

Technische Universität Wien
Fakultät für !nformatik
Cyber-Physical-Systems Group

Biologically-Plausible Large Language Models

Radu Grosu

Natural vs Artificial Neural Networks

Human Brain:

- Synapses: 1 000 000 billion
- Power: 30 W

Chat GPT:

- Synapses: 175 billion
- Power: 30 000 000 W

What is doing nature so much better?

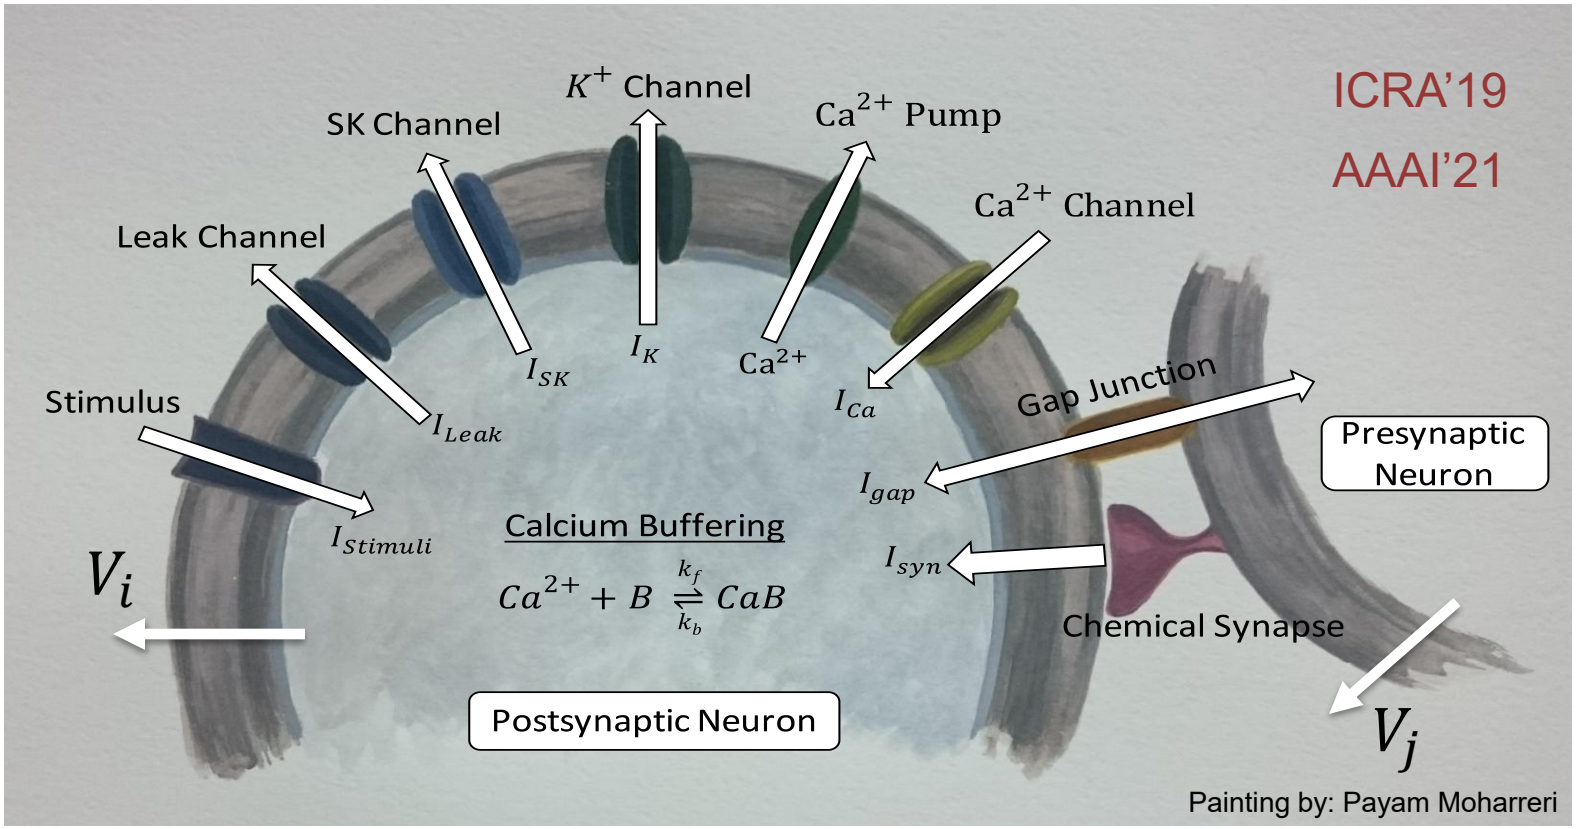
A Reductionist Approach

C.elegans Nematode

- L: 1mm, W: 0.01mm
- 302 nonspiking neurons
- 8000 synapses
- 95 wall muscles
- Associative learning
- Social behavior
- Fully mapped connectom

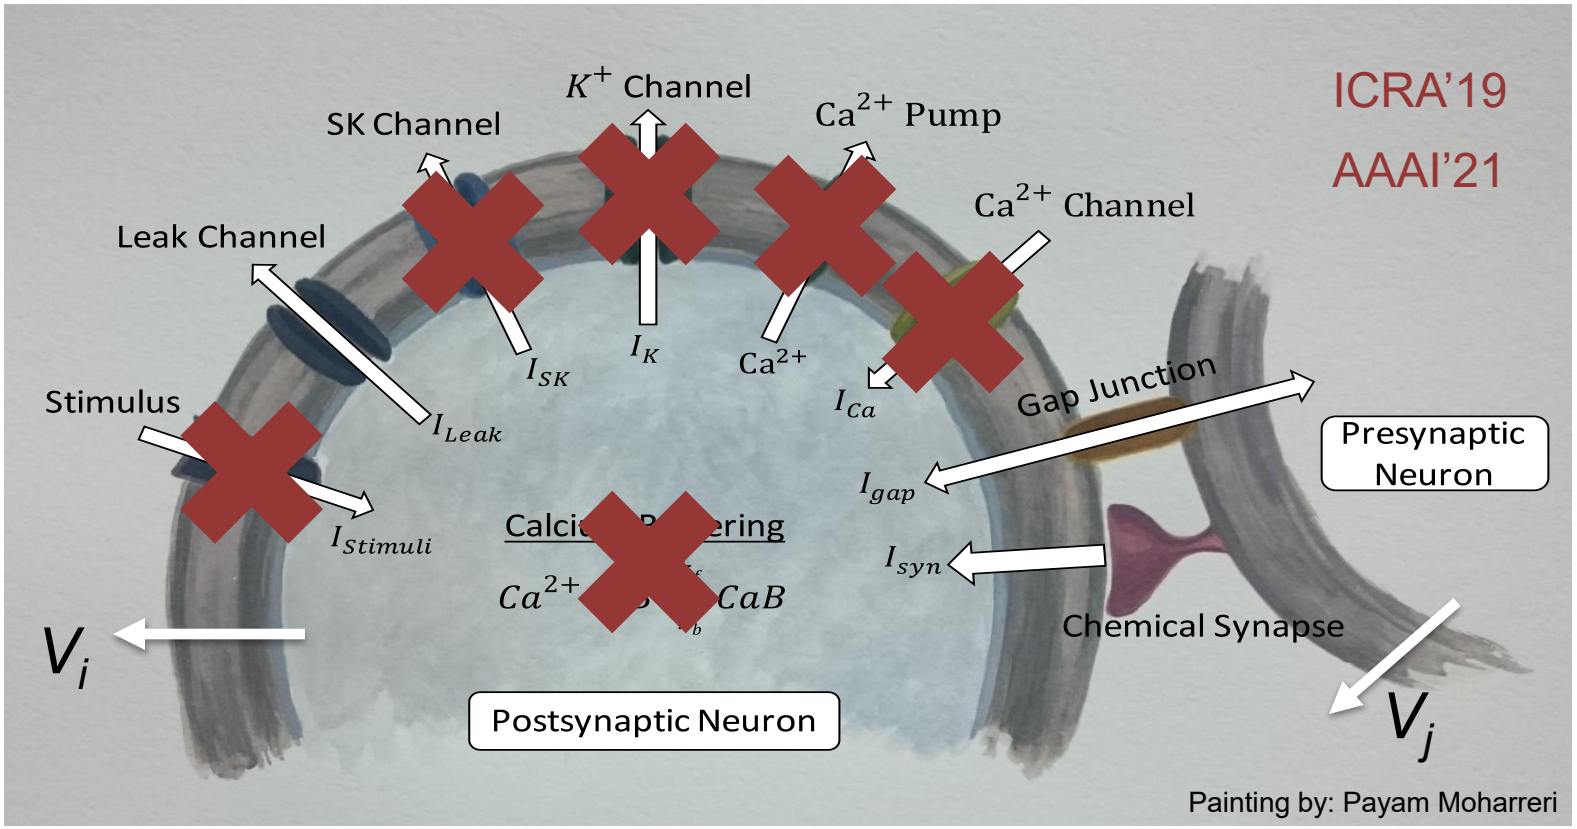


Biophysical Model of a Neuron



$$C_i \dot{V}_i = -\left(I_{Ca} + I_K + I_{SK} + I_{leak,i}\right) + I_{stim,i} + \sum_j I_{syn,ji} + I_{gap,ji}$$

Biophysical Model of a Neuron Used (LTC)



$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$



Currents in Biophysical Model (EEC/LTC)

ICRA'19

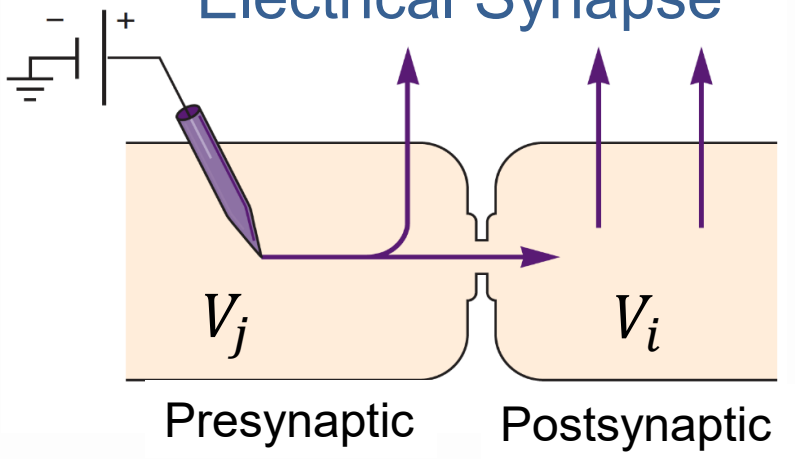
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

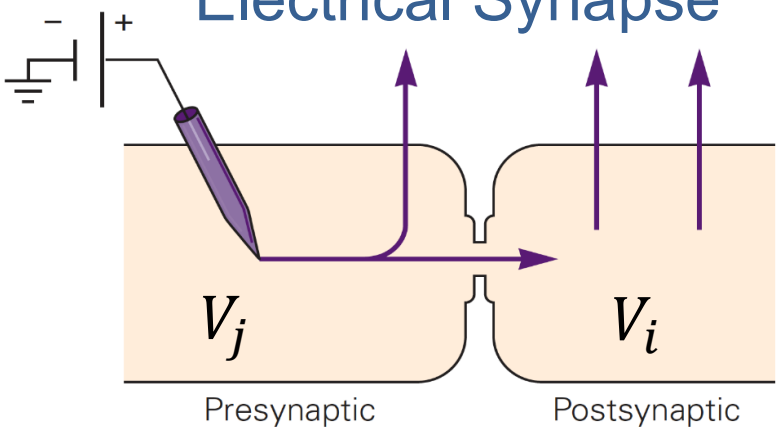
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

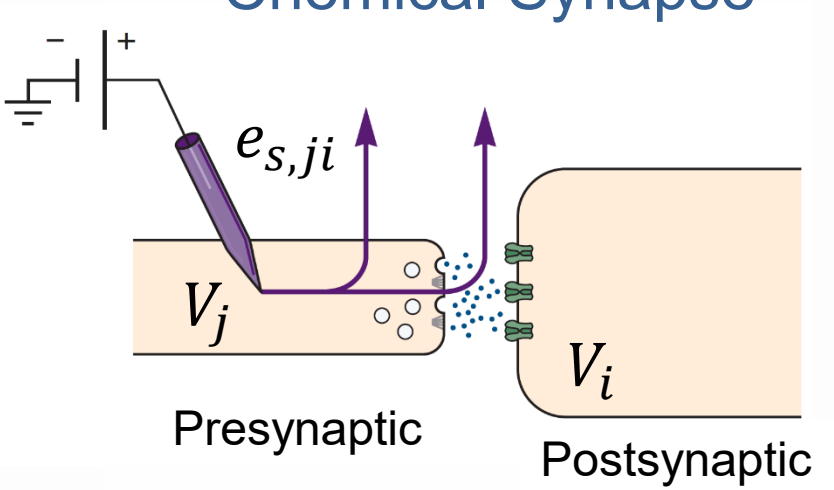
Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



