Rare Events, Real Signals:
Functional Ensembles as a Computational Substrate in Deep Networks

Aditi Aravind* Konstantinos Ladakis* Mario Alexios Savaglio
Stelios M. Smirnakis Maria Papadopouli +

Department of Computer Science, **University of Crete** & Institute of Computer Science, **FORTH**
Archimedes R.U., Athena Research Center
Brigham & Women's Hospital, **Harvard Medical School**

neurons nets https://neurons-nets.eu/

+ Corresponding author (maria@csd.uoc.gr)

Supported by **MSCA** (neurons-nets, #101007926) & **HFRI** (neuron-AD #4085)



Neuronal ensembles as units of cortical computation

- **Neuronal ensembles**: groups of neurons that co-activate more frequently than expected by chance
- **Weak cortical connectivity**: ensembles coordinate to form computational units
- **First order functional connectivity (1FC) ensemble**: neurons with statistically significant pairwise correlations with a target neuron

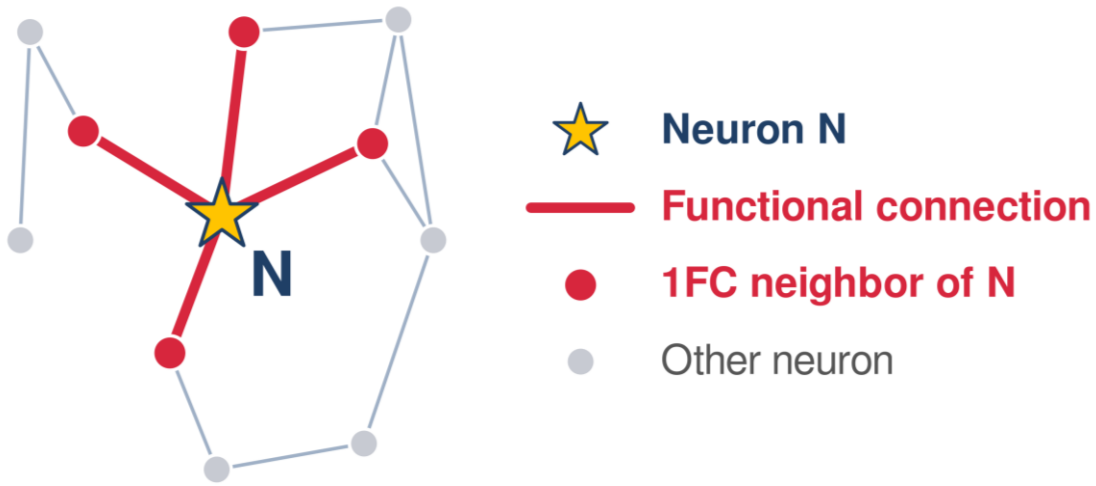

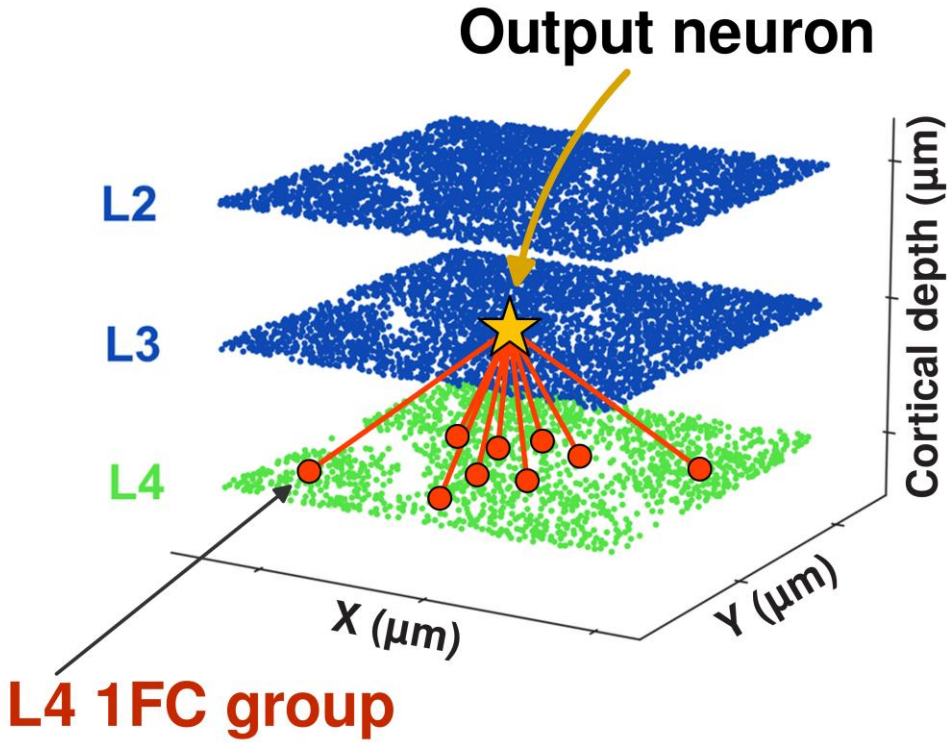

Mouse Visual Cortex

Mouse visual cortex: 1FC ensembles as communication units

- **Weak, statistically significant,** pairwise correlations
- **ReLU-like nonlinearity**, emerging when a large percent of L4-1FC inputs co-fire (~13%)

