

Article

Not peer-reviewed version

The Synaptic Pruning Cliff: Threshold-Like Network Fragility Under Internal Stress and Efficient Recovery in a Computational Model of Depression

[Ngo Cheung](#) *

Posted Date: 13 January 2026

doi: 10.20944/preprints202601.0805.v1

Keywords: MDD; depression; pruning; plasticity; computational biology; threshold; stress



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

The Synaptic Pruning Cliff: Threshold-Like Network Fragility Under Internal Stress and Efficient Recovery in a Computational Model of Depression

Ngo Cheung

MBBS, FHKAM (Psychiatry), Hong Kong SAR, China; info@cheungngomedical.com; Tel.: 98768323

Abstract

Background: Major depressive disorder (MDD) is increasingly viewed through a neuroplasticity lens, with developmental synaptic pruning emerging as a potential core liability. Genetic evidence implicates pruning pathways, while rapid-acting antidepressants like ketamine promote synaptogenesis, suggesting that excessive early elimination leaves circuits vulnerable to later stress. Few computational models, however, capture the specific MDD pattern of latent fragility collapsing under perturbation, followed by recovery via limited plasticity enhancement. **Methods:** An overparameterized feed-forward neural network (~396,000 parameters) was trained on a noisy four-class Gaussian cluster task to represent dense early connectivity. Excessive pruning (95% magnitude-based weight removal, per-layer) simulated adolescent over-elimination. Fragility was assessed under input perturbations and internal neural noise (post-activation Gaussian injections at varying intensities) modeling neuromodulatory disruption. Recovery involved gradient-guided regrowth (50% of pruned connections, prioritized by loss-reduction potential) followed by fine-tuning. Comparisons included random regrowth and a sparsity sweep to identify thresholds. **Results:** The intact network showed robust performance across conditions. Pruning induced sharp collapse (clean accuracy ~51%, standard noisy ~43%), with pronounced sensitivity to internal noise (moderate stress accuracy ~31%) exceeding input noise effects. Gradient-guided regrowth plus fine-tuning restored near-baseline accuracy (clean/standard ~100%) and robustness (combined stress ~97%) despite ~47% persistent sparsity. Targeted regrowth slightly outperformed random under high stress. A critical threshold emerged around 93% sparsity, beyond which combined-stress performance dropped abruptly (>44 percentage points). **Conclusions:** Excessive pruning generates threshold-like intrinsic fragility consistent with stress-triggered MDD relapse, while targeted, limited synaptogenesis efficiently compensates without full density restoration. These findings support a pruning-mediated plasticity deficit as a mechanistic framework for MDD vulnerability and highlight the therapeutic potential of activity-dependent plasticity enhancement. The model provides a testable scaffold for linking polygenic pruning risk to circuit-level decompensation and rapid treatment response.

Keywords: MDD; depression; pruning; plasticity; computational biology; threshold; stress

Introduction

Major depressive disorder (MDD) is one of the most common and disabling psychiatric illnesses. Although twin studies put its heritability near 30–50 percent, genome-wide work shows thousands of small genetic effects that cluster in pathways for neuronal signaling, synaptic upkeep, and immune activity. Developmental synaptic pruning—especially variants in complement component 4 (C4) that boost microglia-mediated synapse removal—has drawn special attention because the same biology also raises risk for schizophrenia [1].

Many researchers now think of MDD less as a monoamine shortage and more as a problem of weakened neuroplasticity. Long-term stress pulls back dendrites, lowers spine numbers, and weakens synapses, changes that mirror the smaller prefrontal and hippocampal volumes seen in

people with depression. Rapid-acting drugs such as ketamine fit this view: within hours they spark bursts of synaptogenesis through brain-derived neurotrophic factor (BDNF) release and mTOR signaling, improving mood even though older structural losses remain [2].

Computational models let us test pruning ideas that would be hard or unethical to study directly. Early simulations focused on schizophrenia-like hallucinations [3]. Later work showed that pruning during adolescence can sharpen but also stiffen recurrent networks [4]. Information-theory studies found that local, activity-based rules keep the connections that matter [5]. Machine-learning projects using sparse subnetworks show that carefully regrowing selected links can restore function [6].

