

A DENDRITIC-INSPIRED NETWORK SCIENCE GENERATIVE MODEL FOR TOPOLOGICAL INITIALIZATION OF CONNECTIVITY IN SPARSE ARTIFICIAL NEURAL NETWORKS

008 **Anonymous authors**

009 Paper under double-blind review

ABSTRACT

014 Artificial neural networks (ANNs) achieve remarkable performance but at the un-
 015 unsustainable cost of extreme parameter density. In contrast, biological networks
 016 operate with ultra-sparse, highly organized structures, where dendrites play a cen-
 017 tral role in shaping information integration. Here we introduce the Dendritic Net-
 018 work Model (DNM), a generative framework that bridges this gap by embedding
 019 dendritic-inspired connectivity principles into sparse artificial networks. Unlike
 020 conventional random initialization, DNM defines connectivity through parametric
 021 distributions of dendrites, receptive fields, and synapses, enabling precise control
 022 of modularity, hierarchy, and degree heterogeneity. This parametric flexibility al-
 023 lows DNM to generate a wide spectrum of network topologies, from clustered
 024 modular architectures to scale-free hierarchies, whose geometry can be charac-
 025 terized and optimized with network-science metrics. Across image classifica-
 026 tion benchmarks (MNIST, Fashion-MNIST, EMNIST, CIFAR-10), DNM consis-
 027 tently outperforms classical sparse initializations at extreme sparsity (99%), in
 028 both static and dynamic sparse training regimes. Moreover, when integrated into
 029 state-of-the-art dynamic sparse training frameworks and applied to Transformer
 030 architectures for machine translation, DNM enhances accuracy while preserving
 031 efficiency. By aligning neural network initialization with dendritic design prin-
 032 ciples, DNM demonstrates that sparse bio-inspired network science modelling is
 033 a structural advantage in deep learning, offering a principled initialization frame-
 034 work to train scalable and energy-efficient machine intelligence.

1 INTRODUCTION

037 Artificial neural networks (ANNs) have demonstrated remarkable potential in various fields; how-
 038 ever, their size, often comprising billions of parameters, poses challenges for both economic viability
 039 and environmental sustainability. Biological neural networks, in contrast, can efficiently process in-
 040 formation using ultra-sparse structures (Drachman, 2005; Walsh, 2013). This efficiency arises from
 041 the brain’s highly structured and evolutionarily optimized network topology. A central component
 042 of this architecture is the dendritic organization, the primary receptive surface of the neuron (Cuntz
 043 et al., 2010). Conventional ANNs omit a crucial component of the brain’s efficiency since they
 044 traditionally depict neurons as simple point-like integrators, mainly ignoring the computing power
 045 inherent in the intricate structure of dendrites.

046 Research has revealed that dendrites are not passive conductors but active computational units ca-
 047 pable of performing sophisticated, nonlinear operations and integration (London & Häusser, 2005;
 048 Larkum, 2013). This insight has motivated theoretical frameworks that model a single neuron as a
 049 multi-layer network, where dendritic branches act as nonlinear subunits that feed a final integrator
 050 at the cell body (Lauditi et al., 2025). As clarified by recent works on dendritic artificial neural
 051 networks (Li et al., 2020; Chavlis & Poirazi, 2025), the dendritic tree’s ability to sample and non-
 052 linearly integrate restricted parts of the input space can be used in neuromorphic physical networks
 053 (Li et al., 2020) or in artificial convolutional layers (Chavlis & Poirazi, 2025). In this paradigm, dis-
 tinct dendritic branches process specific, localized receptive fields without sharing weights, allow-

054 ing for precise, location-specific feature integration. Additional efforts to translate these principles
 055 into neuromorphic systems have confirmed that also dendritic morphology has a significant impact
 056 on spatio-temporal processing and performance(Baek et al., 2024; Jones & Kording, 2021; 2022).
 057 These approaches, however, are often limited to fixed structures that mimic the computational non-
 058 linearity or direct morphology of biological neurons, overlooking the broader rules of connectivity,
 059 which also include non-uniform random network dynamic rewiring by synaptic turnover (Zhang
 060 et al., 2025; Frank et al., 2018).

061 Our understanding of these topological constraints has been revolutionized by the advent of large-
 062 scale functional connectomics (mic, 2025). These studies reveal that biological neural networks
 063 are not uniformly randomly wired; rather, they exhibit specific, non-uniform random connectivity
 064 patterns characterized by "like-to-like" wiring rules and distinct structural motifs across cortical
 065 layers. Translating these high-level connectomic principles, such as modularity, hierarchy, and non-
 066 random receptive field organization, into scalable, topological generative frameworks for designing
 067 the sparse structure of artificial networks remains an open challenge.

068 To address the gap for a flexible, principled framework for generating and testing dendritic topolo-
 069 gies, we introduce a dendritic-inspired network science generative model for sparse topology design
 070 of bipartite layers in deep neural networks: the Dendritic Network Model (DNM). The novelty of
 071 DNM lies in the elaborated mechanism to model the non-uniform topological organization of the
 072 receptive fields on the input layers. The model's parametric approach enables the systematic explo-
 073 ration of the relationship between network structure and computational function. The DNM provides
 074 a principled method for generating sparse network initializations that can be integrated into mod-
 075 ern deep learning frameworks. We demonstrate that this approach can improve performance over
 076 standard sparse initialization techniques and offers a powerful platform for exploring how structural
 077 constraints, inspired by biology, can lead to more efficient and capable artificial neural networks.

078 Our approach can be contrasted with other dendritic-inspired methodologies in the field. For in-
 079 stance, Li et al. (2020) experimentally demonstrated a fully integrated hardware network using
 080 memristor devices, where artificial dendrites provided non-linear integration and filtering to achieve
 081 highly efficient physical networks. Subsequently, Malakasis et al. (2023) utilized bio-realistic
 082 spiking neural networks to show how active dendrites combined with uniformly random synaptic
 083 turnover can optimize learning in binary classification scenarios. These works pave the way for re-
 084 cent advancements like the work of Chavlis & Poirazi (2025), which presents a dendritic-emulating
 085 model that reproduces the nonlinear integrative functions of dendrites. In their framework, a dendrite
 086 is mapped to a node within a tree-like subnetwork, creating a nonlinear computational component
 087 for larger networks. Distinct from these bio-mimetic approaches, the work of Kepner & Robinett
 088 (2019) is rooted in algebraic graph theory. They generate deterministic topologies (RadiX-Nets) us-
 089 ing mixed-radix numeral systems and Kronecker products of adjacency submatrices. Their primary
 090 design goal is to ensure graph-theoretic properties such as constant expansion, path-connectedness,
 091 and symmetry. In contrast, DNM is a generative framework inspired by biological morphology.
 092 Rather than relying on algebraic products, DNM constructs connectivity via parametric distribu-
 093 tions (spatial/non-spatial) of specific biological components: dendritic branches, localized receptive
 094 fields, and synapses. Our focus is on modeling a dendritic-inspired sparse network topology that
 095 allows for the emergence of modular and hierarchical structures rather than enforcing deterministic
 096 symmetry (Figure 1). Thus, while (Li et al., 2020; Chavlis & Poirazi, 2025) focus on the emulation
 097 of nonlinear processing in tree-like dendritic subnetwork structures, our work investigates topolog-
 098 ical dendritic-inspired principles that allows modelling the initial sparse connectivity organization
 099 from a network science perspective.

100 In this article, we describe the Dendritic Network Model in detail and analyse its topology and
 101 geometric characterization. We evaluate its effectiveness with extensive experiments across multiple
 102 architectures and tasks. To assess its basic functionality, we use it to initialize several static and
 103 dynamic sparse training (DST) methods on MLPs for image classification on the MNIST (LeCun
 104 et al., 2002), EMNIST (Cohen et al., 2017), Fashion MNIST (Xiao et al., 2017), and CIFAR-10
 105 (Krizhevsky, 2009) datasets. The results show that DNM clearly outperforms other sparse ini-
 106 tialization methods over all training models tested at 99% sparsity. Next, we extend the tests on
 107 Transformers (Vaswani et al., 2017a) for Machine Translation on the Multi30k en-de (Elliott et al.,
 108 2016), IWSLT14 en-de (Cettolo et al., 2014), and WMT17 en-de (Bojar et al., 2017) benchmarks.
 109 On this architecture, DNM outperforms all topological initialization methods at high sparsity levels.
 110 These findings underscore the potential of DNM in enabling highly efficient and effective network

108 initialization for large-scale sparse neural network training. By analyzing the best-performing DNM
 109 topologies, we can also gain insights into the relationship between network geometry, data structure,
 110 and model performance.
 111

112 2 RELATED WORKS

113 2.1 SPARSE TOPOLOGICAL INITIALIZATION METHODS

116 Dynamic sparse training (DST) trains a neural network with a sparse topology that evolves through-
 117 out the learning process. The initial arrangement of the connections is a critical aspect of this frame-
 118 work. This starting structure determines the initial pathways for information flow and acts as the
 119 foundational scaffold upon which the network learns and evolves. A well-designed initial topology
 120 can significantly improve a model’s final performance and training efficiency, whereas a poor start-
 121 ing point can severely hinder its ability to learn effectively. The principal topological initialization
 122 approaches for dynamic sparse training are grounded in network science theory, where three ba-
 123 sic generative models for monopartite sparse artificial complex networks are the Erdős-Rényi (ER)
 124 model (Erdős & Rényi, 1960), the Watts-Strogatz (WS) model (Watts & Strogatz, 1998), and the
 125 Barabási-Albert (BA) model (Barabási & Albert, 1999). Since the standard WS and BA models are
 126 not directly designed for bipartite networks, they were recently extended into their bipartite coun-
 127 terparts and termed as Bipartite Small-World (BSW) and Bipartite Scale-Free (BSF) (Zhang et al.,
 128 2024b), respectively. BSW generally outperforms BSF for dynamic sparse training (Zhang et al.,
 129 2024b). The Correlated Sparse Topological Initialization (CSTI) (Zhang et al., 2024a) is a feature-
 130 informed topological initialization method that considers the links with the strongest Pearson cor-
 131 relations between nodes and features in the input layer. SNIP (Lee et al., 2018) is a data-informed
 132 pruning method that identifies important connections based on their saliency scores, calculated us-
 133 ing the gradients of the loss function with respect to the weights. Ramanujan graphs (Kalra et al.)
 134 are a class of sparse graphs that exhibit optimal spectral properties, making them suitable for ini-
 135 tializing neural networks with desirable connectivity patterns. The Bipartite Receptive Field (BRF)
 136 network model (Zhang et al., 2025) generates networks with brain-like receptive field connectivity.
 137 This is the first attempt to mimic the structure of brain connections in a sparse network initialization
 138 model. Radix-Nets (Kepner & Robinett, 2019) offer a deterministic approach to "de novo" spar-
 139 sity, utilizing mixed-radix numeral systems and the Kronecker product to construct topologies that
 140 ensure path-connectedness and symmetry while facilitating asymptotic sparsity. Finally, dendritic
 141 Artificial Neural Networks (dANNs) (Chavlis & Poirazi, 2025) introduce a bio-inspired architecture
 142 that mimics the structured connectivity and restricted input sampling of biological dendrites (e.g.,
 143 using Local Receptive Fields). Unlike traditional approaches that strive for class specificity, this
 144 architecture fosters mixed-selective neuronal responses.

145 While the methods discussed above primarily focus on initializing layered or bipartite structures,
 146 a parallel line of research investigates Artificial Neural Networks (ANNs) with general complex
 147 topologies, unconstrained by multipartite restrictions. Monteiro et al. (2016) demonstrated that hy-
 148 brid topologies combining scale-free and small-world properties, inspired by the *C. elegans* connec-
 149 tome, can significantly improve learning curves. Moving beyond manual architecture design, Xie
 150 et al. (2019) utilize random graph models (ER, BA, WS) to generate "randomly wired" networks
 151 that achieve competitive performance in image recognition. To facilitate the translation between
 152 arbitrary graph structures and neural models, Stier & Granitzer (2022) introduced the deepstruct
 153 framework. More recently, Boccato et al. (2024) provided a systematic comparison of these ar-
 154 chitectures, revealing that complex, non-layered topologies can outperform traditional Multilayer
 155 Perceptrons (MLPs) in high-difficulty tasks by potentially exploiting compositional sparsity.

156 3 THE DENDRITIC NETWORK MODEL

157 3.1 BIOLOGICAL INSPIRATION AND PRINCIPLES

158 The architecture of the Dendritic Network Model (DNM) is inspired by the structure of biological
 159 neurons. In the nervous system, neurons process information through complex, branching exten-
 160 sions called dendrites, which act as the primary receivers of synaptic signals. Inspired by this phe-
 161 nomenon, the DNM imposes a structured, dendrite-like organization on how output neurons connect

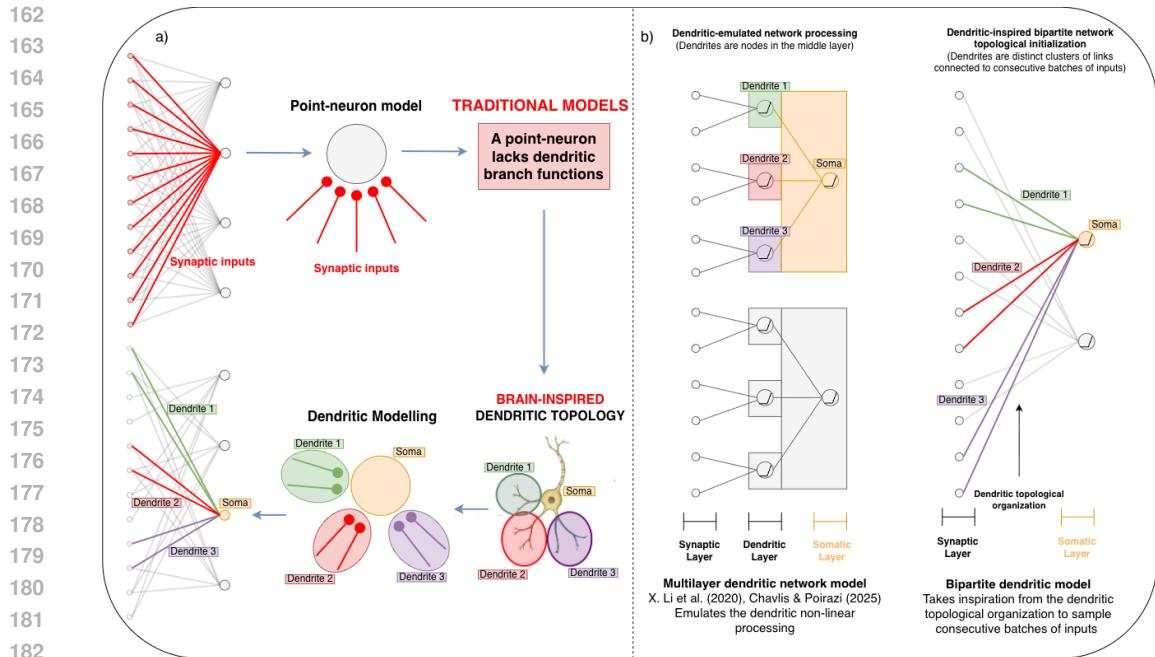


Figure 1: **Comparison of the Dendritic Network Model with traditional point-neurons and existing dendritic architectures.** (a) **From point-neurons to dendritic topology.** Traditional artificial neuron models (top) function as point-integrators, summing all synaptic inputs globally without spatial differentiation. In contrast, the DNM (bottom) introduces a brain-inspired topology where synaptic inputs are organized into distinct dendritic branches. This structure allows the output neuron to process inputs as clustered groups. (b) **Comparison with existing dendritic network models.** The left panel illustrates “Dendritic-emulated network processing” as seen in works like Li et al. (2020) and Chavlis & Poirazi (2025). In these architectures, dendrites are often modeled as explicit computational nodes forming an intermediate layer between inputs and the soma (a multilayer approach). The right panel illustrates the proposed DNM (a dendritic-inspired bipartite network topological initialization). Unlike previous dendritic-inspired models that have a tree-like multilayer structure, DNM embeds dendritic properties directly into the bipartite network topology. It treats dendrites as distinct clusters of links within a bipartite graph, connecting the soma to consecutive batches of inputs. This allows the network to inherit dendritic structural advantages through topological initialization rather than architectural expansion.

to the preceding layer’s inputs (Figure 1). In this work, we refer to the set of connections between two adjacent layers as a bipartite sandwich layer, a term already used in prior literature to denote the bipartite subnetwork of edges that lies between one layer of neurons and the next. We retain this terminology because it captures the specific object our model generates, the pattern of connections, rather than the neurons themselves. By contrast, the term hidden layer refers to the neurons in intermediate layers. Since the DNM defines how neurons connect, but does not generate or modify the neurons themselves, “hidden layer” would not accurately describe the structural entity under consideration. Within each sandwich layer, each output neuron forms multiple dendritic branches, where each branch connects the neuron to a contiguous block of input neurons. These blocks are separated by inactive spaces, segments of the input layer to which the neuron does not connect. All branches belonging to a given output neuron must lie within a predefined local receptive window, resulting in a structured, compartmentalized connectivity pattern. This design moves away from unstructured random sparsity and instead emulates the localized, clustered organization characteristic of biological dendrites.