Mouse Visual Cortex

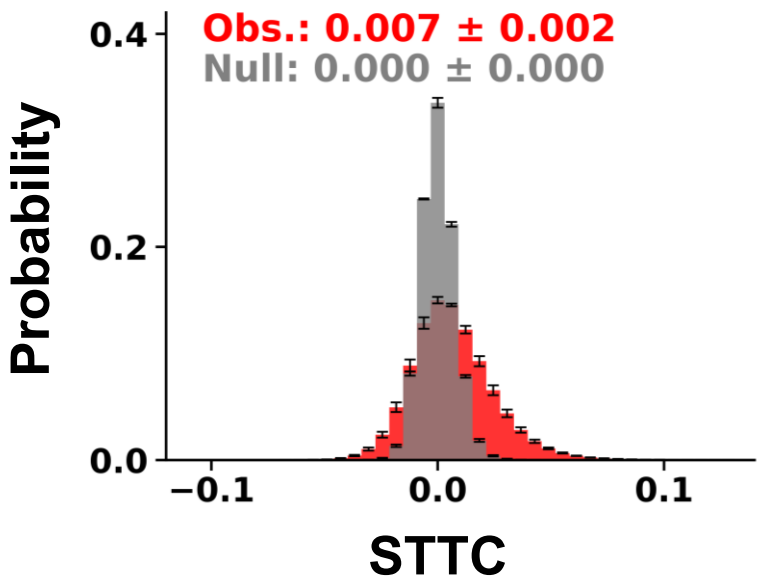

Example Cortical Neuron

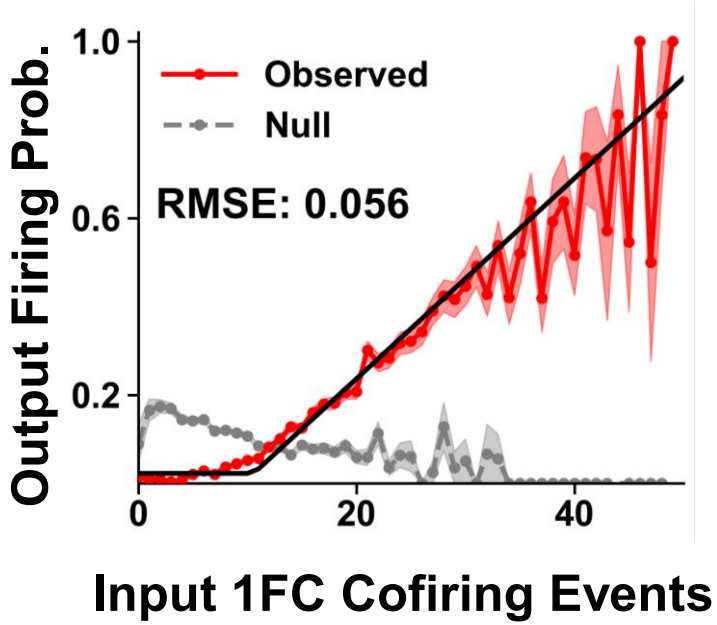

Are these principles preserved in deep networks?

Papadopouli et al. bioarxiv



Innovation: Use of **functional connectivity** for a **local** (node-level) view of ***when*** and ***through which routes*** **information propagates**

Existing methods

whole layer → one number

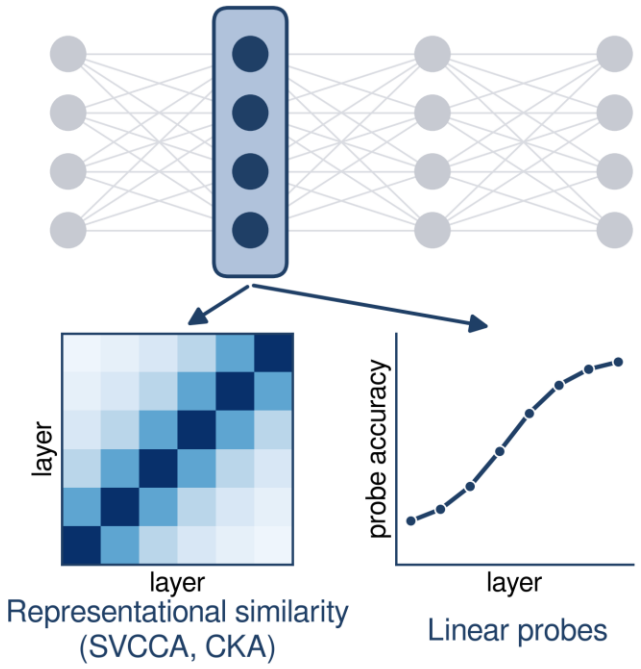

population/layer averages

Our approach

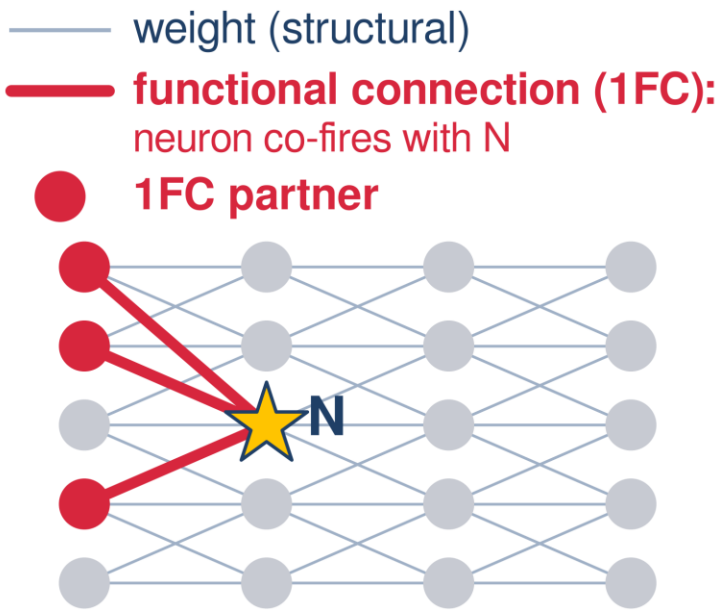

one neuron → its 1FC partners

pairwise functional connectivity (1FC)

- **Systems network neuroscience + information theory** → learned circuit representations

Architecture: Spiking ResNet-18

- Trained on CIFAR-100 (Accuracy 72.18%)

