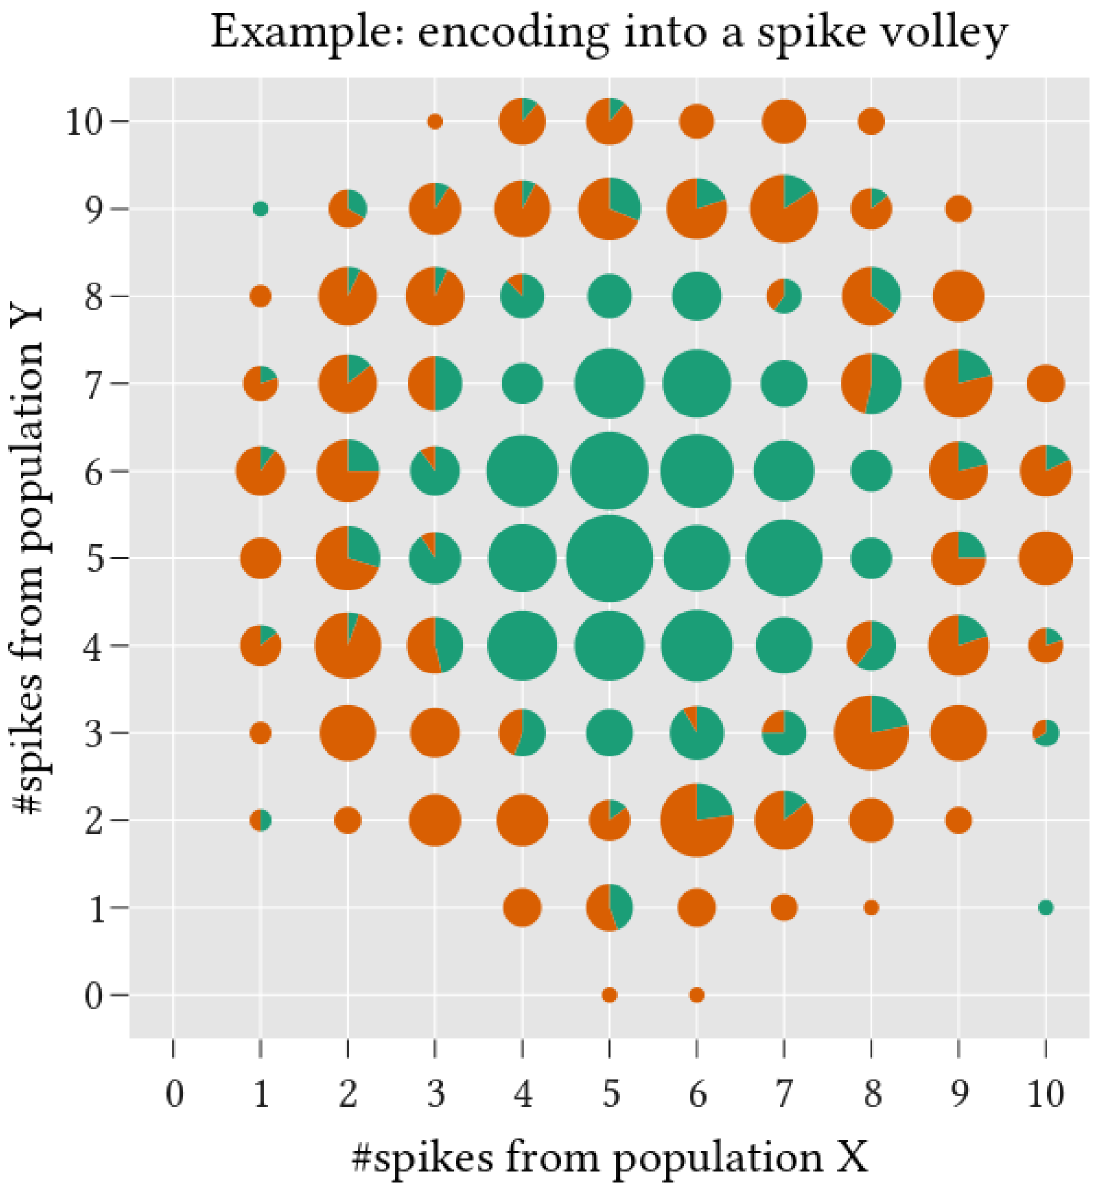

A toy-example:

- Two-dimensional data
- Select **class 1**, ignore **class 2**

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

public

UNIVERSITÄT OSNABRÜCK

Fraunhofer IIS

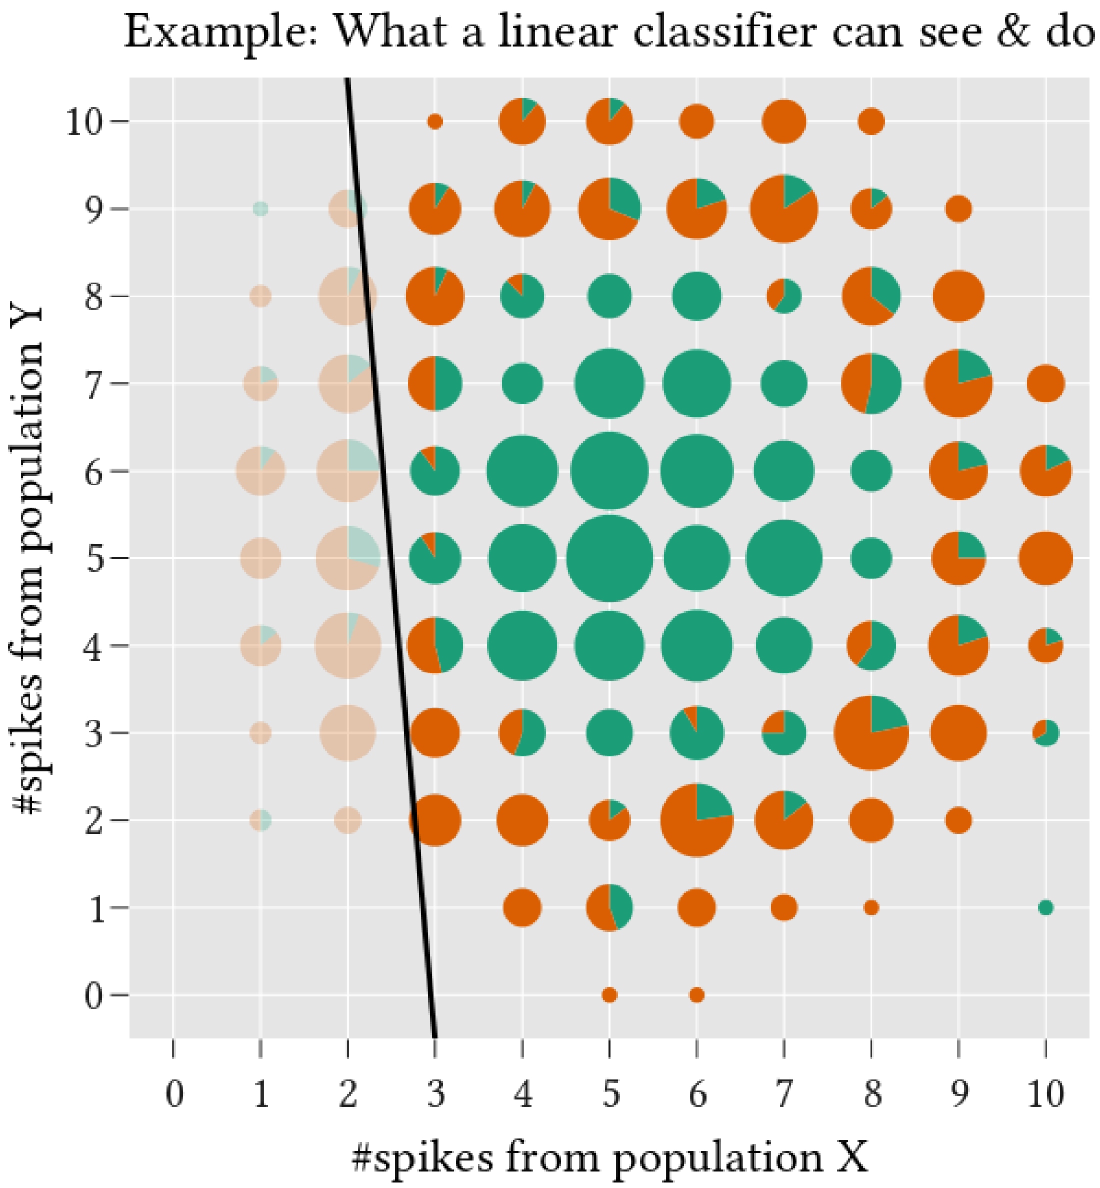

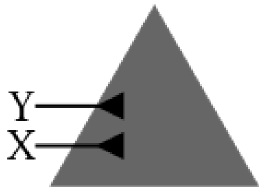

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

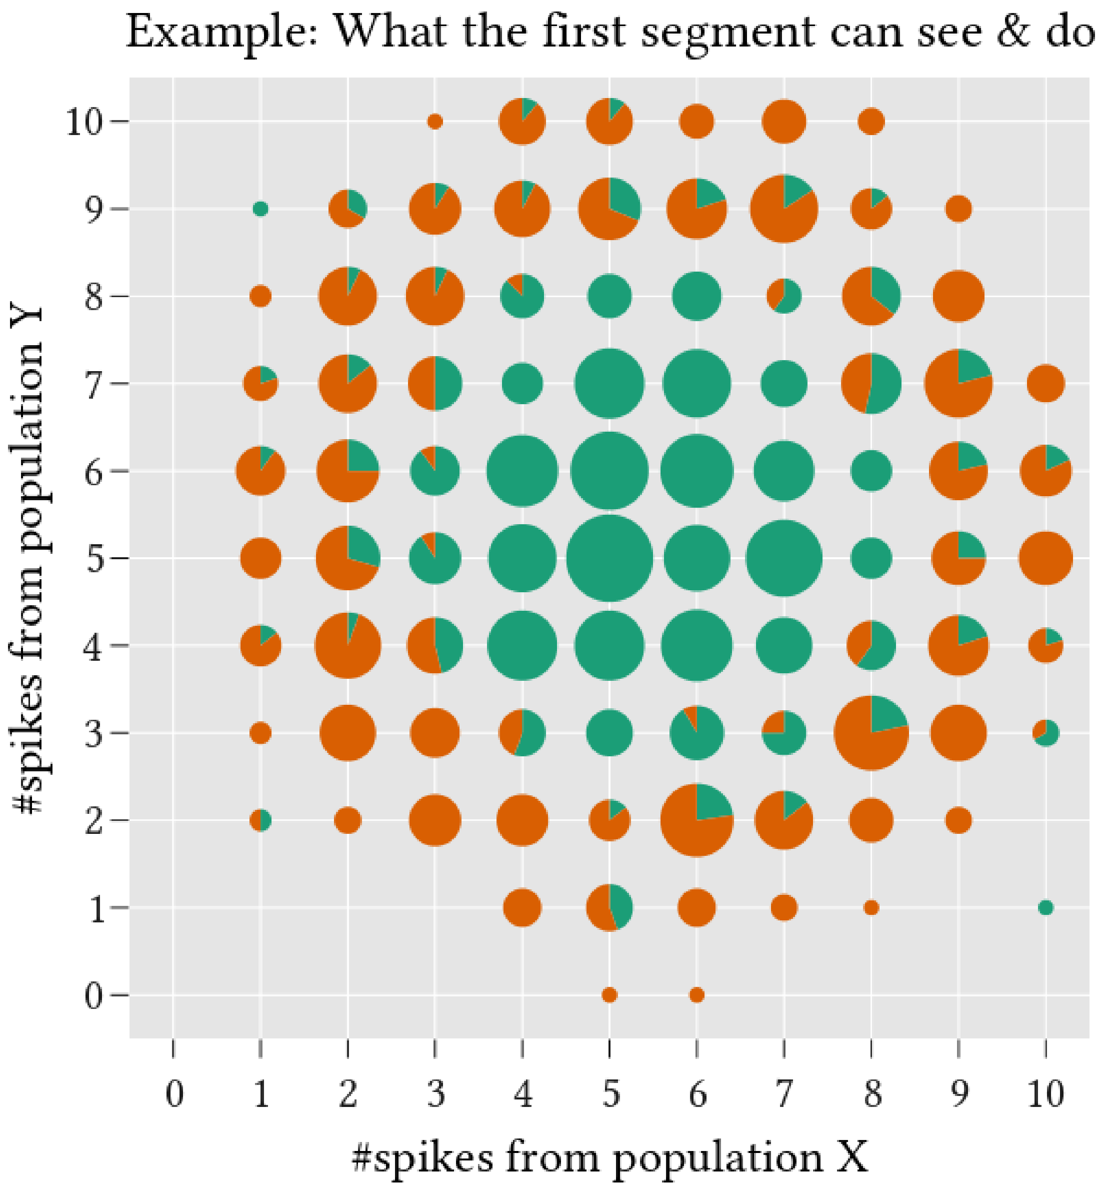

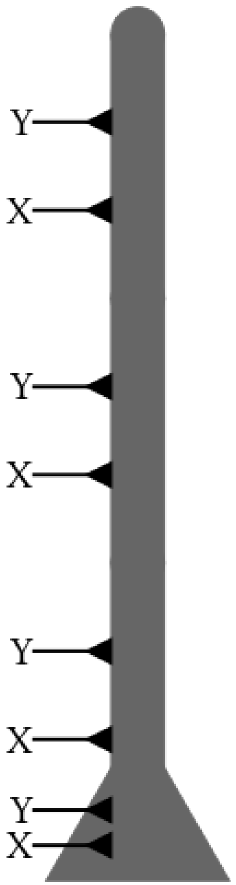

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y)* → *(#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

Multi-compartment neuron:

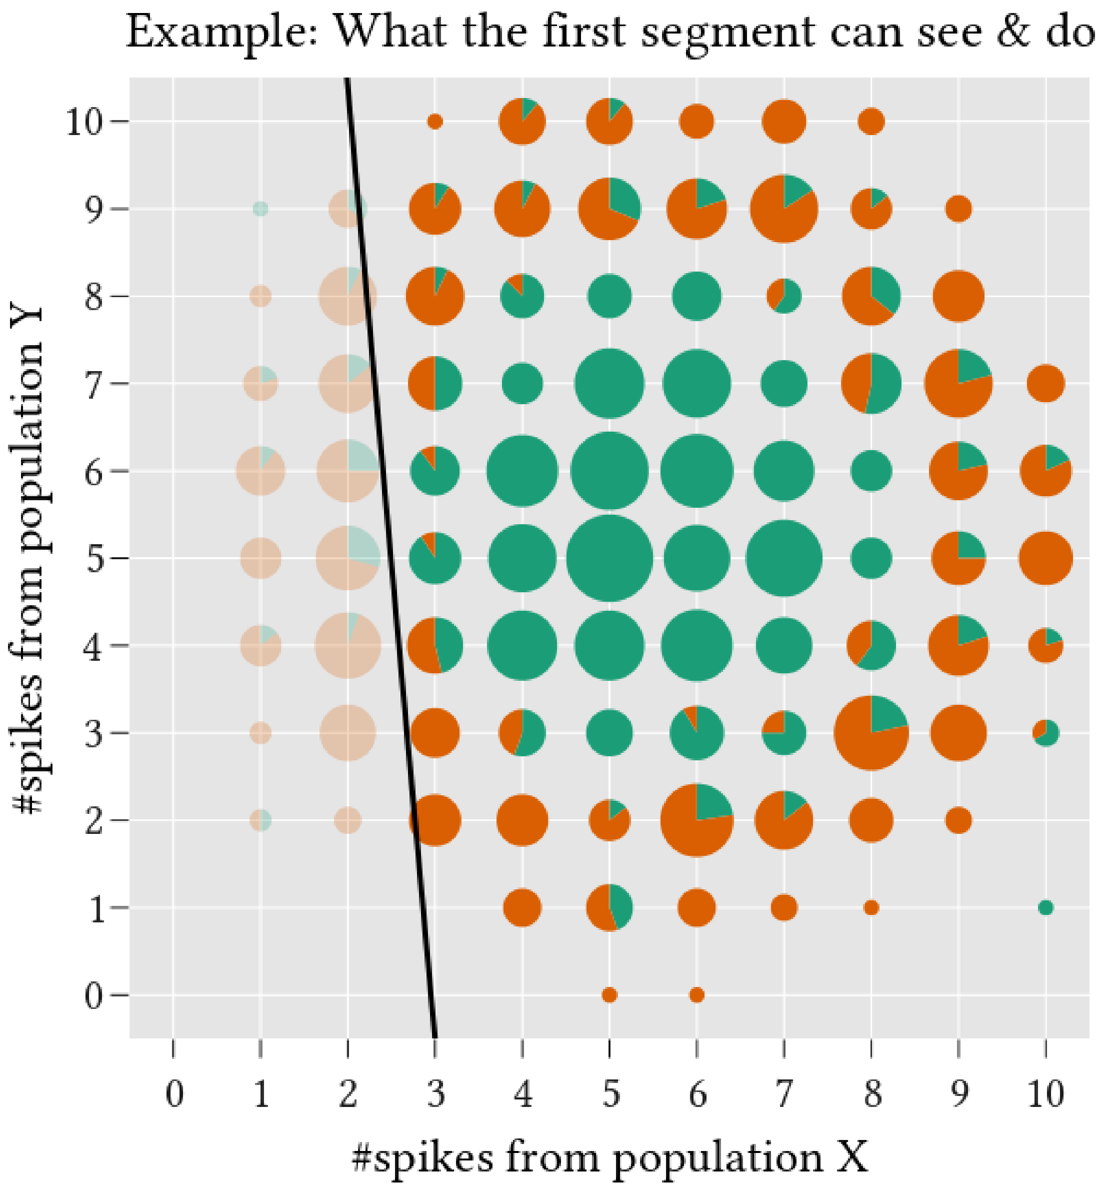

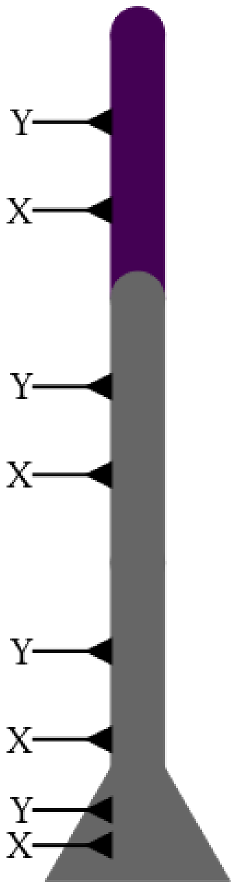

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

Multi-compartment neuron:

- **Iterative** application (~Tree!)