3.2 THE DNM GENERATIVE ALGORITHM

To translate these biological principles into a computational structure, the DNM produces the sparse connectivity matrix of bipartite sandwich layers via a generative algorithm. The process builds connections iteratively for each output neuron j through the following steps: (1) **Degree determi-**

216 **nation:** first, determine the total degree for the output neuron based on a specific degree distribution
 217 strategy; (2) **Receptive field definition:** define the receptive field for the output neuron by topolog-
 218 ically mapping the output neuron’s position to a central point on the input layer and establishing a
 219 receptive window around this center; (3) **Dendritic allocation:** determine the number of dendritic
 220 branches used to connect the output neuron to the input layer; (4) **Dendritic Placement:** place
 221 evenly spread-out dendritic brances within the neuron’s receptive window defined in step 2; (5)
 222 **Synaptic distribution:** distribute the output neuron’s total degree across the dendrites. Appendix F
 223 describes the algorithm in depth.

224

225 3.3 PARAMETRIC SPECIFICATION

226

227 By parametrizing the core biological features of DNM, like the number of dendrites for each output
 228 neuron, the size of the receptive windows, the distribution of synapses across dendrites, and the
 229 degree distribution across output neurons, the DNM provides a flexible framework for generating
 230 network topologies that are sparse, structured, and biologically plausible. Appendix C shows how
 231 the connectivity of an MLP is shaped by the DNM. To apply biological spatial principles to non-
 232 spatial MLPs, we index neurons $i \in \{1, \dots, N\}$, and their physical location x_i is defined linearly
 233 such that the distance between adjacent indices is minimized. This allows us to define "spatial"
 234 distributions where connectivity probabilities depend on the relative distance $|x_i - x_j|$ between
 235 neurons in adjacent layers.

236 **Sparsity (s)** The sparsity parameter (s) defines the percentage of potential connections between
 237 the input and output layers that are absent, controlling the trade-off between computational cost and
 238 representational power.

239
 240 **Dendritic Distribution** The dendritic distribution governs the number of branches that connect
 241 each output neuron to the input layer, which can be seen as the number of distinct input regions a
 242 neuron integrates information from. The central parameter for this is M , which defines the mean
 243 number of dendrites per output neuron. This distribution can be implemented in one of three ways.
 244 The simplest is a fixed distribution, where every output neuron has exactly M dendrites. Alterna-
 245 tively, a non-spatial distribution introduces stochasticity by sampling the number of dendrites for
 246 each output neuron from a probability distribution (e.g., Gaussian or uniform) with a mean of M .
 247 Finally, a spatial distribution makes the number of dendrites for each neuron dependent on its po-
 248 sition within the layer. Using a Gaussian or inverted Gaussian profile, this configuration implies
 249 that some neurons integrate signals from many distinct input regions (a high dendrite count), while
 250 others connect to fewer, more focused regions (a low dendrite count).

251
 252 **Receptive Field Width Distribution** The receptive field of an output neuron j is defined as the
 253 contiguous subset of input neurons to which j *potentially* connects. We define a mapping func-
 254 tion $\phi(j)$ that projects the index of output neuron j to a center coordinate on the input layer. The
 255 receptive field is then the interval $[\phi(j) - \frac{W_j}{2}, \phi(j) + \frac{W_j}{2}]$, where W_j is the receptive field width
 256 determined by the parameter α . The receptive field mirrors the concept of receptive fields in biology.
 257 This process is governed by a mean parameter, α , which specifies the average percentage of consec-
 258 utive input neurons on the input layer from which an output neuron can sample connections. This
 259 distribution can be configured in several ways: a fixed distribution assigns an identical window size
 260 α to all output neurons; a non-spatial distribution introduces variability by drawing each neuron’s
 261 window size from a probability distribution (e.g., Gaussian or uniform) centered on α ; and a spatial
 262 distribution links the window size to the neuron’s position in its layer, allowing for configurations
 263 where receptive windows are, for instance, wider at the center and narrower at the edges.

264
 265 **Degree Distribution** The degree distribution samples the number of incoming connections for
 266 each output neuron. This can be configured using a fixed distribution, where every output neuron
 267 is allocated the same degree. To introduce heterogeneity, a non-spatial distribution can be used to
 268 sample the degree for each neuron from a probability distribution. Finally, a spatial distribution
 269 allows the degree to vary based on the neuron’s position, for instance, by creating highly connected,
 hub-like neurons at the center of the layer. The mean degree is set by the layer size and target
 sparsity.

270 **Synaptic Distribution** Once an output neuron’s total degree is determined, the synaptic distribution
 271 allocates these connections among its various dendritic branches. The allocation can be fixed,
 272 where each dendrite receives an equal number of synapses. Alternatively, a non-spatial distribution
 273 can introduce random variability in synapse counts per dendrite. A spatial distribution can also be
 274 applied, making the number of synapses dependent on a dendrite’s topological location, for example
 275 by assigning more connections to central branches versus outer ones. This distribution has a mean of
 276 $\frac{N_{in} \cdot (1-s)}{M}$, where N_{in} is the size of the input layer. While the degree distribution determines the total
 277 connectivity k_j of an output neuron j , the synaptic distribution governs the partition of k_j across the
 278 neuron’s M_j dendritic branches. Formally, if $s_{j,b}$ is the number of synapses on the b -th dendritic
 279 branch of neuron j , the distribution ensures $\sum_{b=1}^{M_j} s_{j,b} = k_j$. This allocation can be uniform (fixed),
 280 stochastic (non-spatial), or topology-dependent (spatial), allowing specific branches, such as those
 281 in the center of the receptive field, to be more densely connected than distal branches.
 282

283 **Layer Border Wiring Pattern** The DNM includes a setting to control how connections are han-
 284 dled at the boundaries of the input layer. The default behavior is a *wrap-around* topology, where the
 285 input layer is treated as a ring. This means a receptive window for a neuron near one edge can wrap
 286 around to connect to neurons on the opposite edge, ensuring all neurons have a similarly structured
 287 receptive field. Alternatively, a *bounded* pattern can be enforced. In this mode, receptive windows
 288 are strictly confined within the layer’s physical boundaries. If a receptive field extends beyond the
 289 first or last input neuron, it is clamped to the edge. This enforces a more stringent locality, which
 290 we analyze further in Appendix I.
 291

292 3.4 NETWORK TOPOLOGY AND GEOMETRIC CHARACTERIZATION

293 A central hypothesis of this work is that specific topological features, such as modularity and hi-
 294 erarchy, confer distinct inductive biases that facilitate learning in ANNs. To test this hypothesis,
 295 it is essential to demonstrate that the DNM is not limited to a single structural configuration but
 296 rather functions as a flexible generative framework capable of accessing a diverse landscape of
 297 topologies. In this section, we systematically vary the hyperparameters defined in Section 3.3 to
 298 characterize this landscape. Our goal is to show that by tuning the DNM’s parameters, we can
 299 controllably transition the network architecture across three distinct regimes: from unstructured uni-
 300 formly connected random graphs, to input-order-dependent highly modular networks, and finally to
 301 input-order-dependent hierarchical, scale-free (Barabási & Albert, 1999) topologies.¹
 302

303 Figure 2 illustrates this topological diversity by comparing a baseline random network with several
 304 DNM configurations in a 3-layered MLP with 90% sparsity. Each panel displays the network’s
 305 coalescent embedding (Cacciola et al., 2017) in hyperbolic space, its adjacency matrix, and network
 306 science metrics: characteristic path length (L), modularity (Q), structural consistency (σ_c), and the
 307 power-law exponent of the degree distribution (γ).
 308

309 The coalescent embedding maps nodes onto a 2D hyperbolic disk, where the nodes’ angular co-
 310 ordinates are computed via non-linear dimensionality reduction to cluster structurally similar nodes,
 311 while radial coordinates are derived from node popularity (degree). This visualization reveals **hier-
 312 archy** (nodes near the center act as hubs) and **modularity** (angular grouping indicates community
 313 structure). The full algorithmic details are provided in Appendix A.
 314

315 The baseline random network (Figure 2a) lacks structure ($Q = 0.14$, $\sigma_c = 0.04$). Figure 2b rep-
 316 presents a DNM network with $M = 3$, $\alpha = 1$, and all distributions fixed, yielding high modularity
 317 ($Q = 0.64$) and structural consistency ($\sigma_c = 0.76$). A key finding is that by setting a spatial Gaussian
 318 degree distribution (Figure 2c), the DNM generates a hierarchical topology that exhibits scale-free
 319 (Barabási & Albert, 1999) properties. Specifically, the resulting degree distribution follows a power
 320 law $P(k) \sim k^{-\gamma}$ with an exponent $\gamma = 2.30$. Since typical scale-free networks exhibit $2 < \gamma < 3$
 321 (Barabási & Albert, 1999), this confirms that DNM can synthesize architectures with hub-like char-
 322 acteristics and hierarchical organization purely through parametric initialization. Similar measures
 323 are found when setting a spatial Gaussian synaptic distribution (Figure 2d, $Q = 0.54$, $\sigma_c = 0.74$),
 324

325 ¹To facilitate an intuitive exploration of this landscape, we have developed an interactive web application
 326 where readers can adjust the model’s parameters and visualize the resulting network structures. The application
 327 is available at: <https://dendritic-network-model.streamlit.app/>

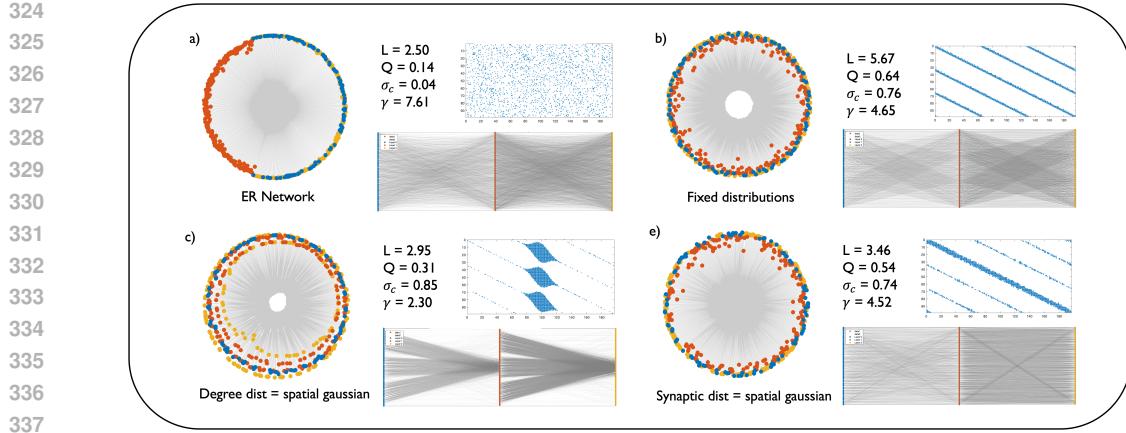


Figure 2: **Geometric and topological characterization of the Dendritic Network Model.** The figure compares a baseline random network (a) with various DNM configurations (b-d) for a 3-layered MLP of size $98 \times 196 \times 196$ with 90% sparsity. Each panel shows a coalescent embedding in hyperbolic space (left), the first layer’s adjacency matrix (top right), a bipartite graph representation (bottom right), and key network science metrics: characteristic path length L , modularity (Q), structural consistency (σ_c), and the power law exponent of the degree distribution (γ). The network in (b) is a standard DNM model, generated using fixed distributions for all parameters, $M = 3$, and $\alpha = 1$. Panels (c-d) modify this standard configuration by switching a single parameter’s distribution to spatial Gaussian: (c) degree distribution, (d) synaptic distribution.

because this configuration does not alter the structure of the network much, as highlighted by the adjacency matrix depicted.

This analysis shows that DNM is a highly flexible framework that can produce a wide spectrum of network architectures. This ability to controllably generate diverse network geometries is fundamental for analyzing the relationship between network structure and computational function in ANNs.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We conduct experiments over two regimes: Static Sparse Training and Dynamic Sparse Training (DST).

Static Sparse Training We first evaluate DNM on Multilayer Perceptrons (MLPs) for image classification tasks. In this regime, the topology remains fixed after initialization to isolate the performance of initial sparse network.

Dynamic Sparse Training (DST) To validate the robustness of DNM as an initialization strategy for evolving topologies, we integrate it into state-of-the-art DST frameworks. We select three DST methods that represent different landscapes of topological evolution mechanisms: SET (Mocanu et al., 2018) utilises random link regrowth; RigL (Evci et al., 2020) adopts gradient-based link regrowth; CHTs (Zhang et al., 2025) uses a network science-based Hebbian-inspired gradient-free link regrowth. A detailed description of these models is provided in Appendix H.2. By testing DNM on these fundamentally different regrowth strategies, we aim to prove that the benefits of the initialization are not limited to a specific training paradigm. Finally, we apply DNM to Transformer (Vaswani et al., 2017b) models for machine translation tasks.

Implementation details For MLP training, we sparsify all layers except the final layer to prevent disconnected output neurons, noting that the final layer has relatively fewer connections compared to previous layers. Comprehensive parameter settings are detailed in Appendix D, and sensitivity tests on the DNM hyperparameters are provided in Appendix E.

378 **Baseline methods** We compare the performance of the DNM initialization against baseline topolo-
 379 gies found in the literature. On static sparse training, we compare DNM with a randomly initialized
 380 network, the Bipartite Small World (BSW) (Zhang et al., 2024b), the Bipartite Receptive Field
 381 (BRF) (Zhang et al., 2025), the Ramanujan graph (Kalra et al.) initialization techniques, the RadiX-
 382 Nets (Kepner & Robinett, 2019) and dANN-R (Chavlis & Poirazi, 2025)², which proved to be the
 383 best-performing dANN variant over our tests. We also include CSTI (Zhang et al., 2024a) and SNIP
 384 (Lee et al., 2018) as the baseline models, noting that their comparison is inherently unfair due to
 385 their data-informed nature. For dynamic sparse training, we compare DNM with a random initial-
 386 ization, BRF, BSW, Ramanujan graph, RadiX-Nets, dANN, SNIP and CSTI. Finally, for tests on
 387 Transformer models, we compare DNM with the BRF initialization, which was proven to be the
 388 state-of-the-art sparse initialization method in previous studies (Zhang et al., 2025).

389 4.2 MLP FOR IMAGE CLASSIFICATION

391 **Static sparse training** As an initial evaluation of DNM’s performance, we compare it to other
 392 topological initialization methods for static sparse training for image classification tasks. On all
 393 benchmarks, DNM outperforms the baseline models, as shown in Table 1. Analyzing the best-
 394 performing DNM networks is crucial to understanding the relationship between network topology
 395 and task performance. This aspect is assessed in Section 5.

397 Table 1: Image classification accuracy of statically trained, 99% sparse MLPs with different initial
 398 network topologies, compared to the fully-connected (FC) model. The scores are averaged over
 399 3 seeds \pm their standard errors. Bold values denote the best performance amongst initialization
 400 methods different from CSTI and SNIP.

401 **Static Sparse Training**

	MNIST	Fashion MNIST	EMNIST	CIFAR10
FC	98.80 \pm 0.00	90.87 \pm 0.02	87.08 \pm 0.04	62.35 \pm 0.13
CSTI	98.11 \pm 0.03	88.55 \pm 0.18	84.74 \pm 0.06	52.60 \pm 0.25
SNIP	98.03 \pm 0.03	88.65 \pm 0.07	85.19 \pm 0.04	61.89 \pm 0.48
Random	95.58 \pm 0.03	86.76 \pm 0.05	78.42 \pm 0.26	54.75 \pm 0.15
BSW	97.27 \pm 0.05	87.87 \pm 0.10	82.92 \pm 0.05	56.26 \pm 0.04
BRF	97.28 \pm 0.03	87.78 \pm 0.14	82.88 \pm 0.02	54.86 \pm 0.08
Ramanujan	96.39 \pm 0.10	86.44 \pm 0.14	81.78 \pm 0.08	54.61 \pm 0.32
RadiX-Nets	97.06 \pm 0.12	88.02 \pm 0.05	82.65 \pm 0.11	50.90 \pm 0.23
dANN-R	96.10 \pm 0.11	86.52 \pm 0.01	80.64 \pm 0.11	51.57 \pm 0.23
DNM	98.07\pm0.09	88.86\pm0.21	85.63\pm0.10	58.71\pm0.28

415 **Dynamic sparse training** We first test DNM on the baseline dynamic sparse training methods,
 416 SET and RigL. The results are shown and discussed in Appendix B, proving that DNM outperforms
 417 the other sparse initialization methods of MLPs (99% sparsity) over the datasets tested. Table 2 shows
 418 the result of the same tests on the state-of-the-art DST method, CHTs. Not only does DNM exhibit
 419 high performance for this task, but it can also surpass the input-informed CSTI method.

