Robustness in sparse artificial neural networks trained with adaptive topology

Bendegúz Sulyok[1,2], Gergely Palla[1,2], Filippo Radicchi[3] and Santo Fortunato[3]

[1]Dept. of Biological Physics, Eötvös Loránd University, Budapest, Hungary
[2]Semmelweis University, Health Services Management Training Centre, Budapest, Hungary
[3]Center for Complex Networks and Systems Research Luddy School of Informatics, Computing, and Engineering Indiana University Bloomington, USA

INTRODUCTION & MOTIVATION

The Efficiency Crisis: Modern Deep Neural Networks (DNNs) rely heavily on massive over-parameterization, creating immense computational, memory and latency footprints.

Partial Solutions: Post-training compression techniques (pruning, quantization, distillation) successfully mitigate inference cost, but cannot lower the overhead at training time.

The solution: We evaluate Dynamic Sparse Training (DST), maintaining 99% sparsity during the entire training cycle.

ARCHITECTURE & DATASETS

General Model Structure: Any regular DNN can be augmented by swapping out Fully Connected (FC) layers for Sparse Linear Layers (SSLs).

Standard visual configuration: We investigated training on three image classification datasets – MNIST, FashionMNIST and EMNIST – using 3 SSLs (1000 hidden neurons each), followed by a single FC readout layer.

Complex convolutional head configuration: We evaluated training on the image classification dataset CIFAR10 by routing the data through a VGG16-BN visual head into 3 SSLs (1000 hidden neuron each) and a FC readout layer.

DYNAMIC OPTIMIZATION

Weight adaptation: Backpropagation optimizes the weights of the links via gradient descent.

Topology Update: The usual training is periodically interrupted and the structure of the SLLs is changed:

a) **Pruning:** A fraction of the weakest links (e.g. zeta=30%) are removed based on the weight magnitude.

b) **Chain Removal:** During an optional second pruning sweep dead ends are removed . This means any link disconnected from the input or output layers.

c) **Link Regrowth:** A number of connections are selected and reactivated in order to preserve the initial sparsity level..