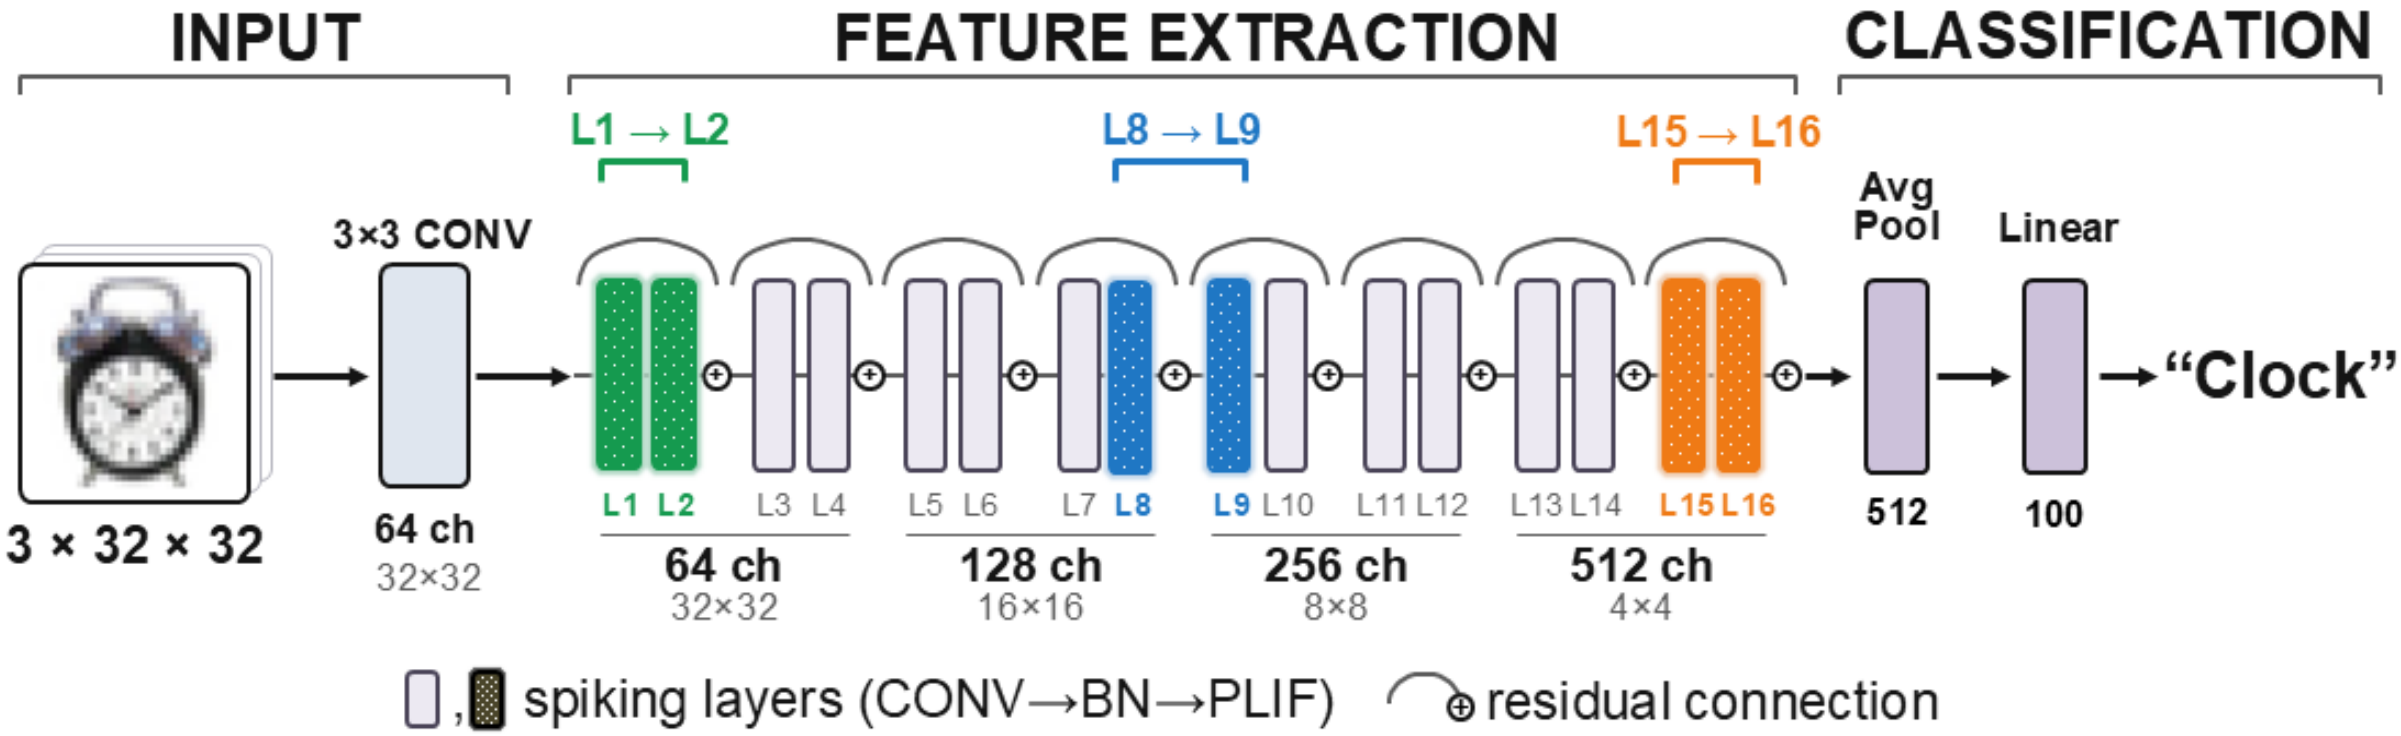

First-order functionally connected (1FC) groups

- **Spiking ResNet-18** on CIFAR-100 (72.18%)
- **1FC group**: layer L−1 neurons with significant positive pairwise functional connectivity (STTC)
- **Functional, not only structural**

Response function:
firing probability vs. aggregate co-firing of the input 1FC group