421 4.3 TRANSFORMER FOR MACHINE TRANSLATION

423 We assess the Transformer’s performance on a machine translation task across three datasets. We
 424 take the best performance of the model on the validation set and report the BLEU score on the test
 425 set. Beam search, with a beam size of 2, is employed to optimize the evaluation process. On the
 426 Multi30k and IWSLT datasets, we conduct a thorough hyperparameter search to find the best settings
 427 for our DNM model. For the WMT dataset, in contrast, we simply use the best settings found in

428 ²For fairness, to perform this comparison, we substituted each of the bipartite sandwich layers in our net-
 429 work with Chavlis and Poirazi’s three-layered subnetwork of sizes x , $2x$, and x respectively, where x is the size
 430 of the input. Then, to compensate for the size difference between the two models, we initialized the dANNs in
 431 a way such that the number of connections between networks is the same, rather than their sparsities. However,
 we also provide tests on the original model published by Chavlis & Poirazi (2025) in Appendix J.

432 Table 2: Image classification on MNIST, Fashion MNIST, EMNIST, and CIFAR10 of the CHTs
 433 model on MLPs with 99% sparsity over various topological initialization methods, compared to the
 434 fully-connected (FC) model. The scores indicate the accuracy of the models, averaged over 3 seeds
 435 \pm their standard errors. Bold values denote the best performance amongst initialization methods
 436 different from CSTI.

	CHTs			
	MNIST	Fashion MNIST	EMNIST	CIFAR10
FC	98.80 \pm 0.00	90.87 \pm 0.02	87.08 \pm 0.04	62.35 \pm 0.13
CSTI	98.70 \pm 0.04	90.56 \pm 0.09	87.47 \pm 0.04	69.59 \pm 0.20
Random	98.46 \pm 0.08	90.02 \pm 0.14	87.04 \pm 0.09	64.62 \pm 0.08
BSW	98.45 \pm 0.03	90.22 \pm 0.07	87.14\pm0.03	67.16 \pm 0.03
BRF	98.52 \pm 0.08	90.55 \pm 0.08	87.09 \pm 0.10	66.72 \pm 0.96
Ramanujan	98.37 \pm 0.04	89.78 \pm 0.12	86.82 \pm 0.09	64.57 \pm 0.10
RadIX-Nets	98.44 \pm 0.05	90.10 \pm 0.18	86.85 \pm 0.06	64.92 \pm 0.11
dANN-R	— \pm	— \pm	— \pm	— \pm
DNM	98.59\pm0.03	90.57\pm0.10	87.14\pm0.09	68.52\pm0.03

450 the previous tests. This approach helps to verify that DNM performs well even without extensive,
 451 dataset-specific tuning. DNM markedly improves the performance of the CHTs algorithm (Table 3,
 452 4).

453 Table 3: Performance comparison of BRF and DNM initialization on Transformer models trained
 454 with CHTs on Multi30k en-de and IWSLT en-de translation tasks with varying sparsity levels (95%
 455 and 90%). BLEU scores (higher is better) are averaged over 3 seeds \pm standard error. Bold indicates
 456 best performance for given sparsity and initialization.

	CHTs			
Initialization	Multi30k		IWSLT	
	0.95	0.90	0.95	0.90
FC	31.38 \pm 0.38		24.48 \pm 0.30	
BRF	28.94 \pm 0.57	29.81 \pm 0.37	21.15 \pm 0.10	21.92 \pm 0.17
DNM	30.54\pm0.42	31.45\pm0.35	22.09\pm0.14	23.52\pm0.24

466 Table 4: Performance comparison of BRF and DNM initialization on Transformer models trained
 467 with CHTs on machine translation tasks across the WMT en-de dataset with varying final sparsity
 468 levels (95% and 90%). Contrary to the BRF model, the DNM model’s parameters were transferred
 469 from the best-performing combinations of previous tests, avoiding any parameter search. Entries are
 470 BLEU scores (higher is better), averaged over 3 seeds \pm standard error. Bold values denote the best
 471 performance for a given sparsity and initialization.

	CHTs	
Initialization	WMT	
	0.95	0.90
FC	25.52	
BRF	20.94 \pm 0.63	22.40 \pm 0.06
DNM	21.34\pm0.20	22.56\pm0.14

481 5 RESULTS ANALYSIS

482 To understand which network structures are inherently best suited for specific tasks, we analyze the
 483 topologies of the top-performing models from our static sparse training experiments. Static sparse
 484 training is ideal for this analysis because its fixed topology allows us to link network structure to task

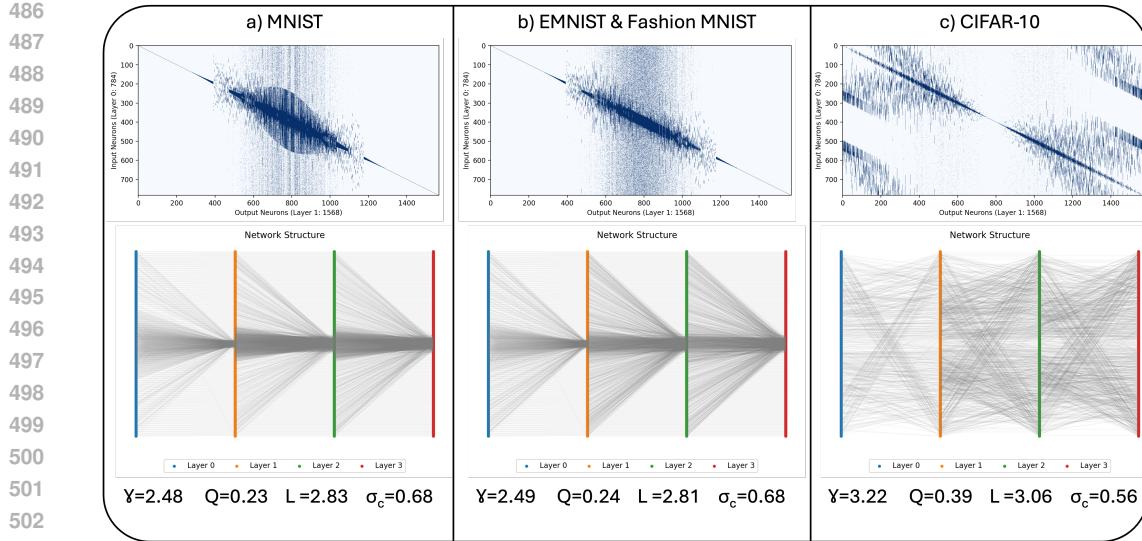


Figure 3: **Representation of the best performing DNM models on image classification.** The figure compares the best performing DNM architectures on MNIST (a), Fashion MNIST and EMNIST (b), and CIFAR10 (c). Each panel shows the network’s adjacency matrix (top) and the network’s layerwise representation (bottom). Furthermore, each panel exhibits the network’s topological measures: characteristic path length L , modularity (Q), structural consistency (σ_c), and the power law exponent of the degree distribution (γ).

performance directly, isolating it as a variable in a way that is impossible with dynamic methods. Figure 3 shows the adjacency matrices of these models, their direct bipartite graph representations, and their key metrics in network science. For image classification on Fashion MNIST and EMNIST, the optimal network’s topology is identical, and very similar to that on MNIST. These networks are scale-free ($\gamma \leq 3$) (Barabási & Albert, 1999) and exhibit a small characteristic path. Finally, we obtain contrasting results when assessing the network adopted for CIFAR-10 classification. Its higher γ parameter indicates that this network lacks hub nodes, possibly hinting that for more complex datasets like CIFAR-10, a more distributed and less hierarchical connectivity pattern is advantageous. Such a topology might promote the parallel processing of localized features across the input space, which is critical for natural image recognition, where object location and context vary significantly. In Appendix N, we expand our analysis by examining the topologies of the best-performing and worst-performing models on each of the dataset, and Appendix E gives a more detailed analysis of the best parameter combinations for each of the tests performed.

Overall, this analysis reveals a compelling relationship between task complexity and optimal network topology. While simpler, more structured datasets like MNIST and EMNIST benefit from scale-free, hierarchical architectures that can efficiently integrate global features through hub neurons, the more complex CIFAR-10 dataset favors a flatter, more distributed architecture. This underscores the potential of the DNM: its parametric flexibility allows it to generate these distinct, task-optimized topologies, moving beyond a one-size-fits-all approach to sparse initialization.

6 CONCLUSION

In this work, we introduced the Dendritic Network Model (DNM), a novel generative framework for initializing sparse neural networks inspired by the structure of biological dendrites. We have shown that the DNM is a highly flexible tool capable of producing a wide spectrum of network architectures, from modular to hierarchical and scale-free, by systematically adjusting its core parameters.

Our extensive experiments across multiple architectures demonstrate the effectiveness of our approach. At extreme sparsity levels, DNM consistently outperforms alternative topological initialization methods in both static and dynamic sparse training regimes, sometimes exceeding the performance of the data-informed CSTI and SNIP.

540 REFERENCES

541

542 Functional connectomics spanning multiple areas of mouse visual cortex. *Nature*, 640(8058):435–
543 447, 2025.

544 Eunhye Baek, Sen Song, Chang-Ki Baek, Zhao Rong, Luping Shi, and Carlo Vittorio Cannistraci.
545 Neuromorphic dendritic network computation with silent synapses for visual motion perception.
546 *Nature Electronics*, 7(6):454–465, 2024.

547 Mukund Balasubramanian and Eric L Schwartz. The isomap algorithm and topological stability.
548 *Science*, 295(5552):7–7, 2002.

549

550 Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286
551 (5439):509–512, 1999.

552 Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. Deep rewiring: Training
553 very sparse deep networks. *arXiv preprint arXiv:1711.05136*, 2017.

554

555 Tommaso Boccato, Matteo Ferrante, Andrea Duggento, and Nicola Toschi. Beyond multilayer
556 perceptrons: Investigating complex topologies in neural networks. *Neural Networks*, 171:215–
557 228, 2024.

558 Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian
559 Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. Findings of the 2017
560 conference on machine translation (wmt17). Association for Computational Linguistics, 2017.

561

562 Alberto Cacciola, Alessandro Muscoloni, Vaibhav Narula, Alessandro Calamuneri, Salvatore Nigro,
563 Emeran A Mayer, Jennifer S Labus, Giuseppe Anastasi, Aldo Quattrone, Angelo Quartarone, et al.
564 Coalescent embedding in the hyperbolic space unsupervisedly discloses the hidden geometry of
565 the brain. *arXiv preprint arXiv:1705.04192*, 2017.

566

567 Carlo Vittorio Cannistraci and Alessandro Muscoloni. Geometrical congruence, greedy navigability
568 and myopic transfer in complex networks and brain connectomes. *Nature Communications*, 13
(1):7308, 2022.

569

570 Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report
571 on the 11th IWSLT evaluation campaign. In Marcello Federico, Sebastian Stüker, and François
572 Yvon (eds.), *Proceedings of the 11th International Workshop on Spoken Language Translation:
573 Evaluation Campaign*, pp. 2–17, Lake Tahoe, California, December 4–5 2014. URL <https://aclanthology.org/2014.iwslt-evaluation.1>.

574

575 Spyridon Chavlis and Panayiota Poirazi. Dendrites endow artificial neural networks with accurate,
576 robust and parameter-efficient learning. *Nature communications*, 16(1):943, 2025.

577

578 Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist
579 to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pp.
2921–2926. IEEE, 2017.

580

581 Hermann Cuntz, Friedrich Forstner, Alexander Borst, and Michael Häusser. One rule to grow them
582 all: a general theory of neuronal branching and its practical application. *PLoS computational
583 biology*, 6(8):e1000877, 2010.

584

585 David A Drachman. Do we have brain to spare?, 2005.

586

587 Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. Multi30k: Multilingual english-
588 german image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pp.
589 70–74. Association for Computational Linguistics, 2016. doi: 10.18653/v1/W16-3210. URL
<http://www.aclweb.org/anthology/W16-3210>.

590

591 Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publication of the Mathematical
592 Institute of the Hungarian Academy of Sciences*, 5:17–60, 1960.

593

594 Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery:
595 Making all tickets winners. In *International conference on machine learning*, pp. 2943–2952.
596 PMLR, 2020.

594 AC Frank, S Huang, M Zhou, A Gdalyahu, G Kastellakis, TK Silva, E Lu, X Wen, P Poirazi,
 595 JT Trachtenberg, et al. Hotspots of dendritic spine turnover facilitate clustered spine addition and
 596 learning and memory. *nat commun* 9: 422, 2018.

597

598 Ole-Christoffer Granmo. The tsetlin machine—a game theoretic bandit driven approach to optimal
 599 pattern recognition with propositional logic. *arXiv preprint arXiv:1804.01508*, 2018.

600 Siddhant Jayakumar, Razvan Pascanu, Jack Rae, Simon Osindero, and Erich Elsen. Top-kast: Top-k
 601 always sparse training. *Advances in Neural Information Processing Systems*, 33:20744–20754,
 602 2020.

603

604 Ilenna Simone Jones and Konrad Paul Kording. Might a single neuron solve interesting machine
 605 learning problems through successive computations on its dendritic tree? *Neural Computation*,
 606 33(6):1554–1571, 2021.

607

608 Ilenna Simone Jones and Konrad Paul Kording. Do biological constraints impair dendritic compu-
 609 tation? *Neuroscience*, 489:262–274, 2022.

610

611 Suryam Arnav Kalra, Arindam Biswas, Pabitra Mitra, and Biswajit Basu. Sparse neural architectures
 via deterministic ramanujan graphs. *Transactions on Machine Learning Research*.

612

613 Jeremy Kepner and Ryan Robinett. Radix-net: Structured sparse matrices for deep neural networks.
 In *2019 IEEE international parallel and distributed processing symposium workshops (IPDPSW)*,
 614 pp. 268–274. IEEE, 2019.

615

616 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
 617 *arXiv:1412.6980*, 2014.

618

619 TN Kipf. Semi-supervised classification with graph convolutional networks. *arXiv preprint*
 620 *arXiv:1609.02907*, 2016.

621

622 Alex Krizhevsky. Learning multiple layers of features from tiny images. pp. 32–33, 2009. URL
<https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.

623

624 Matthew Larkum. A cellular mechanism for cortical associations: an organizing principle for the
 625 cerebral cortex. *Trends in neurosciences*, 36(3):141–151, 2013.

626

627 Mike Lasby, Anna Golubeva, Utku Evci, Mihai Nica, and Yani Ioannou. Dynamic sparse training
 628 with structured sparsity. *arXiv preprint arXiv:2305.02299*, 2023.

629

630 Clarissa Lauditti, Enrico M Malatesta, Fabrizio Pittorino, Carlo Baldassi, Nicolas Brunel, and Ric-
 631 cardo Zecchina. Impact of dendritic nonlinearities on the computational capabilities of neurons.
PRX Life, 3(3):033003, 2025.

632

633 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to
 634 document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 2002.

635

636 Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning
 637 based on connection sensitivity. *arXiv preprint arXiv:1810.02340*, 2018.

638

639 Xinyi Li, Jianshi Tang, Qingtian Zhang, Bin Gao, J Joshua Yang, Sen Song, Wei Wu, Wenqiang
 640 Zhang, Peng Yao, Ning Deng, et al. Power-efficient neural network with artificial dendrites.
Nature Nanotechnology, 15(9):776–782, 2020.

641

642 Sean Lie. Harnessing the power of sparsity for large gpt ai models. *Cerebras Systems Blog*, 2022.

643

644 Michael London and Michael Häusser. Dendritic computation. *Annu. Rev. Neurosci.*, 28(1):503–
 645 532, 2005.

646

647 Linyuan Lü, Liming Pan, Tao Zhou, Yi-Cheng Zhang, and H Eugene Stanley. Toward link pre-
 648 dictability of complex networks. *Proceedings of the National Academy of Sciences*, 112(8):
 649 2325–2330, 2015.

650

651 Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):
 652 261–277, 1988.

648 Nikos Malakasis, Spyridon Chavlis, and Panayiota Poirazi. Synaptic turnover promotes efficient
 649 learning in bio-realistic spiking neural networks. In *2023 57th Asilomar Conference on Signals,*
 650 *Systems, and Computers*, pp. 942–949. IEEE, 2023.

651

652 Adam Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families i: Bipartite ramanujan
 653 graphs of all degrees. In *2013 IEEE 54th Annual Symposium on Foundations of computer science*,
 654 pp. 529–537. IEEE, 2013.

655

656 Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H Nguyen, Madeleine Gibescu,
 657 and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connec-
 658 tivity inspired by network science. *Nature communications*, 9(1):2383, 2018.

659

660 Roberto LS Monteiro, Tereza Kelly G Carneiro, José Roberto A Fontoura, Valéria L da Silva,
 661 Marcelo A Moret, and Hernane Borges de Barros Pereira. A model for improving the learning
 662 curves of artificial neural networks. *PloS one*, 11(2):e0149874, 2016.