Chemical Synapse



$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

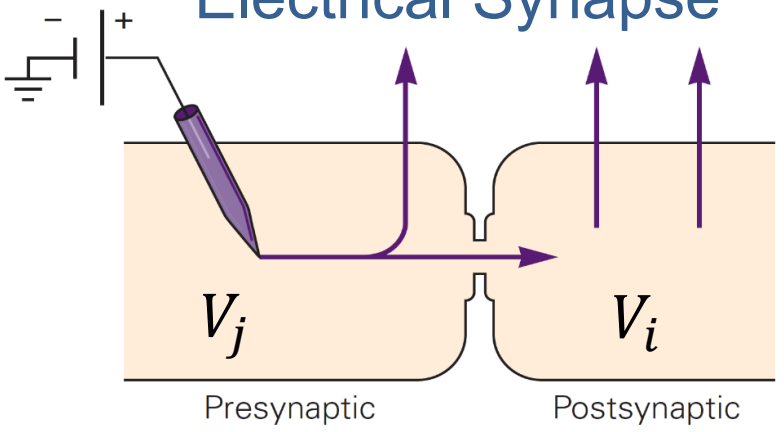
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

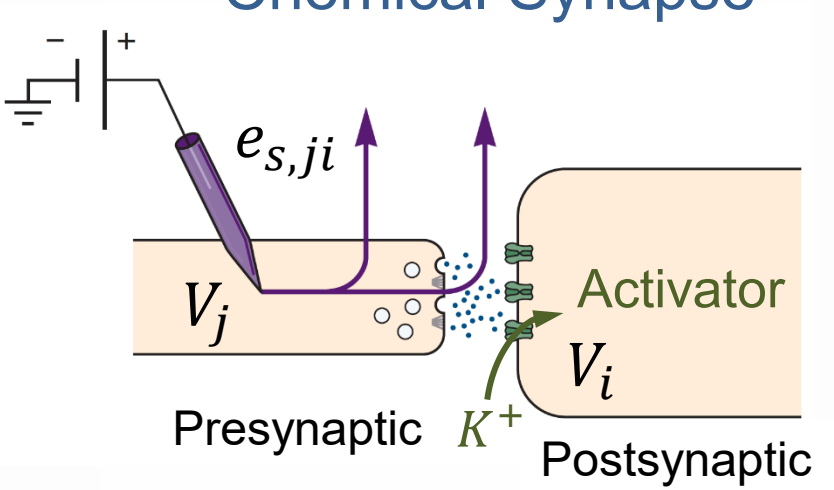
Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



Chemical Synapse



$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

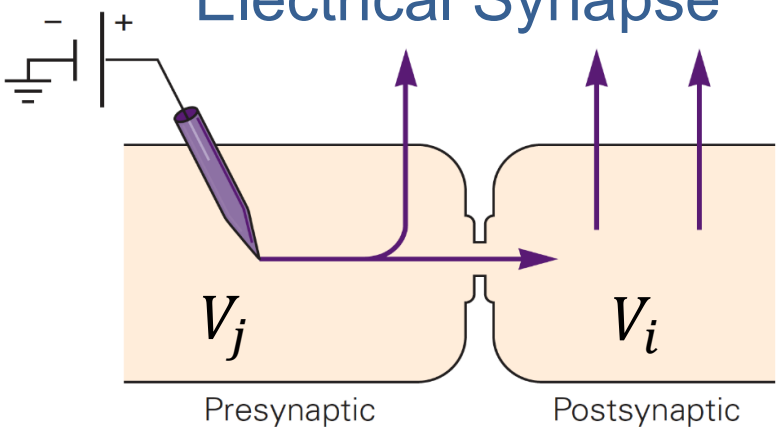
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

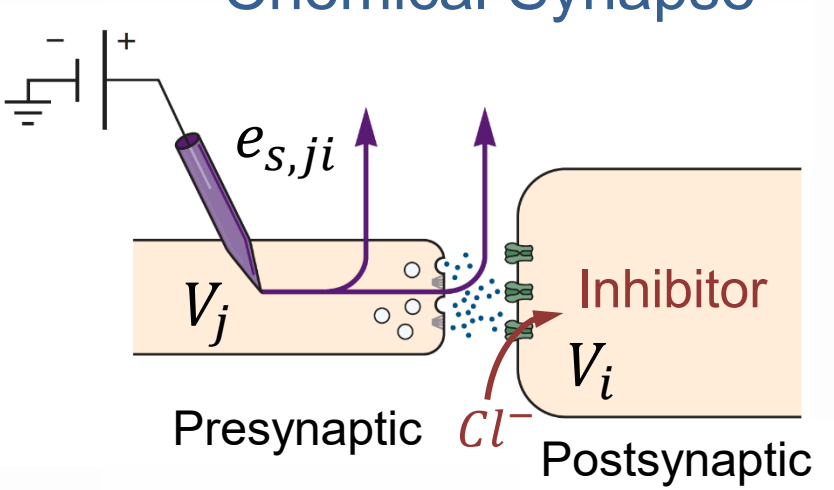
Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



Chemical Synapse



$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

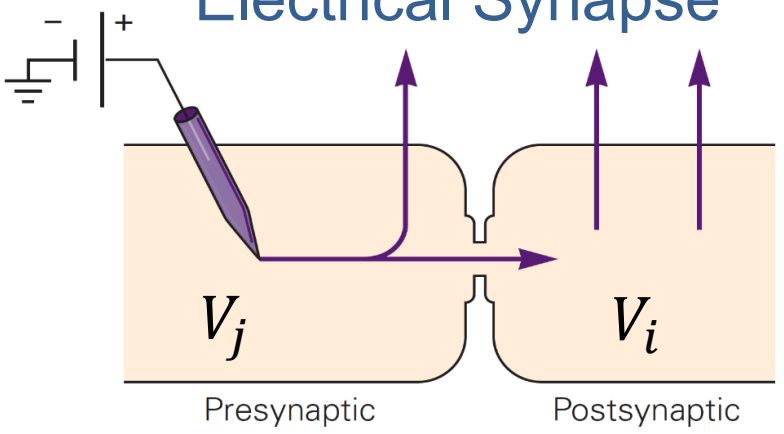
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

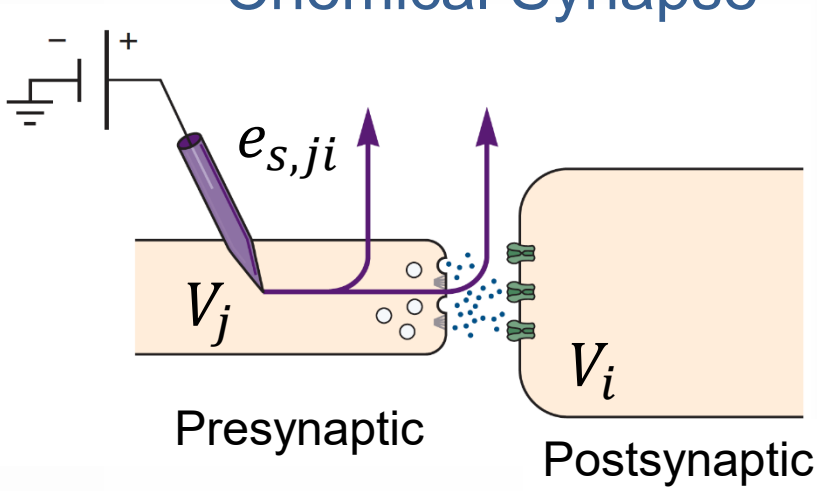
Leaking Current

$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



Chemical Synapse



$$I_{syn,ji}(t) = g_{s,ji}\sigma(V_j(t), \mu_j)(E_{s,ji} - V_i(t))$$

$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

Currents in Biophysical Model (EEC/LTC)

ICRA'19

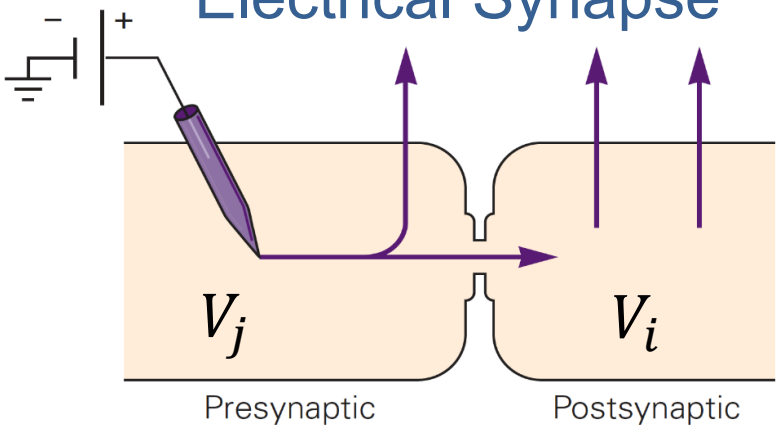
AAAI'21

$$C_i\dot{V}_i(t) = I_{leak,i}(t) + \sum_j (I_{syn,ji}(t) + I_{gap,ji}(t))$$

Leaking Current

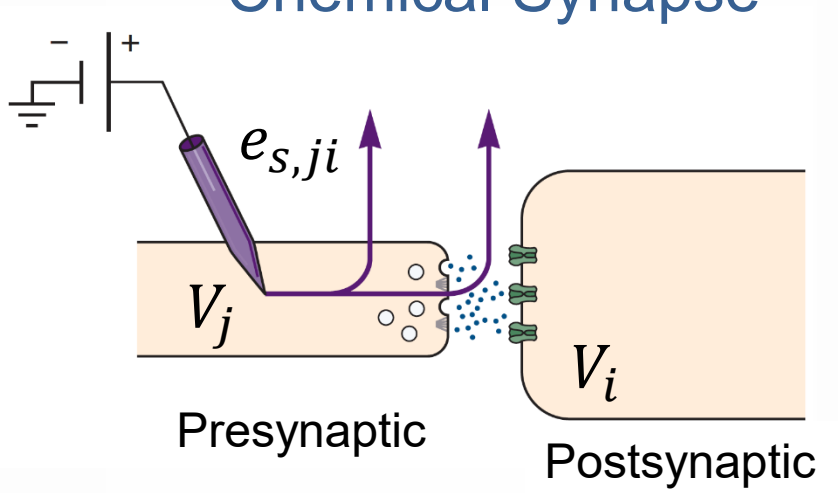
$$I_{leak,i}(t) = g_{l,i}(E_{l,i} - V_i(t))$$

