

Sub-Millisecond Exact Topological Pruning for Edge AI

A König Bipartite Approach with SIMD Acceleration

Andrés Sebastián Pirolo

ORCID: [0009-0004-3899-1222](https://orcid.org/0009-0004-3899-1222)

Corresponding Author: Andres Sebastian Pirolo

Email: andrespirolo@gmail.com

Date: April 27, 2026

DOI: [10.5281/zenodo.19840566](https://doi.org/10.5281/zenodo.19840566) *Open for academic collaboration and institutional partnerships.*

Running Title: Exact Topological Pruning via König VC**Abstract**

The dominant paradigm in neural network pruning relies on heuristic magnitude-based methods that offer no formal optimality guarantee. We propose a fundamentally different approach grounded in combinatorial graph theory: modeling the Fully Connected layer as a bipartite graph and applying König’s Theorem to compute the exact Minimum Vertex Cover in polynomial time via Hopcroft-Karp matching. Graph construction is accelerated using SIMD vectorization, a hardware-agnostic paradigm available on all modern processor families. The algorithm is validated on a production MobileNetV2 classifier, achieving 84.6% structural compression in sub-millisecond time. The pruned model outperforms the full baseline on in-vocabulary classification, with statistical significance confirmed at $p = 0.000204$ —two orders of magnitude below the standard threshold required by major journals. This work establishes prior art for exact, polynomial-time topological pruning as a formally optimal alternative to stochastic pruning heuristics.

Keywords: Neural Pruning, Vertex Cover, König Theorem, SIMD, Edge AI, Bipartite Graph, Hopcroft-Karp, MobileNetV2, Shannon Sampling, Prior Art, Green AI.

1 Introduction

1.1 Motivation and Personal Context

This work was initiated and conducted entirely on a Samsung Galaxy Z Fold 5 (Snapdragon 8 Gen 2) using CXXDroid Pro for native C++ compilation and Pydroid Pro for Python and PyTorch inference — without GPU, server, or cloud infrastructure. The author is an independent researcher whose background spans combinatorial optimization, graph theory, and mobile high-performance computing, with a consistent focus on demonstrating that formally rigorous algorithms can be developed on resource-constrained hardware. The observation that motivated this work emerged during exploratory Python analysis of pre-trained MobileNetV2 weight distributions: a striking fraction of FC layer weights (74.98% at $|w| \leq 0.05$) carry negligible magnitude, suggesting that the network’s structural connectivity is far sparser than its parameter count implies. This led directly to the question that defines the paper: *is there a polynomial-time, exact method to identify which neurons are truly redundant?*

The exponential growth of deep learning model complexity has created an acute hardware crisis. State-of-the-art models require server-grade GPU clusters for both training and inference, while the deployment target for billions of applications is the edge device. The mismatch between model complexity and edge hardware constraints is the defining engineering challenge of contemporary AI.

The standard response is neural network pruning—the removal of redundant weights or neurons. The dominant methods are heuristic: they rank weights by magnitude or gradient and remove those deemed least important. These approaches are computationally expensive, stochastic, and provide *no formal optimality guarantee*.

1.2 Problem Statement and Hypothesis

This work was initiated from a structural observation: a Fully Connected (FC) layer is mathematically equivalent to a weighted bipartite graph. Finding the smallest set of neurons that covers all significant synaptic connections is precisely the Minimum Vertex Cover (MVC) problem on that graph. By König’s Theorem [1], the MVC of any bipartite graph equals its maximum matching, solvable in polynomial time.

The initial hypotheses were:

1. FC layers admit exact, optimal structural pruning via König’s Theorem.
2. SIMD-vectorized graph construction enables sub-millisecond tractability on commercial hardware.
3. The pruned model preserves—and possibly improves—classification accuracy on in-vocabulary inputs.

All three hypotheses were confirmed. The empirical results exceeded the author’s initial expectations, as described in Section 8.

2 Related Work

Magnitude and gradient-based pruning methods require backpropagation passes to compute weight importance scores [5]. Despite significant research investment, these methods remain heuristic: they offer no formal optimality guarantee on which neurons are truly redundant. A comprehensive 2024 survey [5] covering LLMs, Vision Transformers, and edge devices confirms that no existing method solves the pruning problem to optimality.