663

664 Alessandro Muscoloni, Josephine Maria Thomas, Sara Ciucci, Ginestra Bianconi, and Carlo Vittorio Cannistraci. Machine learning meets complex networks via coalescent embedding in the
 665 hyperbolic space. *Nature communications*, 8(1):1615, 2017.

666

667 Alessandro Muscoloni, Ilyes Abdelhamid, and Carlo Vittorio Cannistraci. Local-community net-
 668 work automata modelling based on length-three-paths for prediction of complex network struc-
 669 tures in protein interactomes, food webs and more. *BioRxiv*, pp. 346916, 2018.

670

671 Alessandro Muscoloni, Umberto Michieli, Yingtao Zhang, and Carlo Vittorio Cannistraci. Adaptive
 672 network automata modelling of complex networks. 2022.

673

674 Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national
 675 academy of sciences*, 103(23):8577–8582, 2006.

676

677 Julian Stier and Michael Granitzer. Deepstruct–linking deep learning and graph theory. *Software
 678 Impacts*, 11:100193, 2022.

679

680 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 681 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-
 682 tion processing systems*, 30, 2017a.

683

684 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 685 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-
 686 tion processing systems*, 30, 2017b.

687

688 Christopher A Walsh. Peter huttenlocher (1931–2013). *Nature*, 502(7470):172–172, 2013.

689

690 Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393
 691 (6684):440–442, 1998.

692

693 Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmark-
 694 ing machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

695

696 Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring randomly wired neural
 697 networks for image recognition. In *Proceedings of the IEEE/CVF international conference on
 698 computer vision*, pp. 1284–1293, 2019.

699

700 Geng Yuan, Xiaolong Ma, Wei Niu, Zhengang Li, Zhenglun Kong, Ning Liu, Yifan Gong, Zheng
 701 Zhan, Chaoyang He, Qing Jin, et al. Mest: Accurate and fast memory-economic sparse training
 702 framework on the edge. *Advances in Neural Information Processing Systems*, 34:20838–20850,
 703 2021.

704

705 Yingtao Zhang, Jialin Zhao, Wenjing Wu, and Alessandro Muscoloni. Epitopological learning and
 706 cannistraci-hebb network shape intelligence brain-inspired theory for ultra-sparse advantage in
 707 deep learning. In *The Twelfth International Conference on Learning Representations*, 2024a.

702 Yingtao Zhang, Jialin Zhao, Wenjing Wu, Alessandro Muscoloni, and Carlo Vittorio Cannistraci.
703 Epitopological learning and cannistraci-hebb network shape intelligence brain-inspired theory for
704 ultra-sparse advantage in deep learning. In *The Twelfth International Conference on Learning*
705 *Representations*, 2024b. URL <https://openreview.net/forum?id=iayEcORsGd>.
706
707 Yingtao Zhang, Diego Cerretti, Jialin Zhao, Wenjing Wu, Ziheng Liao, Umberto Michieli, and
708 Carlo Vittorio Cannistraci. Brain network science modelling of sparse neural networks enables
709 transformers and llms to perform as fully connected. *arXiv preprint arXiv:2501.19107*, 2025.
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756 A GLOSSARY OF NETWORK SCIENCE
757758 In this section, we introduce the basic notions of network science mentioned in this article.
759760 **Scale-Free Network** A Scale-Free Network (Barabási & Albert, 1999) is characterized by a highly
761 uneven distribution of degrees amongst the nodes. A small number of nodes, called hubs, have a
762 very high degree, and a large number of nodes have very few connections. The degree distribution
763 of scale-free networks follows a power law trend $P(k) \sim k^\gamma$, where γ is a constant smaller than 3.
764 In contrast, nodes in random networks are distributed following a Binomial distribution.
765766 **Watts-Strogatz Model and Small-World Network** A Small-World Network (Watts & Strogatz,
767 1998) is characterized by a small average path length. This property implies that any two nodes can
768 communicate through a short chain of connections. The Watts-Strogatz (Watts & Strogatz, 1998)
769 model is well known for its high clustering and short path lengths. This network is modelled by a
770 parameter β between 0 and 1 that can determine its level of clustering. When β takes low values
771 ($\beta \approx 0$), the WS network is a highly clustered lattice. On the other hand, when β approaches 1, the
772 network becomes a random small-world graph. Intermediate values of β can generate a clustered
773 network that maintains small-world connectivity.774 Formally, a network is small-world when the path of length L between two randomly chosen nodes
775 grows proportionally to the logarithm of the number of nodes (N) in the network, that is:
776

777
$$L \propto \log N. \quad (1)$$

778 **Structural Consistency** Structural consistency (Lü et al., 2015) is an index based on the first-
779 order matrix perturbation of the adjacency matrix, which represents the predictability of the network
780 structure. A perturbation set ΔE is randomly sampled from the original link set E . Identifying as
781 E^L the links ranked as the top L according to the structural perturbation method (Lü et al., 2015),
782 with $L = |\Delta E|$, the structural consistency σ_c is calculated as:
783

784
$$\sigma_c = \frac{|E^L \cap \Delta E|}{\Delta E}. \quad (2)$$

785

786 **Modularity** Modularity (Newman, 2006) quantifies the tendency of nodes in the network to form
787 distinct communities (or clusters). This measure ranges from -1 to 1. A high modularity score
788 (close to 1) hints at the presence of dense connections between nodes within communities, but
789 sparse connections between nodes belonging to different communities. A modularity score close
790 to 0, in contrast, suggests that the network lacks any community organization and the interaction
791 between nodes is essentially uniform. When modularity approaches -1, the network exhibits an
792 anticomunity structure. This means that nodes are strongly connected across the network, and
793 there is little differentiation into separate groups. In other words, a negative modularity represents a
794 cohesive network. The formula to compute the modularity (Q) is:
795

796
$$Q = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j), \quad (3)$$

797

798 where A represents the network's adjacency matrix, and k_i and k_j are the degrees of nodes i and j ,
799 respectively. $\delta(c_i, c_j)$ is the Kronecker delta function, which equals one if i and j are in the same
800 community, else it equals 0.
801802 **Characteristic path length** The characteristic path length is computed as the average node-pairs
803 length in the network; it is a measure associated with the network's small-worldness (Cannistraci &
804 Muscoloni, 2022). The characteristic path length (L) is derived by:
805

806
$$L = \frac{1}{n(n-1)} \cdot \sum_{i,j} d(i, j), \quad (4)$$

807

808 where n is the number of nodes in the network, and $d(i, j)$ is the shortest path length between node i
809 and node j .
810

810
Coalescent Embedding Coalescent embedding (Muscoloni et al., 2017) is a class of machine
 811 learning algorithms used for unsupervised dimensionality reduction and embedding complex net-
 812 works in a geometric space, often hyperbolic. This method maps high-dimensional information on
 813 a low-dimensional embedding while maintaining the essential topological features of the network.
 814 This embedding reveals latent structures of the system, like hierarchical and scale-free structures.
 815 In this article, coalescent embedding maps the networks that have latent hyperbolic geometry onto
 816 the two-dimensional hyperbolic space. The approach involves 4 steps: 1) links are pre-weighted
 817 with topological rules that approximate the underlying network geometry (Cannistraci & Muscoloni,
 818 2022); 2) non-linear dimensionality reduction; 3) generation of angular coordinates; 4) generation
 819 of radial coordinates.

820 This process is illustrated in Figure 2, which showcases the results of applying a specific coalescent
 821 embedding pipeline to four different synthetic networks. The embeddings shown were generated
 822 without any initial link pre-weighting (step 1). For the non-linear dimensionality reduction (step
 823 2), the Isomap (Balasubramanian & Schwartz, 2002) algorithm was used. Finally, the angular co-
 824 ordinates (step 3) were determined using Equidistant Adjustment (EA), a process that preserves the
 825 relative order of the nodes while arranging them at perfectly uniform angular intervals.

827 B EXPERIMENTS ON BASELINE DST METHODS

828 In this section, we provide the results of our experiments on the baseline dynamic sparse training
 829 methods, SET and RigL. The results are shown in Tables 5 and 6, proving that DNM outperforms
 830 the other sparse initialization methods of MLPs (99% sparsity) over the datasets tested. The model’s
 831 performance is comparable to the input-informed CSTI and SNIP, highlighting that DNM’s high
 832 degree of freedom can match a topology induced by data features.

833
 834
 835 Table 5: Image classification on MNIST, Fashion MNIST, EMNIST, and CIFAR-10 of the SET
 836 model on MLPs with 99% sparsity over various topological initialization methods, compared to the
 837 fully-connected (FC) model. The scores indicate the accuracy of the models, averaged over 3 seeds
 838 \pm their standard errors. Bold values denote the best performance amongst initialization methods
 839 different from CSTI and SNIP.

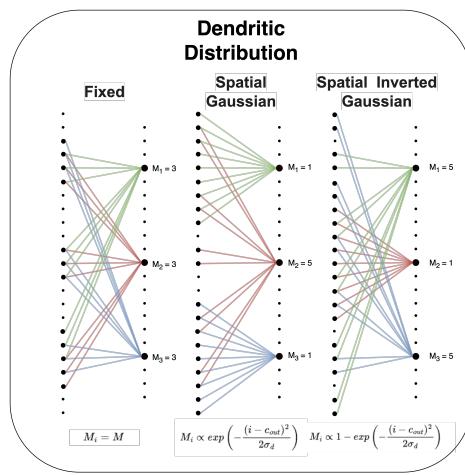
840
SET

	MNIST	Fashion MNIST	EMNIST	CIFAR10
FC	98.80 \pm 0.00	90.87 \pm 0.02	87.08 \pm 0.04	62.35 \pm 0.13
CSTI	98.40 \pm 0.06	89.96 \pm 0.07	86.70 \pm 0.10	65.31 \pm 0.16
SNIP	98.66 \pm 0.02	90.43 \pm 0.08	87.13 \pm 0.02	63.45 \pm 0.14
Random	98.16 \pm 0.06	89.17 \pm 0.15	86.03 \pm 0.12	62.80 \pm 0.24
BSW	98.22 \pm 0.03	89.28 \pm 0.09	86.21 \pm 0.02	64.13 \pm 0.11
BRF	98.56 \pm 0.03	89.58 \pm 0.11	86.21 \pm 0.11	64.40 \pm 0.25
Ramanujan	98.08 \pm 0.03	88.72 \pm 0.11	85.89 \pm 0.04	62.28 \pm 0.15
RadiX-Nets	98.37 \pm 0.08	89.33 \pm 0.08	86.15 \pm 0.09	55.91 \pm 0.13
dANN-R	97.95 \pm 0.10	88.91 \pm 0.02	85.47 \pm 0.09	57.44 \pm 0.09
DNM	98.67\pm0.04	89.66\pm0.05	87.32\pm0.11	64.47\pm0.17

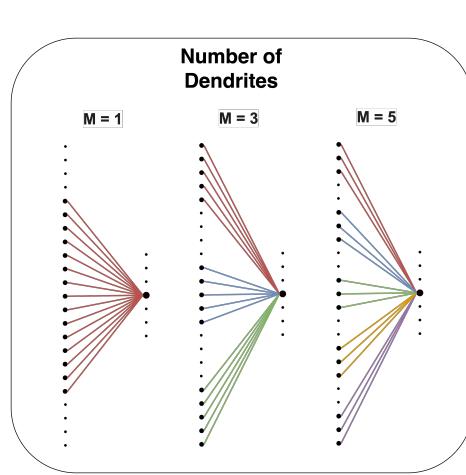
855 C IMPACT OF DNM PARAMETERS ON NETWORK SCIENCE

856 In this section, we provide a visual breakdown of how the topological structure of the Dendritic Neu-
 857 ral Model (DNM) adapts when individual control parameters are varied. While the mathematical
 858 definitions of these distributions are provided in the main text, visualizing the resulting connectivity
 859 patterns offers greater insight into the network’s plasticity and pruning capabilities. Figure 4 illus-
 860 trates the effects of varying key DNM parameters on the resulting network topology. Each subfigure
 861 isolates a single parameter change while holding all others constant, allowing for a clear view of its
 862 specific influence.

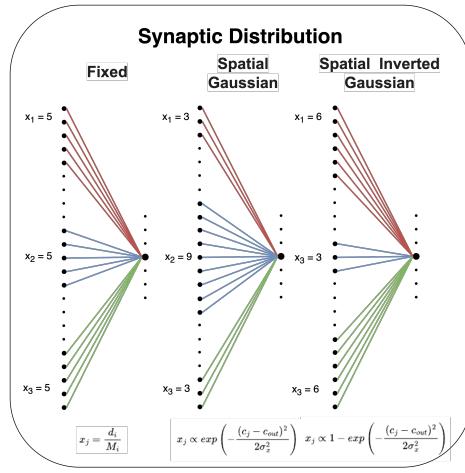
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878



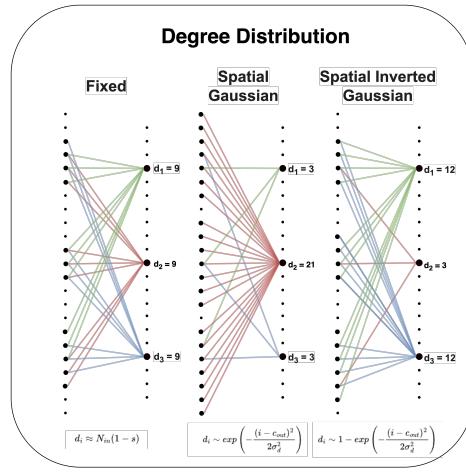
(a) Effect of the dendritic distribution on the network topology.



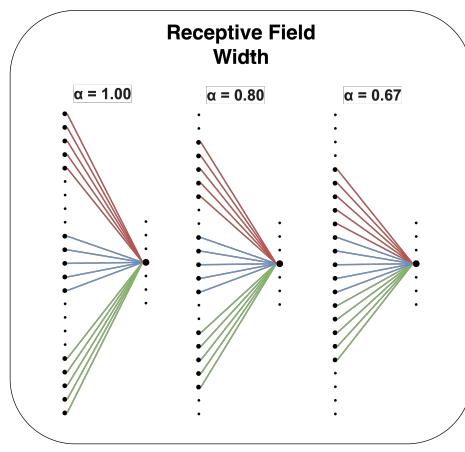
(b) Effect of the number of dendrites on the network topology.



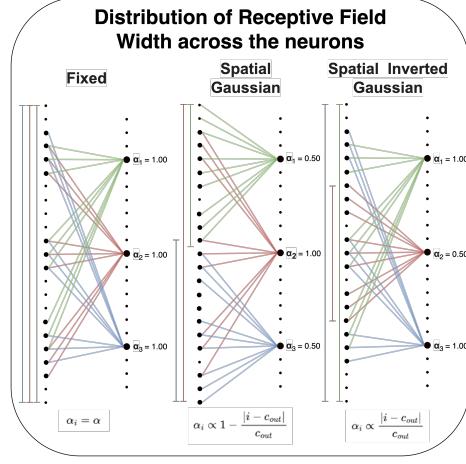
(c) Effect of the synaptic distribution on the network topology.



(d) Effect of the degree distribution on the network topology.



(e) Effect of the receptive field width on the network topology.

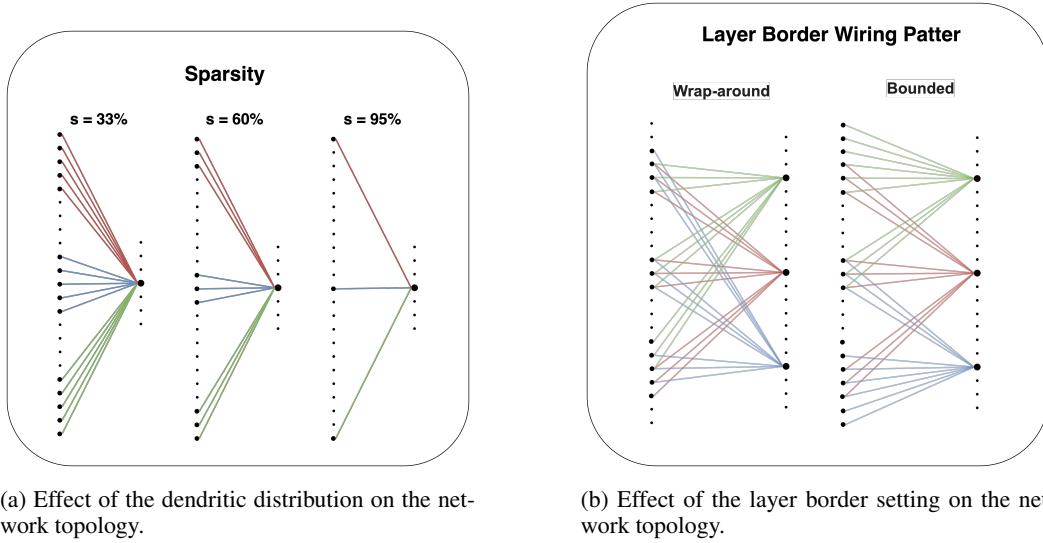


(f) Effect of the receptive field width distribution on the network topology.