Electrical Synapse



$$I_{gap,ji}(t) = g_{g,ji}(V_j(t) - V_i(t))$$

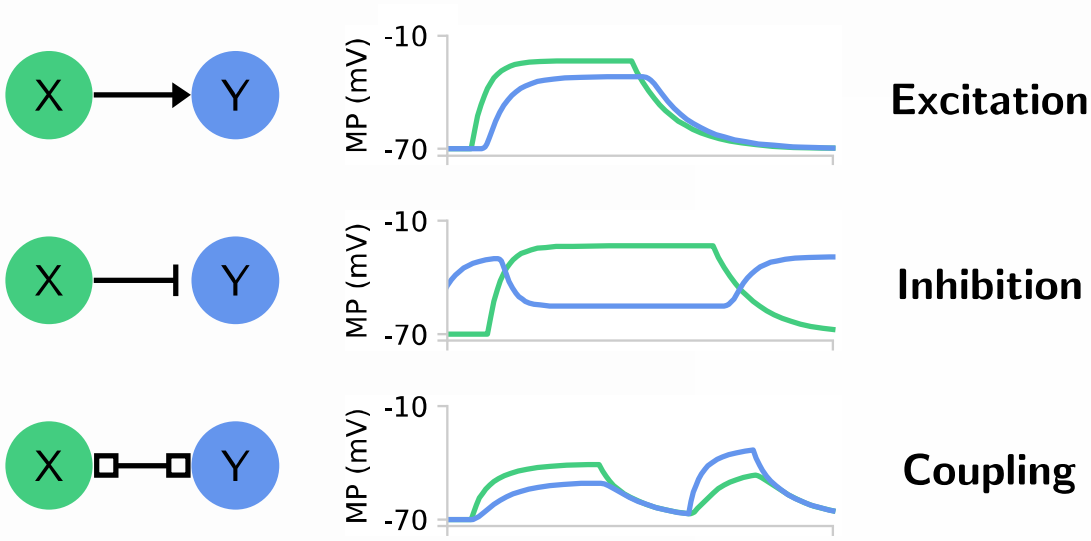
Chemical Synapse



$$I_{syn,ji}(t) = g_{s,ji}\sigma(V_j(t), \mu_j)(E_{s,ji} - V_i(t))$$

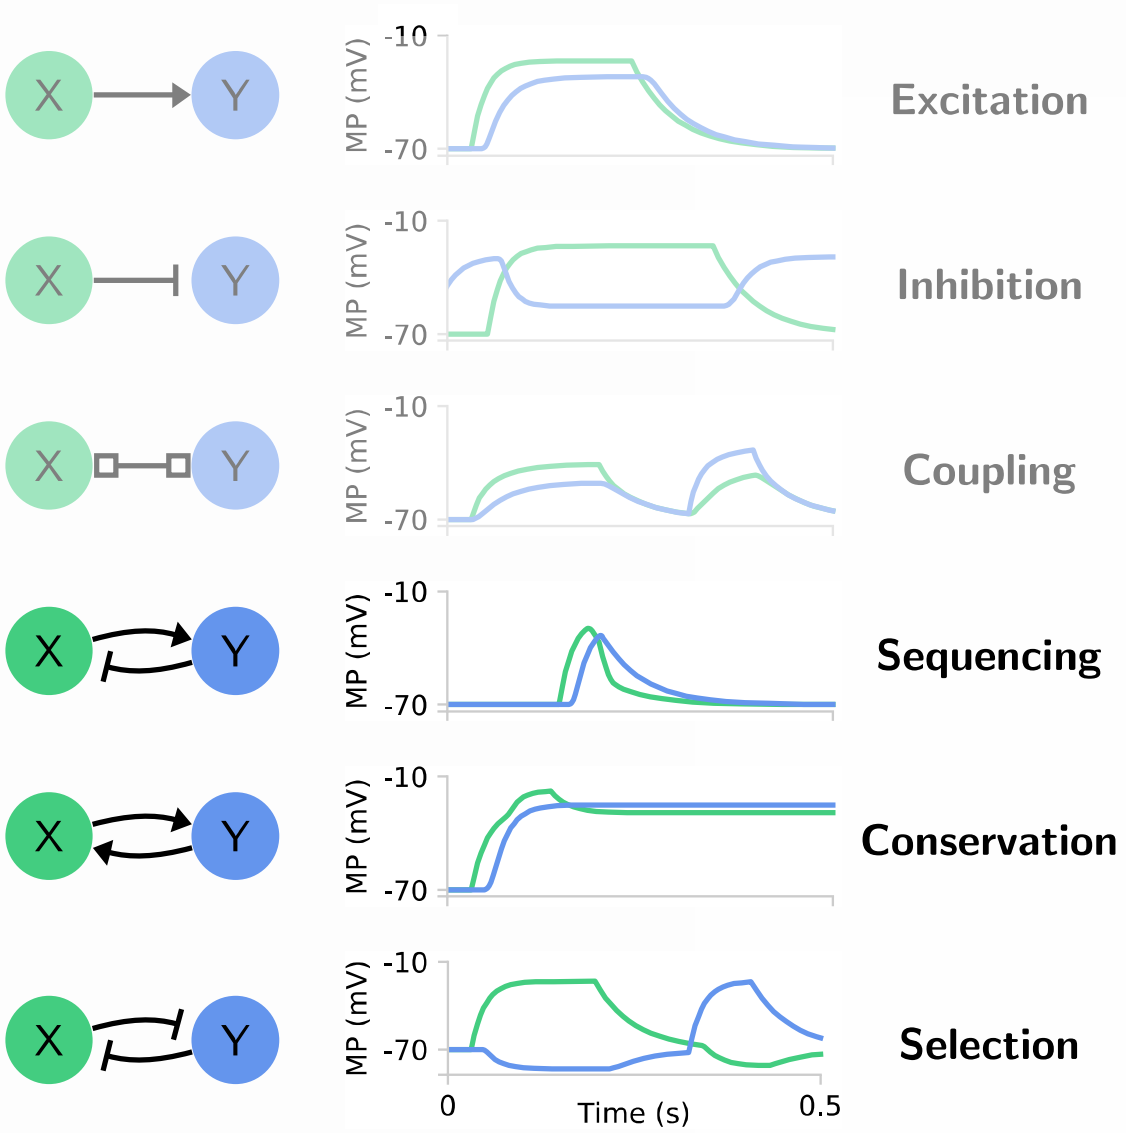
$$\sigma(V_j(t), \mu_j) = \frac{1}{1 + e^{-(V_j(t) - \mu_j)}}$$

Primitive Control Motifs (NCPs)



ICRA'19

Feedback Control Motifs (NCPs)



