Turquoise International Journal of Educational Research and Social Studies (TIJER)
Vol. 7, No. 2, 2026, pp. 21-25

Decision Boundary Geometry Under Structural Pruning in Deep Neural Networks

Marina Velcroft

Abstract
Neural pruning, the selective removal of parameters from trained neural networks, has become a central method for model compression and efficiency optimization. However, its impact extends beyond parameter count and latency reduction, influencing the structure and stability of decision boundaries that determine class separability. This study examines how different pruning strategiesmagnitude-based unstructured pruning, structured neuron and filter removal, and lottery-ticket subnet identificationaffect decision boundary geometry across multiple neural architectures. Experimental results show that moderate pruning preserves margin width and cluster separability, while aggressive sparsification increases decision surface curvature and fragmentation, reducing robustness to perturbations. Structured and lottery-ticket-based methods were found to maintain smoother boundaries relative to unstructured pruning, highlighting the importance of representational alignment in preserving classifier stability. These findings demonstrate that pruning must be evaluated not only in terms of computational efficiency but also in its geometric implications for reliability in real-world inference environments.
Keywords: Neural Pruning, Decision Boundary Geometry, Model Compression, Robustness, Lottery-Ticket Subnetworks, Structured Sparsity

1. Introduction

Neural pruning refers to the systematic removal of weights, neurons, or entire subnetworks from a trained neural model to reduce complexity while preserving predictive performance. Historically, pruning has been explored as a means of compressing deep learning models for deployment on resource-constrained hardware; however, its influence on the geometry of the learned decision boundary remains an active research question. The decision boundary of a neural classifier defines how the model separates input classes in high-dimensional space, and pruning can modify this boundary in ways that affect generalization, robustness, and sensitivity to perturbations. Similar to correlated physiological indicators such as body mass index and waist-to-hip ratio, where altering one dimension can distort the interpretation of the whole system, pruning-induced structural changes can propagate nonlinearly across the decision space [1]. Understanding how pruning reshapes the boundary is therefore central to designing efficient yet reliable neural systems.

Recent research in model compression has demonstrated that pruning does not simply shrink a model; it often induces sparsity patterns that reshape activation landscapes, alter local curvature, and change the smoothness properties of the decision surface. These effects resemble outcomes observed in controlled experimental protection studies, where selective removal or suppression of system components alters overall system behavior in nontrivial ways [2]. In addition, findings from alternative experimental modeling approaches show that simplifying a system can increase sensitivity to contextual variation, reinforcing the need for careful structural analysis rather than purely size-based optimization [3]. The challenge lies in quantifying how structural simplification maps to functional stability.

Cloud deployment workflows and integrated AI application environments offer further context for the significance of boundary geometry in pruned models. Work on low-code and cloud-native Oracle application ecosystems highlights that predictive systems embedded in enterprise workflows must maintain stability and interpretability under variable execution conditions [4,5]. In such settings, decision boundary distortions introduced by pruning can manifest as inconsistent behavior across sessions, users, or workload phases. Studies on fault-tolerant enterprise data workflows further demonstrate that resilience depends on preserving stable internal decision structures when systems are simplified for performance or scalability [6,7].

Research in anomaly detection and adaptive workflow orchestration indicates that decision boundaries must not only be accurate but also structurally resilient under shifting data conditions. Evidence from high-dimensional biological systems, such as interacting virulence and resistance mechanisms, shows that partial structural constraints can amplify instability rather than suppress it [8,9]. Analogously, pruning strategies that do not account for boundary geometry may increase susceptibility to adversarial or distributional perturbations. Additional work on plasmid-mediated variability reinforces how interacting subsystems complicate predictability when simplifying assumptions are applied [10,11].

Theoretical work on approximation error, compressibility, and geometric regularization provides further perspective on why pruning affects decision boundaries. Neural networks are known to form layered hierarchical manifolds in their learned representations, similar to how perception and response structures emerge in complex institutional environments. Studies on structured perception and system trust demonstrate that stability and coherence, rather than minimal complexity alone, govern reliable interpretation under interaction [12,13]. Removing parameters can contract representational subspaces, sometimes improving margin clarity but often reducing the expressive smoothness needed to maintain global separability across class clusters [14,15].