Few models, however, tackle a pattern that looks more like depression: circuits that work fine until stress pushes them past a tipping point, and that can recover with only partial regrowth. The pruning-mediated plasticity-deficit hypothesis argues that too much trimming in adolescence leaves networks thin and fragile; later stress then exposes that weakness, while limited adult plasticity blocks repair [7]. Shared genetic signals with schizophrenia make this plausible.

In the study that follows, we build an over-connected feed-forward network and train it on a noisy classification job to mimic rich early wiring. Heavy pruning represents adolescent over-elimination, creating vulnerability—especially to internal noise meant to model neuromodulatory shifts. We then allow targeted regrowth guided by loss gradients, standing in for BDNF/mTOR-driven synaptogenesis, and ask whether that selective rebuilding restores resilience without returning to the original density.

By capturing both stress-sensitive collapse and efficient rescue at lasting sparsity, the simulation puts the pruning-plasticity idea for MDD to the test. It adds a depression-specific lens—fragility to intrinsic, not just external, disturbance—and offers a framework for tailoring synaptogenic therapies to individuals who carry high pruning risk.

Methods

Computational Environment and Reproducibility

All experiments were coded in Python with PyTorch. NumPy and PyTorch random seeds were fixed at 42. Because some GPU operations are non-deterministic, all runs used CPU only. The project is organised around a single configuration file so that every hyper-parameter can be reproduced or altered systematically. The algorithm was demonstrated in Figure 1.

Task and Dataset

We designed a four-class classification problem that captures the need for robust decision boundaries in noisy settings. Two-dimensional inputs were drawn from four Gaussian clouds centred at $(-3, -3)$, $(3, 3)$, $(-3, 3)$ and $(3, -3)$. The training set contained 12 000 samples with additive Gaussian noise of standard deviation 0.8. Evaluation used three disjoint splits: a standard noisy set (4 000 samples, $\sigma = 0.8$), a clean set (2 000 samples, $\sigma = 0.0$) and a higher-noise set introduced later for stress tests. Seeds 100, 200 and 300 generated the three splits, preventing overlap with training data. Mini-batch sizes were 128 for learning and 1 000 for testing.

Network Architecture

The initial model was intentionally oversized, mirroring exuberant early connectivity. It consisted of an input layer with 2 units, hidden layers of 512, 512 and 256 ReLU units, and a 4-unit softmax output. This layout required roughly 396 500 trainable weights. During some tests, Gaussian noise was injected after hidden activations to simulate internal stress.

Baseline Training

The dense network was trained for 20 epochs with Adam (learning rate 0.001) and cross-entropy loss on the noisy training set. No internal stress was applied while learning the baseline weights.

Pruning Procedure

After baseline training, 95 % of the weights in every layer were removed by magnitude pruning. Within each matrix, weights were ranked by absolute value and the smallest 95 % were set to zero. Masks recorded the pruned locations and remained fixed during later phases. This single-step elimination follows evidence that stronger weights usually mark frequently used synapses [5].

Plasticity Restoration

To mimic drug-induced synaptogenesis, half of the previously pruned connections were reinstated. In the main condition, gradients were accumulated for 30 batches at masked sites; the 50 % with the largest absolute gradients were restored and initialised with small values drawn from a normal distribution ($\sigma = 0.03$). A comparison condition chose the same number of sites uniformly at random. The network was then fine-tuned for 15 epochs with Adam (learning rate 0.0005). Masks were re-applied after each update so overall sparsity stayed constant.

Stress Tests and Evaluation

Performance was measured on: (a) the clean test set, (b) the standard noisy set, and (c) several stress conditions. External stress increased input noise to $\sigma = 1.0$ or 2.0 . Internal stress added Gaussian noise after hidden activations with $\sigma = 0.3$ (mild), 0.5 (moderate), 1.0 (high) or 1.5 (severe). A combined challenge used input $\sigma = 1.0$ and internal $\sigma = 0.5$. Classification accuracy was recorded for each condition. Additional sweeps pruned 0–99 % of weights to locate failure points, and compared gradient-guided versus random regrowth.