ICRA'19

Policy Motifs: While Loop

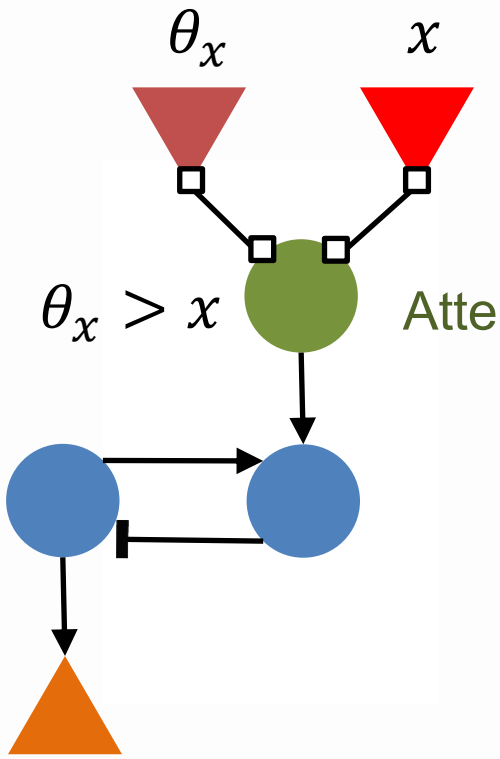
ICRA'19

Sensory Neurons

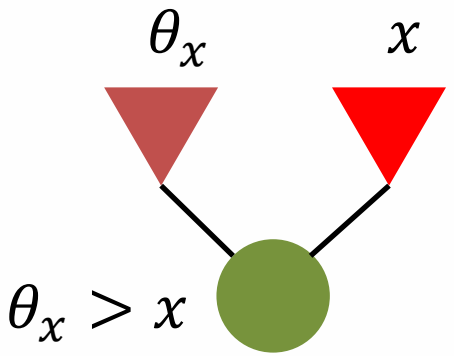
Inter Neurons

Command Neurons

Motor Neurons



Attention



while $(\theta_x \leq x)$

goBack

Parallel-Parking NCP: Composition

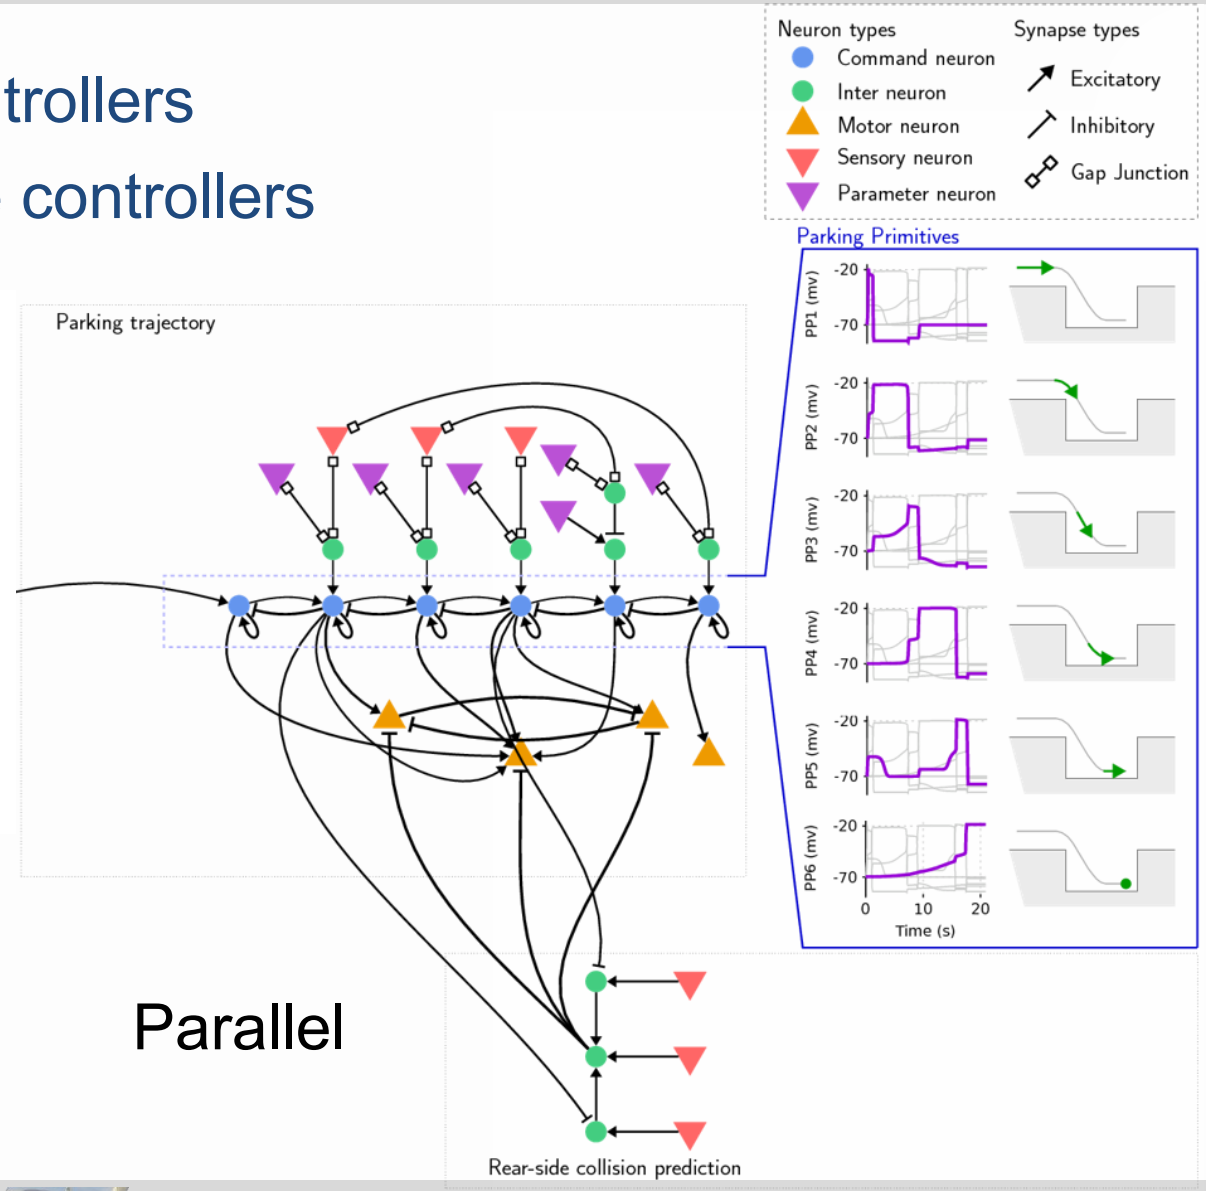
Motifs: Local controllers

NCP: Compose controllers

Sequential

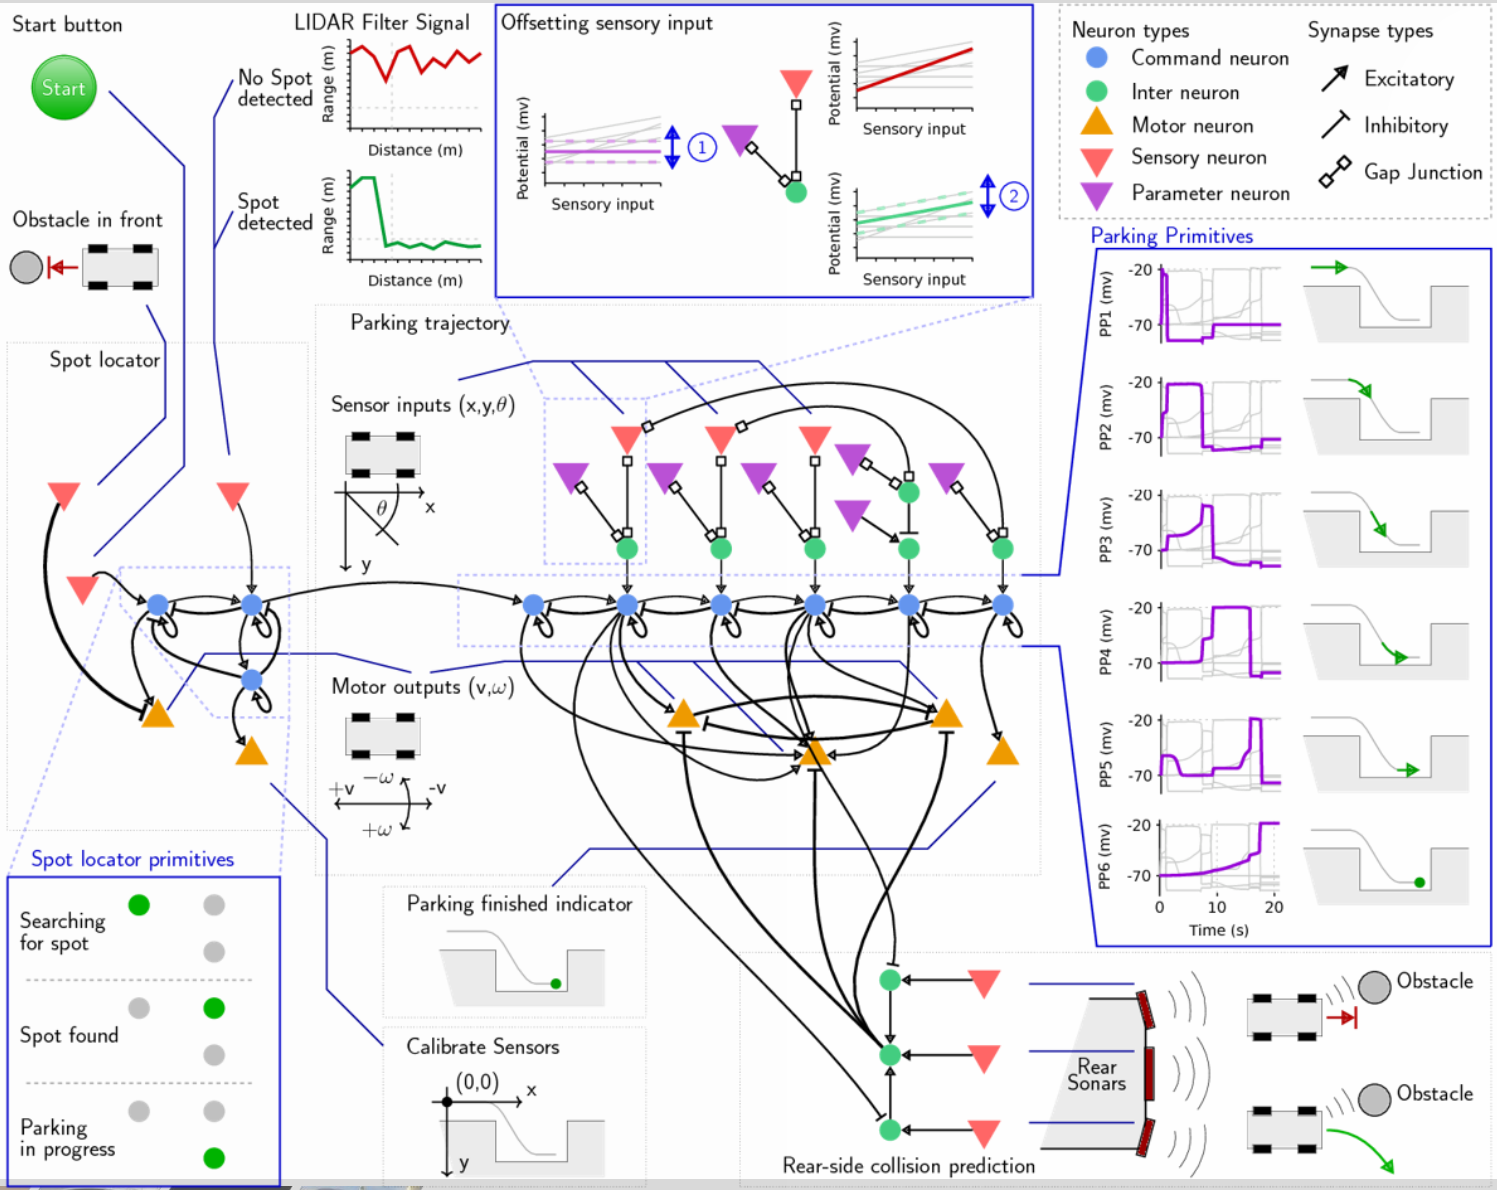
ICRA 2019

Parallel



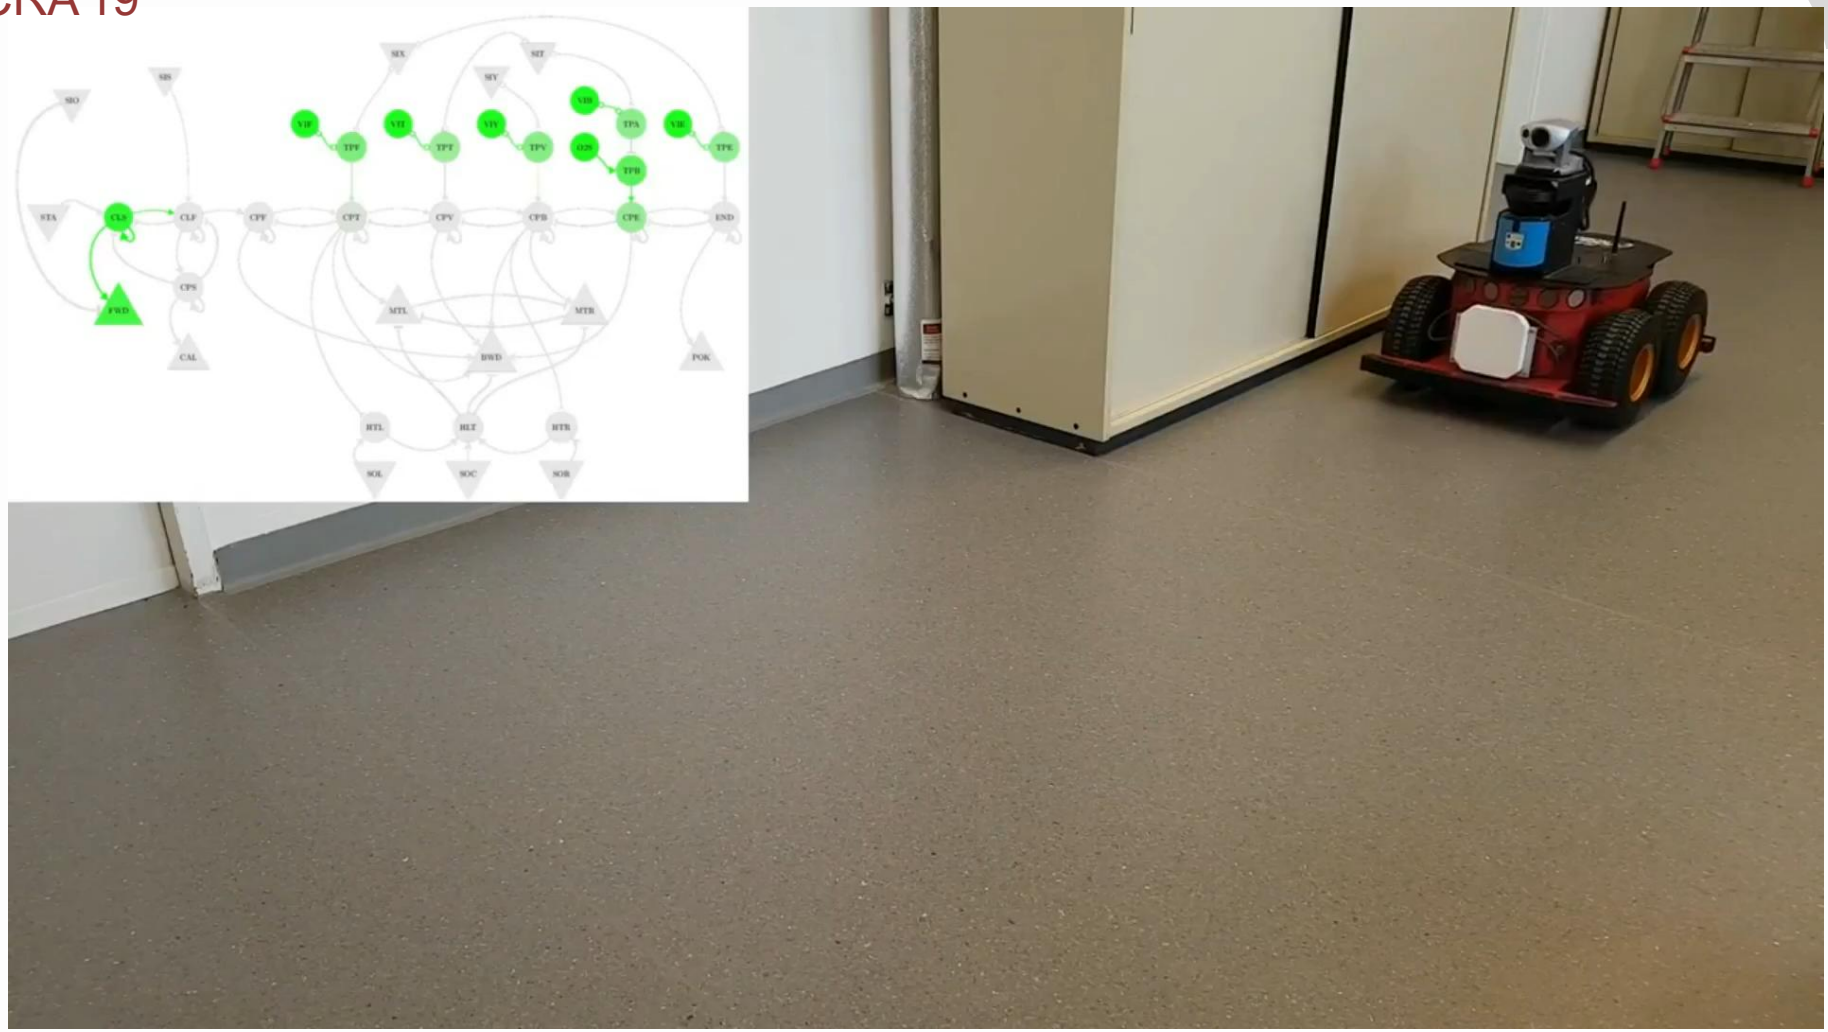
Parallel-Parking NCP: Hierarchy

ICRA'19



Parallel-Parking NCP

ICRA'19



Lane Keeping in Autonomous Driving