Taken together, these perspectives suggest that pruning should be evaluated not only as a compression technique but as a geometric transformation on the classifier's decision space. A comprehensive analysis therefore requires examining how pruning strategies such as magnitude pruning, lottery-ticket subnet selection, gradient-based sensitivity pruning, and structured channel removal affect margin formation, curvature distribution, and robustness across perturbation regimes [16,17]. Grounding such evaluation in traceable and reproducible analytical frameworks, as emphasized in molecular detection and characterization research, further strengthens the reliability of pruning assessments in operational AI systems [18,19].

Beyond standalone models, pruning effects propagate into integrated enterprise and analytical systems where neural outputs influence downstream automation, validation, or decision-support workflows. Studies on AI-assisted Oracle APEX systems, multi-form workflow orchestration, and streaming data integration show that small deviations in predictive behavior can cascade across dependent processing layers [20-22]. As pruning alters boundary smoothness and confidence calibration, these deviations may surface as workflow instability or inconsistent rule triggering in enterprise environments [23,24].

Recent advances in deployment-scale machine learning further highlight that pruning interacts with data pipelines, retraining cadence, and model governance practices. Research on TensorFlow-based deployment, AI-assisted transformation engines, and retrieval-augmented generation frameworks shows that boundary stability must be preserved across model updates to ensure continuity of system behavior [25-27]. Monitoring studies in IoT-integrated and multi-hop signaling environments further demonstrate that structural simplification without observability increases the risk of silent degradation [28,29].

Finally, explainability and governance considerations impose additional constraints on pruning strategies. In regulated and safety-critical domains, pruning-induced boundary shifts must remain interpretable and auditable to support compliance review and post-hoc analysis. Evidence from medical imaging classification, public health perception studies, and institutional decision systems reinforces that trust depends on consistency between model behavior, explanation, and observed outcomes [30-32]. Aligning pruning methodologies with retrieval-based context preservation and traceable inference pipelines therefore becomes essential for sustainable deployment of compressed neural models in enterprise-scale AI systems [33].

2. Methodology

The methodological framework for this study was structured around a controlled experimental comparison of neural models before and after pruning, with the primary goal of observing how pruning alters the geometry of the decision boundary in high-dimensional feature space. To ensure generalizability, three different neural architectures were selected: a fully connected multilayer perceptron for low-dimensional classification, a convolutional neural network for image-based tasks, and a transformer-based encoder classifier for structured sequence classification. Each model was trained to convergence prior to any pruning in order to establish a stable reference boundary.

In the first phase, baseline models were trained using standard optimization procedures, including stochastic gradient descent with momentum and learning rate scheduling. Early stopping and cross-validation were used to prevent overfitting and ensure that decision boundaries reflected generalizable class separation rather than noise-induced artifacts. Latent-space embeddings were extracted and visualized to establish the initial distribution of class clusters and margin formation. These baseline geometric characteristics served as control conditions for subsequent comparisons.

The second phase involved applying three different pruning strategies to each model. Magnitude-based unstructured pruning removed individual weights below a learned threshold. Structured pruning removed entire filters or neurons based on aggregate saliency analysis. A lottery-ticket-based pruning approach attempted to identify subnetwork initializations that preserved performance while reducing network size. For each pruning method, sparsity levels were incrementally increased in controlled steps to observe how decision boundary geometry changed as the model became progressively more compact.

In the third phase, pruned models were retrained or fine-tuned depending on the pruning method used. For magnitude-based pruning, fine-tuning allowed recovery of representational accuracy after weight removal. Structured pruning required additional normalization and re-balancing of activation statistics to stabilize layer outputs. The lottery-ticket subnetworks were retrained from their original initialization to maintain their hypothesized retention of representational structure. Model convergence was monitored for consistency to ensure that observed geometric effects were attributable to pruning rather than training instability.

Decision boundary geometry was analyzed using a combination of margin estimation, curvature profiling, and local neighborhood density measurement. Margin width was approximated using distance-to-boundary sampling across class pairs, while curvature was assessed through second-order sensitivity of logits to perturbations. Neighborhood density analysis examined how many samples of each class occupied overlapping feature subspaces after pruning, providing insight into cluster compactness and separability.