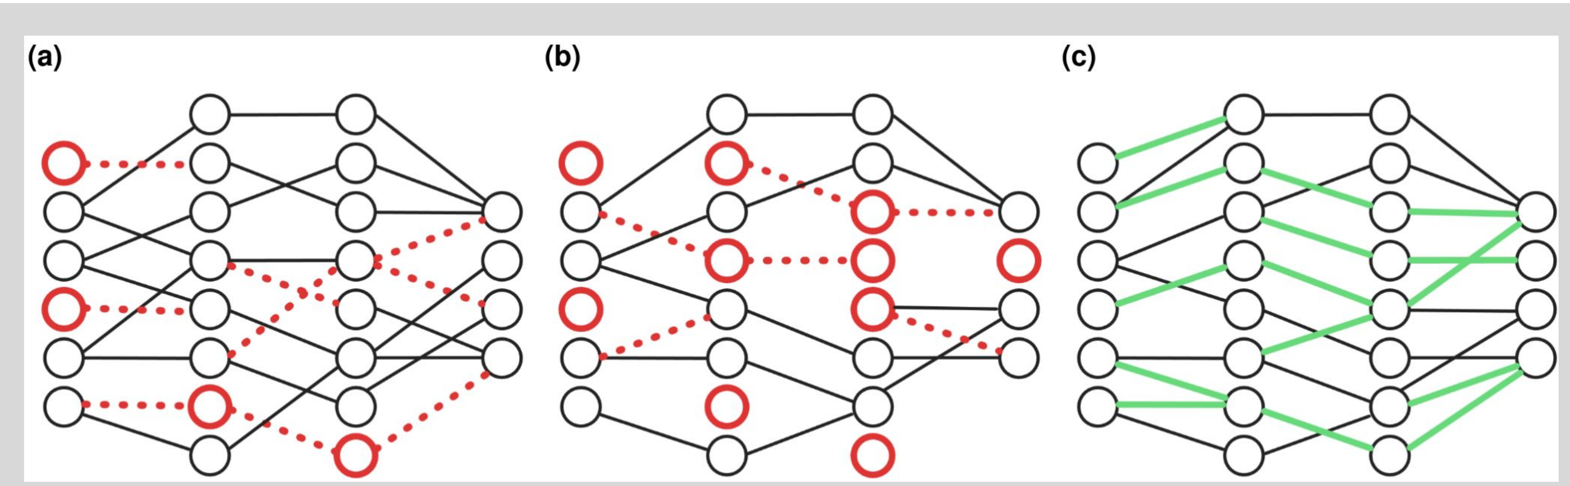

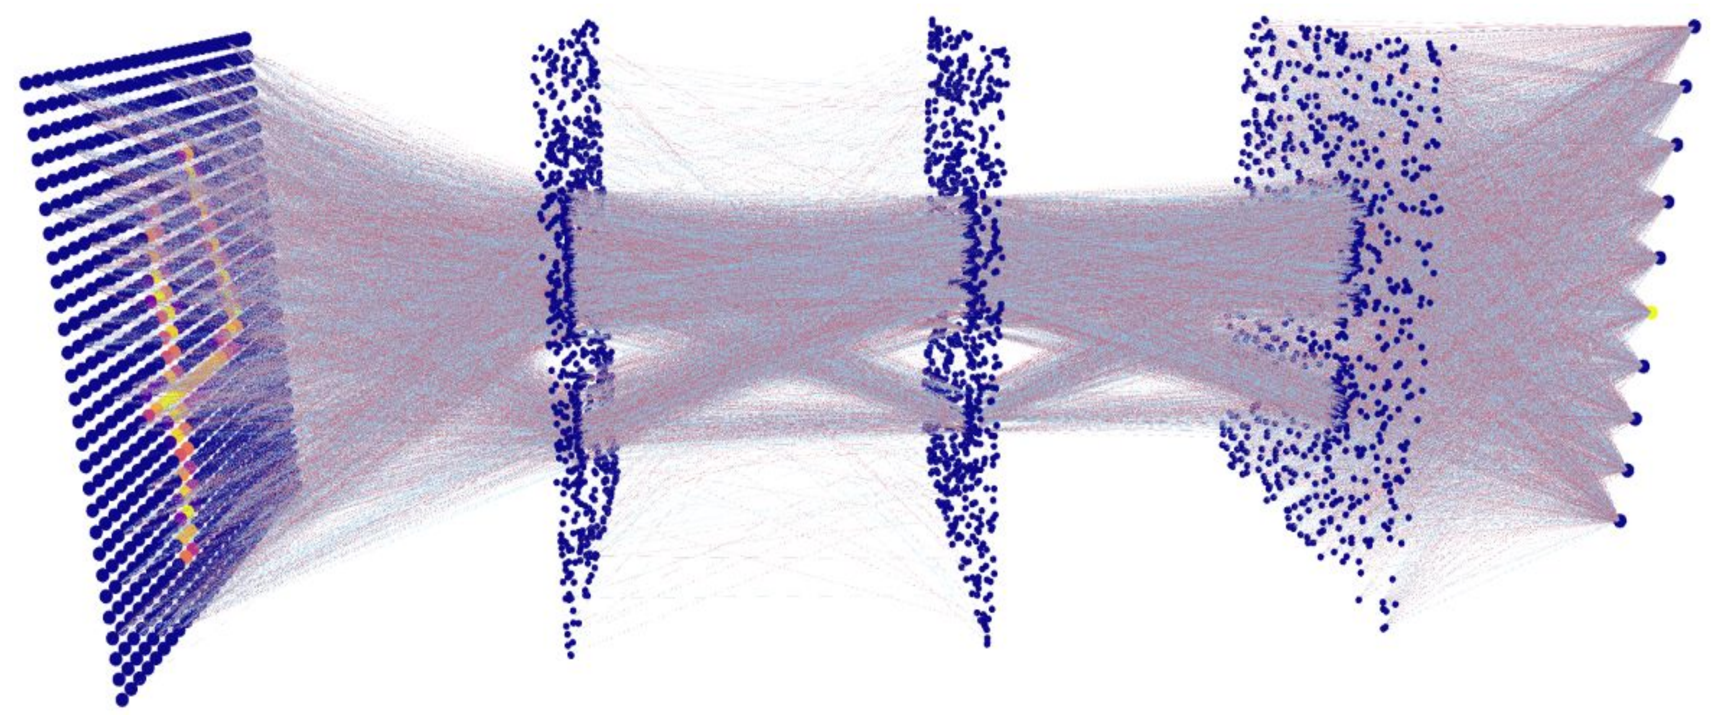

REGROWTH METHODS

Random Link Regrowth[1] (RLR): New links are introduced uniformly at random. This method does not benefit from Chain Removal.

Epitopological Growth via[2] (CH3L3): A network automaton heuristic for link prediction. CH3L3 evaluates local structural clustering and preferentially grows missing links within dense communities.

Chain Removal: An optional second pruning sweep executed before growing new links; used to remove dead ends and unused connections. Particularly useful for CH3L3.