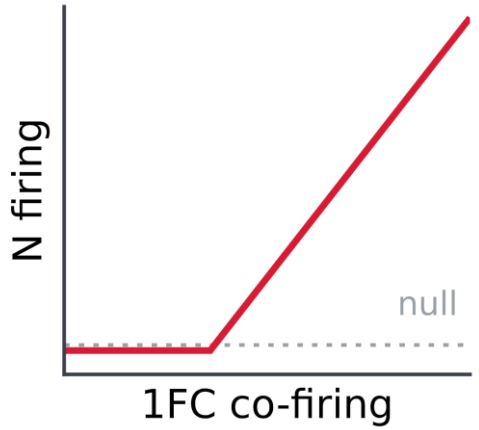

Neuron response

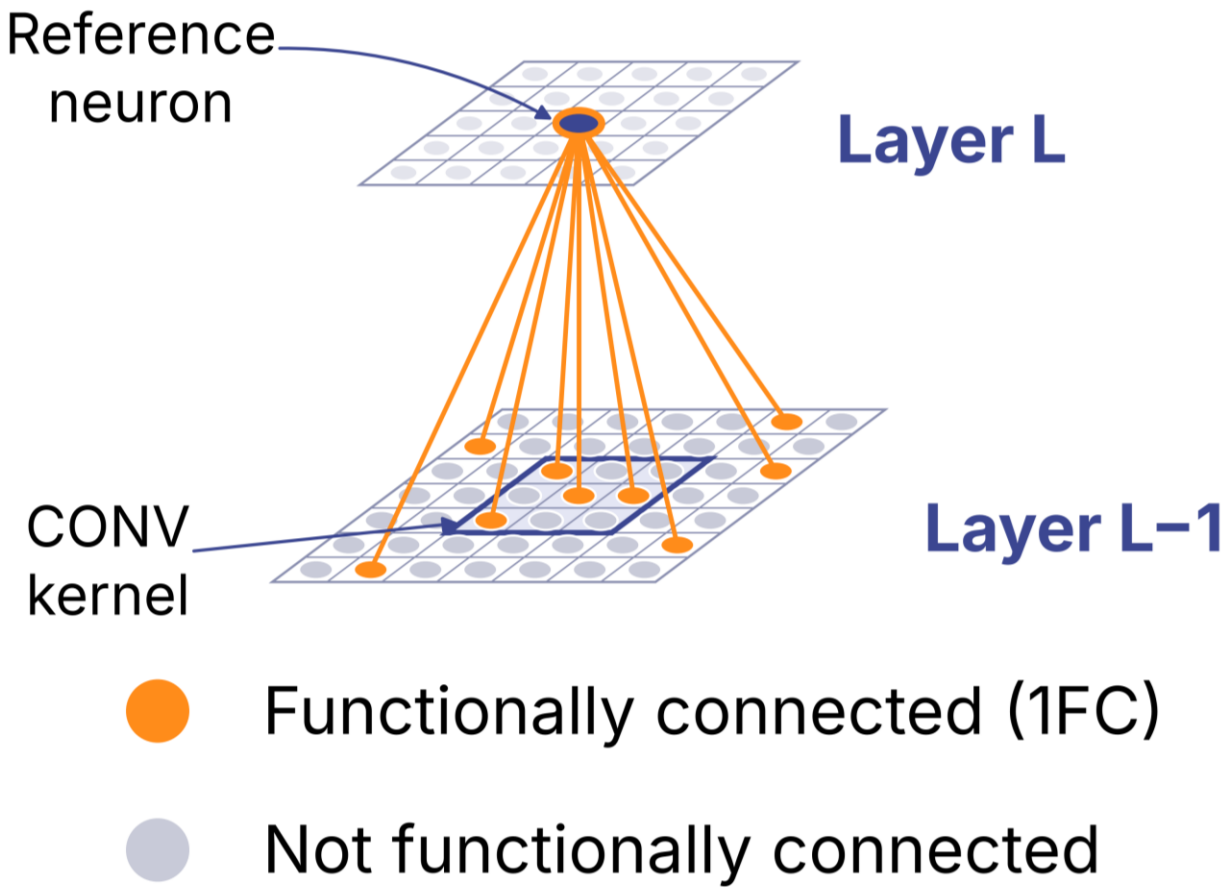

Weak correlations; ReLU-like response

- Statistically significant but weak pairwise correlations
- Response well fitted by ReLU; gain scales with ensemble size