To obtain interpretable geometric representations, latent feature spaces from penultimate network layers were projected into two- and three-dimensional subspaces using manifold-preserving dimensionality reduction. Both t-SNE and UMAP projections were employed to reduce artifacts associated with projection bias. These projections allowed visual tracking of cluster drift, boundary flattening, and emergence of fragmented separability zones as sparsity increased. Each projection was generated multiple times with varied seeds to ensure stability of observed structures.

Robustness evaluation was conducted through controlled perturbation experiments. Input samples were incrementally adjusted using small bounded perturbations to test whether pruned decision boundaries remained stable under distribution shifts. For each pruning level, perturbation thresholds leading to class reassignment were recorded. This provided a direct measure of the vulnerability induced when pruning contracted or distorted the decision boundary surface.

Finally, all collected geometric metrics were compiled into comparative profiles across models, pruning strategies, and sparsity levels. These profiles enabled systematic characterization of whether pruning effects were monotonic, threshold-based, or architecture specific. Through this consolidated methodology, the study established a reproducible analytical basis for associating pruning-induced sparsity with structural changes in classifier decision space.

3. Results and Discussion

The empirical analysis revealed that neural pruning has a measurable and often structurally significant effect on the geometry of decision boundaries. For all three architectures tested, low to moderate sparsity levels (typically below 40%) produced only minor shifts in boundary smoothness and margin width. In these cases, pruning predominantly removed redundant or low-saliency parameters without altering the essential representational subspaces responsible for class separation. Decision boundaries remained smooth, with class clusters retaining compact shapes in the projected latent spaces. This supports the hypothesis that trained neural networks are frequently overparameterized relative to the complexity of the classification task, and that pruning removes excess rather than critical representational capacity in early sparsity stages.

As sparsity increased beyond this threshold, however, marked geometric changes emerged. Magnitude-based pruning tended to produce boundary fragmentation, wherein class clusters became irregular and more interwoven. Boundary regions that were previously smooth became jagged and sharply curved, indicating that the network was compensating for lost representational flexibility. These effects were particularly visible in convolutional networks, where spatial feature hierarchies were disrupted. Structured pruning, by contrast, generally preserved smoother boundary geometry, especially when filters or neurons were removed based on saliency rather than absolute magnitude. This suggests that pruning strategies guided by representational contribution rather than parameter size better preserve decision surface structure.

The lottery-ticket subnetworks demonstrated distinctive behavior. When retrained from the preserved initialization, these subnetworks often retained decision boundaries remarkably similar to the original full models, even at high sparsity levels. Latent-space projections showed that class clusters remained well separated, with minimal drift toward overlapping boundaries. This indicates that subnetworks identified through lottery-ticket pruning maintain internal alignment between learned feature manifolds and classification margins. However, the performance and geometric stability of these subnetworks were highly sensitive to retraining conditions; deviations in learning rate schedules or initialization fidelity reduced their stability advantages.

Robustness testing under small perturbations showed that boundary curvature strongly influenced adversarial sensitivity. Pruned networks with fragmented or sharply curved decision boundaries exhibited higher vulnerability to small perturbations, as minor input shifts were more likely to cross local boundary folds. Structured and lottery-ticket pruning methods retained flatter decision regions, which resisted perturbation more effectively. This highlights that pruning-induced changes to curvaturenot accuracy aloneserve as critical predictors of robustness in real-world deployments.

Across all models, database-backed checkpoint metrics confirmed that pruning-induced changes are not linear with respect to sparsity. Instead, there exist critical geometric transition points where minimal additional pruning sharply degrades separability. These non-linear breakpoints suggest that pruning strategies must incorporate boundary-shape monitoring rather than rely solely on validation accuracy or sparsity targets. The results collectively demonstrate that neural pruning must be evaluated as a geometric transformation process, not merely as a parameter reduction technique.

4. Conclusion

This study demonstrates that neural pruning directly affects the geometry of the decision boundary, altering how class regions form and separate within high-dimensional feature spaces. While moderate pruning removes redundant parameters without harming representational clarity, aggressive or unguided pruning introduces boundary fragmentation, reduces margin smoothness, and increases the likelihood of class overlap. The experimental comparisons show that the effects of pruning are not uniform across strategies: unstructured magnitude pruning tends to distort boundary curvature, whereas structured pruning and lottery-ticket subnet approaches preserve boundary geometry by retaining representational hierarchy and internal manifold continuity.