914
915
916
917
Figure 4: Representations of the network's topology obtained by varying a DNM parameter while
918
keeping all others fixed.

918
 919 Table 6: Image classification on MNIST, Fashion MNIST, EMNIST, and CIFAR-10 of the RigL
 920 model on MLPs with 99% sparsity over various topological initialization methods, compared to the
 921 fully-connected (FC) model. The scores indicate the accuracy of the models, averaged over 3 seeds
 922 \pm their standard errors. Bold values denote the best performance amongst initialization methods
 923 different from CSTI and SNIP.

	RigL			
	MNIST	Fashion MNIST	EMNIST	CIFAR10
FC	98.80 ± 0.00	90.87 ± 0.02	87.08 ± 0.04	62.35 ± 0.13
SNIP	98.76 ± 0.05	90.50 ± 0.06	87.30 ± 0.04	63.31 ± 0.25
CSTI	98.77 ± 0.02	90.19 ± 0.03	87.28 ± 0.06	60.59 ± 0.46
Random	98.66 ± 0.27	89.88 ± 0.04	87.18 ± 0.07	64.13 ± 0.11
BSW	98.74 ± 0.03	90.12 ± 0.06	87.28 ± 0.10	65.19 ± 0.23
BRF	98.18 ± 0.03	89.79 ± 0.02	87.05 ± 0.14	63.55 ± 0.78
Ramanujan	98.37 ± 0.04	89.78 ± 0.12	86.82 ± 0.09	64.57 ± 0.10
Radix-Nets	98.44 ± 0.05	90.10 ± 0.18	86.85 ± 0.06	64.57 ± 0.10
dANN-R	98.54 ± 0.05	89.44 ± 0.05	86.81 ± 0.04	62.03 ± 0.06
DNM	98.74 ± 0.06	90.22 ± 0.02	87.35 ± 0.15	65.58 ± 0.13



952
 953 Figure 5: Representations of the network’s topology obtained by varying a DNM parameter while
 954 keeping all others fixed.

958 D HYPERPARAMETER SETTINGS AND IMPLEMENTATION DETAILS

960 Our experimental setup is designed to replicate the conditions in Zhang et al. (2025). Configurations
 961 are assessed on validation sets before being tested on separate test sets. All reported scores are the
 962 average of three runs using different random seeds, presented with their corresponding standard
 963 errors.

965 D.1 MLP FOR IMAGE CLASSIFICATION

967 Models are trained for 100 epochs using Stochastic Gradient Descent (SGD) with a learning rate
 968 of 0.025, a batch size of 32, and a weight decay of 5×10^{-4} . All sparse models are trained
 969 at a 99% sparsity level. For dynamic methods, we used SET, RigL, and CHTs. The regrowth
 970 strategy for CHTs is CH2_L3n (Muscoloni et al., 2018). For our DNM, we conduct a grid search
 971 over its key hyperparameters. We tested a mean dendrite count (M) of 3. For the dendrite, degree,
 972 receptive field width, and synaptic distributions, we searched across fixed, spatial Gaussian, and

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 7: Hyperparameters of MLP on Image Classification Tasks

Hyper-parameter	MLP
Hidden Dimension	1568 (3072 for CIFAR10)
# Hidden layers	3
Batch Size	32
Training Epochs	100
LR Decay Method	Linear
Start Learning Rate	0.025
End Learning Rate	$2.5e^{-4}$
ζ (fraction of removal)	0.3
Update interval (for DST)	1 epoch
Momentum	0.9
Weight decay	$5e^{-4}$

spatial inverted Gaussian options. The mean receptive field width (α) was fixed at 1.0. For the BSW baseline, the rewiring probability is searched in the set $\{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. For the BRF baseline, we searched the randomness parameter r over the same set of values, and also tested both fixed and uniform degree distributions.

D.2 TRANSFORMER FOR MACHINE TRANSLATION

We use a standard 6-layer Transformer architecture with 8 attention heads and a model dimension of 512. The dimension of the feed-forward network is set to 1024 for Multi30k and 2048 for IWSLT14 and WMT17. All models are trained using the Adam optimizer (Kingma & Ba, 2014) with the noam learning rate schedule. Dataset-specific training parameters are as follows:

- Multi30k: Trained for 5,000 steps with a learning rate of 0.25, 1000 warmup steps, and a batch size of 1024.
- IWSLT14: Trained for 20,000 steps with a learning rate of 2.0, 6000 warmup steps, and a batch size of 10240.
- WMT17: Trained for 80,000 steps with a learning rate of 2.0, 8000 warmup steps, and a batch size of 12000.

We evaluated models at final sparsity levels of 90% and 95%. For Multi30k and IWSLT14, we performed a comprehensive hyperparameter search. The search for IWSLT14 included a mean dendrite count $M \in \{3, 7, 21\}$ and various combinations of fixed and spatial distributions for all DNM parameters. For WMT17, to assess generalization, we did not perform a new search. Instead, we directly applied the best-performing DNM configuration identified from the IWSLT14 experiments. This configuration used $M=7$ with a fixed dendritic distribution, alongside spatial Gaussian or inverse-Gaussian distributions for degree, receptive field width, and synapses.

E SENSITIVITY TESTS

We provide sensitivity tests for DNM hyperparameters. First, we focus on the analysis of CHTs on MLPs for image classification at 99% sparsity. Next, we study the parametric configurations for the CHTs model on the Multi30k translation task at 90% sparsity. For each task, we vary one parameter at a time, keeping the others fixed to a specific configuration. We calculate the coefficient of variation (CV) of the scores to quantify the sensitivity of the model to each parameter. A low CV indicates that the model’s performance is relatively stable across different settings of that parameter, suggesting low sensitivity. Conversely, a high CV suggests that the model’s performance is more variable and sensitive to changes in that parameter. We average each parameter’s CV across various parametric configurations to obtain a robust measure of sensitivity. Next, we analyse the top 5% best-performing configurations for each task to understand the commonalities in the optimal settings. This method not only helps us understand which parameters are most influential but also guides future configurations for similar tasks.

1026
1027 **Dynamic Sparse Training for Image Classification** We analyse the sensitivity of DNM parame-
1028 ters for the initialization of CHTs on MLPs for image classification at 99% sparsity. The analysis is
1029 performed over MNIST, Fashion MNIST, EMNIST, and CIFAR-10. The sensitivity analysis, sum-
1030 marized in Table 8, evaluates the impact of DNM initialization parameters for CHTs at 99% sparsity.
1031 Across all benchmarks, the Degree Distribution consistently emerges as the most critical parameter,
1032 highlighting the paramount importance of initial network connectivity. Following in descending or-
1033 der of influence are the Receptive Field Width, Dendritic, and Synaptic distributions. Analyzing the
1034 top 5% best-performing configurations, we observe similar trends across various datasets (Figure 6).
1035 The most relevant findings is that a spatial gaussian receptive field width distribution is constantly
1036 preferred.

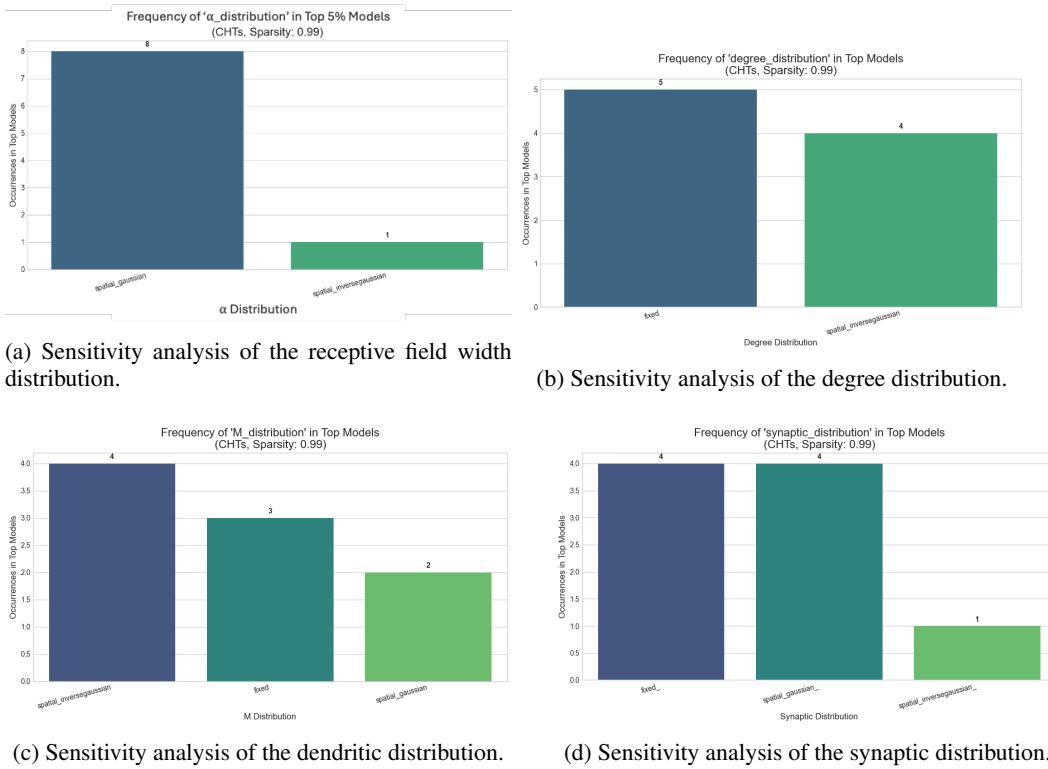


Table 8: Sensitivity Analysis Results for CHTs on MNIST, EMNIST, Fashion MNIST, CIFAR-10 for Image Classification. The table presents the average coefficient of variation (CV) for each DNM parameter across different configurations at 99% sparsity level. A higher CV indicates greater sensitivity of the model’s performance to changes in that parameter.

	MNIST	EMNIST	Fashion MNIST	CIFAR-10
Degree Dist	0.001538	0.001419	0.003379	0.017639
Rec Field Width Dist	0.000506	0.001585	0.001205	0.005381
Dendritic Dist	0.000525	0.001265	0.001136	0.004983
Synaptic Dist	0.000464	0.000989	0.000964	0.004306

1076 **Dynamic Sparse Training for Machine Translation** We focus on the analysis of DNM for the
1077 initialization of the CHTs model on Multi-30k for machine translation. Both at 90% and 95% spar-
1078 sity, the calculated coefficients of variation do not surpass 0.01. This indicates that the model’s
1079 performance is relatively stable across different settings of each parameter, suggesting low sensi-
tivity (Table 9). At both sparsity levels, the dendritic distribution appears to be the most sensitive,

1080 whereas M and α are the most stable. Analyzing the top 5% best-performing configurations, we
 1081 observe that spatial distributions generally outperform fixed distributions (Figure 7).
 1082

1083 Table 9: Sensitivity Analysis Results for CHTs on Multi30k for Machine Translation. The table
 1084 presents the average coefficient of variation (CV) for each DNM parameter across different config-
 1085 urations at 90% and 95% sparsity levels. A higher CV indicates greater sensitivity of the model’s
 1086 performance to changes in that parameter.

Parameter	Average Coefficient of Variation (CV)	
	90%	95%
Dendritic distribution	0.008665	0.010520
Degree distribution	0.010910	0.014196
Receptive Field Width Distribution	0.008511	0.009607
Synaptic distribution	0.009444	0.011526
M	0.008821	0.009938
α	0.007576	0.009607

1095 *Note: Higher CV indicates greater impact on performance.*

1098 F MODELLING DENDRITIC NETWORKS

1100 The Dendritic Network Model (DNM) generates the sparse connectivity matrix of bipartite sand-
 1101 which layers by iteratively building connections for each output neuron based on the principles of
 1102 dendritic branching and localized receptive fields. The generation process can be broken down into
 1103 the following steps:

- 1104 1. Determine the degree of each output neuron in the layer based on one of the three dis-
 1105 tribution strategies (fixed, non-spatial, spatial). A probabilistic rounding and adjustment
 1106 mechanism ensures that no output neuron is disconnected and the sampled degrees sum
 1107 precisely to the target total number of connections of the layer.
- 1108 2. Next, for each output neuron j , define a receptive field. This is done by topologically
 1109 mapping the output neuron’s position at a central point in the input layer and establishing a
 1110 receptive window around this center. The size of this window is controlled by the parameter
 1111 $\alpha_j \in [0, 1]$, which determines the fraction of the input layer that the neuron can connect to.
 1112 α_j itself can be fixed or sampled from a spatial or non-spatial distribution.
- 1113 3. For each output neuron, determine the number of dendritic branches, M_j to be used to con-
 1114 nect it to the input layer. Again, M_j is determined based on one of the three configurations
 1115 (fixed for all neurons, or sampled from a distribution that could depend on the neuron’s
 1116 position in the layer).
- 1117 4. Place the M_j dendrites as dendritic centers within the neuron’s receptive window, spacing
 1118 them evenly across the window.
- 1119 5. The neuron’s total degree, obtained from step 1, is distributed across its M_j dendrites ac-
 1120 cording to a synaptic distribution (fixed, spatial, or non-spatial). For each dendrite, connec-
 1121 tions are formed with the input neurons that are spatially closest to its center.
- 1122 6. Finally, the process ensures connection uniqueness and adherence to the precise degree
 1123 constraints.

1126 G DYNAMIC SPARSE TRAINING (DST)

1128 Dynamic sparse training (DST) is a subset of sparse training methodologies that allows for the evolu-
 1129 tion of the network’s topology during training. Sparse Evolutionary Training (SET) (Mocanu et al.,
 1130 2018) is the pioneering method in this field, which iteratively removes links based on the absolute
 1131 magnitude of their weights and regrows new connections randomly. Subsequent developments have
 1132 expanded upon this method by refining the pruning and regrowth steps. One such advancement was
 1133 proposed by Deep R (Bellec et al., 2017), a method that evolves the network’s topology based on
 stochastic gradient updates combined with a Bayesian-inspired update rule. RigL (Evci et al., 2020)

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

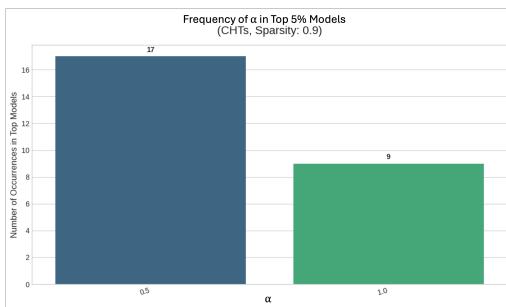
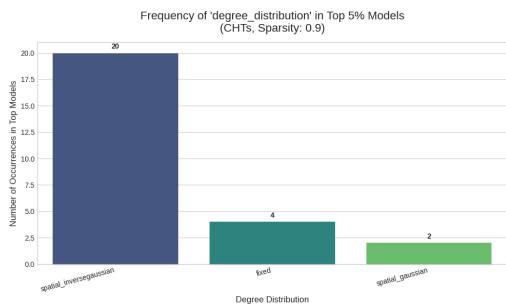
1170

1171

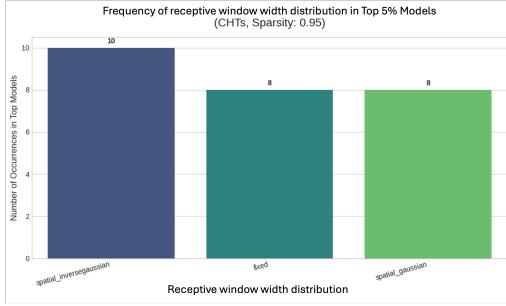
1172

1173

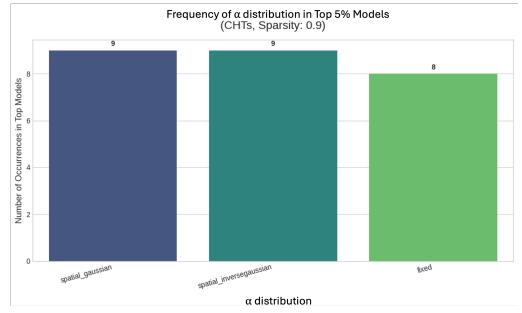
1174

(a) Sensitivity analysis of α .

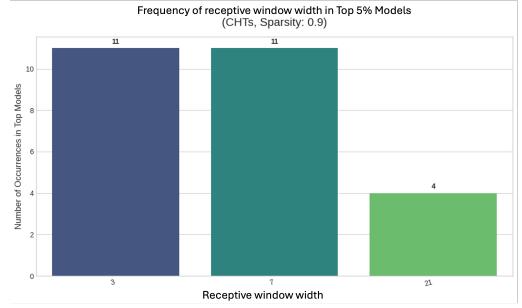
(c) Sensitivity analysis of the degree distribution.



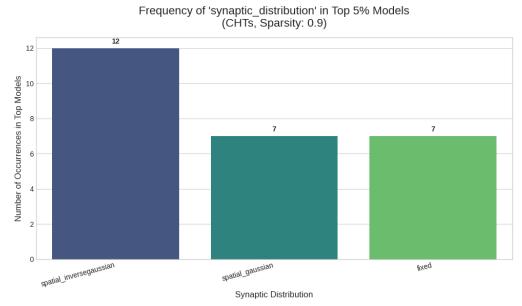
(e) Sensitivity analysis of the dendritic distribution.