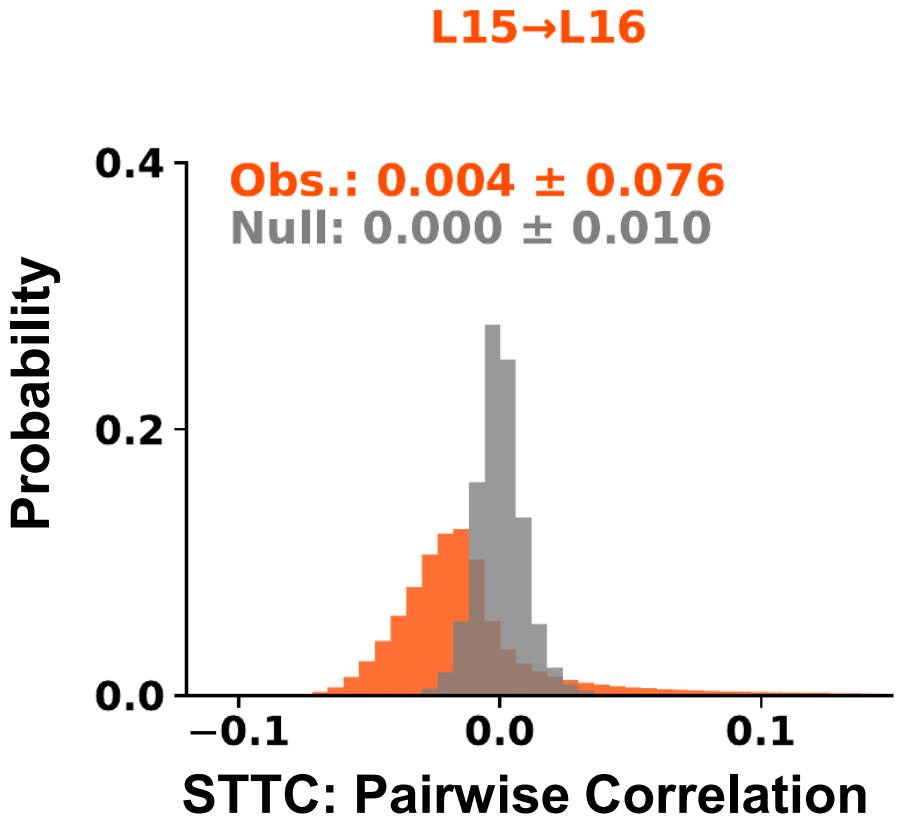

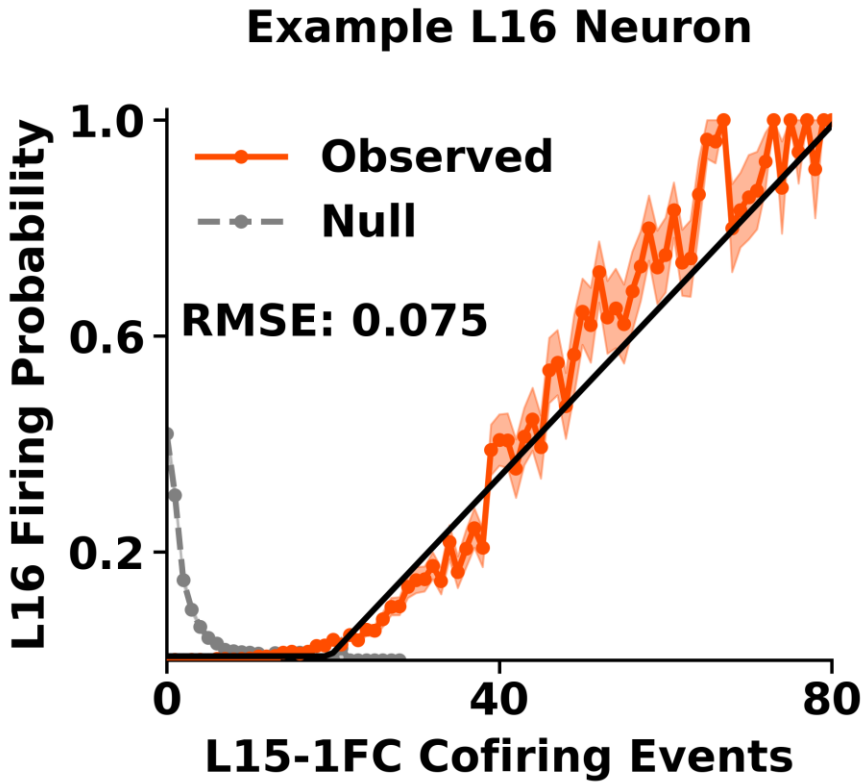

Example: L15 → L16 · Similar in early and intermediate layers

- **Two response regimes**: Firing of the target neuron rises sharply only when ≥ 5–16% of the 1FC group co-fires
- **Partner identity** adds little beyond aggregate co-firing in predicting the firing event of target neuron