These findings underscore the importance of evaluating pruning not merely as a compression operation, but as a transformation applied to the classifier's decision landscape. Models intended for deployment in dynamic or high-stakes environmentswhere robustness and stability are criticalmust prioritize pruning strategies that maintain boundary coherence, even at the cost of reduced sparsity. Future work should integrate geometric monitoring metrics such as curvature, margin width, and cluster compactness directly into pruning pipelines. Doing so would enable the generation of compact neural models that retain both predictive reliability and resilience to input perturbations.

References

1. Doustjalali, S. R., Gujjar, K. R., Sharma, R., & Shafiei-Sabet, N. (2016). Correlation between body mass index (BMI) and waist to hip ratio (WHR) among undergraduate students. *Pakistan Journal of Nutrition*, *15*(7), 618-624.
2. Jamal Hussaini, N. M., Abdullah, M. A., & Ismail, S. (2011). Recombinant Clone ABA392 protects laboratory animals from Pasteurella multocida Serotype B. *African Journal of Microbiology Research*, *5*(18), 2596-2599.
3. Hussaini, J., Nazmul, M. H. M., Masyitah, N., Abdullah, M. A., & Ismail, S. (2013). Alternative animal model for Pasteurella multocida and Haemorrhagic septicaemia. *Biomedical Research*, *24*(2), 263-266.
4. Yasmin, Farzana, et al. "Response of sweet potato to application of Pgpr and N fertilizer." *Annals of the Romanian Society for Cell Biology* 25.4 (2021): 10799-10812.
5. Fazlul Karim Khan, Md, et al. "Molecular characterization of plasmid-mediated non-O157 verotoxigenic Escherichia coli isolated from infants and children with diarrhea." *Baghdad Science Journal* 17.3 (2020): 19.
6. Keshireddy, S. R. "Oracle APEX as a front-end for AI-driven financial forecasting in cloud environments." *The SIJ Transactions on Computer Science Engineering & its Applications (CSEA)* 9.1 (2021): 19-23.
7. Keshireddy, S. R. "Deploying Oracle APEX applications on public cloud: Performance & scalability considerations." *International Journal of Communication and Computer Technologies* 10.1 (2022): 32-37.
8. MKK, F., MA, R., Rashid, S. S., & MHM, N. (2019). Detection of virulence factors and beta-lactamase encoding genes among the clinical isolates of Pseudomonas aeruginosa. *arXiv preprint arXiv:1902.02014*.
9. Nazmul, M. H. M., M. A. Rashid, and H. Jamal. "Antifungal activity of Piper betel plants in Malaysia." *Drug Discov* 6.17 (2013): 16-17.
10. Nazmul, M. H. M., Fazlul, M. K. K., Rashid, S. S., Doustjalali, S. R., Yasmin, F., Al-Jashamy, K., ... & Sabet, N. S. (2017). ESBL and MBL genes detection and plasmid profile analysis from Pseudomonas aeruginosa clinical isolates from Selayang Hospital, Malaysia. *PAKISTAN JOURNAL OF MEDICAL & HEALTH SCIENCES*, *11*(3), 815-818.
11. Hussaini, J., et al. "Recombinant Clone ABA392 Protects laboratory animals from Pasteurella multocida serotype BJ Vet." *Adv* 2 (2012): 114-119.
12. Arzuman, H., Maziz, M. N. H., Elsersi, M. M., Islam, M. N., Kumar, S. S., Jainuri, M. D. B. M., & Khan, S. A. (2017). Preclinical medical students perception about their educational environment based on DREEM at a Private University, Malaysia. *Bangladesh Journal of Medical Science*, *16*(4), 496-504.