(b) Sensitivity analysis of the receptive field width distribution.



(d) the receptive window width.



(f) Sensitivity analysis of the synaptic distribution.

Figure 7: Sensitivity analysis of the DNM parameters for CHTs initialization over Transformer models for machine translation on Multi30k at 90% sparsity.

advanced the field further by leveraging the gradient information of non-existing links to guide the regrowth of new connections. MEST (Yuan et al., 2021) is a method that exploits information on the gradient and the weight magnitude to selectively remove and randomly regrow new links, similarly to SET. MEST introduces the EM&S technique that gradually decreases the density of the network until it reaches the desired sparsity level. Top-KAST (Jayakumar et al., 2020) maintains a constant sparsity level through training, iteratively selecting the top K weights based on their magnitude and applying gradients to a broader subset of parameters. To avoid the model being stuck in a suboptimal sparse subset, Top-KAST introduces an auxiliary exploration loss that encourages ongoing adaptation of the mask. A newer version of RigL, sRigL (Lasby et al., 2023), adapts the principles of the original model to semi-structured sparsity, speeding up the training from scratch of vision models. CHT (Zhang et al., 2024b) is the state-of-the-art (SOTA) dynamic sparse training framework that adopts a gradient-free regrowth strategy that relies solely on topological information (network shape intelligence). This model suffers from two main drawbacks: it has time complexity $\mathcal{O}(N \cdot d^3)$ (N node network size, d node degree), and it rigidly selects top link prediction scores, which causes suboptimal link removal and regrowth during the early stages of training. For this reason, this model

1188 was evolved into CHTs (Zhang et al., 2025), which adopts a flexible strategy to sample connections
 1189 to remove and regrow, and reduces the time complexity to $\mathcal{O}(N^3)$. The same authors propose a
 1190 sigmoid-based gradual density decay strategy, namely CHTss (Zhang et al., 2025), which proves to
 1191 be the state-of-the-art dynamic sparse training method over multiple tasks. CHTs and CHTss can
 1192 surpass fully connected MLP, Transformer, and LLM models over various tasks using only a small
 1193 fraction of the networks’ connections.

H BASELINE METHODS

1197 We describe in detail the models compared in our experiments.

H.1 SPARSE NETWORK INITIALIZATION

1201 **Bipartite Scale-Free (BSF)** The Bipartite Scale-Free network (Zhang et al., 2024b) is an exten-
 1202 sion of the Barabási-Albert (BA) model (Barabási & Albert, 1999) to bipartite networks. We detail
 1203 the steps to generate the network. 1) Generate a BA monopartite network consisting of $m+n$ nodes,
 1204 where m and n are the numbers of nodes of the first and second layer of the bipartite network, re-
 1205 spectively. 2) Randomly select m and n nodes to assign to the two layers. 3) Count the number
 1206 of connections between nodes within the same layer (frustrations). If the two layers have an equal
 1207 number of frustrations, match each node in layer 1 with a frustration to a node in layer 2 with a
 1208 frustration, randomly. Apply a single rewiring step using the Maslov-Sneppen randomization (MS)
 1209 procedure for every matched pair. If the first layer counts more frustrations, randomly sample a sub-
 1210 set of layer 1 with the same number of frustrations, and repeat step 1. For each remaining frustration
 1211 in layer 1, sequentially rewire the connections to the opposite layer using the preferential attachment
 1212 method from step 1. If the second layer has more frustrations than the first, apply the opposite pro-
 1213 cedure. The resulting network will be bipartite and exhibit a power-law distribution with exponent
 1214 $\gamma = 2.76$.

1215 **Bipartite Small-World (BSW)** The Bipartite Small-World network (Zhang et al., 2024b) is an
 1216 extension of the Watts-Strogatz model to bipartite networks. It is modelled as follows: 1) Build a
 1217 regular ring lattice with a number of nodes $N = \#L_1 + \#L_2$, with $L_1 > L_2$, where L_1 and L_2
 1218 represent the nodes in the first and second layers of the bipartite network, respectively. 2) Label the
 1219 N nodes in a way such that for every L_1 node positioned in the network, $\#L_1/\#L_2$ nodes from L_2
 1220 are placed at each step. Then, at each step, establish a connection between an L_1 node and the $K/2$
 1221 closest L_2 neighbours in the ring lattice. 3) For every node, take every edge connecting it to its $K/2$
 1222 rightmost neighbors, and rewire it with a probability β , avoiding self-loops and link duplication.
 1223 When $\beta = 1$, the generated network corresponds to a random graph.

1224 **Correlated Sparse Topological Initialization (CSTI)** The Correlated Sparse Topological Initial-
 1225 ization (CSTI) (Zhang et al., 2024b) initializes the topology of the layers that interact directly with
 1226 the input features. The construction of CSTI follows four steps. 1) **Vectorization:** Denoting as n
 1227 the number of randomly sampled input data from the training set and as M the number of valid
 1228 features with variance different from zero among these samples, we build an $n \times M$ matrix. 2)
 1229 **Feature selection:** We perform feature selection by calculating the Pearson Correlation for each
 1230 feature. Hence, we construct a correlation matrix. 3) **Connectivity selection:** Next, we construct a
 1231 sparse adjacency matrix, with entries “1” corresponding to the top-k% values from the correlation
 1232 matrix (where the value of k depends on the desired sparsity level). A scaling factor \times determines
 1233 the dimension of the hidden layer. 4) **Assembling topologically hubbed network blocks:** Finally,
 1234 the selected adjacency matrix masks the network to form the initialized topology for each bipartite
 1235 sandwich layer.

1236 **SNIP** SNIP (Lee et al., 2018) is a static sparse initialization method that prunes connections based
 1237 on their sensitivity to the loss function. The sensitivity of a connection is defined as the absolute
 1238 value of the product of its weight and the gradient of the loss with respect to that weight, evaluated
 1239 on a small batch of training data. Connections with the lowest sensitivity are pruned until the desired
 1240 sparsity level is reached. This method allows for the identification of important connections before
 1241 training begins, enabling the training of sparse networks from scratch.

Ramanujan Graphs Ramanujan Graphs (Lubotzky et al., 1988) are a class of optimal expander graphs that exhibit excellent connectivity properties. They are characterized by their spectral gap, which is the difference between the largest and second-largest eigenvalues of their adjacency matrix. A larger spectral gap indicates better expansion properties, meaning that the graph is highly connected and has a small diameter. Ramanujan graphs are constructed using deep mathematical principles from number theory and algebraic geometry. They are known for their optimal expansion properties, making them ideal for applications in network design, error-correcting codes, and computer science. In this article, we use bipartite Ramanujan graphs as a sparse initialization method for neural networks.

In our experiments, we built bipartite Ramanujan graphs as a theoretically-grounded initialization method, following the core principles outlined by Kalra et al.. Drawing inspiration from the findings of Marcus et al. (2013), which prove the existence of bipartite Ramanujan graphs for all degrees and sizes, we constructed these graphs as follows:

- Generate d random permutation matrices, where d is the desired degree of the graph. Each permutation matrix represents a perfect matching, which is a set of edges that connects each node in one layer to exactly one node in the other layer without any overlaps.
- Iteratively combine these matchings. In each step, deterministically decide whether to add or subtract the successive matching to the current adjacency matrix. This decision is made by minimizing a *barrier function* that ensures that the eigenvalues remain within the Ramanujan bounds.

Bipartite Receptive Field (BRF) The Bipartite Receptive Field (BRF) Zhang et al. (2025) is the first sparse topological initialization model that generates brain-network-like receptive field connectivity. The BRF directly generates sparse adjacency matrices with a customized level of spatial-dependent randomness according to a parameter $r \in [0, 1]$. A low value of r leads to less clustered topologies. As r increases towards 1, the connectivity patterns tend to be generated uniformly at random. Specifically, when r tends to 0, BRF builds adjacency matrices with links near the diagonal (adjacent nodes from the two layers are linked), whereas when r increases 1, this structure tends to break.

Mathematically, consider an $N \times M$ bipartite adjacency matrix $M_{i,j=1,\dots,M, j=1,\dots,N}$, where M represents the input size and n represents the output size. Each entry $m_{i,j}$ of the matrix is set to 1 if node i from the input layer connects to node j of the output layer, and 0 otherwise. Define the scoring function

$$S_{i,j} = d_{ij}^{\frac{1}{1-r}}, \quad (5)$$

where

$$d_{ij} = \min\{|i - j|, |(i - M) - j|, |i - (j - N)|\} \quad (6)$$

is the distance between the input and output neurons. $S_{i,j}$ represents the distance of an entry of the adjacency matrix from the diagonal, raised to the power of $\frac{1-r}{r}$. When the parameter r tends to zero, the scoring function becomes more deterministic; when it tends to 1, all scores $S_{i,j}$ become more uniform, leading to a more random adjacency matrix.

The model is enriched by the introduction of the degree distribution parameter. The Bipartite Receptive Field with fixed sampling (BRF) sets the degree of all output neurons to be fixed to a constant value. The Bipartite Receptive Field with uniform sampling, on the other hand, samples the degrees of output nodes from a uniform distribution.

RadiX-Nets RadiX-Nets (Kepner & Robinett, 2019) represent a family of sparse deep neural network topologies initialized "de novo," meaning they are constructed with a sparse structure from the outset rather than being pruned from a dense parent network. This approach addresses the limitations of hardware capacity by deterministically generating topologies that are much more diverse than previous sparse implementations (such as X-Nets) while preserving critical graph-theoretic properties. The construction relies on mixed-radix numeral systems and the Kronecker product to achieve connectivity. Formally, a mixed-radix topology is defined by an ordered set of integers

$$\mathcal{N} = (N_1, N_2, \dots, N_L), \quad (7)$$

which implicitly defines a numeral system representing integers in the range 0 to $N' - 1$, where $N' = \prod_{i=1}^L N_i$. The final RadiX-Net topology is constructed via the Kronecker product of adjacency submatrices from these mixed-radix systems and dense adjacency submatrices. This construction guarantees two essential properties:

- Path-connectedness: There exists at least one path between any input and output node, ensuring that information can flow through the network.
- Symmetry: Each input node has the same number of paths to each output node, promoting uniformity in information distribution.

For a RadiX-Net defined by a mean radii μ and d radices (with $d \sim \log_\mu(N')$), the density of the graph scales asymptotically as $\frac{1}{\mu^{d-1}}$, allowing for significant sparsity while maintaining expressive power.

dANNs Dendritic Artificial Neural Networks (dANNs) (Chavlis & Poirazi, 2025) incorporate the structural connectivity and restricted sampling properties of biological dendrites to improve parameter efficiency and robustness⁹. Unlike the point-neuron model of standard ANNs, a dANN splits the computational unit into two layers: a dendritic layer and a somatic layer. Input data is first fed into the dendritic layer via sparse connections (synaptic weights), where each dendrite performs a linear weighted sum followed by a nonlinearity. These dendritic activations are then weighted (cable weights) and summed at the soma before passing through a second nonlinearity. A defining feature of dANNs is their restricted input sampling, inspired by the receptive fields (RFs) of the visual cortex¹². The authors propose three specific sampling strategies:

- Random Sampling (dANN-R): Input features are selected randomly for each dendrite.
- Local Receptive Fields (dANN-LRF): Inputs are sampled from a spatially restricted neighborhood (e.g., a 4×4 pixel grid), preserving spatial locality.
- Global Receptive Fields (dANN-GRF): Sampling is restricted per soma, but dendrites belonging to that soma sample from distributed locations around a central point.

The dANN models exhibit mixed selectivity, where nodes respond to multiple classes. This property results in networks that are more robust to noise and overfitting, matching or outperforming dense networks while using orders of magnitude fewer trainable parameters.

H.2 DYNAMIC SPARSE TRAINING (DST)

SET (Mocanu et al., 2018) At each training step, SET removes connections based on weight magnitude and randomly regrows new links.

RigL (Evci et al., 2020) At each training step, RigL removes connections based on weight magnitude and regrows new links based on gradient information.

CHTs (Zhang et al., 2025) Cannistraci-Hebb Training (CHT)(Muscoloni et al., 2022) is a brain-inspired gradient-free link regrowth method. It predicts the existence and the likelihood of each nonobserved link in a network. The rationale is that in complex networks that have a local-community structure, nodes within the same community tend to activate simultaneously ("fire together"). This co-activation encourages them to form new connections among themselves ("wire together") because they are topologically isolated. This isolation, caused by minimizing links to outside the community, creates a barrier that reinforces internal signaling. This strengthened signaling, in turn, promotes the creation of new internal links, facilitating learning and plasticity within the community. CHTs enhances this gradient-free regrowth method by incorporating a soft sampling rule and a node-based link-prediction mechanism.

I ANALYSIS OF THE BOUNDED LAYER BORDER WIRING PATTERN

To investigate the impact of the layer border wiring pattern setting, we compare the default wrap-around topology with the bounded topology. We tested the bounded configuration on the same

1350 image classification tasks as the default, at 99% sparsity using static sparse training (Table 10), SET
 1351 (Table 11), and CHTs (Table 12). Overall, the performance of the bounded model is comparable to
 1352 that of the default wrap-around model, with some variations across datasets and training methods.
 1353 The bounded model outperforms the default on Fashion MNIST and CIFAR10 when using static
 1354 sparse training and SET, while the default has a slight edge on MNIST and EMNIST. For CHTs, the
 1355 results are mixed, with each model excelling in different datasets. These findings suggest that while
 1356 strict locality can be beneficial in certain contexts, the flexibility of wrap-around connections may
 1357 provide advantages in others. The choice of wiring pattern should thus be informed by the specific
 1358 characteristics of the task and dataset at hand.

1359 Table 10: Comparison of wrap-around and bounded topology performances for static sparse training
 1360 of MLPs. The entries represent the accuracy for image classification over different datasets at 99%
 1361 sparsity, averaged over 3 seeds \pm their standard error.

Static Sparse Training				
	MNIST	EMNIST	Fashion MNIST	CIFAR10
Bounded	97.64 \pm 0.10	84.00 \pm 0.06	89.19\pm0.01	61.63\pm0.18
Wrap-around	97.82\pm0.03	84.76\pm0.13	88.47 \pm 0.03	59.04 \pm 0.17

1370 Table 11: Comparison of wrap-around and bounded topology performances for SET initialization
 1371 on MLPs. The entries represent the accuracy for image classification over different datasets at 99%
 1372 sparsity, averaged over 3 seeds \pm their standard error.

SET				
	MNIST	EMNIST	Fashion MNIST	CIFAR10
Bounded	98.40\pm0.02	86.52\pm0.02	89.78\pm0.09	65.67\pm0.18
Wrap-around	98.36 \pm 0.05	86.50 \pm 0.06	89.75 \pm 0.04	64.81 \pm 0.01

1380 Table 12: Comparison of wrap-around and bounded topology performances for CHTs initialization
 1381 on MLPs. The entries represent the accuracy for image classification over different datasets at 99%
 1382 sparsity, averaged over 3 seeds \pm their standard error.

CHTs				
	MNIST	EMNIST	Fashion MNIST	CIFAR10
Bounded	98.66\pm0.03	87.35 \pm 0.00	90.68\pm0.09	68.03 \pm 0.14
Wrap-around	98.62 \pm 0.01	87.40\pm0.04	90.62 \pm 0.16	68.76\pm0.11

J COMPARISON OF DNM WITH OTHER DENDRITIC NETWORKS

1393 Previous work, like that of Chavlis & Poirazi (2025), has explored the integration of dendritic struc-
 1394 tures into artificial neural networks, demonstrating improvements in parameter efficiency and ro-
 1395 bustness. Our Dendritic Network Model (DNM) distinguishes itself through its comprehensive ap-
 1396 proach to modeling dendritic connectivity. While Chavlis & Poirazi (2025) primarily focuses on
 1397 the functional aspects of dendrites within a tree-like multilayered structure (dendritic and somatic
 1398 layers), DNM embeds dendritic properties directly into the bipartite network topology. DNM treats
 1399 dendrites as distinct clusters of links within a bipartite graph, connecting the soma to consecutive
 1400 batches of inputs. This allows the network to inherit dendritic structural advantages through topo-
 1401 logical initialization rather than architectural expansion 1.