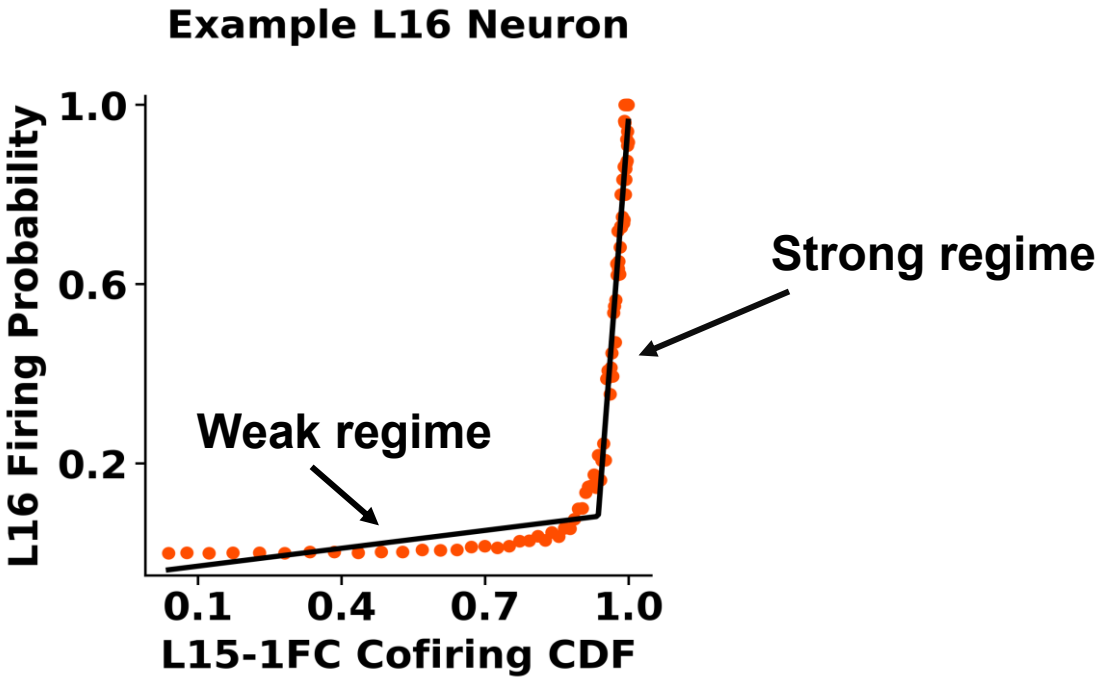

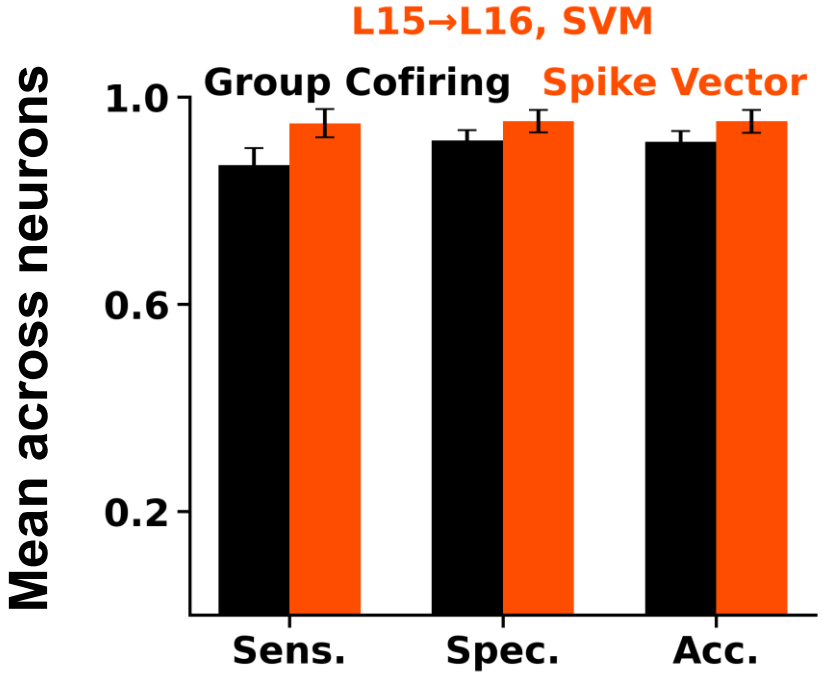

Example: L15 → L16 · Similar in early and intermediate layers

Representational significance

- **Larger 1FC co-firing** → more reliable encoding of the presented image
- **Rare but highly coordinated** activity patterns
- **Hub neurons**: high degree of connectivity, high mutual information