TRAINING CONVERGENCE

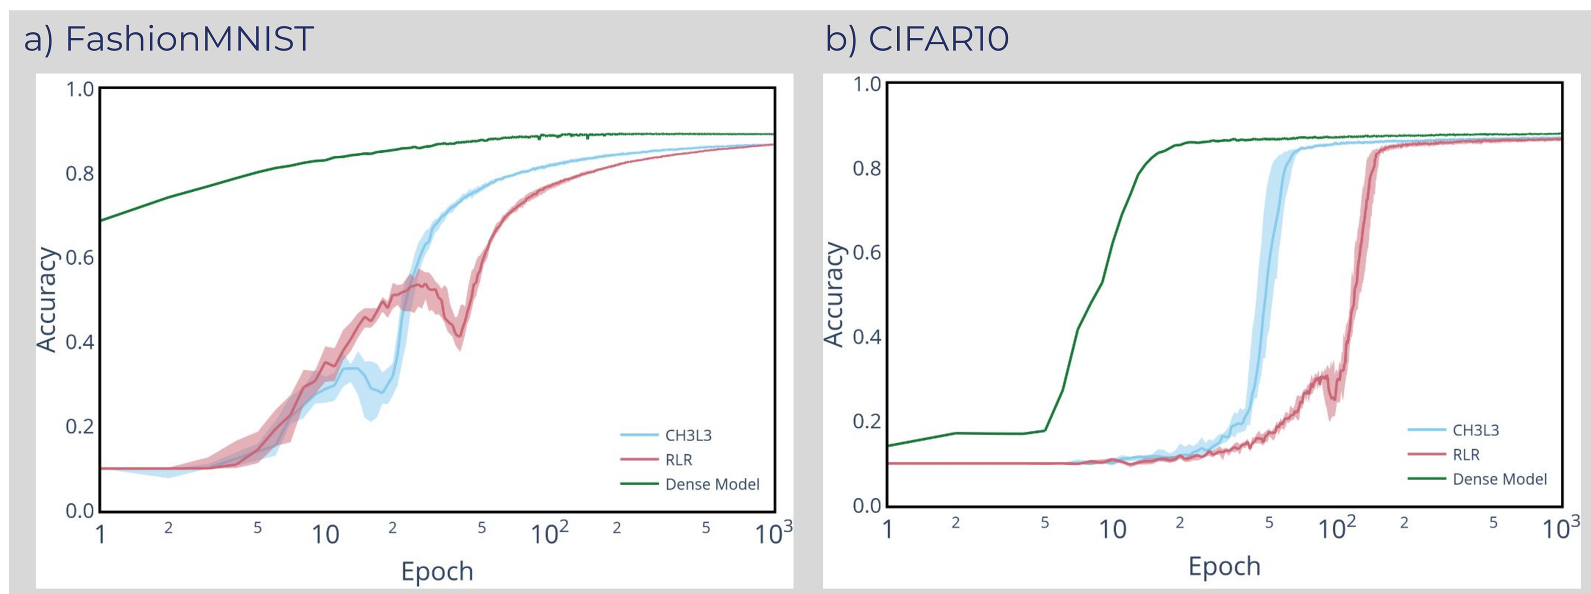

Dense Model: FC networks of the same shape were trained to serve as baseline for the structural analysis.

Final Accuracy: Across all tested image datasets both methods eventually converge to the same classification accuracy as the FC network of the same shape.

Convergence Speed: The epitopological framework (CH3L3) exhibits vastly superior early-stage learning velocity over the simpler RLR method.

PERTURBATION ANALYSIS

Trained models were subjected to varied structural and weight-based perturbations – without any fine-tuning – in order to gauge performance degradation.

- **Weight Order Pruning:** Dropping the strongest links first results in catastrophic loss of performance across all configurations.
- **Random Pruning:** All models show a slow and steady decline when removing links uniformly at random. Models trained with CH3L3 are slightly more robust.
- **Reverse Weight Order Pruning:** The RLR configurations show significantly higher resilience against removing the weak links first, suggesting these models still have high levels of redundancy.
- **Weight Shuffling:** During this perturbation type we shuffled weights in bins. Models trained with CH3L3 tend to be more robust.
- **Weight Modification:** The RLR configurations are more robust against adding random Gaussian noise to the weight values.