Figure 1. Schematic overview of the developmental pruning and plasticity simulation pipeline. The model progresses through five distinct stages: (1) Initial training of a fully connected dense network representing childhood development; (2) Magnitude-based pruning to 95% sparsity, simulating aggressive adolescent synaptic elimination; (3) Stress testing via input perturbation and internal neural noise to evaluate vulnerability; (4) Therapeutic recovery modeling, where gradient accumulation identifies and restores high-utility synapses (analogous to BDNF-dependent plasticity); and (5) Final outcome analysis comparing regrowth strategies and identifying critical sparsity thresholds.

Results

Baseline Performance

After 20 training epochs the fully connected model mastered the four-class task (Table 1). Accuracy reached 100 percent on both the noise-free test set and the standard noisy set ($\sigma = 0.8$). The network was also resilient to stronger perturbations: accuracy fell only to 97.8 percent when input noise was raised to $\sigma = 1.0$ and to 83.6 percent at $\sigma = 2.0$. Injecting internal Gaussian noise after every hidden layer hardly mattered; even the most severe setting ($\sigma = 1.5$) left performance above 99 percent. The intact system therefore combined perfect baseline accuracy with ample safety margins under stress while employing 396 548 trainable parameters (zero sparsity).

Consequences of 95 Percent Pruning

Removing the smallest 95 percent of weights cut the active parameter count to about 21 050 (94.7 percent sparsity) and revealed a dramatic brittleness. Clean-set accuracy dropped to 50.8 percent and fell further on the standard noisy set (43.1 percent). Extra input noise now produced pronounced failures: 41.2 percent accuracy at $\sigma = 1.0$ and 39.6 percent at $\sigma = 2.0$. Internal noise was even more damaging; moderate disruption ($\sigma = 0.5$) reduced accuracy to 31.2 percent, and high disruption ($\sigma = 1.0$) to 29.0 percent. A combined challenge (input $\sigma = 1.0$ plus internal $\sigma = 0.5$) yielded 32.0 percent. Hence excessive pruning did not just lower overall capacity – it erased the network's tolerance to variability.

Recovery Through Gradient-Guided Regrowth

We next reinstated one-half of the pruned connections (about 187 700 weights) at small random strengths, choosing sites with the largest accumulated gradients. After 15 fine-tuning epochs the network, still 47.3 percent sparse, almost fully regained its earlier abilities. Accuracy returned to 100 percent on clean data and 99.9 percent on the standard noisy set. Robustness also rebounded: 97.3 percent at input $\sigma = 1.0$ and 84.0 percent at $\sigma = 2.0$. Internal stress accuracies climbed to 99.8 percent for $\sigma = 0.5$, 99.0 percent for $\sigma = 1.0$, and 95.0 percent for $\sigma = 1.5$. Under the combined challenge the model reached 96.9 percent. Thus limited, targeted synaptogenesis restored function without returning to the original density.

Table 1
Performance Metrics Across Developmental Stages: Baseline, Pruned, and Recovered Networks

Metric / Condition	Baseline (0% Sparse)	Pruned (94.7% Sparse)	Recovered (47.3% Sparse)
Standard Performance			
Clean Accuracy	100.0%	50.8%	100.0%
Standard Accuracy	100.0%	43.1%	99.9%
Input Perturbation Resilience			
Input Noise (+1.0)	97.8%	41.2%	97.3%
Input Noise (+2.0)	83.6%	39.6%	84.0%
Internal Neural Stress (Noise)			
Mild ($\sigma=0.3$)	100.0%	35.5%	99.8%
Moderate ($\sigma=0.5$)	100.0%	31.2%	99.8%
High ($\sigma=1.0$)	99.9%	29.0%	99.0%
Severe ($\sigma=1.5$)	99.9%	30.4%	95.0%
Composite Stress			
Combined (Input=1.0, Int=0.5)	98.0%	32.0%	96.9%

Note. The Pruned stage represents the fragile state following excessive adolescent synaptic elimination. The Recovered stage utilizes gradient-guided plasticity to restore function despite retaining 47.3% sparsity.