1402 Section 4.2 presents a direct comparison between DNM and the best performing dANN model.
 1403 The results indicate that DNM consistently outperforms dANNs across various image classification
 1404 tasks at extreme sparsity levels (99%). These tests were performed by substituting each bipartite

sandwich layer in our original network with dANN’s three-layered subnetwork, ensuring that the total number of trainable parameters remained constant for a fair comparison. We attribute our model’s superior performance to a structural limitation in the dANN framework: the inclusion of dendritic integration layers increases network depth, which in turn forces each bipartite sandwich layer to be significantly sparser. For completeness, we perform additional tests on the network originally proposed by Chavlis & Poirazi (2025). In particular, we take a network instantiated with three sequential dendritic blocks, following the structure:

$$Input \rightarrow (Dendrite \rightarrow Soma) \times 3 \rightarrow Output.$$

We set the width of the dendritic layers to twice the size of the input layer ($N_{dendrite} = 2 \times N_{in}$). Conversely, the somatic integration layers are fixed at a width $N_{soma} = 256$, and the synapses parameter is set to 128. We compare the three models introduced by Chavlis & Poirazi (2025), dANN-R, dANN-LRF, and dANN-GRF, against a modified dANN in which we replace the connectivity to the dendritic layers with a DNM topology that maintains the same sparsity levels. The parameters chosen for the DNM topology were extracted from the best performing configuration found in Section 4.2 for EMNSIT at 99% sparsity. The results, summarized in Table 13, indicate that the DNM-initialized model outperforms dANN, which is the model that we found to perform the best out of those proposed by Chavlis & Poirazi (2025). This further underscores the effectiveness of DNM’s topological approach in enhancing network performance.

Table 13: Comparison of the dANN models (Chavlis & Poirazi, 2025) with a modified version where DNM replaces the connectivity patterns in the dendritic layer (dANN-DNM).

Model	MNIST	EMNIST	Fashion MNIST
dANN-R	98.49 \pm 0.05	85.96 \pm 0.09	89.77 \pm 0.08
dANN-LRF	98.51 \pm 0.00	85.60 \pm 0.15	89.39 \pm 0.07
dANN-GRF	98.53 \pm 0.06	86.26 \pm 0.05	89.77 \pm 0.08
dANN-DNM	98.70\pm0.05	86.64\pm0.07	90.09\pm0.04

K TRANSFERABILITY OF OPTIMAL TOPOLOGIES

A critical question for our generative model of network topology is whether its principles are generalizable across different tasks. To investigate this, we conducted a comprehensive transfer learning experiment to assess if an optimal topology discovered on one task could be effectively applied to others. This tests the hypothesis that the DNM can capture fundamental structural priors beneficial for a class of problems, thereby reducing the need for extensive hyperparameter searches on every new dataset.

Experimental Design We identified the best-performing DNM hyperparameter configuration from the static sparse training experiments for each of the four datasets: MNIST, Fashion MNIST, EMNIST, and CIFAR-10. We then performed a cross-transfer analysis where the optimal configuration for a source dataset was used to initialize MLP models for the other three target datasets. We compared these transferred topologies against baseline initialization methods and against the DNM models specifically tuned for the target task (“DNM (Fine-tuned)”).

Results The results, summarized in Table 14, reveal two distinct trends in topological transferability. First, we observe high transferability among the simpler grayscale datasets. Regardless of whether the parameters are transferred from MNIST, EMNIST, or Fashion MNIST, the resulting DNM models consistently outperform the baseline initialization methods on the other grayscale targets. In these cases, the transferred performance is often comparable to the task-specific fine-tuned models.

However, a significant performance gap emerges when transferring topologies from simpler tasks to the more complex CIFAR-10 dataset. As shown in Table 14, configurations optimized for MNIST, EMNIST, or Fashion MNIST yield poor performance when applied to CIFAR-10 (approximately 47-48% accuracy compared to 58.71% for the fine-tuned model). This suggests that the structural

priors learned from simpler inputs are insufficient for the complexity of natural images. Consequently, for practical scenarios involving image classification with MLPs, we recommend utilizing the parametric configuration derived from CIFAR-10. This configuration acts as a more robust baseline for complex tasks. The specific parameters for this recommended configuration are detailed in Table 15.

Table 14: Performance of transferred DNM topologies on static sparse training at 99% sparsity. We evaluate the transferability of the single best hyperparameter configuration found on each source dataset applied to the other targets. Scores represent accuracy averaged over 3 seeds \pm standard error. Bold values denote the best performance among transfer and baseline methods.

	Static Sparse Training			
	MNIST	Fashion MNIST	EMNIST	CIFAR10
FC	98.80 \pm 0.00	90.87 \pm 0.02	87.08 \pm 0.04	62.35 \pm 0.13
CTSI	98.11 \pm 0.03	88.55 \pm 0.18	84.74 \pm 0.06	52.60 \pm 0.25
SNIP	98.03 \pm 0.03	88.65 \pm 0.07	85.19 \pm 0.04	61.89 \pm 0.48
Random	95.58 \pm 0.03	86.76 \pm 0.05	78.42 \pm 0.26	54.75 \pm 0.15
BSW	97.27 \pm 0.05	87.87 \pm 0.10	82.92 \pm 0.05	56.26 \pm 0.04
BRF	97.28 \pm 0.03	87.78 \pm 0.14	82.88 \pm 0.02	54.86 \pm 0.08
Ramanujan	96.39 \pm 0.10	86.44 \pm 0.14	81.78 \pm 0.08	54.61 \pm 0.32
Radix-Nets	97.06 \pm 0.12	88.02 \pm 0.05	82.65 \pm 0.11	50.90 \pm 0.23
dANN-R	96.10 \pm 0.11	86.52 \pm 0.01	80.64 \pm 0.11	51.57 \pm 0.23
DNM (Fine-tuned)	98.07 \pm 0.09	88.86 \pm 0.21	85.63 \pm 0.10	58.71 \pm 0.28
Transferred from MNIST	98.07 \pm 0.09	88.92 \pm 0.09	85.80 \pm 0.04	48.24 \pm 0.25
Transferred from Fashion MNIST	98.05 \pm 0.09	88.86 \pm 0.21	85.63 \pm 0.10	47.52 \pm 0.39
Transferred from EMNIST	98.05 \pm 0.09	88.86 \pm 0.21	85.63 \pm 0.10	47.52 \pm 0.39
Transferred from CIFAR-10	97.48 \pm 0.38	83.48 \pm 0.02	88.62 \pm 0.11	58.71 \pm 0.28

Discussion This experiment strongly suggests that the structural principles identified by DNM as optimal for MNIST serve as a powerful and generalizable prior for other image classification tasks. The ability to transfer a high-performing topology with minimal performance loss has significant practical implications, as it can drastically reduce the computational cost associated with architecture search for new applications. This finding reinforces the idea that bio-inspired, structured initialization is not merely a task-specific trick but a robust strategy for building efficient sparse networks.

L LIMITATIONS

While the Dendritic Network Model demonstrates significant improvements in accuracy at extreme sparsity levels, we acknowledge a practical limitation regarding current hardware acceleration. Most contemporary GPUs are highly optimized for dense matrix operations; consequently, unstructured sparsity does not always translate into wall-clock training speedups without specialized software support. However, the hardware landscape is rapidly evolving to address this bottleneck. Emerging technologies are specifically designed to efficiently support unstructured sparsity (Lie, 2022). Our method is designed to be future-proof, positioned to fully leverage dynamic sparse training and inference as this specialized hardware becomes widely accessible.

M PRACTICAL GUIDELINES FOR DNM INITIALIZATION

To facilitate the adoption of the Dendritic Network Model (DNM) without the need for extensive hyperparameter search, we provide recommended baseline configurations for MLPs and Transformers. These settings are derived from our most robust experimental results: the CIFAR-10 configuration for image classification and the IWSLT14-to-WMT17 transferred configuration for machine translation.

1512 M.1 MLP FOR IMAGE CLASSIFICATION
1513

1514 For Multi-Layer Perceptrons (MLPs) applied to image classification tasks, particularly those involving
1515 complex visual features, we recommend the configuration detailed in Table 15. This setup was
1516 found to be the most effective for the CIFAR-10 dataset, and proved to stably outperform other
1517 initialization baselines on diverse datasets (K).

1518 Table 15: Recommended DNM parameter configuration for practical image classification scenarios.
1519 This configuration corresponds to the optimal settings identified for CIFAR-10.
1520

1521 Parameter	1522 Value / Distribution
1523 Sparsity (s)	0.99
1524 Mean Dendrites	3
1525 Dendritic Distribution	Fixed
1526 Receptive window width	1.0
1527 Receptive window width distribution	Spatial Inverse-Gaussian
1528 Degree Distribution	Spatial Inverse-Gaussian
1529 Synaptic Distribution	Spatial Inverse-Gaussian
1530 Layer Border Wiring Pattern	Wrap-around

1531 M.2 TRANSFORMERS FOR MACHINE TRANSLATION
1532

1533 For Transformer architectures applied to machine translation, we recommend the configuration that
1534 successfully transferred from IWSLT14 to the large-scale WMT17 dataset (4.3). As shown in Table
1535 16, this configuration uses a higher dendrite count ($M = 7$) compared to MLPs and employs spatial
1536 distributions to organize connectivity.

1537 Table 16: Recommended DNM configuration for Transformers on Machine Translation tasks (de-
1538 rived from IWSLT14 to WMT17 transfer).
1539

1541 Parameter	1542 Value / Setting
1543 Sparsity (s)	0.90
1544 Mean Dendrites	7
1545 Dendritic Distribution	Fixed
1546 Receptive window width	1.0
1547 Receptive window width distribution	Spatial Gaussian
1548 Degree Distribution	Spatial Inverse-Gaussian
1549 Synaptic Distribution	Spatial Inverse-Gaussian
1550 Layer Border Wiring Pattern	Wrap-around

1551 N RELATIONSHIP BETWEEN TOPOLOGY AND PERFORMANCE
1552

1553 In this section, we analyze how different topological properties of DNM-initialized networks cor-
1554 relate with their performance on image classification tasks. We focus on four key graph-theoretic
1555 metrics: powerlaw distribution, average path length, clustering coefficient, and degree distribution.
1556 Understanding these relationships can provide insights into why certain configurations yield better
1557 results.
1558

1559 To investigate the structural drivers of performance in DNM-initialized networks, we analyzed the
1560 correlation between key graph-theoretic metrics and test accuracy across the MNIST, Fashion-
1561 MNIST, EMNIST, and CIFAR-10 datasets. For each dataset, we isolated the top 10, middle 10,
1562 and bottom 10 performing models from our hyperparameter search to visualize how topological
1563 variance dictates model efficacy. The results are summarized in Figure 8.

1564 **Powerlaw gamma** Regarding the degree distribution exponent (γ), we see Fashion-MNIST align-
1565 ing with CIFAR-10. Both datasets exhibit a positive correlation between accuracy and γ ($r = 0.31$

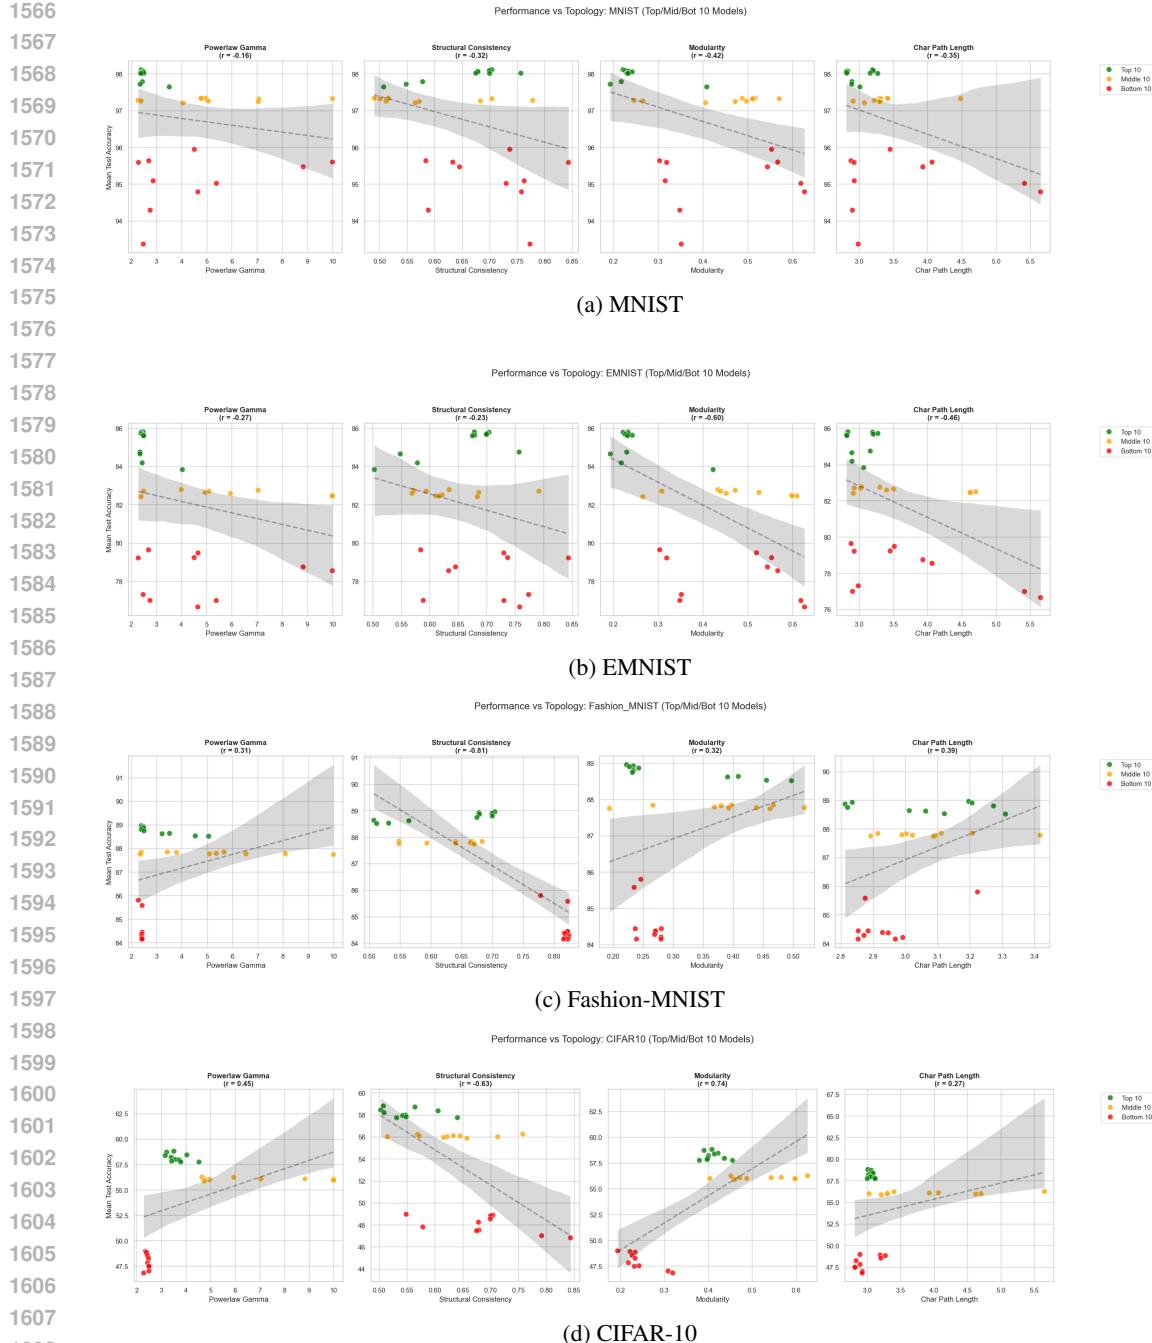


Figure 8: **Performance vs. Topology Analysis.** Scatter plots showing the relationship between test accuracy and four topological metrics: Powerlaw Gamma (γ), Structural Consistency (σ_c), Modularity (Q), and Characteristic Path Length (L). Models are color-coded by performance tier: Top 10 (Green), Middle 10 (Orange), and Bottom 10 (Red). The trend lines and Pearson correlation coefficients (r) highlight task-specific topological preferences.

for Fashion-MNIST, $r = 0.45$ for CIFAR-10). Since a higher γ implies a steeper decay in the degree distribution (fewer extreme hubs), this indicates that complex tasks prefer a more distributed connectivity where information processing is shared among many nodes rather than concentrated in a few central hubs. Conversely, MNIST and EMNIST show negative correlations ($r = -0.16$

1620 and $r = -0.27$), suggesting they benefit from the strong, centralized integration provided by heavy-tailed, hub-dominated topologies.
 1621
 1622

1623 **Structural Consistency** Across all four datasets, we observe a consistent negative correlation
 1624 between Structural Consistency (σ_c) and accuracy. Fashion-MNIST exhibits the strongest negative
 1625 correlation ($r = -0.81$), reinforcing the finding that strict structural predictability is detrimental to
 1626 initialization performance, regardless of the dataset complexity.
 1627

1628 **Modularity** We observe that higher modularity (Q) correlates positively with performance across
 1629 CIFAR-10. The best-performing topologies exhibit high modularity, suggesting that complex natural
 1630 image recognition benefits from distinct, specialized communities of neurons that process local fea-
 1631 tures independently before integration. Conversely, for the simpler MNIST and EMNIST datasets,
 1632 lower modularity appears advantageous, indicating that a more integrated network structure suffices
 1633 for less complex tasks.
 1634