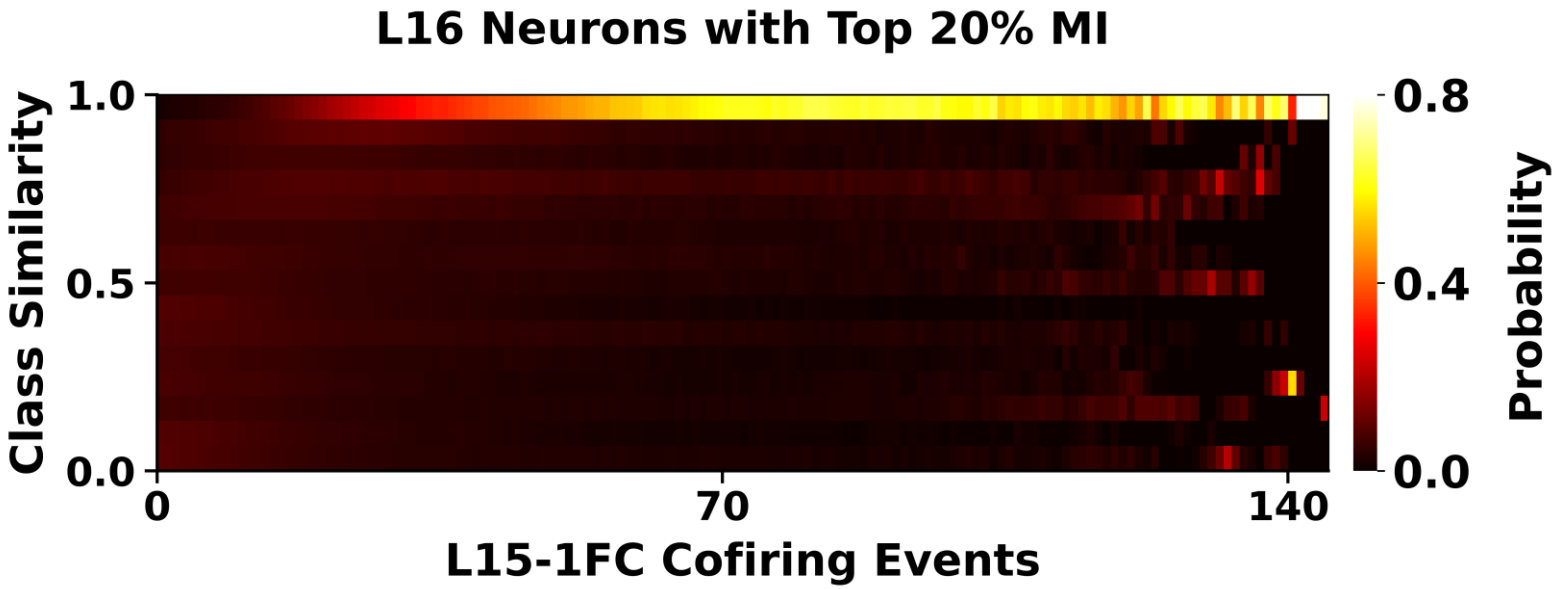

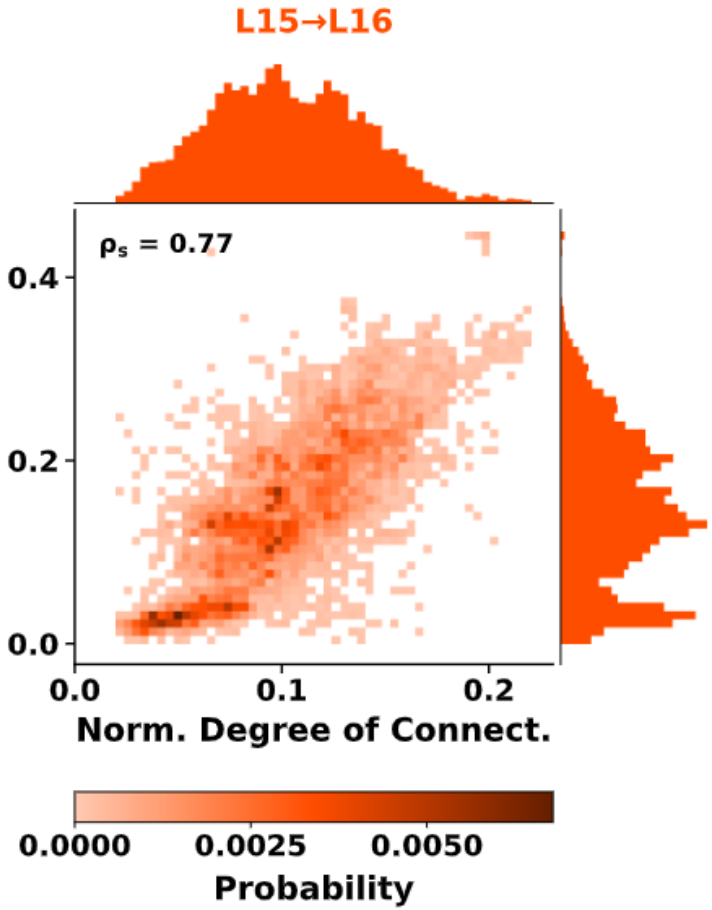

Top-20% MI L16 neurons · brighter = co-firing matches the neuron's preferred class