Random Versus Targeted Regrowth

A control experiment repeated the same 50 percent reinstatement but selected sites uniformly at random. Clean and standard accuracies again reached 100 percent, and moderate stresses were handled equally well. At higher load the gradient-guided method held a narrow edge (99.4 percent versus 98.8 percent at high internal noise), suggesting that biologically informed targeting confers a small but measurable benefit.

Sparsity Threshold

Varying the pruning fraction revealed a sharp tipping point. Performance stayed above 96 percent (combined stress) until sparsity exceeded roughly 90 percent; between 90 and 93 percent accuracy collapsed by more than 40 points, and beyond 95 percent it hovered near chance (25–31 percent). The model therefore requires a minimum synaptic density to preserve reliable computation, and dropping below that critical level triggers a sudden loss of reserve (Table 2).

Table 2
Critical Threshold Analysis: Performance Collapse at High Sparsity Levels

Sparsity Level	Clean Accuracy	Standard Accuracy	Stress Resilience (Moderate)	Combined Stress
0%	100.0%	100.0%	100.0%	97.8%
50%	100.0%	100.0%	100.0%	97.5%
70%	100.0%	99.9%	99.9%	97.3%
80%	100.0%	99.8%	99.6%	96.5%
90%	100.0%	99.8%	83.5%	79.3%
93%	76.0%	57.9%	34.6%	35.1%
95%	76.0%	67.4%	30.9%	31.5%
97%	23.9%	25.0%	30.6%	29.7%
99%	23.9%	25.0%	25.4%	24.5%

Note. A critical "cliff" in performance is observed between 90% and 93% sparsity, particularly under stress conditions, indicating the loss of computational reserve required for robust processing.

Discussion

Interpretation of Results

Our updated simulation clarifies how severe synaptic pruning may set the stage for major depressive disorder (MDD) and how stress can then tip the system into failure (Figure 2). When the model kept most of its original connections, it stayed accurate even when inputs were noisy or when we added internal "neural" noise. This resilience resembles the stable thinking and mood of healthy people. After we removed 95 % of the synapses, however, the network's accuracy crashed to chance levels. Internal noise—used here as a stand-in for stress-related changes such as altered HPA-axis activity or inflammation—was particularly damaging, dropping accuracy below 32 %. Clinically, this mirrors how many patients with depression feel slowed or unsettled even in calm settings.

A key finding was how well the system recovered after we let only half of the lost synapses grow back in a targeted, gradient-guided way. Even with roughly 47 % sparsity, performance returned to baseline in easy conditions and to 95–99 % of baseline under stress. Targeted regrowth worked slightly better than random regrowth when stress was highest, implying that where new synapses form matters more than how many return.

We also saw a sharp tipping point: once pruning exceeded about 93 %, performance under stress plunged by more than 44 %. Below that level, the network coped well; above it, accuracy fell rapidly. Such nonlinear behavior suggests that modest genetic or environmental pushes past a critical density could trigger sudden clinical vulnerability, a pattern consistent with the polygenic architecture of MDD.

New Insights from the Improved Model

Earlier pruning studies often focused on external noise or on schizophrenia. By adding internal noise after hidden-layer activations, we captured a different, depression-relevant weakness: networks that were over-pruned fell apart even on clean inputs, whereas dense or recovered networks stayed stable. This echoes evidence that MDD involves a reduced signal-to-noise ratio inside prefrontal and limbic circuits [2].




Figure 2. The Two-Hit Model of Synaptic Pruning and Depression. A schematic representation of the simulation results demonstrating the progression from a healthy neural state to Major Depressive Disorder (MDD) and subsequent recovery. Phase 1 depicts a healthy, dense network with high resilience to noise. Phase 2 illustrates the "First Hit," where developmental pruning exceeds a critical tipping point (approx. 93%), resulting in a fragile sparse network. Phase 3 shows the "Second Hit," where internal noise—modeling HPA-axis stress or inflammation—causes a catastrophic drop in performance, mirroring clinical MDD symptoms. Phase 4 and Phase 5 demonstrate therapeutic rescue via Ketamine-induced plasticity. The model indicates that targeted, gradient-guided regrowth restores function to near-baseline levels more effectively than random regrowth.

