Hardware-aware spatially-embedded networks generalize better and faster.

Gabriel Béna*, Jimmy Weber*, Yassine Taoudi-Benchekroun*, Dan F. M. Goodman†, Melika Payvand*
*Institute of Neuroinformatics, UZH / ETH Zurich, Switzerland
†Department of Electrical and Electronic Engineering, Imperial College London, United Kingdom

A central challenge for any intelligent system is *compositional generalization*: recombining a finite set of learnt primitives into a much larger space of novel structures. Despite years of progress, standard neural networks still fail to generalize systematically on standard compositional benchmarks [1, 2], and the strongest current solutions largely rely on heavily engineered, hand-crafted architectures [3, 4]. Biological brains, by contrast, solve such problems routinely while operating under severe spatial, metabolic and communication constraints. A growing body of work now reframes these very constraints—long regarded as mere limitations—as *performance-driving inductive biases* that promote modularity, efficiency, and structured computation [5, 6, 7]. Inspired by this reframing, we systematically study how spatial priors can serve as useful inductive biases for compositional learning, proposing (i) controlled benchmarks whose axes of compositional difficulty can be independently tuned; and (ii) modelling frameworks that allow spatial priors to be injected without sacrificing fair comparisons against unconstrained baselines.
Specifically, for (i), we introduce the **2D-PVR meta-dataset**, a 2D grid-navigation reformulation of Pointer Value Retrieval (PVR) [8] that disentangles *combinatorial width* (size of the atomic function library) from *sequential depth* (length of the reasoning chain), while keeping local input statistics invariant. For (ii), we develop **Spatial DEEP-R**, a hardware–software co-design extension of the Deep Rewiring algorithm [9, 10] that trains sparse recurrent networks under strict connectivity budgets. By embedding neurons in different metric spaces and penalizing connections by distance, the networks are driven toward locally clustered, energy-efficient topologies, while non-spatial baselines retain strict identical computational power.

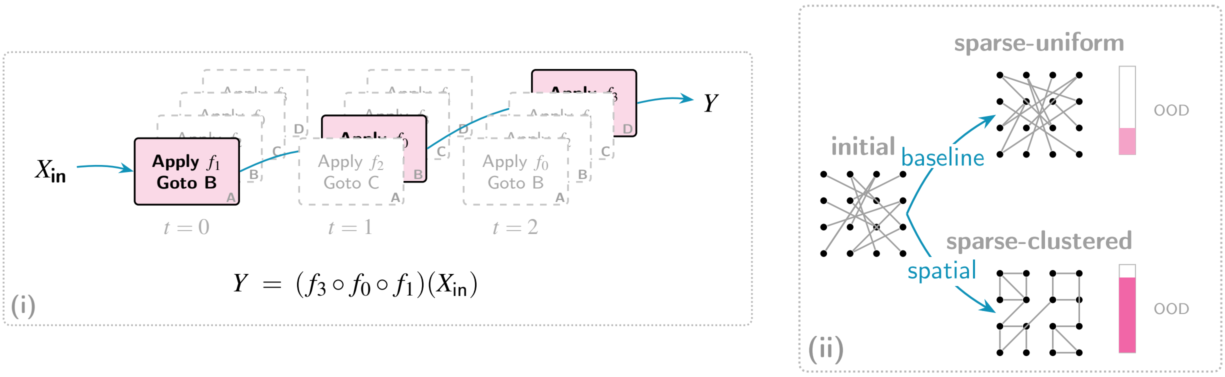

Figure 1: **Random vs spatial network architectures (ii) on compositional benchmark (i)**

Using a rigorous statistical protocol, we find that spatially embedded networks generalize both *better* and *earlier* than their non-spatial counterparts on out-of-distribution data (OOD). The advantage is most pronounced in tasks with *high-width*, i.e., with a large library of primitives, matching the intuition than modular structures are well-equipped to tackle highly compositional problems. Strikingly, a non-Euclidean connectivity prior inspired by the neuromorphic memristor-based *Mosaic* architecture [10, 11] also significantly outperforms unconstrained baselines, showing that well-designed hardware-friendly connectivity can itself become an *active driver* of compositional performance, not just a constraint to be tolerated for physical realizability. Finally, by causally lesioning structural and functional neuron clusters, we uncover a *structure–function gap*: while spatial constraints reliably induce modular physical topologies, the underlying computation seems to remain distributed, meaning generalization could be driven through other organizational properties than clean isolation of functional sub-routines. Ultimately, by bridging neuromorphic principles with systematic benchmarking, this framework opens new avenues for understanding compositional intelligence in artificial systems.

[1] B. Lake & M. Baroni, arXiv:1711.00350, 2018.
[2] N. Kim & T. Linzen, EMNLP, 9087–9105, 2020.
[3] J. Andreas et al., arXiv:1511.02799, 2017.
[4] X. Chen et al., arXiv:2008.06662, 2020.
[5] J. Achterberg et al., Nat. Mach. Intell., 5(12), 1369–1381, 2023.
[6] M. Khona et al., Neural Comput., 35(11), 1850–1869, 2023.
[7] X.-J. Zhang et al., PNAS Nexus, 4(1), pgae580, 2025.
[8] C. Zhang et al., arXiv:2107.12580, 2022.
[9] G. Bellec et al., arXiv:1711.05136, 2018.
[10] J. Weber et al, IEEE, 2024
[11] T. Dalgaty et al., Nat. Commun., 15, 142, 2024.