Alexander Amini



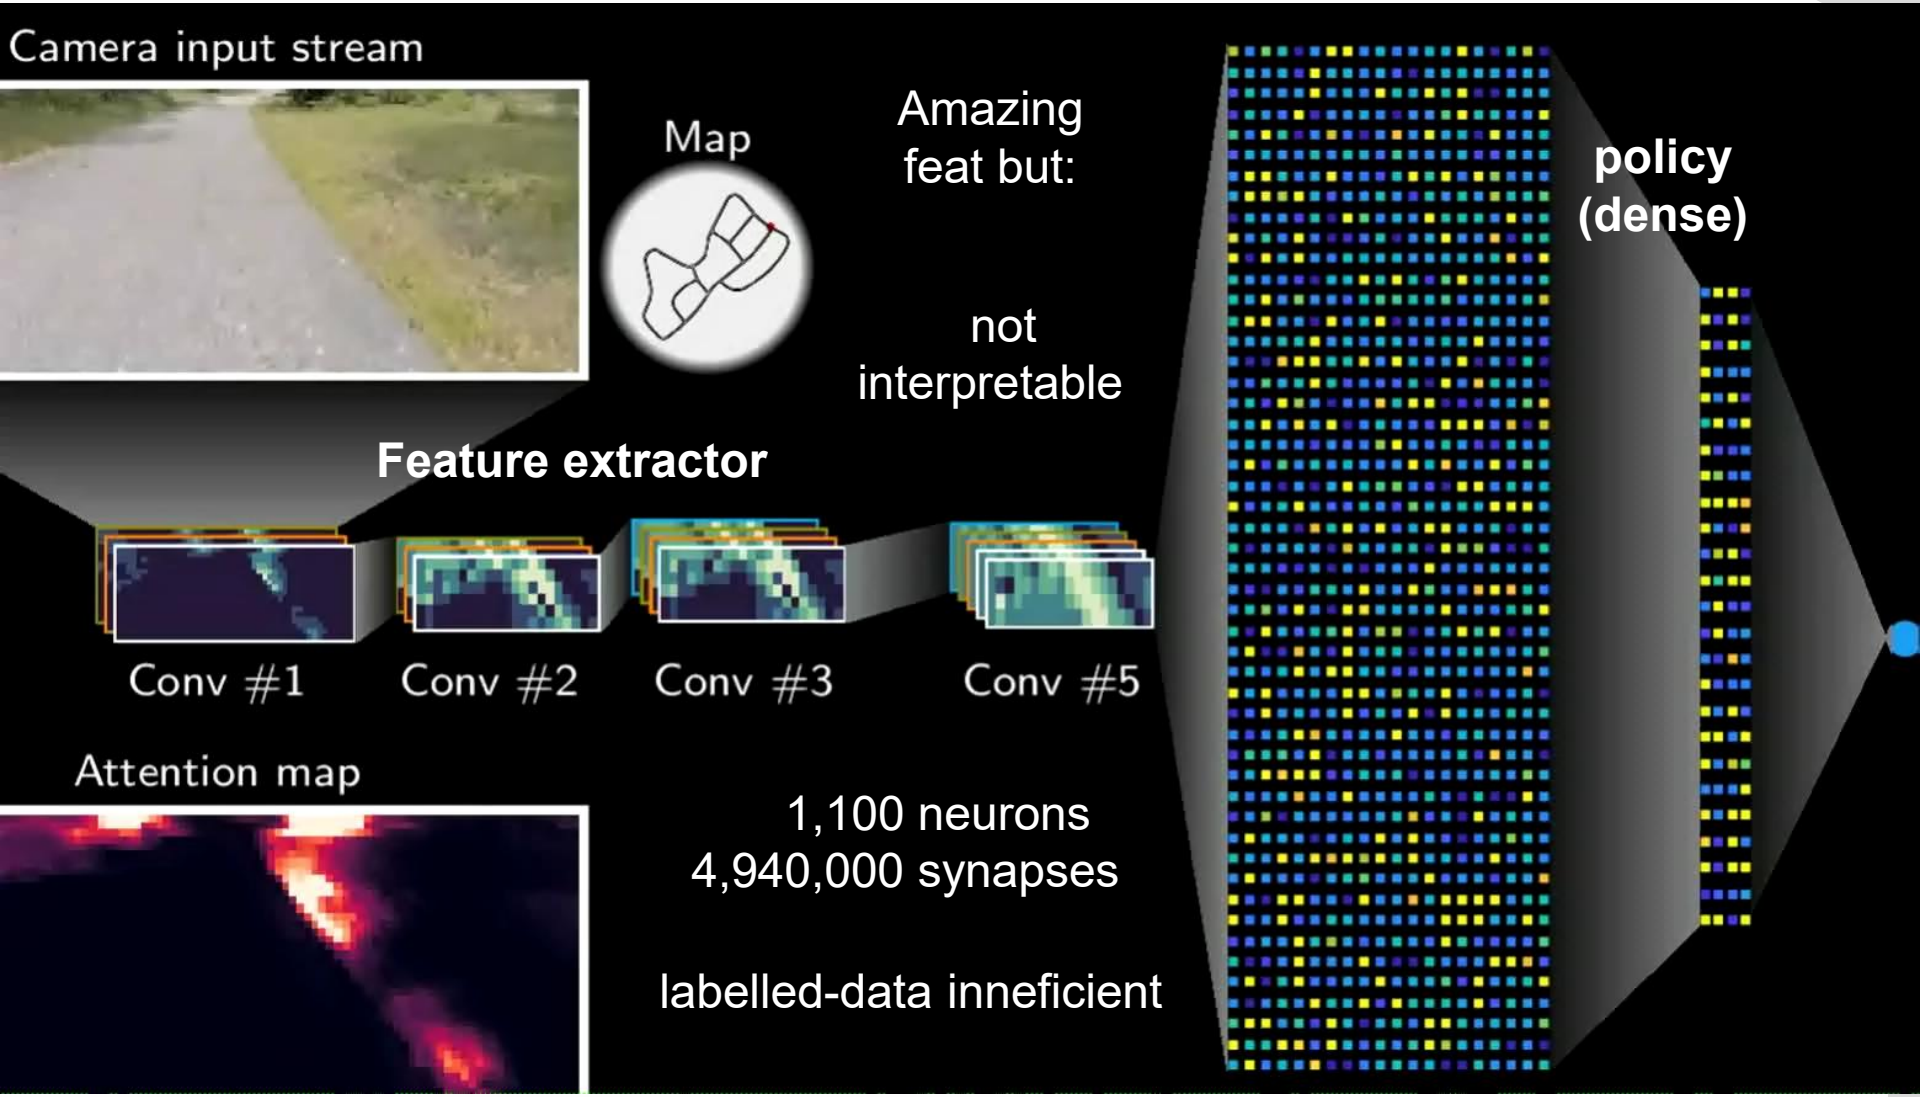
Daniela Rus

MIT-Toyota: A Fully Data-Driven Approach

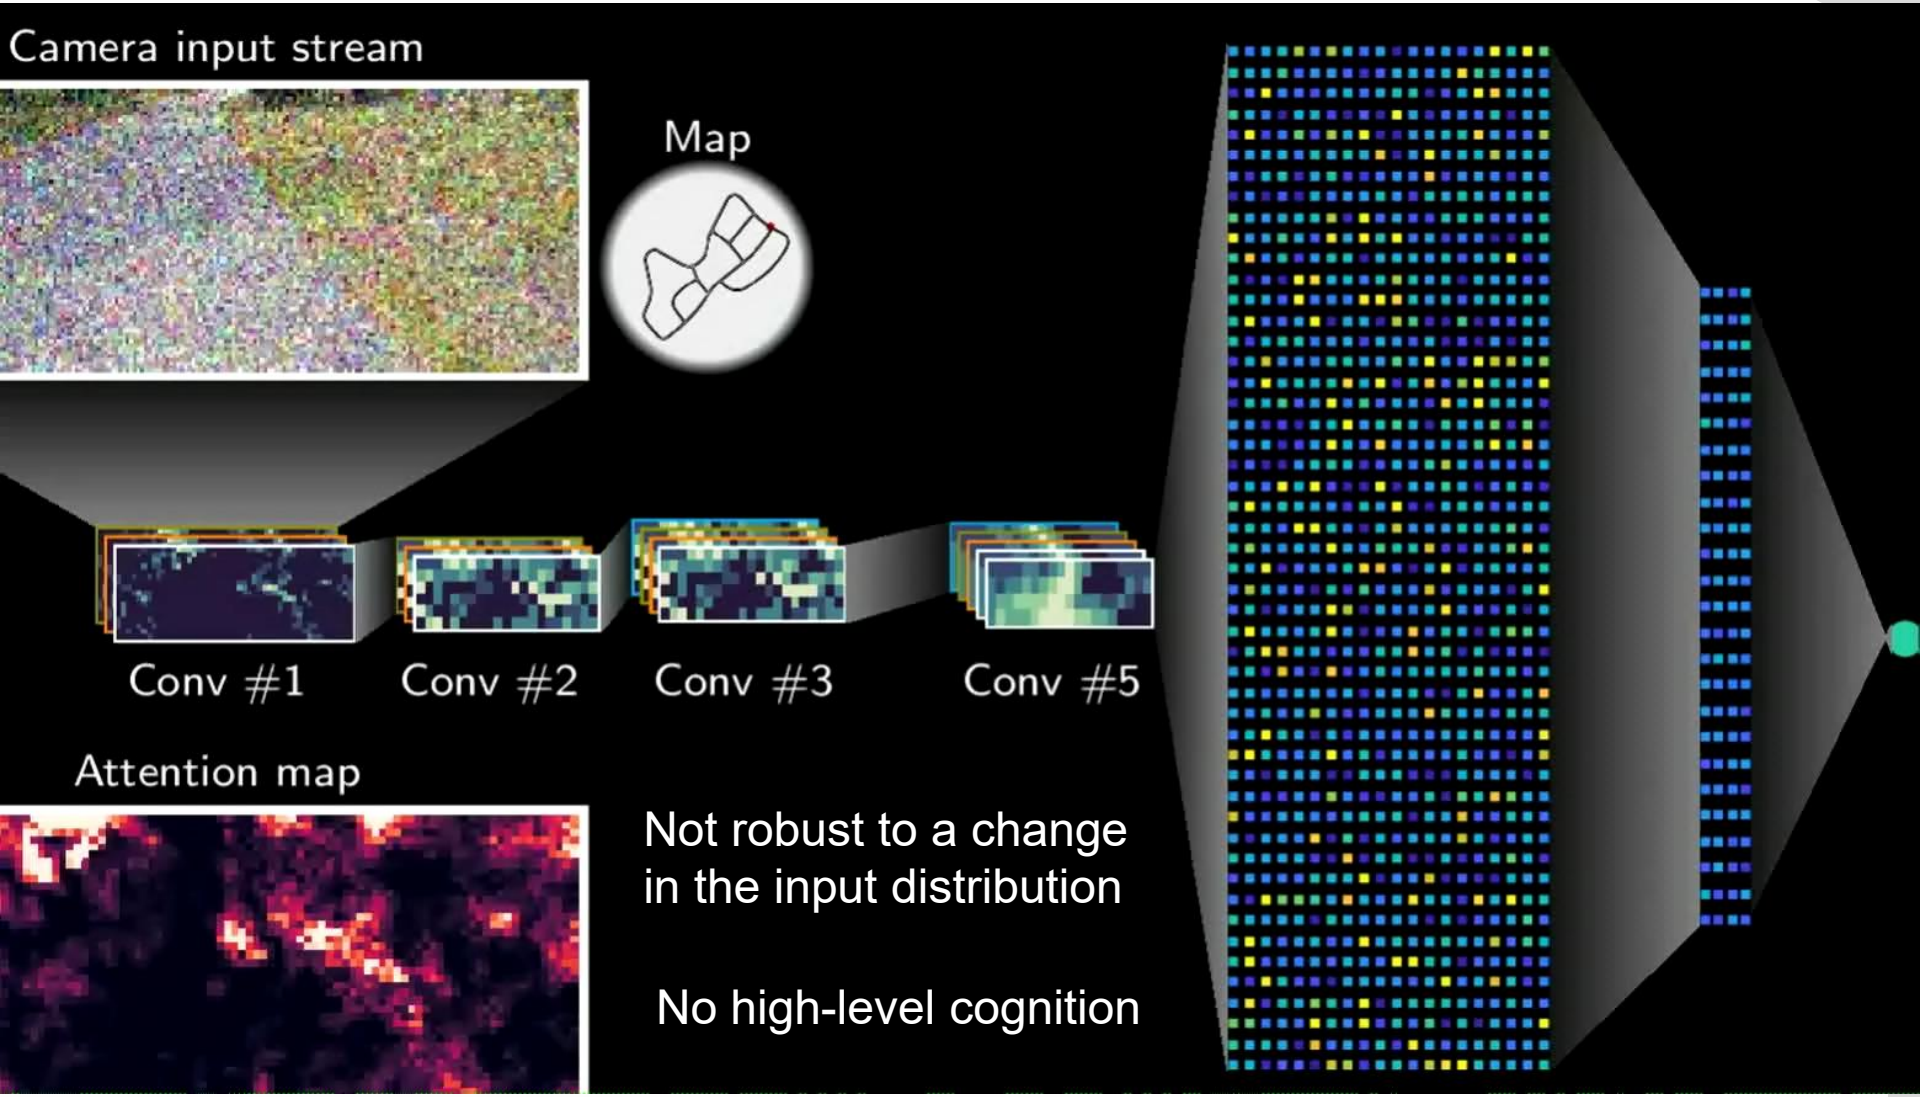
Supervised Learning: Learning from (state, action) Sequences

End2end: Learn one CNN from Sensors to Actuators

Lane Keeping: MIT End-To-End Solution



MIT End-To-End Solution while Raining



NCP Group TUW-MIT-IST

TECHNISCHE UNIVERSITÄT WIEN



Radu Grosu



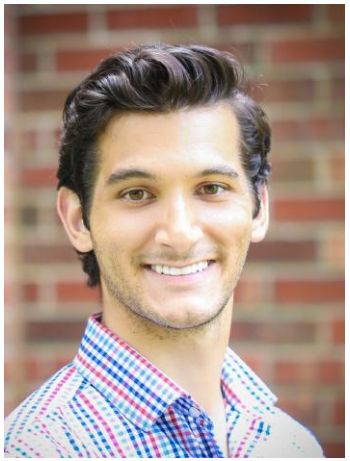
Daniela Rus



Tom Henzinger

MIT CSAIL

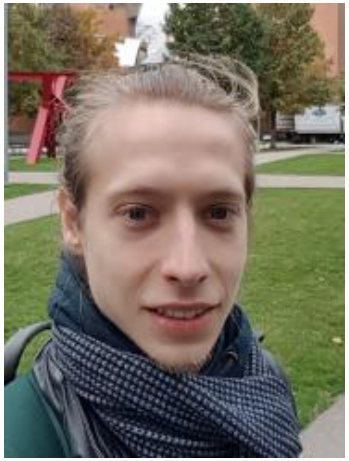
FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics



Alexander Amini



Ramin Hasani



Mathias Lechner

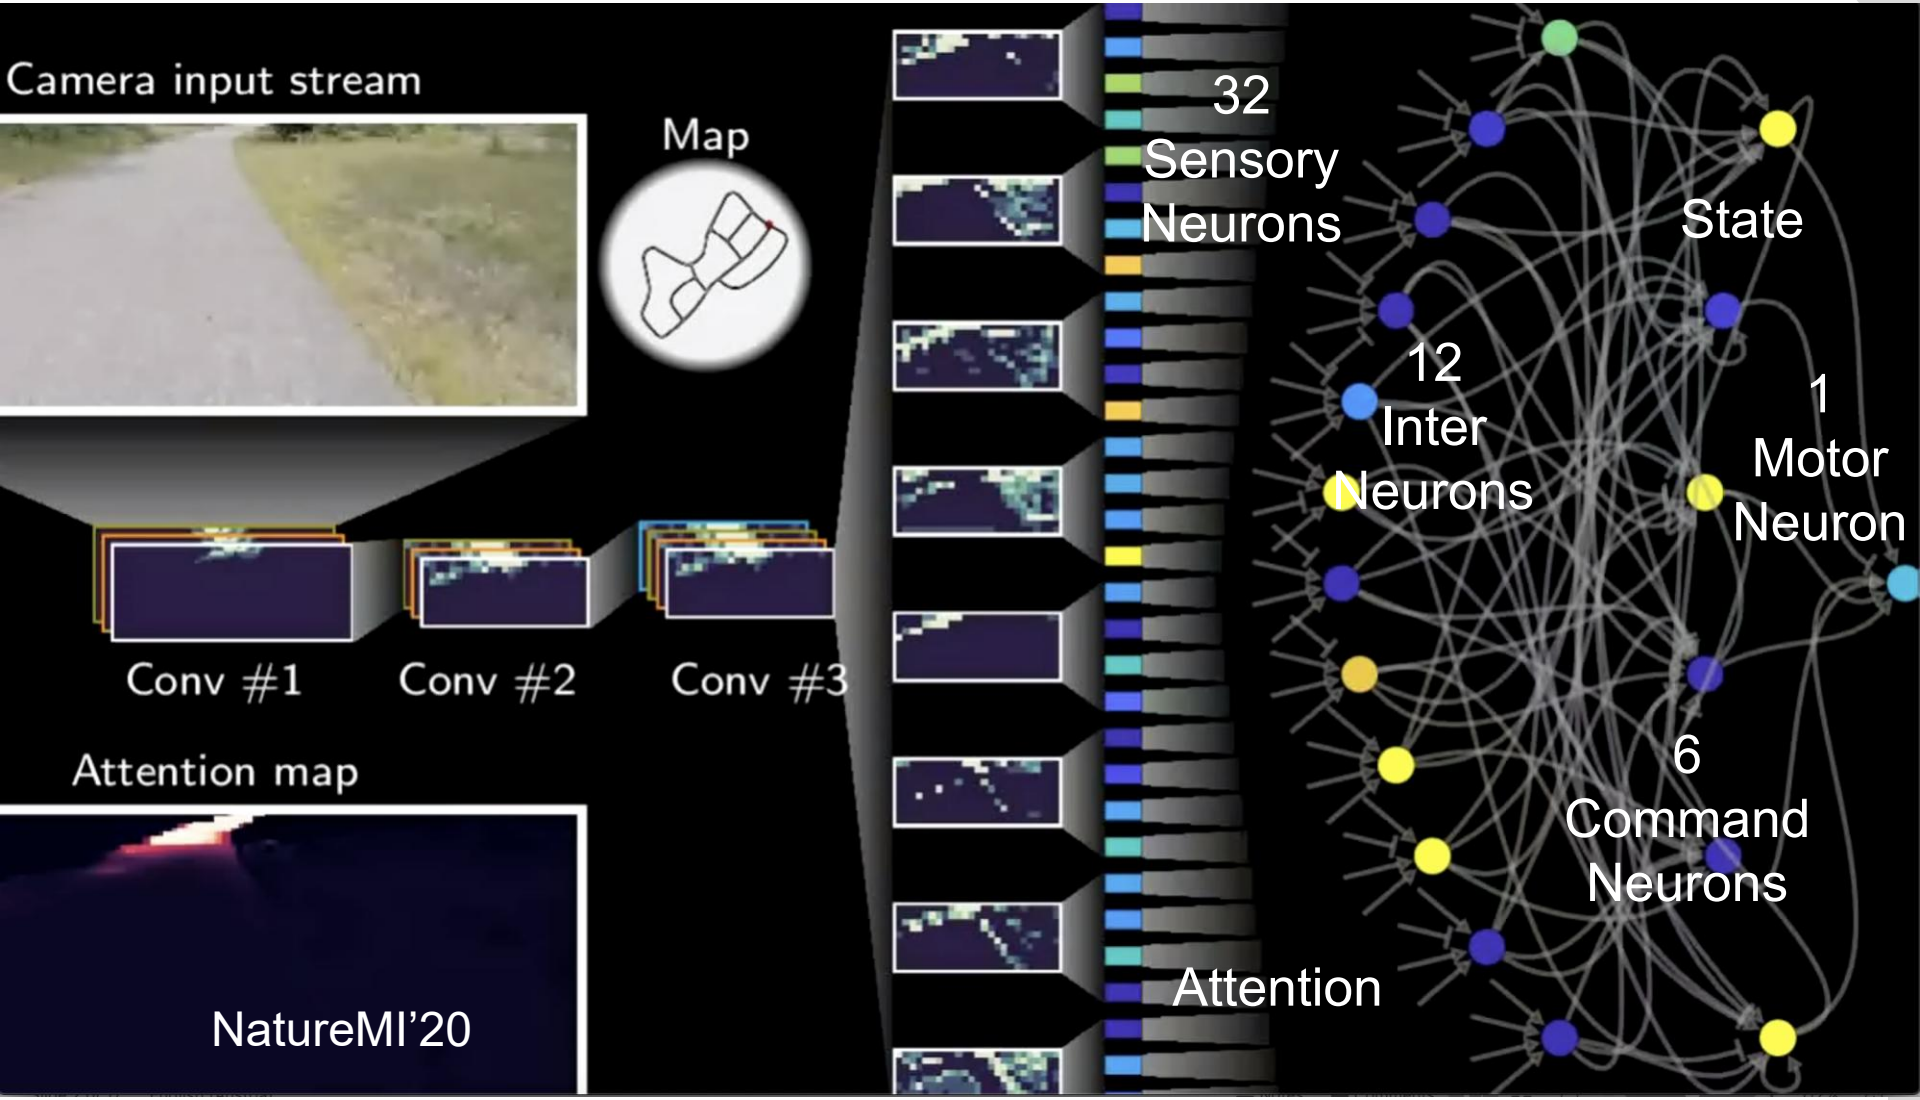
IST AUSTRIA
Institute of Science and Technology