Second, the benefit of gradient-guided regrowth bolsters the notion that activity-dependent mechanisms—such as BDNF release and mTOR-related protein synthesis—direct new spines to the most advantageous locations [8]. Random regrowth is easier to code, but it's not as likely to happen in real life. In vivo, ketamine-induced synaptogenesis happens near active synapses, not at random.

Implications for the Pruning-Mediated Plasticity Deficit Hypothesis

These findings reinforce the perspective that excessive developmental pruning constitutes an initial "hit," while constrained adult plasticity serves as a subsequent "hit" in response to stress [7]. The tipping-point behavior aligns with genetic research that associates complement and microglial pruning pathways with the risk of major depressive disorder (MDD) (Figure 3). Our successful rescue with partial, targeted regrowth is similar to how quickly ketamine works in the clinic; it restores function in hours—too quickly for full regrowth—by causing bursts of spine formation and homeostatic scaling [9]. The fact that recovery happened even with only half the original number of synapses shows that treatments can make up for permanent developmental losses.

Practically, patients with high genetic loading in pruning pathways might benefit most from treatments that reopen plasticity windows. The model's sensitivity to internal noise also highlights intrinsic circuit instability as a core feature of depression, distinct from sensory processing deficits seen in other disorders.

In summary, our simulation supports a two-hit framework: developmental over-pruning creates a fragile network, and later stress exposes that fragility. Pharmacologically enhancing targeted plasticity can largely reverse the damage, underscoring the promise of neuroplasticity-centered treatments and of integrated computational, genetic, and imaging work to refine this hypothesis.




Figure 3. The Pruning-Mediated Plasticity Deficit Hypothesis. A conceptual framework illustrating the "Two-Hit" model of depression. Hypothesis Core Framework: The model posits that excessive developmental pruning

(Hit 1) creates a latent vulnerability, which is unmasked by stress and constrained plasticity in adulthood (Hit 2). Biological Mechanisms: Genetic loading in complement and microglial pathways drives the initial over-pruning, resulting in intrinsic circuit instability and heightened sensitivity to internal noise. Clinical Intervention: Treatments like Ketamine function by reopening plasticity windows. This induces bursts of spine formation and homeostatic scaling, allowing for rapid functional rescue via targeted regrowth, effectively compensating for the permanent developmental synaptic loss.

Novelty of the Present Model

Our simulation adds several fresh angles to earlier work on synaptic pruning and mood disorders. Previous models often studied pruning as it relates to developmental efficiency or to schizophrenia-like hyper-activity. Here, we aimed instead at a depression-relevant pattern: a network that works well in quiet conditions but falls apart when stressed. We did this by adding "internal" noise after each hidden-layer activation, treating it as a stand-in for neuromodulatory changes that accompany stress in major depressive disorder (MDD). That choice shifts attention from external sensory noise toward the brain's own fluctuating state, which is closer to how many patients experience sudden dips in mood or cognition without obvious triggers.

A second point of novelty is how we let the network recover. Rather than growing synapses at random, we restored half of the lost connections that carried the largest gradients—an abstract proxy for activity-dependent processes such as BDNF/mTOR signaling that follow rapid-acting antidepressant treatment. The method produced only a slight edge over random regrowth, yet the consistent benefit under heavy stress hints that where new synapses form can matter more than sheer numbers. To our knowledge, few pruning studies have made that direct comparison.

Finally, by mapping accuracy across a full sparsity sweep, we found a sharp tipping point near 93 % loss. Performance stayed strong below that value but slipped fast once the threshold was crossed. Such non-linear behavior fits the idea that modest genetic or environmental pushes can flip a resilient brain into vulnerability.

Potential Impact and Translational Implications

If similar dynamics appear in more realistic networks, they could help make sense of clinical observations. The steep threshold suggests that pruning-related genetic risk might show up only after synapse loss passes a critical line, offering one explanation for the mixed penetrance of risk alleles in MDD. The rapid rescue we saw—despite the network remaining about 47 % sparse—mirrors how ketamine can lift mood within hours without fully restoring adolescent-level connectivity. That finding supports treatments that aim for targeted, not wholesale, synaptogenesis.