Example: What the second segment can see & do

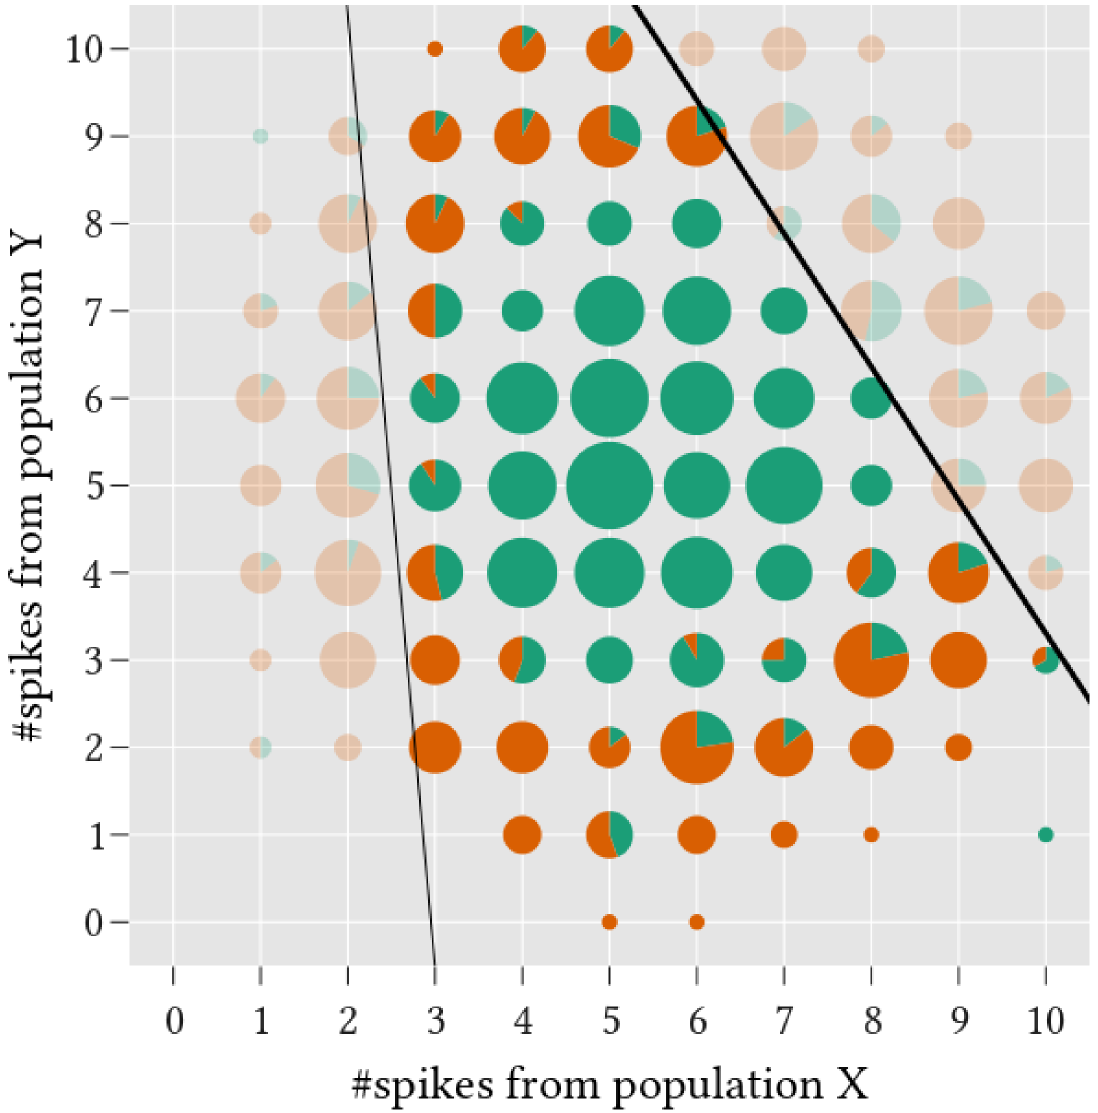

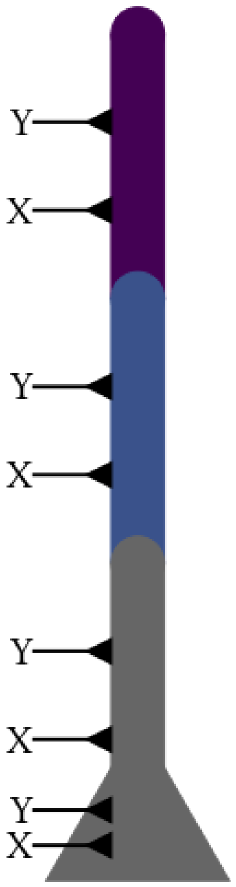

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

Multi-compartment neuron:

- **Iterative** application (~Tree!)

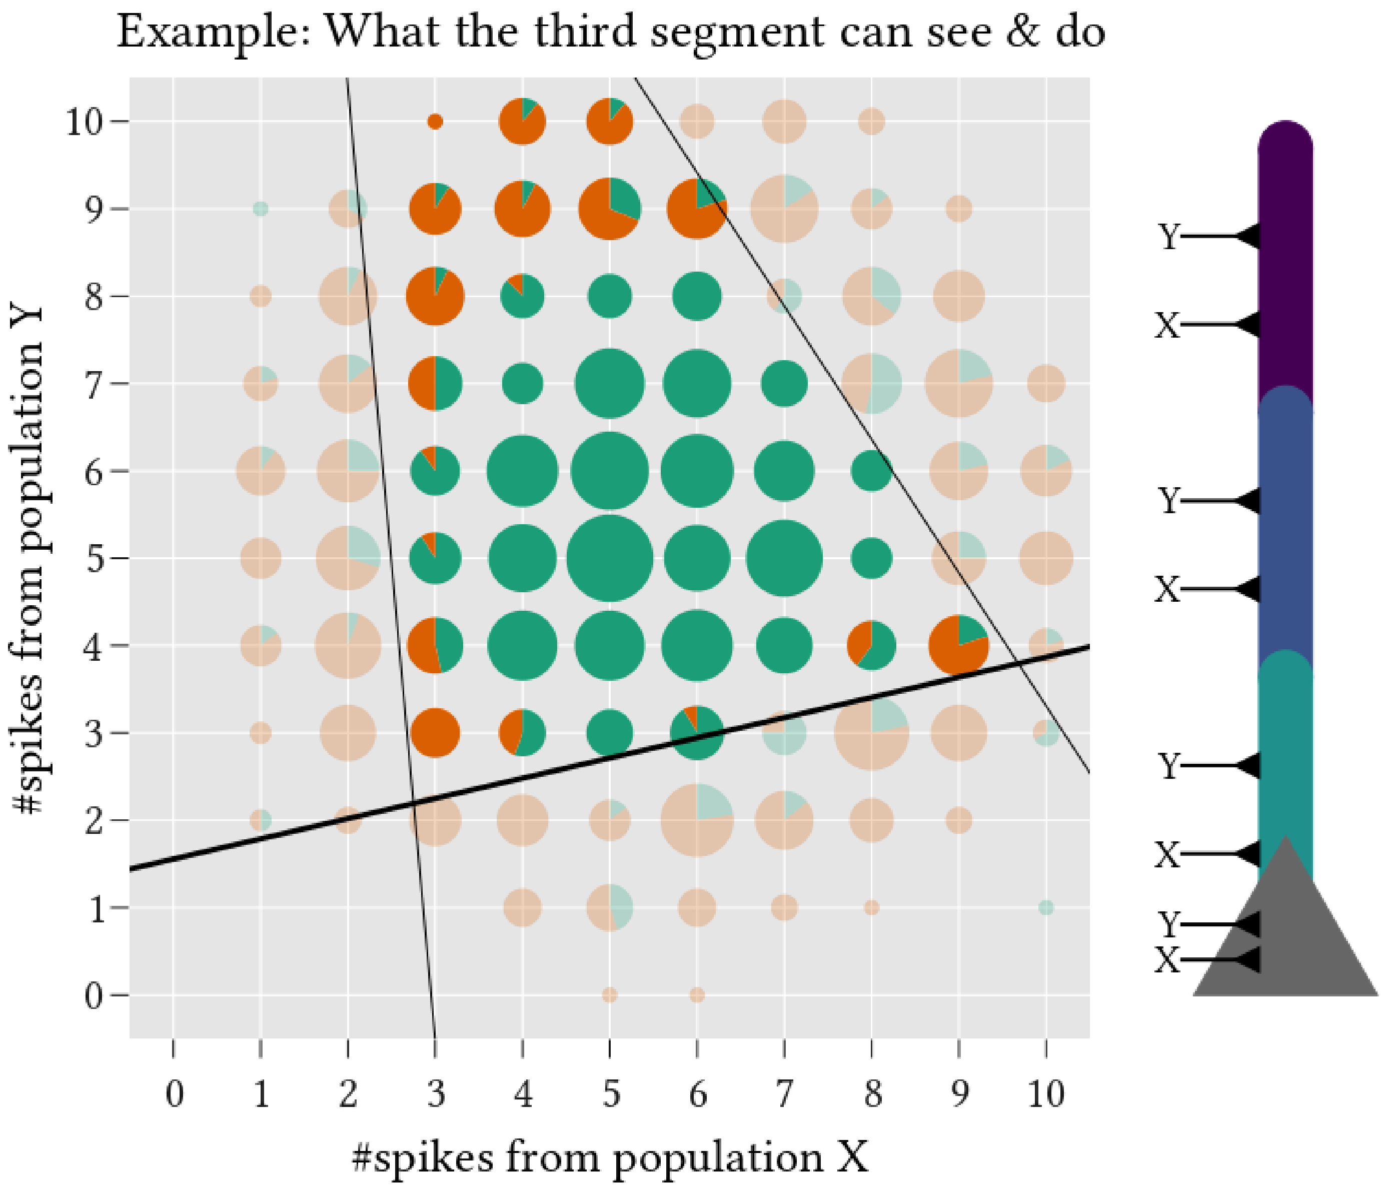

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

Multi-compartment neuron:

- **Iterative** application (~Tree!)