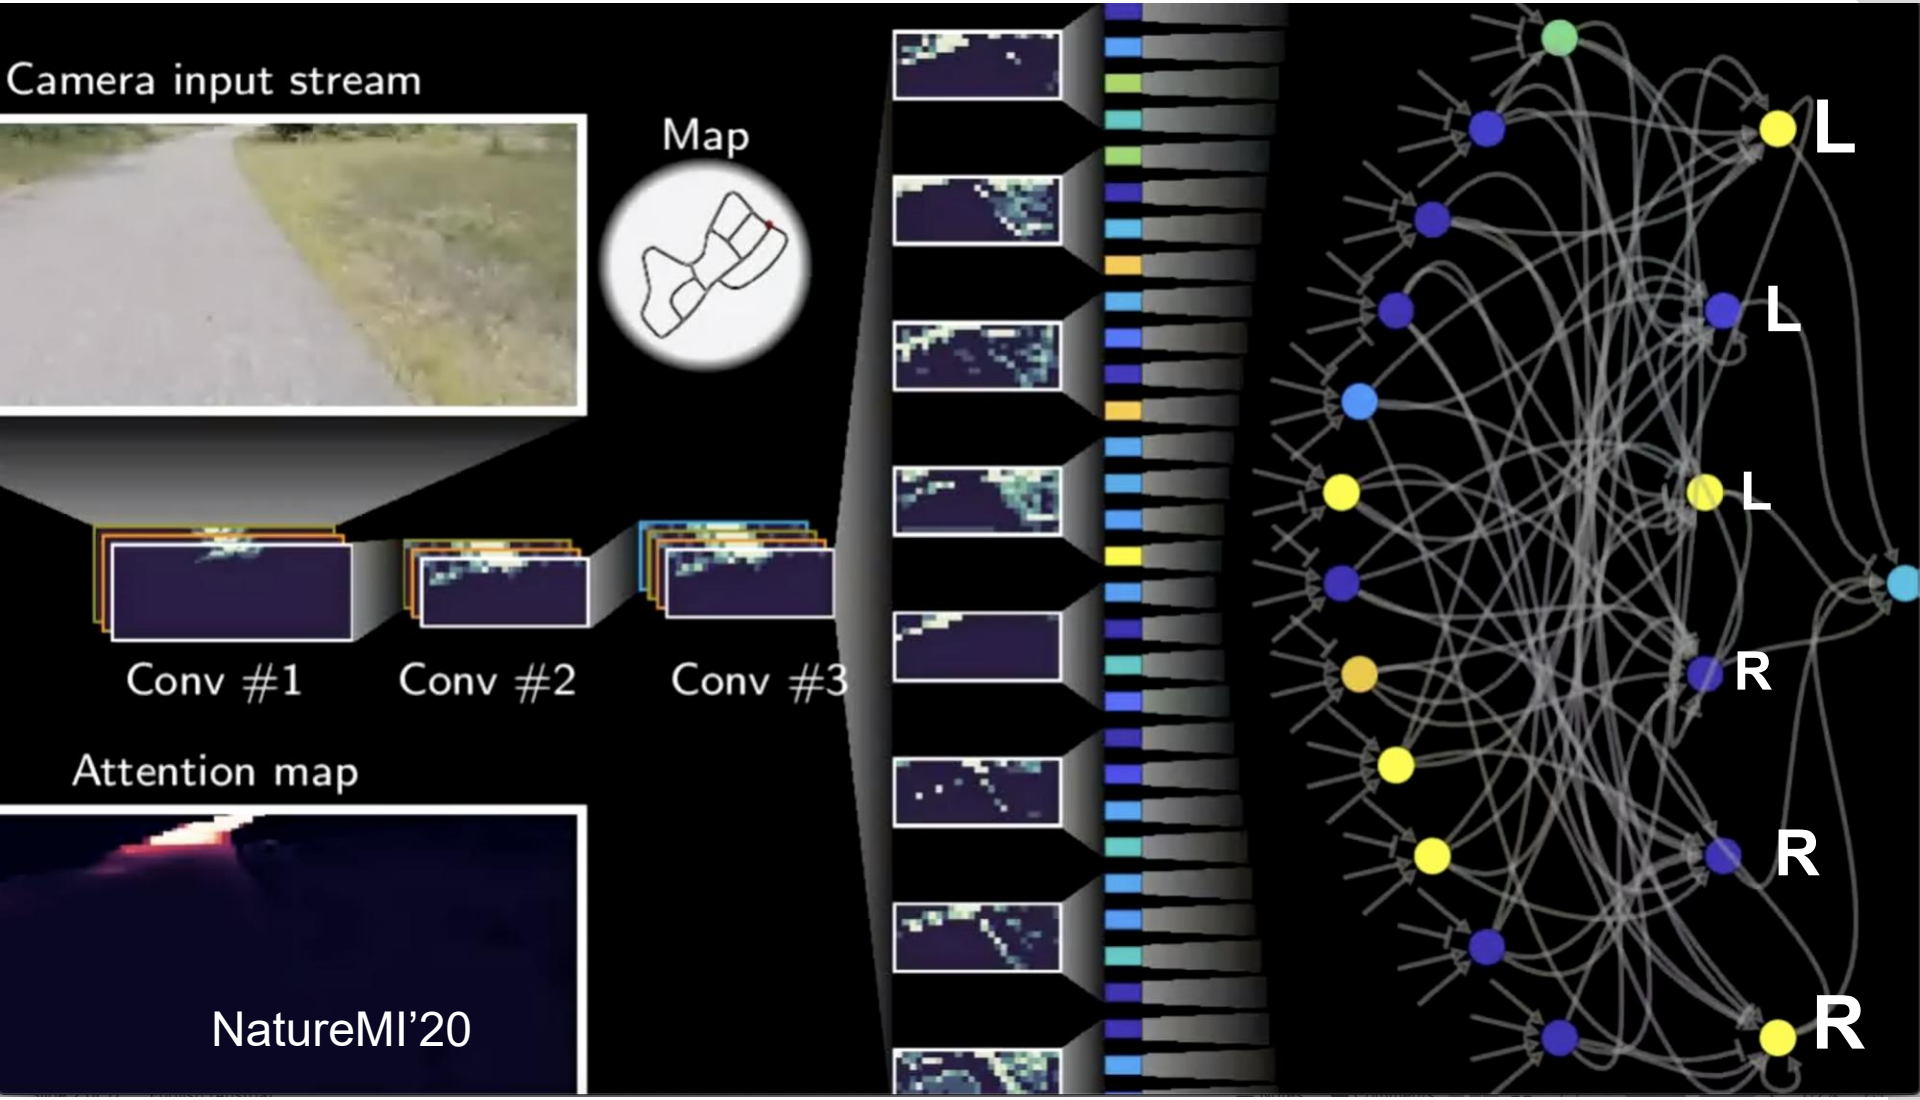
TU WIEN

Cyber-Physical-Systems Group

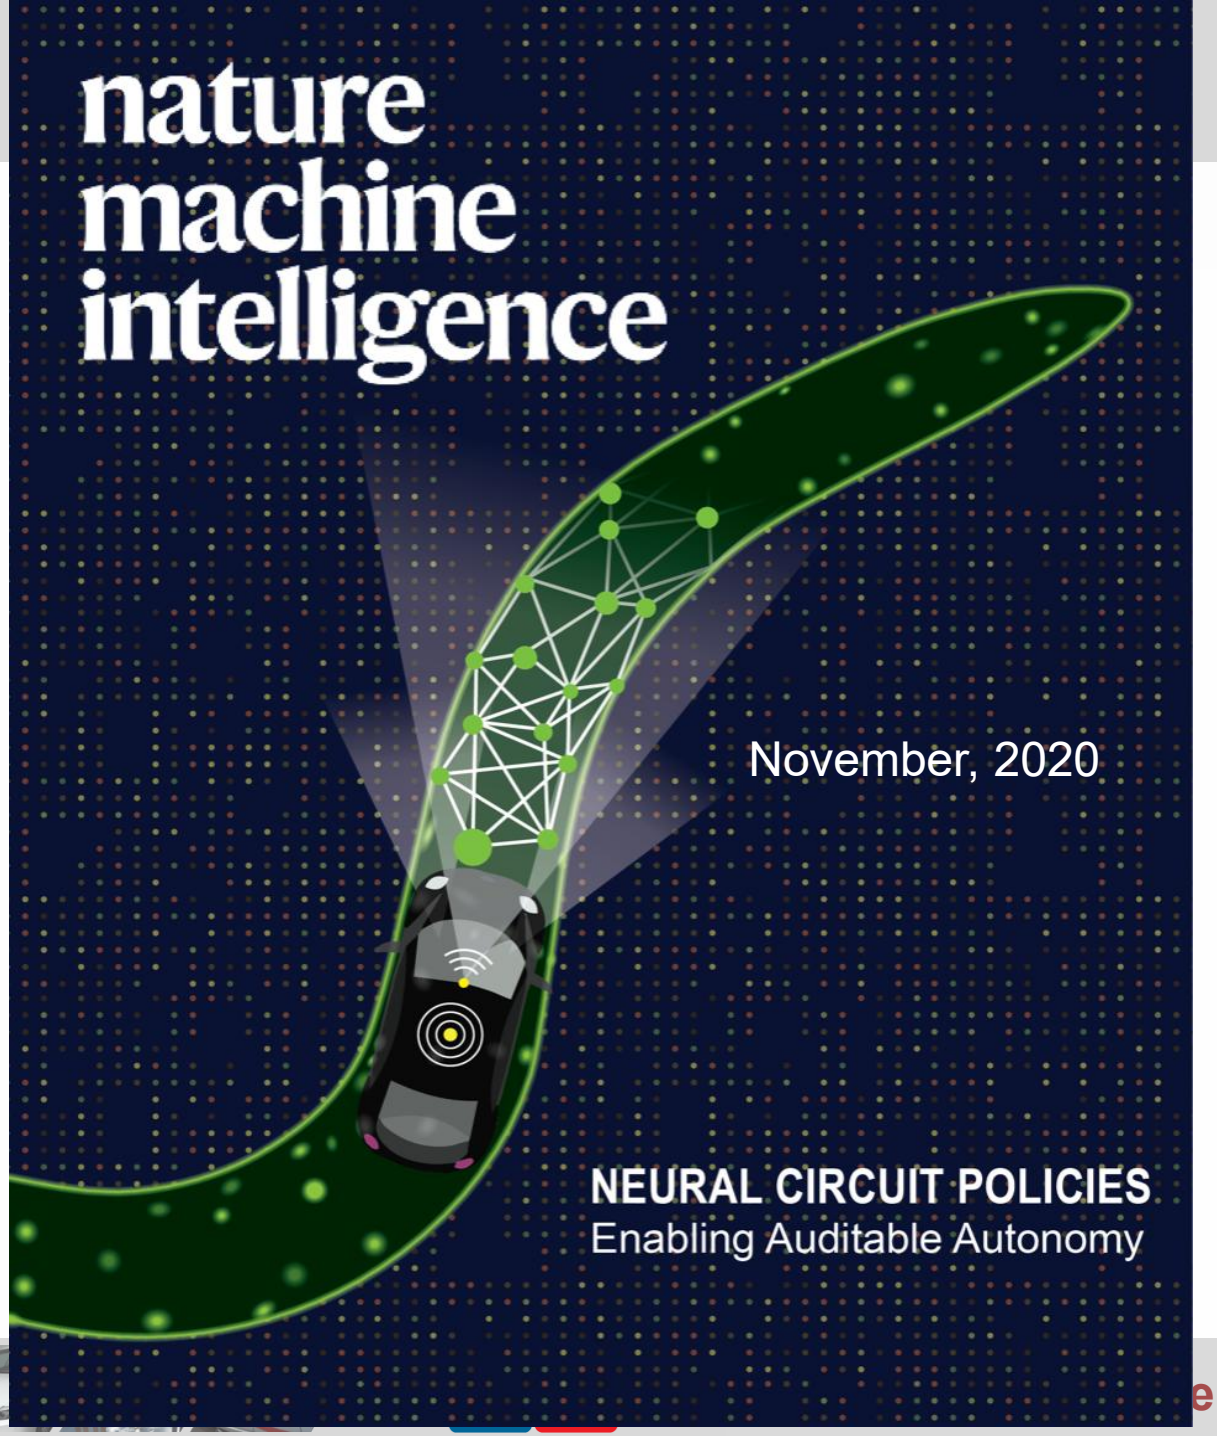
NCP Architecture for Lane Keeping



NCP Architecture for Lane Keeping



NCP Architecture in Action


nature
machine
intelligence
November, 2020
NEURAL CIRCUIT POLICIES
Enabling Auditable Autonomy

Netexplo Innovation Forum 2021



14 April 2021 - The Innovation Forum 2021 will showcase 10 breakthrough digital innovations with the potential of profound and lasting impact on the digital society. The innovations are spotted worldwide through a year-long work by the Netexplo University Network, made up of globally renowned universities in the technology domain. Some among the most promising global innovators in digital technologies will present their views and discuss about the trends that are shaping our futures.

10. WORM BRAIN DRIVER

Austria

Frugal AI inspired by primitive brains is enough to drive an autonomous vehicle.

Taking inspiration from the tiny brains of threadworms, researchers developed a frugal AI system capable of controlling a self-driving car. Only a few neurons were needed, compared with the millions used in most networks. Using only 19 neurons in the control circuit and just 75,000 parameters, the system successfully drove a vehicle. By using nature to improve deep learning, the team believes their new method can reduce training time and make it possible to implement AI in relatively simple systems.

https://en.unesco.org/news/unesco-netexplo-forum-2021
https://github.com/mlech26l/keras-ncp


The Innovation
Excellence Award
For the Outstanding Application of HPC
Hyperion
Research
The Innovation
Excellence
Award
HPC
USER FORUM
Global • 2022
Presented to:
Ramin Hasani (CSAIL, MIT), Mathias Lechner (IST Austria),
Alexander Amini (CSAIL MIT), Thomas Henzinger (IST Austria),
Radu Grosu (TU Wien), Daniela Rus (CSAIL, MIT)
For the Outstanding Application of HPC for
Business and Scientific Achievements
2022

LiquidAI Valued at Hundreds of Millions



A Problem However Remained (LRCs)

Integrating neural differential equations is hard:

$$C_i \dot{V}_i(t) = \sum_j I_{ji}(V(t), U(t))$$

Problem: equations are very stiff and hard to integrate.

The solution is as simple as it is unexpected:

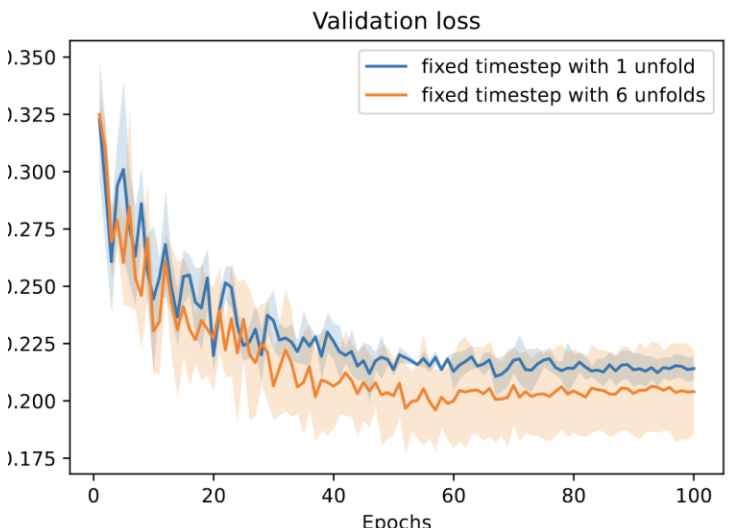
$$\dot{V}_i(t) = \sum_j I_{ji}(V(t), U(t)) / C_i(V(t), U(t))$$

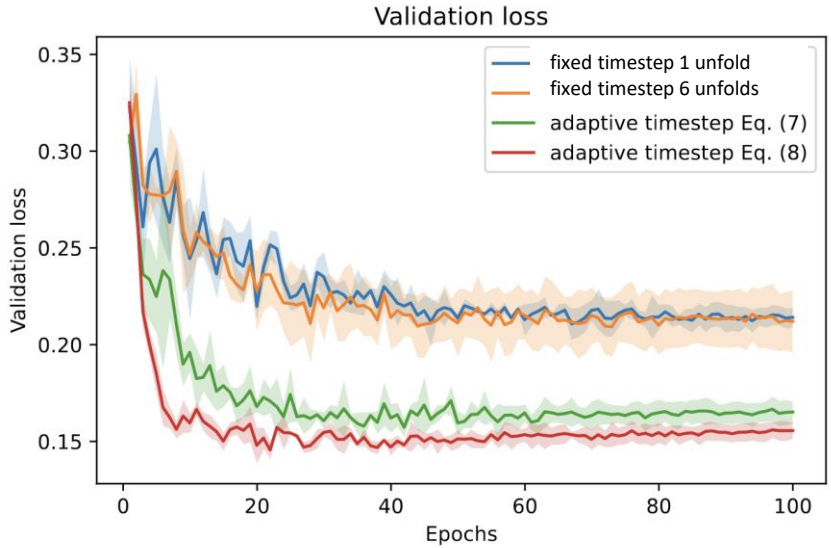
Optimal conductance can be simultaneously learned! LRCs

ICRA'24, NeurIPS'25

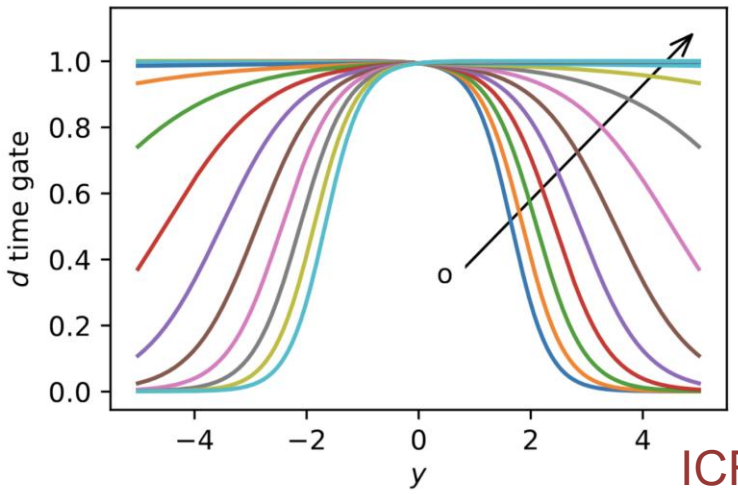


The Difference Made by the Time Gate (GCU)





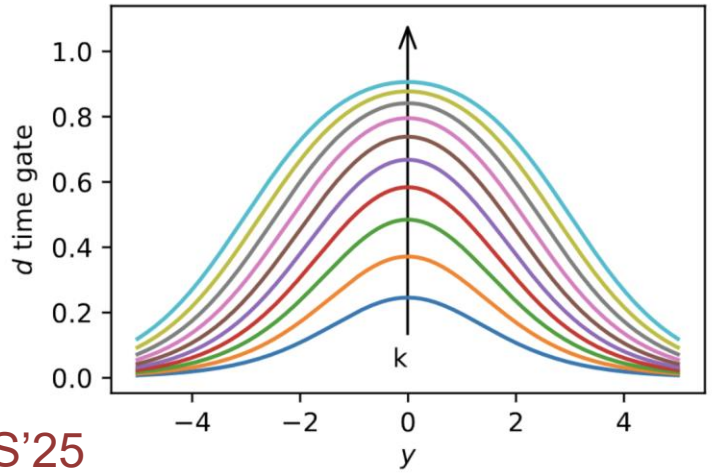
(7) $\delta_t = \sigma(w_t)$



$y_t = [V_{t-1}, U_t]$

$w_t = Oy_t + p$

(8) $\delta_{i,t} = \sigma(w_t + k) + \sigma(w_t - k)$



ICRA'24, NeurIPS'25



Why are LrcSSMs Sustainable?