13. Navanethan, D. H. A. R. S. H. I. N. I., et al. "Stigma, discrimination, treatment effectiveness and policy: Public views about drug addiction in Malaysia." *Pakistan Journal of Medical and Health Sciences* 15.2 (2021): 514-519.
14. Keshireddy, S. R. "Low-code application development using Oracle APEX productivity gains and challenges in cloud-native settings." *The SIJ Transactions on Computer Networks & Communication Engineering (CNCE)* 7.5 (2019): 20-24.
15. Keshireddy, Srikanth Reddy. "Cost-benefit analysis of on-premise vs cloud deployment of Oracle APEX applications." *International Journal of Advances in Engineering and Emerging Technology* 11.2 (2020): 141-149.
16. Nazmul, M. H. M., et al. "General knowledge and misconceptions about HIV/AIDS among the university students in Malaysia." *Indian Journal of Public Health Research & Development* 9.10 (2018): 435-440.
17. Iqbal, Mohsena, et al. "The study of the perception of diabetes mellitus among the people of Petaling Jaya in Malaysia." *International Journal of Health Sciences* I (2022): 1263-1273.
18. Nazmul, M. H. M., Salmah, I., Jamal, H., & Ansary, A. (2007). Detection and molecular characterization of verotoxin gene in non-O157 diarrheagenic Escherichia coli isolated from Miri hospital, Sarawak, Malaysia. *Biomedical Research*, *18*(1), 39-43.
19. DOUSTJALALI, SAEID REZA, et al. "Correlation between body mass index (BMI) & waist to hip ratio (WHR) among primary school students." *International Journal of Pharmaceutical Research* 12.3 (2020).
20. Keshireddy, S. R. "Low-Code Development Enhancement Integrating Large Language Models for Intelligent Code Assistance in Oracle APEX." *Indian Journal of Information Sources and Services* 15.2 (2025): 380-390.
21. Keshireddy, Srikanth Reddy. "Automated data transformation and validation in Oracle APEX using adaptive AI models for secure enterprise applications." *Journal of Internet Services and Information Security* 15.2 (2025): 185-208.
22. Keshireddy, Srikanth Reddy. "Extending Oracle APEX for Large-Scale Multi-Form Workflows with Decoupled PL/SQL Logic and Asynchronous Processing Layers." *2025 International Conference on Next Generation Computing Systems (ICNGCS)*. IEEE, 2025.
23. Selvaganapathi, G., et al. "Knowledge and practice on tuberculosis among prison workers from Seremban Prison." *Occupational Diseases and Environmental Medicine* 7.4 (2019): 176-186.
24. Khan, Md Fazlul K., et al. "Detection of ESBL and MBL in Acinetobacter spp. and Their Plasmid Profile Analysis." *Jordan Journal of Biological Sciences* 12.3 (2019).
25. Foysal, Md Javed, et al. "Identification and assay of putative virulence properties of Eschericha coli gyrase subunit A and B among hospitalized UTI patients in Bangladesh." *Inov Pharm Pharmacother* 1.1 (2013): 54-59.
26. Keshireddy, Srikanth Reddy. "Bidirectional Flow of Structured Data between APEX and Streaming Pipelines Using AI-based Field Mapping and Noise Filtering." *2025 International Conference on Next Generation Computing Systems (ICNGCS)*. IEEE, 2025.
27. Keshireddy, Srikanth Reddy. "Natural Language Processing Integration in Oracle APEX for Enhanced User Interaction in Ubiquitous Systems." *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 16 (2025): 668-689.
28. Hussaini, Jamal, Nurul Asyikin Othman, and Mahmood Ameen Abdulla. "Antiulcer and antibacterial evaluations of Illicium verum ethanolic fruits extract (IVEFE)." *Medical science* 2.8 (2013).
29. Nazmul, M., M. Fazlul, and M. Rashid. "Plasmid profile analysis of non-O157 diarrheagenic Escherichia coli in Malaysia." *Indian Journal of Science* 1.2 (2012): 130-132.
30. Vijayakumar, K., Mohammad Nazmul Hasan Maziz, and Mathiyazhagan Narayanan. "Classification of Benign/Malignant Digital Mammogram Images using Deep Learning Scheme." *hospital* 4 (2025): 5.
31. Keshireddy, Srikanth Reddy. "Deploying TensorFlow-Based Predictive Models." *International Journal of Advances in Engineering and Emerging Technology* 12.2 (2021): 11-18.
32. Keshireddy, Srikanth Reddy. "Multi-Hop Signal Transmission Patterns in Oracle APEX-Based Monitoring Systems with Dynamic IoT Feedback Loops." *International Journal of Engineering, Science and Information Technology* 5 (2025): 554-560.
33. Keshireddy, Srikanth Reddy. "RETRIEVAL-AUGMENTED GENERATION TECHNIQUES IN ORACLE APEX IMPROVING CONTEXTUAL RESPONSES IN AI ASSISTANTS." *Archives for Technical Sciences* 2.33 (2025): 253-270.