1635 **Characteristic Path** Like CIFAR-10, Fashion-MNIST models show a positive correlation ($r =$
 1636 0.39), indicating a benefit from longer path lengths that preserve local processing. This stands in
 1637 contrast to MNIST and EMNIST ($r < 0$), which favor "small-world" architectures with short global
 1638 integration paths.
 1639

1640 O LEARNING RATE SENSITIVITY

1641 This section analyzes the impact of learning rate variations on model stability (Standard Error)
 1642 and predictive performance (Mean Test Accuracy) across three distinct datasets: Fashion MNIST,
 1643 EMNIST, and MNIST. The plots illustrate the trade-off between convergence speed and stability as
 1644 the learning rate is increased from 0.01 to 0.1 (Figure 9). Across all three datasets, a learning rate
 1645 in the range of 0.025 to 0.05 appears to be the optimal operating window. This range consistently
 1646 provides the best balance of maximizing test accuracy while avoiding the instability and divergence
 1647 (high standard error) associated with rates approaching 0.10.
 1648

1649

1650 P COMPARISON WITH DENSE EQUIVALENT ARCHITECTURES

1651

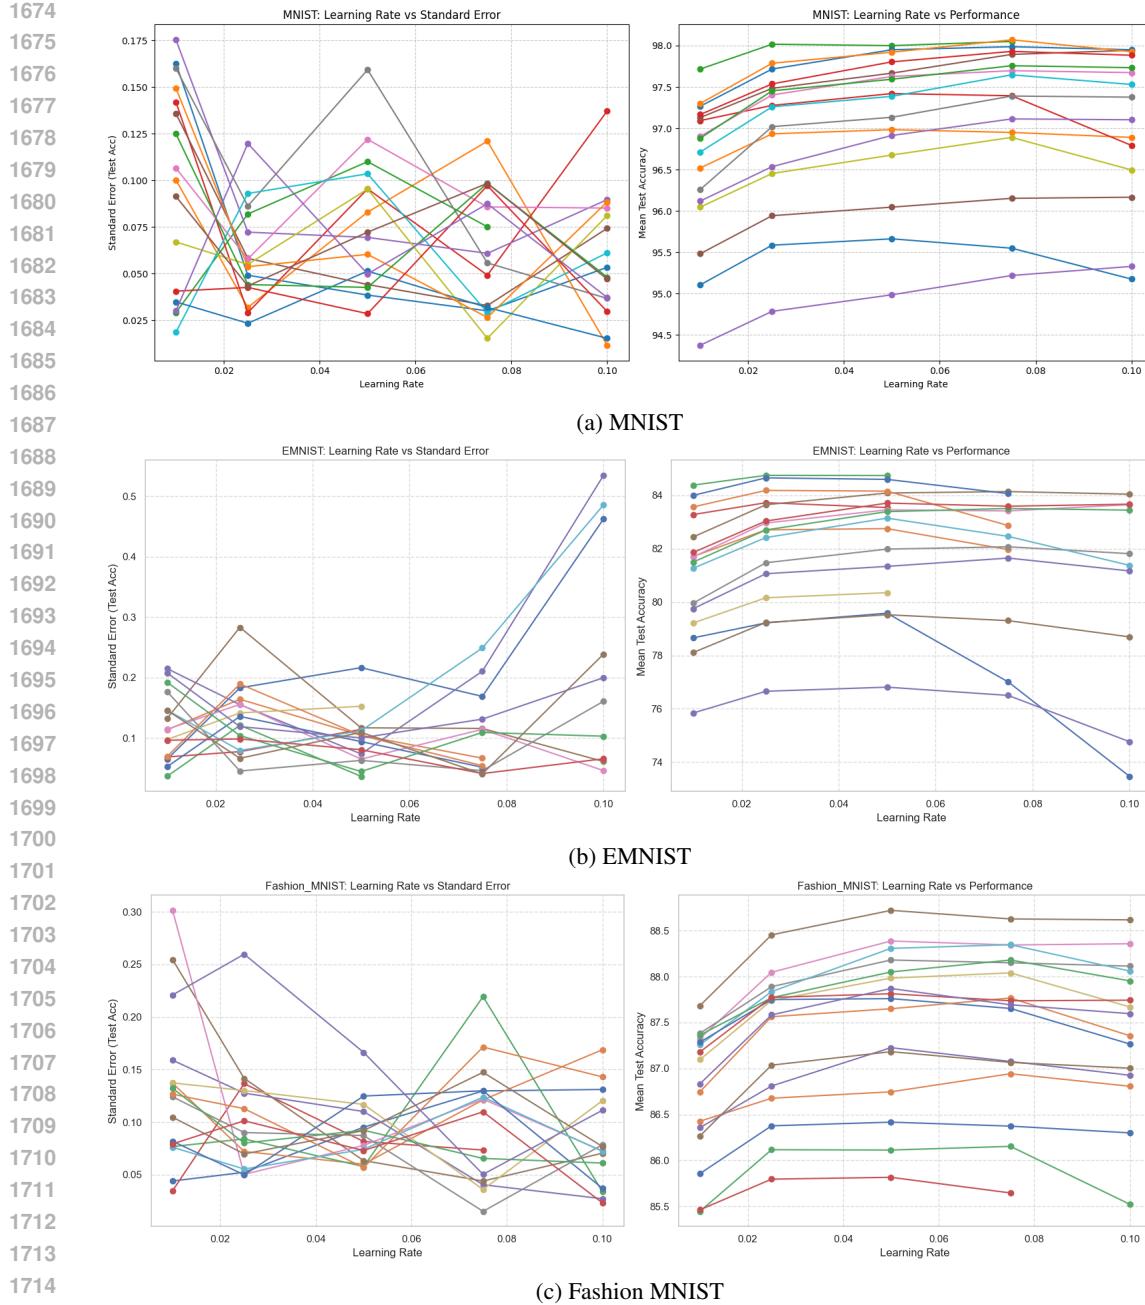
1652 To rigorously validate the structural advantage of the Dendritic Network Model (DNM), it is neces-
 1653 sary to decouple the benefits of network topology from the benefits of parameter efficiency. While
 1654 Table 1 compares DNM against the over-parameterized Fully Connected (FC) network, a more rig-
 1655 orous baseline is the *Dense Equivalent* (DE) architecture.
 1656

1657

1658 **Motivation** At 99% sparsity, a sparse network retains only 1% of the parameters of the original FC
 1659 model. A critical question arises: *Does the sparse topology provide an inductive bias that facilitates*
 1660 *learning, or could a dense network with the same tiny parameter budget perform equally well?* The
 1661 Dense Equivalent model is constructed by reducing the hidden dimension size H of the MLP such
 1662 that the total number of trainable parameters equals the non-zero parameter count of the sparse DNM
 1663 model. If DNM outperforms the DE baseline, it confirms that the *arrangement* of connections (the
 1664 dendritic topology) is superior to a fully connected structure of equal capacity. Conversely, if the
 1665 DE model performs comparably, it would suggest that the performance is governed solely by the
 1666 parameter count, negating the need for complex sparse initialization.
 1667

1668

1669 **Results** Table 17 presents the comparison between the Dense Equivalent baseline and the sparse
 1670 initializations at 99% sparsity for image classification on the CIFAR-10 dataset. To match the ex-
 1671 treme parameter reduction of 99% sparsity, the layer widths of the dense network must be drastically
 1672 reduced. This severely limits the width of the representation that can be propagated through the net-
 1673 work. In contrast, other sparse initialization methods maintain the original wide layer dimensions
 but sparsify the connections. As shown in the results, DNM significantly outperforms the Dense
 Equivalent baseline across all datasets.



1728 Table 17: Comparison of Static Sparse Training vs. Dense Equivalent (DE) architectures on the
 1729 CIFAR-10 dataset. Models utilize the same number of trainable parameters (approx. 1% of the
 1730 original FC network). The DE model uses reduced hidden dimensions to match the parameter count
 1731 of the 99% sparse DNM. Results are averaged over 3 seeds \pm standard error.

Static Sparse Training	
CIFAR10	
FC	62.35 \pm 0.13
CSTI	52.60 \pm 0.25
SNIP	61.89 \pm 0.48
DE	56.37 \pm 0.31
Random	54.75 \pm 0.15
BSW	56.26 \pm 0.04
BRF	54.86 \pm 0.08
Ramanujan	54.61 \pm 0.32
RadiX-Nets	50.90 \pm 0.23
dANN-R	51.57 \pm 0.23
DNM	58.71\pm0.28

1748 fixed a parametric configuration for DNM (all distributions fixed, 3 dendrites) without performing
 1749 any hyperparameter tuning.

1751 **Pattern Recognition Bandit (Granmo, 2018)** We evaluated DNM on a "PatternBandit" task,
 1752 where an agent must detect a specific local signal pattern ([1, 1, 1, 1]) hidden within a noisy state
 1753 vector to receive a reward (Range: -1.0 to 1.0). This task tests the network's ability to maintain local
 1754 feature integration capabilities under sparsity. We compared a Dense baseline against Random and
 1755 DNM initializations using CHTs at 90% and 95% sparsity. Table 18 proves that DNM is capable of
 1756 outperforming both the dense and the randomly initialized sparse networks, both at 90% and 95%
 1757 sparsities.

1758 **Synthetic Node Classification** We applied DNM to the linear transformation layers of a Graph
 1759 Convolutional Network (GCN) (Kipf, 2016) on a synthetic dataset with high informative feature
 1760 density (85% informative features). We tested extreme sparsity regimes (99%, 95%, and 90%) to
 1761 see if DNM prevents topological collapse. The results are shown in Table 19, and show that the
 1762 principles of DNM generalize effectively to non-visual and non-language domains, outperforming
 1763 fully connected networks at extreme sparsity levels.

R DNM'S INTERACTION WITH TRAINING DYNAMICS

1768 We agree that the primary focus of our work is to introduce and validate DNM as a novel sparse
 1769 topological initialization framework. Our central hypothesis is that a bio-inspired, structured initial
 1770 topology provides a significant structural advantage over conventional random or other sparse ini-
 1771 tializations. To test this, our primary investigation focused on the most critical dynamic related to
 1772 our contribution: the dynamics of sparsity itself. To address the lack of exploration of how DNM in-
 1773 teracts with other training dynamics beyond sparsity, we compared our model with a random sparse
 1774 initialization for static sparse training on MLPs for image classification on CIFAR-10. For these
 1775 tests, we selected the best performing parametric configuration for image classification on CIFAR-
 1776 10. DNM demonstrates superior training dynamics compared to random initialization, as shown
 1777 in Figure 10. This figure compares the accuracy, loss, and average gradient norm through training
 1778 epochs of the DNM-initialized network compared to the random one.

1779 **Gradient Norms** In the Random initialization, the gradient norm for deeper layers (e.g., Layer 2,
 1780 represented by the blue line) remains suppressed and stagnant around 0.3–0.4. In contrast, the DNM
 1781 initialization facilitates significantly healthier gradient flow, with Layer 2 norms rising steadily to
 about 0.6. This suggests that the dendritic topology mitigates the vanishing gradient problem in

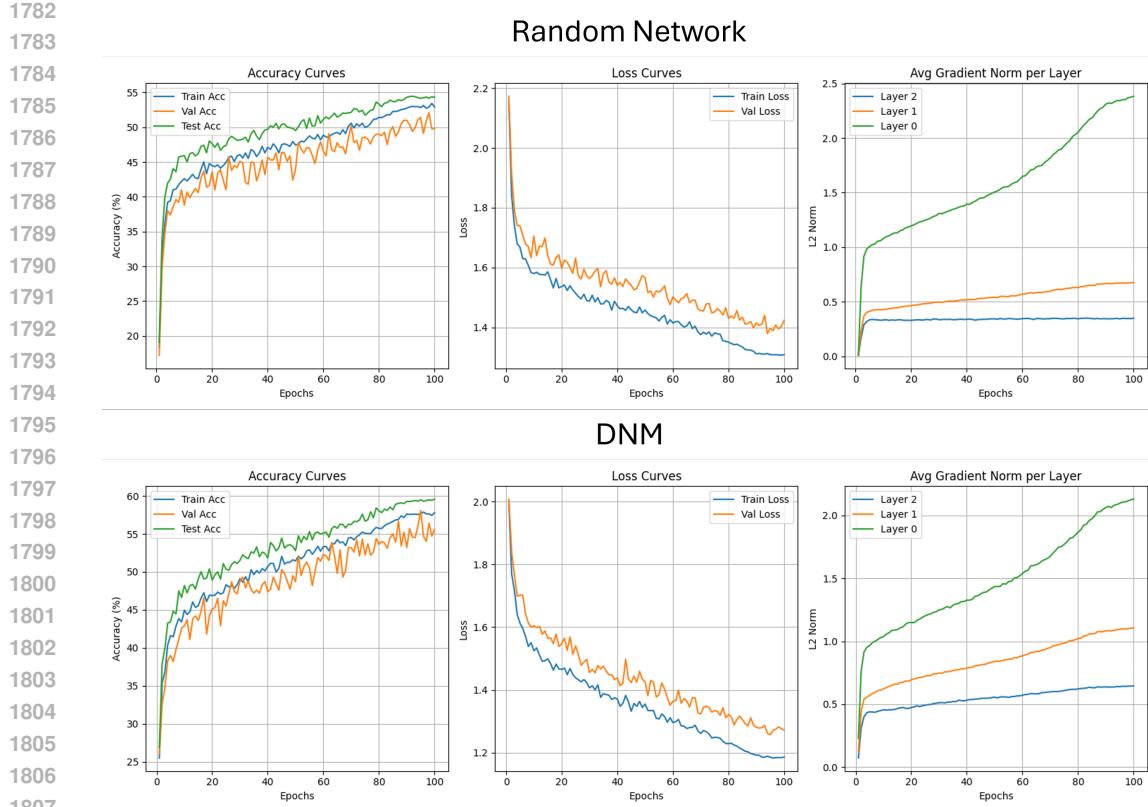


Figure 10: **Comparison of training dynamics between DNM and Random initialization.** The plots illustrate the evolution of test accuracy, training/validation loss, and gradient norms (Layer 2) over 100 epochs on CIFAR-10. DNM exhibits faster convergence, higher final accuracy, and healthier gradient flow compared to random initialization.

sparse networks better than random connections, allowing for more effective weight updates in downstream layers.

Loss and Accuracy The DNM model achieves a lower final training and validation loss compared to the Random model. Consequently, DNM reaches a higher final testing accuracy ($\approx 60\%$) compared to the Random baseline ($\approx 55\%$).

Convergence Speed Analysis To quantify the convergence speed, we utilize the Area Across Epochs (AAE) metric (?). AAE integrates the area between the ideal accuracy (100%) and the learning curve over the training duration E . A lower AAE integrates the area underneath the learning curve over the training epochs E . A higher AAE indicates that the model reached higher accuracy levels faster and sustained them throughout the training process. The DNM model achieves a 53.32 AAE measure, while the random model achieves 49.68. DNM yields a higher AAE by a margin of 3.64, confirming that it converges significantly faster.

S COMBINATION OF DNM WITH PRUNING AND REGULARIZATION TECHNIQUES

DNM performs well as an initialization method of DST frameworks, which inherently involves pruning links. DNM can also be combined effectively with regularization techniques, as proved by the tests on image classification, where models were trained with weight decay, which is one of the most popular forms of L2 regularization.

1836 Table 18: Performance comparison on Mean Reward (averaged over 3 seeds). DNM outperforms
 1837 both dense and random baselines at high sparsity.
 1838

1839	Sparsity	Model Configuration	Mean Reward	vs Random Baseline
1840	95%	Dense (Baseline)	0.71	-
		Random + CHTs	0.65	-
		DNM + CHTs	0.72	+10.77%
1841	90%	Dense (Baseline)	0.71	-
		Random + CHTs	0.65	-
		DNM + CHTs	0.75	+15.39%

1842 Table 19: GCN Synthetic Node Classification Accuracy. Tests performed on extreme sparsity
 1843 regimes to evaluate topological collapse.
 1844

1845	Sparsity	Model Configuration	Mean Accuracy	vs Random Baseline
1846	99%	Dense (Baseline)	0.7544	-
		Random Sparse	0.6433	-
		DNM + CHTs	0.7294	+13.38%
1847	95%	Dense (Baseline)	0.7544	-
		Random Sparse	0.7572	-
		DNM + CHTs	0.7594	+0.29%
1848	90%	Dense (Baseline)	0.7544	-
		Random Sparse	0.7561	-
		DNM + CHTs	0.7628	+0.89%

1863 T REPRODUCTION STATEMENT

1864 All experiments were conducted on NVIDIA A100 80GB GPUs. MLP and Transformer models
 1865 were trained using a single GPU. The code to reproduce the experiments will be made publicly
 1866 available upon publication.
 1867

1869 U CLAIM OF LLM USAGE

1870 The authors declare that Large Language Models (LLMs) were used in the writing process of this
 1871 manuscript. However, the core idea and principles of the article are entirely original and were not
 1872 generated by LLMs.
 1873