Causal evidence from ablation

- **1FC-ranked ablation** degrades accuracy faster than |w| or gradient saliency

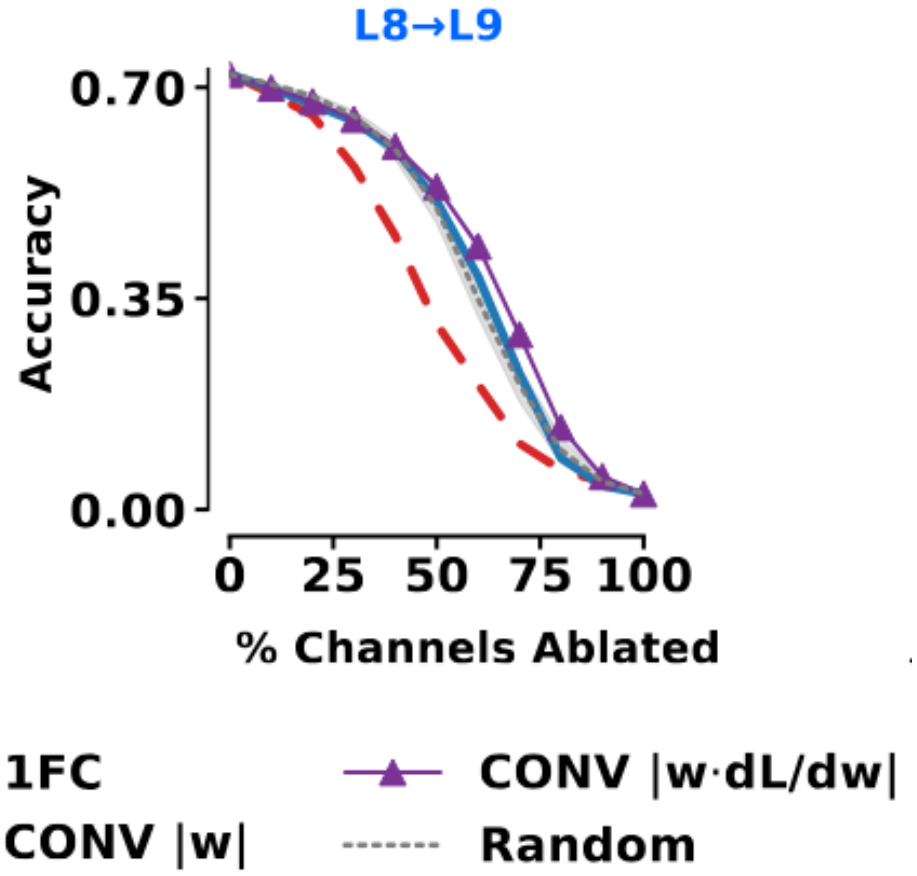

Noise disrupts responses; Adversarial does not

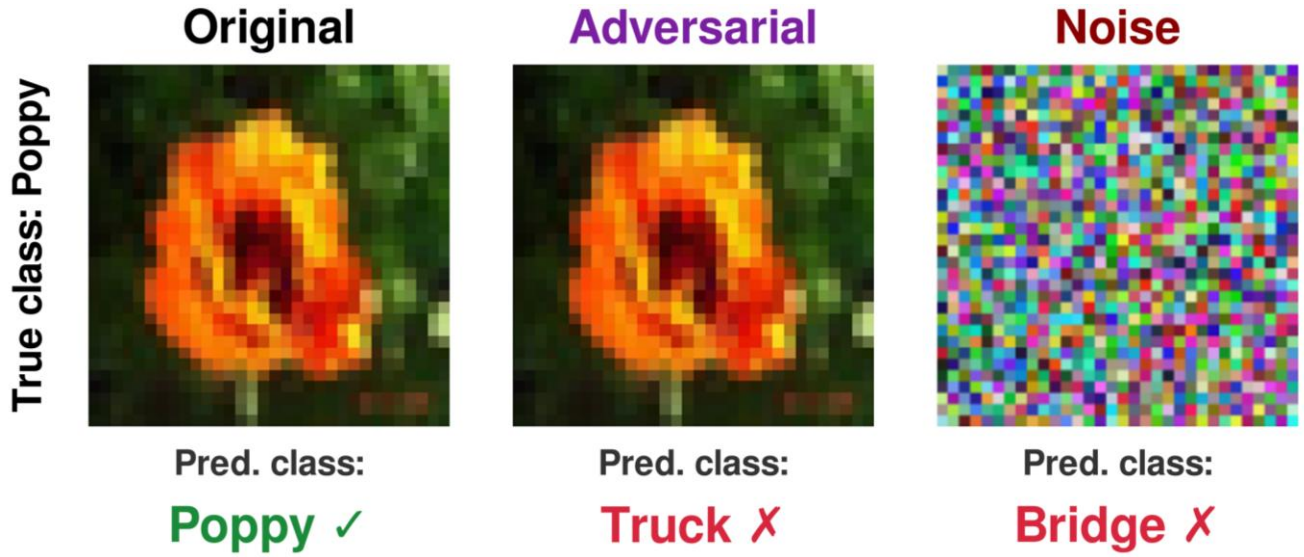

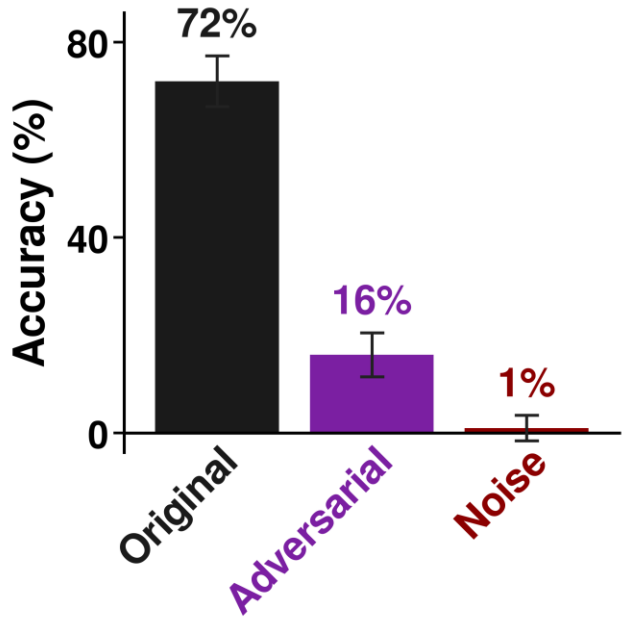

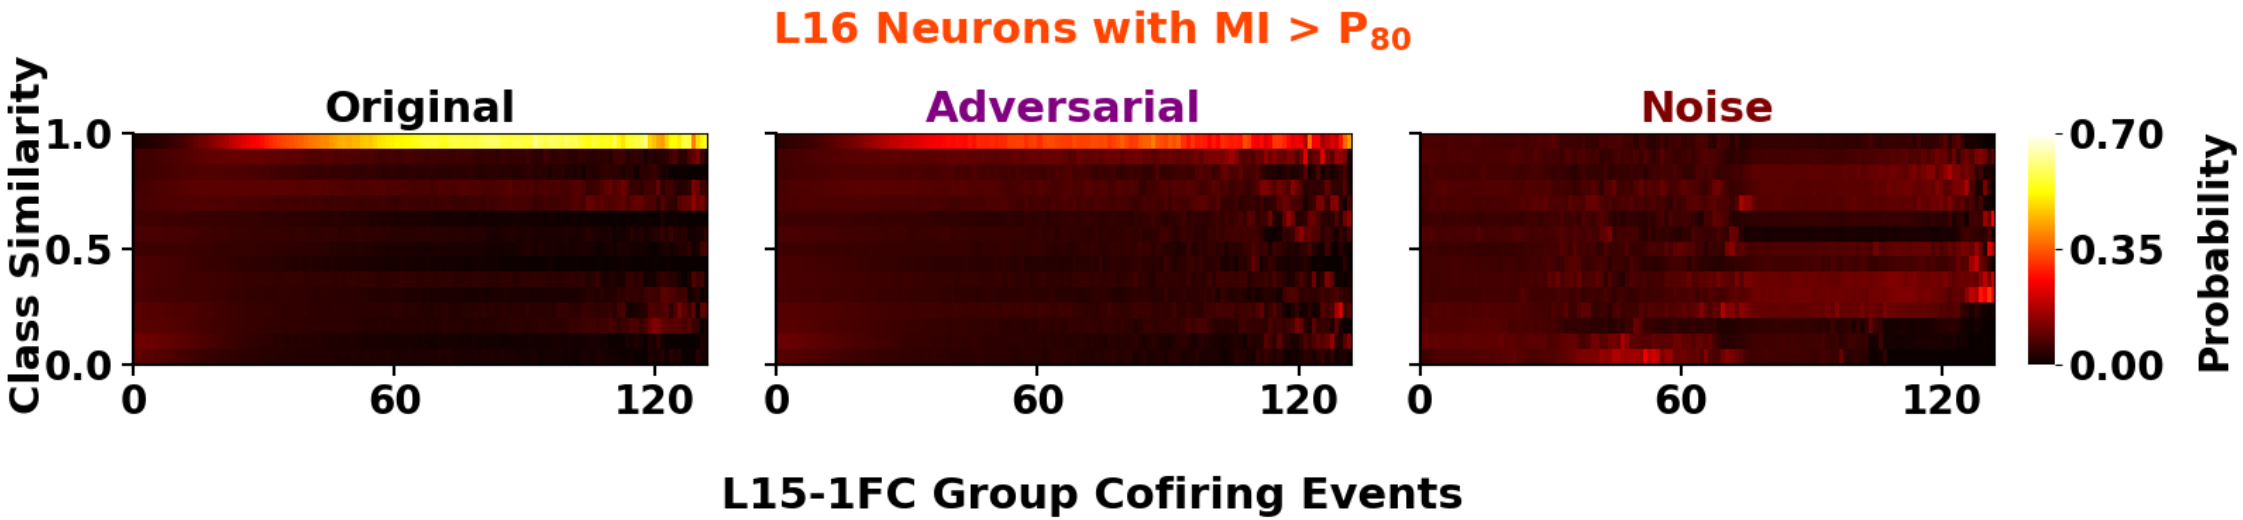

Conclusions

- Principles of cortical functional connectivity are preserved in deep spiking networks.
- Neurons respond to their input 1FC ensemble through a ReLU-like function, and its gain scales with ensemble size.
- Class information is concentrated in ***rare*** but ***highly coordinated co-firing*** events.
- This structure is ***shaped by learning***: uniform noise disrupts it, while it is largely preserved under adversarial input.
- These properties are largely conserved in other architectures and datasets: a spiking ResNet-14 on DVS Gesture and in an ANN ResNet-18.

Check out our paper!