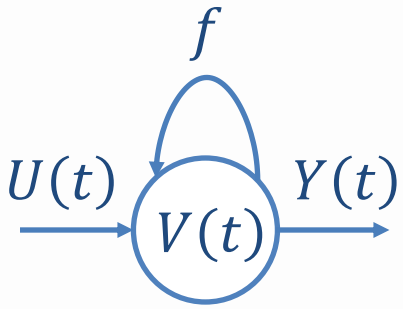
Because of Recursion

NeurIPS'25

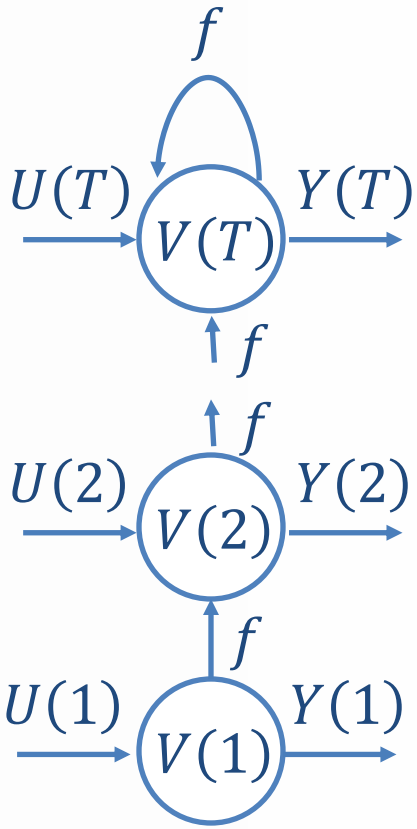
$$\dot{V}(t) = f(V(t), U(t))$$

$$Y(t) = g(V(t))$$

$$V(t+1) = V(t) + f\big(V(t), U(t+1)\big)dt$$



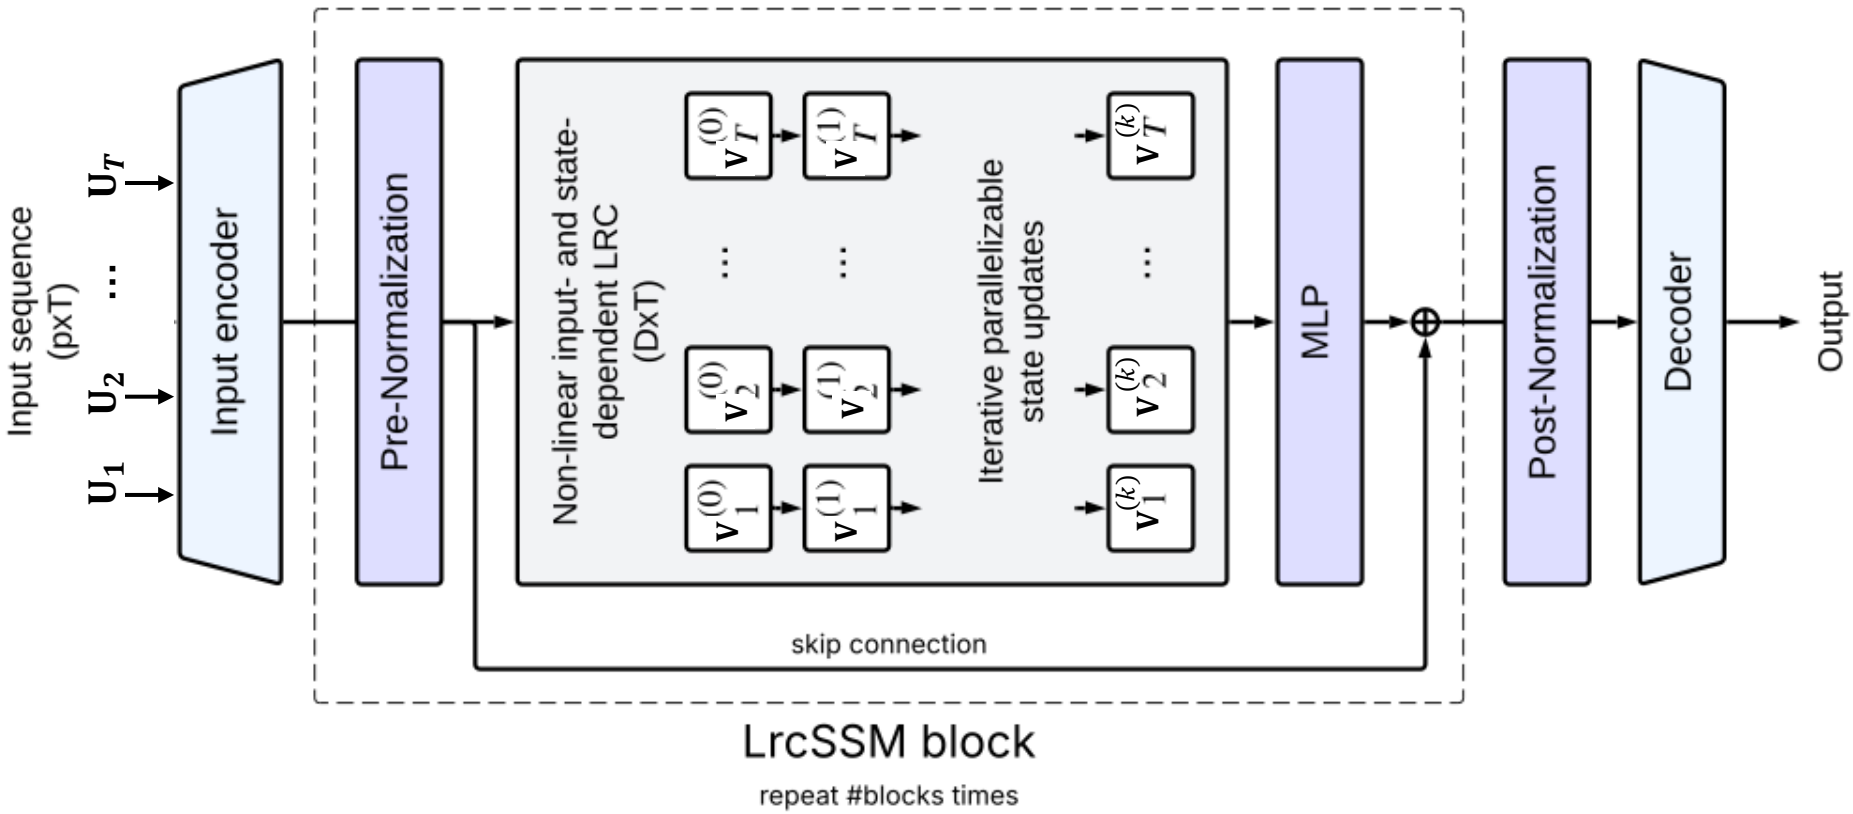
Inference



Virtual Unfolding During Learning

Problem: f^n

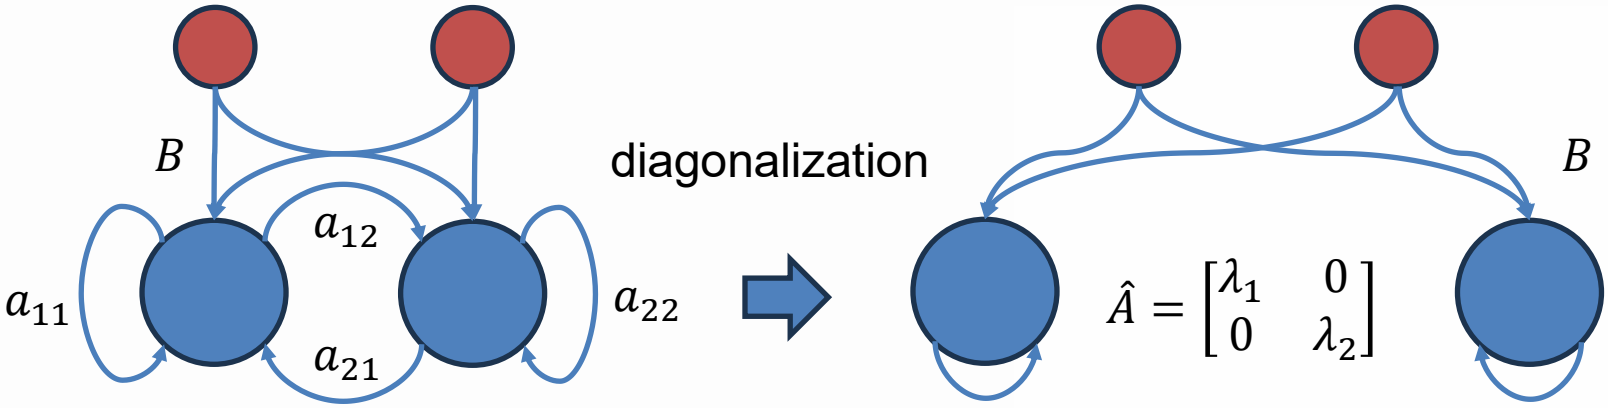
Scaling Up to an LLM (LrcSSM)



$$\dot{V}_i(t) = \underbrace{I_{ii}\big(V_i(t), U(t)\big)/C_i(V_i(t), U(t))}_{\text{diagonal in } V(t)} + \sum_j I_{ji}(U_t)/C_i(U(t))$$

ICRA'24, NeurIPS'25

Diagonalization



$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

$$\lambda_1 = a_{11}^* + a_{12}a_{22}^*a_{21} \qquad \lambda_2 = a_{22}^* + a_{21}a_{11}^*a_{12}$$

$$\lambda_1 = \frac{1}{1-a_{11}} + a_{12}\frac{1}{1-a_{22}}a_{21} \qquad \lambda_2 = \frac{1}{1-a_{22}} + a_{21}\frac{1}{1-a_{11}}a_{12}$$

$$\det(A - \lambda I) = \begin{vmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{vmatrix} = (a_{11} - \lambda)(a_{22} - \lambda) - a_{21}a_{12} = 0$$

ICRA'24, NeurIPS'25

NeurIPS 2025 Area-Chair Review

The authors present a novel deep SSM architecture that has nonlinear dynamics within each layer. The authors leverage recent advances in parallelization of nonlinear recursions to render their architecture computationally tractable.
As Reviewer dhK1 said:

This innovative paper has opened up new perspectives for me on the parallelization of nonlinear recurrent architectures. It has already begun to affect how I think about future research directions. Papers like these should definitely be discussed at conferences like NeurIPS.

I agree with this assessment and I congratulate the authors on their nice work.

NeurIPS'25



Conclusions

I presented LRCs, a path towards sustainable AI, that is:

- Biologically plausible, as the state of LRC neurons is their membrane potential, like in the EEC and LTC models;
- Biologically more faithful then EEC/LTCs, as LRCs capture saturation phenomena and state-input dependent capacitance;
- Extremely efficient, as LRCs converge considerably faster and to a considerably lower loss then LTCs do;
- Sustainable, as LRCs scale up very smoothly to the structured-state-space-model approach to LLMS.