Example: What the soma can see & do

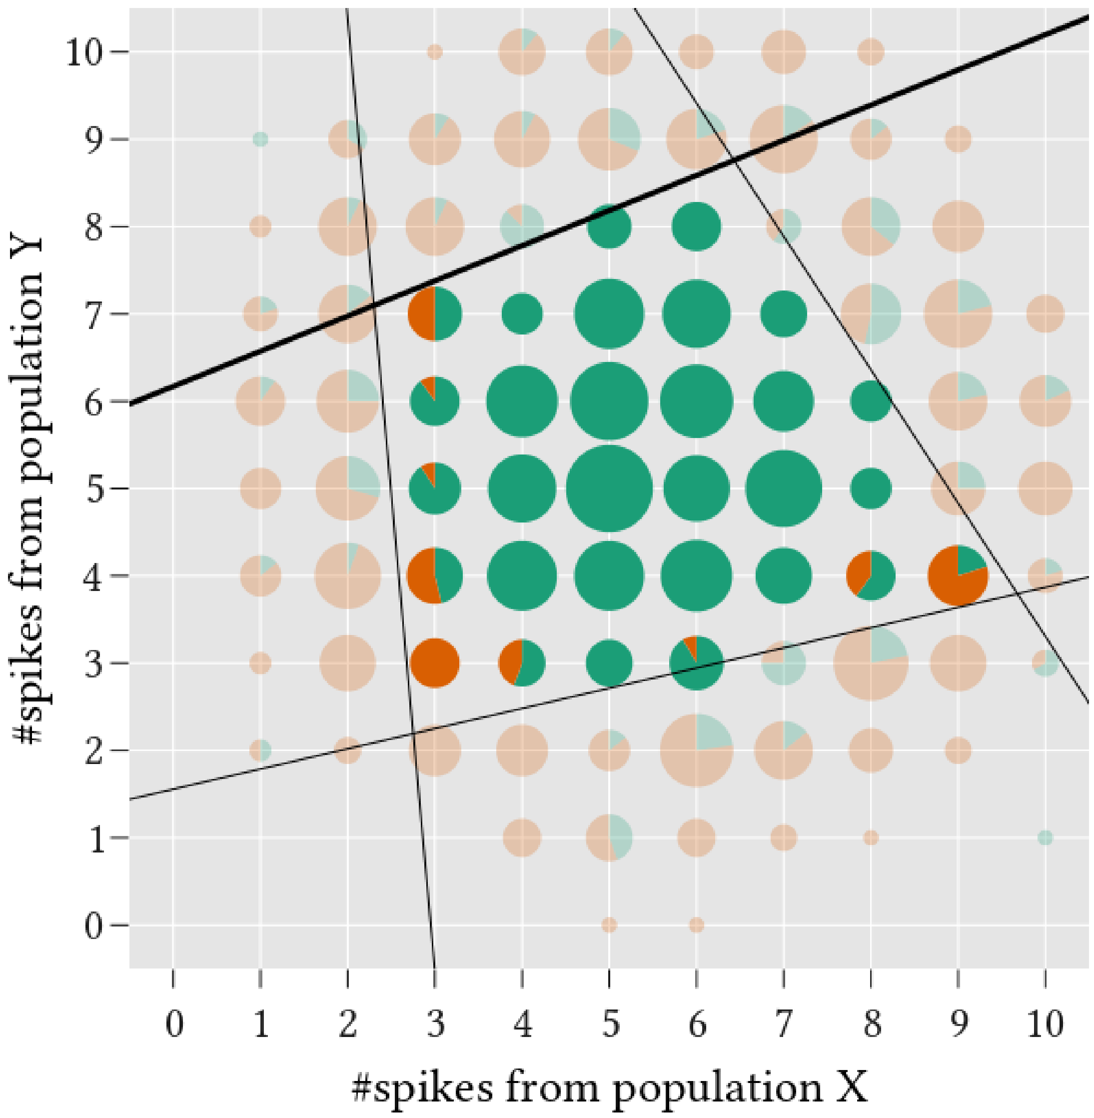

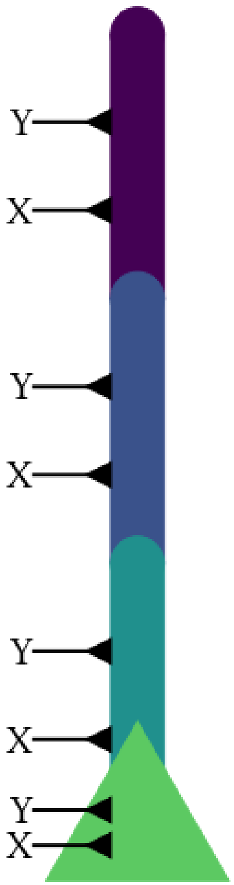

A toy-example:

- Two-dimensional data
- Select class 1, ignore class 2

Encoding:

- Encode X, Y each by one **population of 10 neurons**
- Each sample becomes a **volley of coincident spikes**, i.e.: *(x,y) → (#x spikes, #y spikes)*

Linear classifier (point neuron):

- Optimal split (e.g. info-gain)
- Problem: combining neurons

Multi-compartment neuron:

- **Iterative** application (~Tree!)
- Finds NICE **sparse** solution!