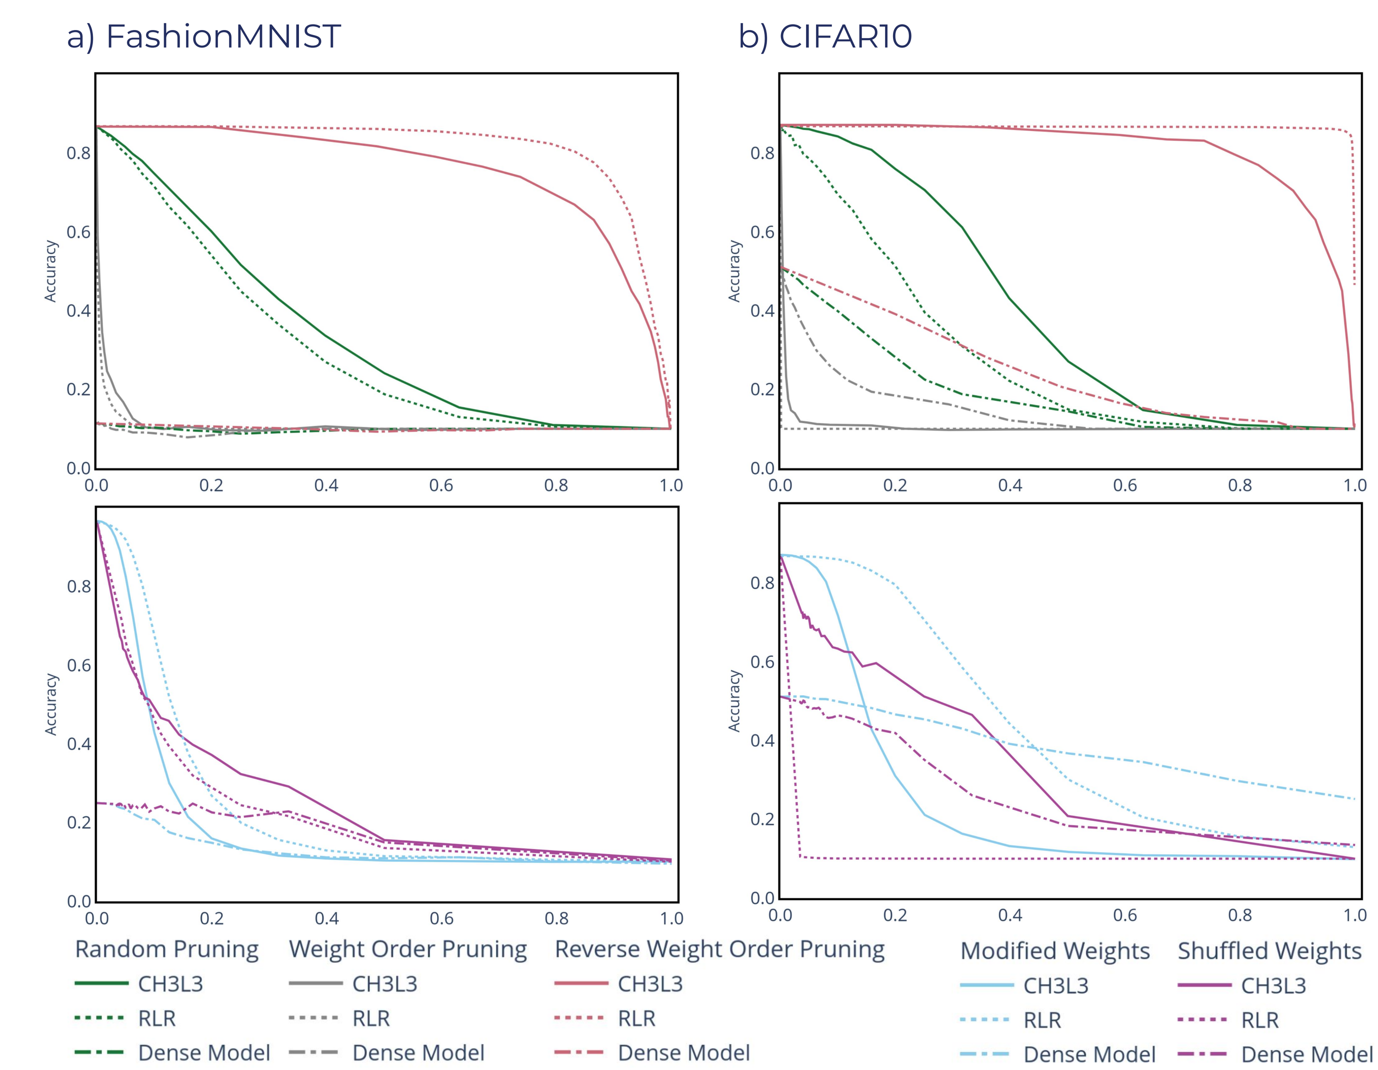

CONCLUSION

Success in Training: Model optimization remains highly effective even under the strict 99% sparsity constraint.

Network Science Principles: CH3L3 compresses the artificial neural pathways into tightly packed, highly effective local clusters, accelerating convergence and these models maintain strong structural robustness against certain perturbations.

Structural Trade-off: While showing slower convergence, RLR naturally spreads out links more evenly across the network, resulting in higher number of active neurons and possibly further post-training compression.

[1] Mocanu, D. C., Mocanu, E., Stone, P., Nguyen, P. H., Gibescu, M., & Liotta, A. (2018). Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. Nature communications, 9(1), 2383.
[2] Zhang, Y., Zhao, J., Wu, W., Muscoloni, A., & Cannistraci, C. V. (2024, May). Epitopological learning and cannistraci-hebb network shape intelligence brain-inspired theory for ultra-sparse advantage in deep learning. In The Twelfth International Conference on Learning Representations.

We acknowledge the support of the AccelNet-MultiNet program, a project of the National Science Foundation (award #1927425 and #1927418)