The model's special sensitivity to internal noise also points to possible biomarkers. Measures of resting-state signal-to-noise, whether from EEG or fMRI, might flag people whose circuits sit close to the pruning threshold and who could benefit most from plasticity-enhancing drugs. In the long run, preventive efforts that soften adolescent pruning, perhaps through anti-inflammatory or stress-buffering strategies, might keep vulnerable brains safely below that critical point.

Limitations

Several shortcuts limit how tightly we can link the simulation to biology. The four-class Gaussian task is a far cry from real-world emotional regulation. We used a feed-forward network, so it lacks the feedback loops that could model rumination or persistent negative bias. Our pruning method removes weights strictly by magnitude and ignores biological cues like complement tagging or microglial action. Likewise, gradient-guided regrowth captures some spirit of BDNF-driven plasticity but skips details such as local dendritic signaling or astrocytic support. Stress entered the model simply as additive Gaussian noise, omitting slow hormonal changes or inflammatory signaling. Finally, we treated pruning in isolation, without layering on other risks such as chronic stress or genetic diversity.

Conclusions

Despite those gaps, the work offers clear computational support for a pruning-based vulnerability in depression. Excessive early synapse loss can leave a network that looks fine until stress pushes it over the edge—much like stress-triggered episodes in MDD. Yet a modest, well-placed burst of synaptogenesis can restore stability, echoing the fast relief seen with ketamine and related drugs. By highlighting a sharp threshold and the value of targeted plasticity, the model encourages both prevention—keeping pruning below the danger line—and personalized rescue—reopening plasticity in those who need it. Future studies should move toward recurrent, multitask networks, include glial and inflammatory factors, and pair modeling with genetic and imaging data to test these ideas more directly.

Funding Declaration: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. .

Ethics Declaration: Not applicable. .

Conflicts of Interest: None declared.

References

1. Sekar A, Bialas AR, de Rivera H, et al. Schizophrenia risk from complex variation of complement component 4. *Nature*. 2016;530(7589):177–183. <https://doi.org/10.1038/nature16549>
2. Duman RS, Aghajanian GK. Synaptic dysfunction in depression: Potential therapeutic targets. *Science*. 2012;338(6103):68–72. <https://doi.org/10.1126/science.1222939>
3. Hoffman RE, Dobscha SK. Cortical pruning and the development of schizophrenia: A computer model. *Schizophrenia Bulletin*. 1989;15(3):477–490. <https://doi.org/10.1093/schbul/15.3.477>
4. Averbek BB. Pruning recurrent neural networks replicates adolescent changes in working memory and reinforcement learning. *Proceedings of the National Academy of Sciences*. 2022;119(22):e2121331119. <https://doi.org/10.1073/pnas.2121331119>
5. Scholl C, Rule ME, Hennig MH. The information theory of developmental pruning: Optimizing global network architectures using local synaptic rules. *PLoS Computational Biology*. 2021;17(10):e1009458. <https://doi.org/10.1371/journal.pcbi.1009458>
6. Liu S, Chen T, Chen X, et al. Sparse training via boosting pruning plasticity with neuroregeneration. In *Advances in Neural Information Processing Systems* (Vol. 34, pp. 9908–9922). <https://doi.org/10.48550/arXiv.2103.01600>
7. Cheung N. From Pruning to Plasticity: Refining the Etiological Architecture of Major Depressive Disorder Through Causal and Polygenic Inference. *Preprints*. <https://doi.org/10.20944/preprints202601.0601.v1>
8. Monteggia LM, Kavalali ET. Synaptic basis of rapid antidepressant action. *European Archives of Psychiatry and Clinical Neuroscience*. 2024;275:1539–1546. <https://doi.org/10.1007/s00406-024-01898-6>
9. Moda-Sava RN, Murdock MH, Parekh PK, et al. Sustained rescue of prefrontal circuit dysfunction by antidepressant-induced spine formation. *Science*. 2019;364(6436):eaat8078. <https://doi.org/10.1126/science.aat8078>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.