Summary

- active dendrites are an **established fact in neuroscience,** but mostly **absent in machine learning models**

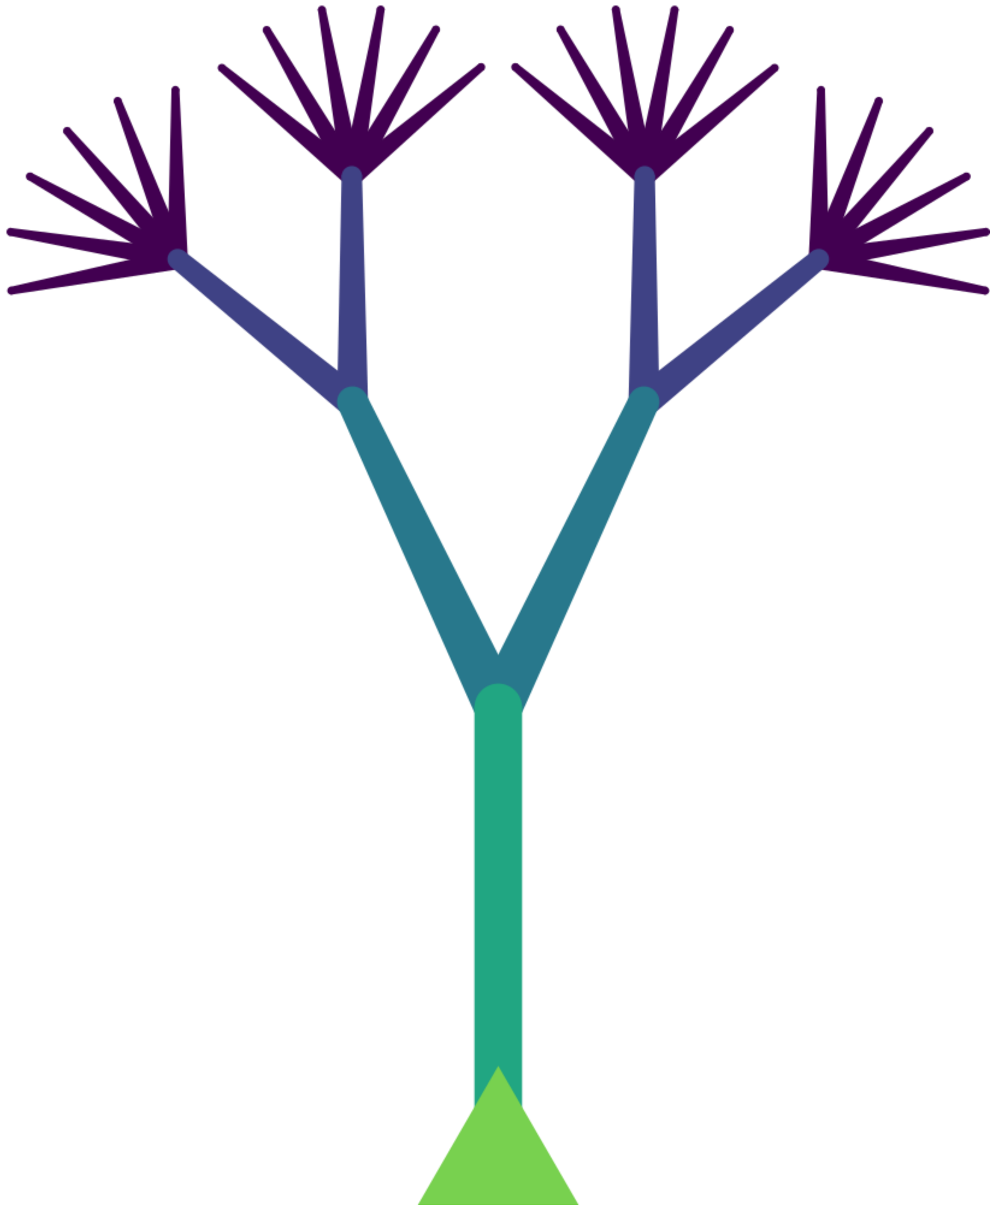

UNIVERSITÄT OSNABRÜCK

Fraunhofer IIS

Summary

- active dendrites are an **established fact in neuroscience,** but mostly **absent in machine learning models**

- active dendrites could realize **interesting computations**
  - **temporal pattern detection**
  - **timing invariance**
  - longer **memory traces** (>100ms)
  - **nonlinear computation** within each neuron

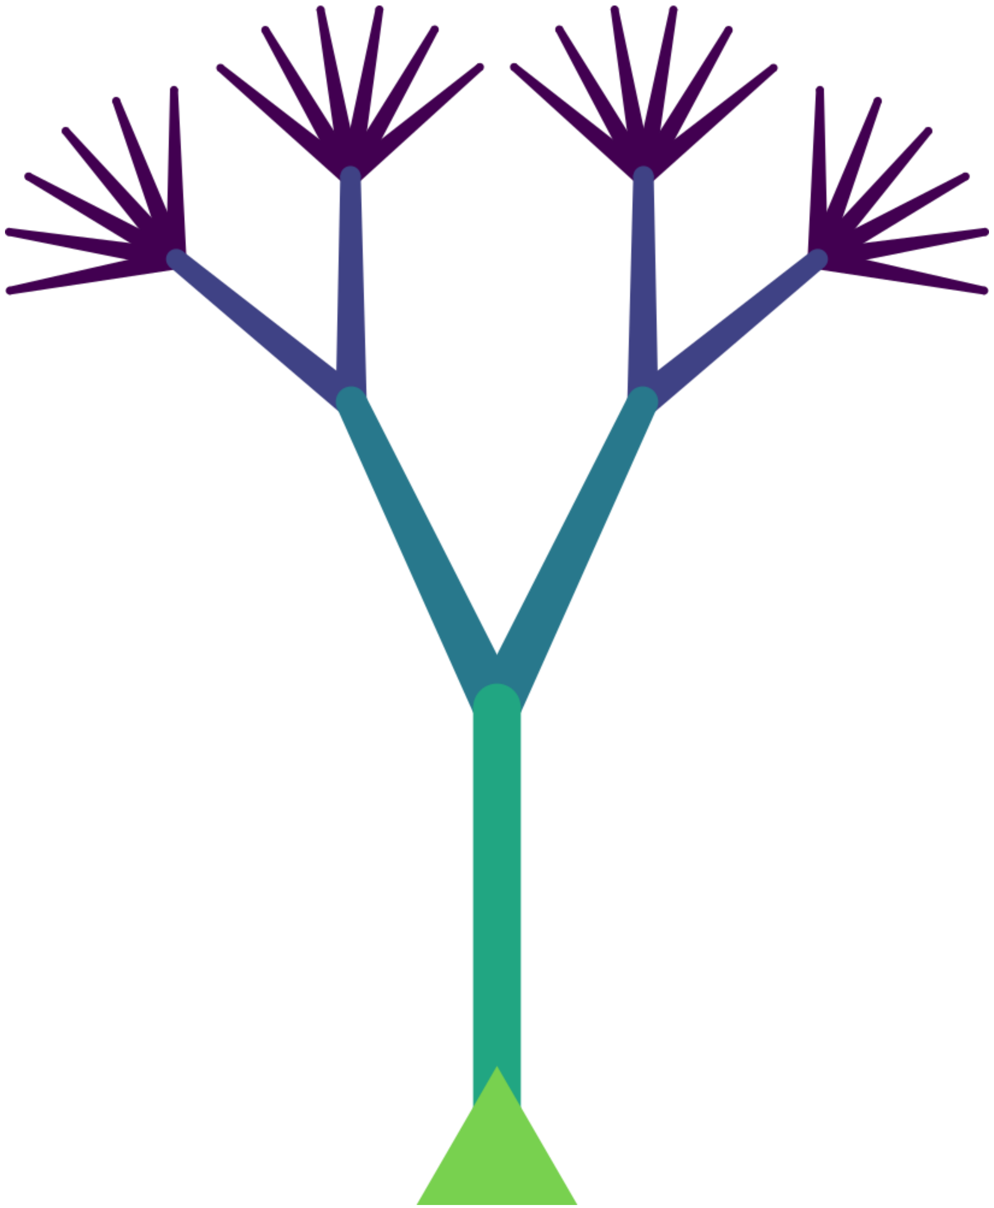

public

UNIVERSITÄT OSNABRÜCK

Fraunhofer IIS

Summary

- active dendrites are an **established fact in neuroscience,** but mostly **absent in machine learning models**

- active dendrites could realize **interesting computations**
  - **temporal pattern detection**
  - **timing invariance**
  - longer **memory traces** (>100ms)
  - **nonlinear computation** within each neuron

- **stochasticity** & **delay** have a purpose in this framework

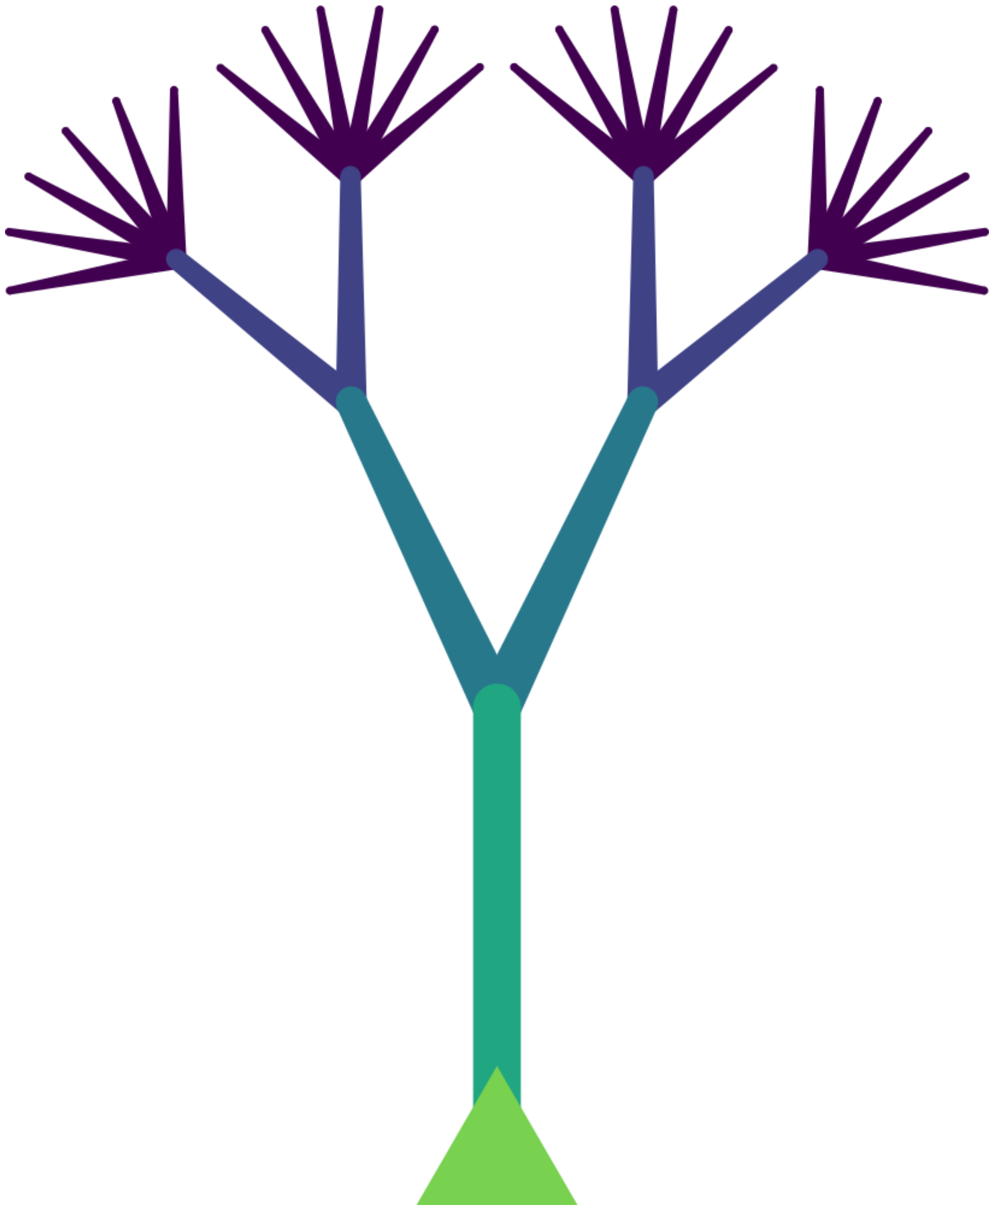

Summary

- active dendrites are an **established fact in neuroscience**, but mostly **absent in machine learning models**
- active dendrites could realize **interesting computations**
  - **temporal pattern detection**
  - **timing invariance**
  - longer **memory traces** (>100ms)
  - **nonlinear computation** within each neuron
- **stochasticity** & **delay** have a purpose in this framework
- Entire model only makes use of **binary** states

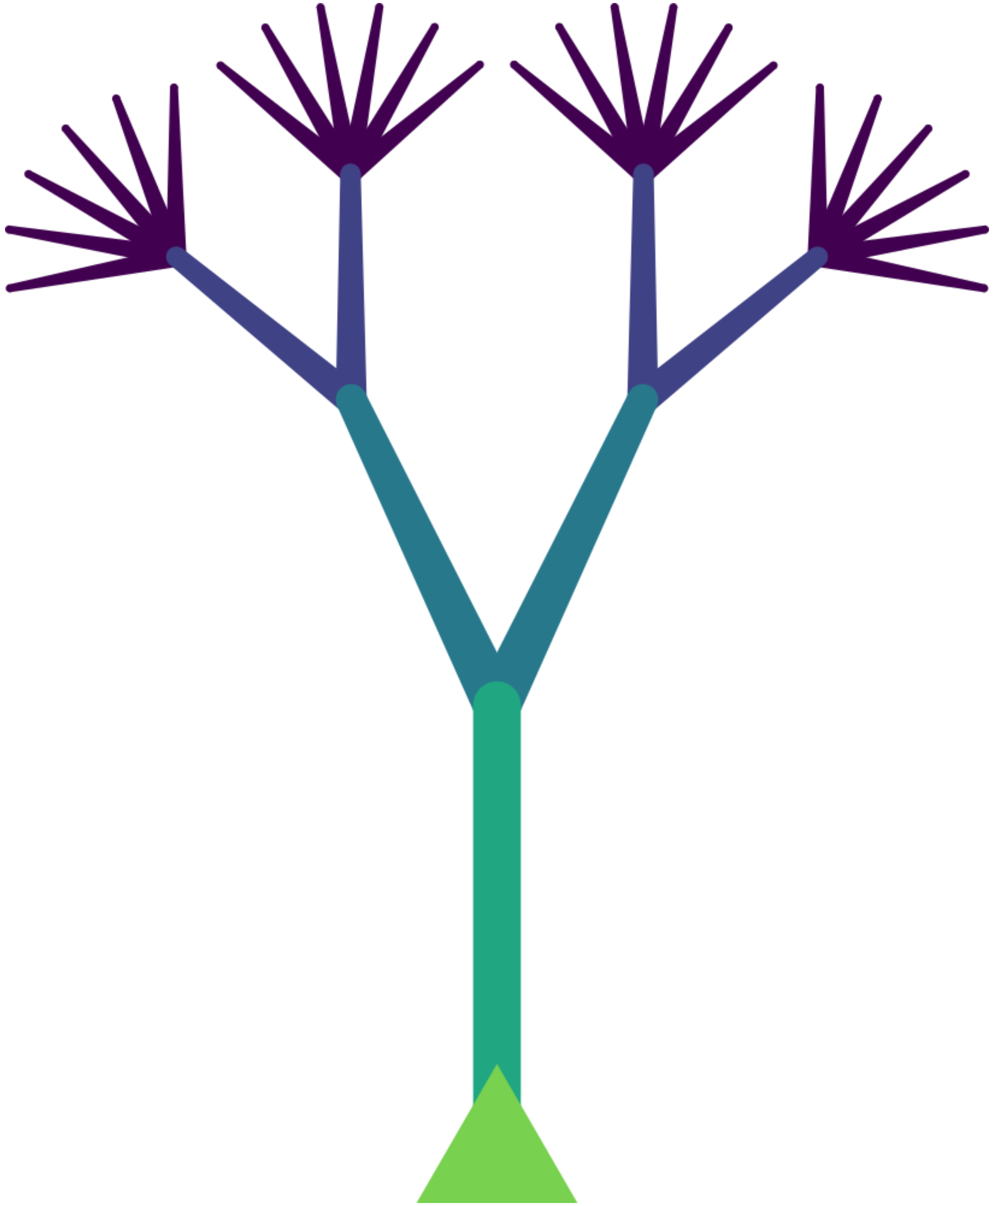

Summary

- active dendrites are an **established fact in neuroscience,** but mostly **absent in machine learning models**

- active dendrites could realize **interesting computations**
  - **temporal pattern detection**
  - **timing invariance**
  - longer **memory traces** (>100ms)
  - **nonlinear computation** within each neuron

- **stochasticity** & **delay** have a purpose in this framework

- Entire model only makes use of **binary** states

→ **Attractive for neuromorphic hardware community?**

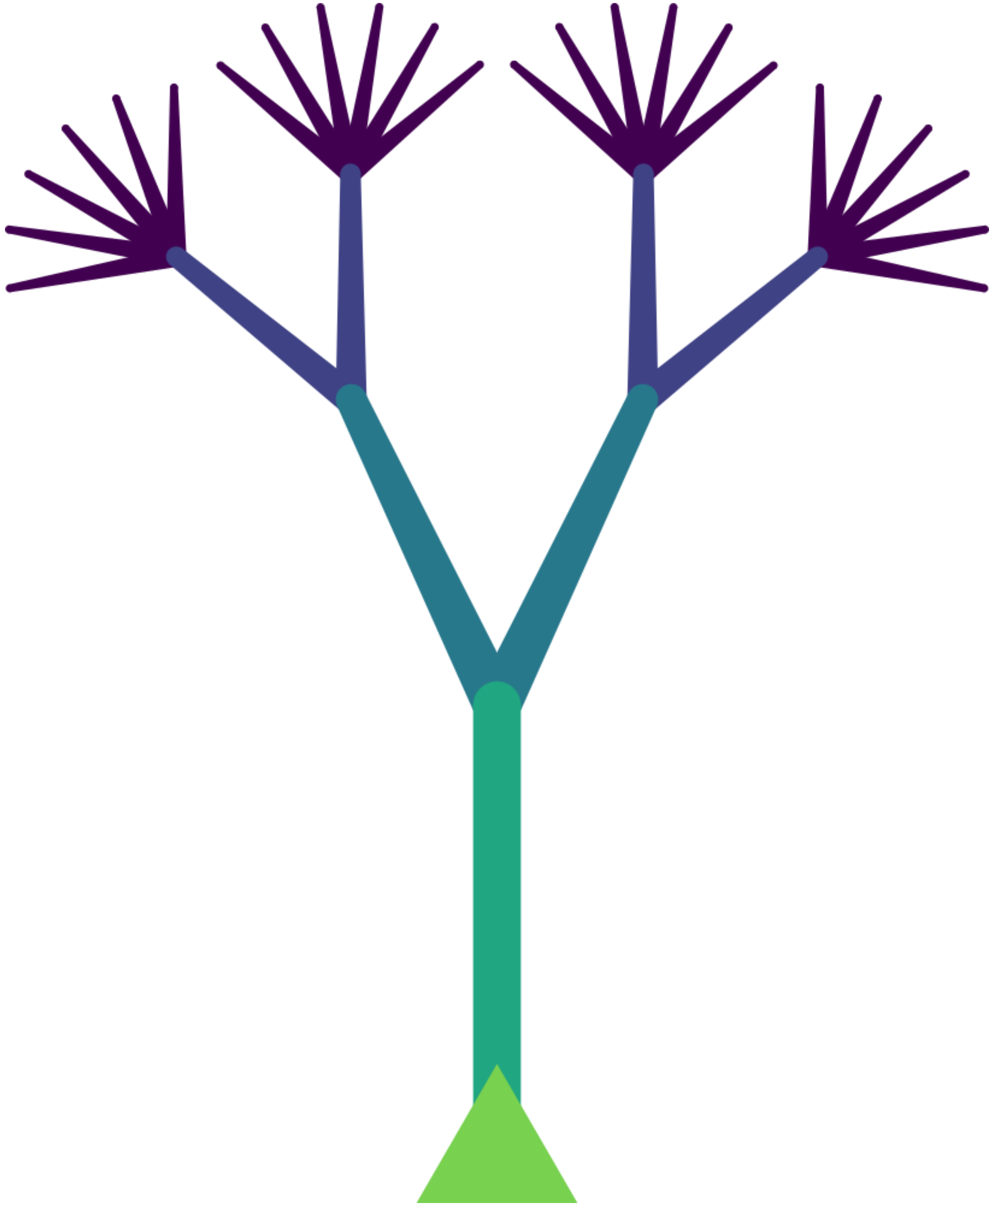

Further information & contacts

"Event-based pattern detection in active dendrites"

Preprint: https://www.biorxiv.org/content/10.1101/690792v3

⚠ Work in progress – revision upcoming ⚠

"Making spiking neurons more succinct with multi-compartment models"
https://dl.acm.org/doi/epdf/10.1145/3381755.3381763

Osnabrück University:

Pascal Nieters
pnieters@uni-osnabrueck.de

Prof. Gordon Pipa
gpipa@uni-osnabrueck.de

Fraunhofer IIS:

Gunter Föttinger (Head of *Embedded AI*, IIS)
gunter.foettinger@iis.fraunhofer.de

Alexander Jaschke (BizDev for *Com. Systms.*, IIS)
alexander.jaschke@iis.fraunhofer.de

Johannes Leugering (Expert in *Embedded AI*, IIS)
johannes.leugering@iis.fraunhofer.de

Dr. Marco Breiling (Chief Scientist for Com. Systms., IIS)
marco.breiling@iis.fraunhofer.de

Fin.