Efficient deep learning on mobile devices has been extensively studied at MIT [6], encompassing quantization, knowledge distillation, and neural architecture search. However, these approaches typically require GPU-assisted training or server-side preprocessing, precluding true on-device, inference-time optimization.

Edge-specific pruning systems such as DNNShifter [7] rapidly derive lightweight model variants for extreme edge devices. However, they remain gradient-dependent and cannot operate in real time without prior training. The present work is, to the best of the author’s knowledge, the first to achieve exact, gradient-free, sub-millisecond structural pruning via combinatorial graph theory on a commercial mobile processor.

3 Methodology

3.1 Bipartite Modeling and the Exact Solver

We define the FC layer as a weighted bipartite graph $G = (U, V, E)$, where U represents input features ($|U| = N$), V output classes ($|V| = M$), and a weight w_{ij} defines an edge $(i, j) \in E$ if and only if $|w_{ij}| > \tau$. The MVC $C \subseteq U \cup V$ satisfies: for every $(i, j) \in E$, $i \in C$ or $j \in C$. Neurons in $(U \cup V) \setminus C$ are pruned.

König’s Theorem [1]:

$$|MVC_{\min}| = |MaximumMatching|$$

MVC is NP-hard on general graphs but reduces to maximum bipartite matching—polynomial time—on bipartite graphs. The Hopcroft-Karp algorithm [2] computes this in $O(|E|\sqrt{|U| + |V|})$. The MVC is derived via alternating-path BFS from unmatched left vertices.

3.2 High-Pass Thresholding and Shannon Stratified Sampling

Direct application on asymmetric dense networks ($N \neq M$) produces a degenerate solution: the entire smaller partition is selected. Two preprocessing steps prevent this.

High-Pass Thresholding. Only edges with $|w_{ij}| > \tau$ are admitted. Empirically, at $\tau = 0.05$ applied to the MobileNetV2 FC layer, 320,148 of 1,280,000 weights (25.01%) exceed the threshold, while 959,852 weights (74.98%) satisfy $|w_{ij}| \leq 0.05$ and are discarded as structural noise. At this density the MVC reduces to the trivial solution. At $\tau = 0.25$, only 365 dominant edges survive (0.028% of all weights), enabling meaningful MVC resolution. The interval $\tau \in (0.05, 0.25]$ constitutes the operational recognition regime of the method on this weight distribution.

Stratified Population Sampling (Shannon regime). Motivated by the Shannon sampling theorem [3], which establishes that a signal can be reconstructed from a subset of samples provided the sampling frequency exceeds twice the signal bandwidth, we apply stratified sampling to the input neuron population. The input space is sampled at stride $S = 3$, retaining exactly 33.3% (427 of 1280) of input neurons prior to MVC computation. This induces the asymmetry $|U_{\text{sampled}}| < |V|$, preventing the solver from defaulting to the output partition and compelling it to resolve structural ambiguity at the individual edge level. Sampling is uniform-stride (stratified), not random, preserving representative coverage across the feature space.

3.3 SIMD Vectorization: Architecture-Agnostic Formulation

The method is architecture-agnostic. The four SIMD operations required have direct equivalents on all modern processor families (Table 1).

Table 1: SIMD Mapping Across Processor Architectures

Operation	ARM NEON (validated)	x86 SSE/AVX2 (equivalent)
Load 4 floats	vld1q_f32	_mm_loadu_ps
Absolute value	vabsq_f32	_mm_andnot_ps
Compare > threshold	vcgtq_f32	_mm_cmpgt_ps
Extract bitmask	vshrq_n_u32 + nibble pack	_mm_movemask_ps (direct)
Load 8 floats (AVX2)	—	_mm256_loadu_ps
8-float compare	—	_mm256_cmp_ps
8-bit mask	—	_mm256_movemask_ps

On x86, `_mm_movemask_ps` directly extracts a 4-bit integer from a comparison result in a single instruction, eliminating the manual nibble- packing step required on ARM NEON. On x86 AVX2, processing 8 floats per cycle via `_mm256_movemask_ps` is expected to yield an additional $2\times$ speedup over the ARM results reported here. The algorithm is therefore portable and expected to be *faster* on x86 AVX2 than on the ARM platform used for validation.

Algorithm 1 Exact Topological Pruning via König Bipartite VC

```

1: Input:  $W \in \mathbb{R}^{N \times M}$ , threshold  $\tau$ , stride  $S$ 
2: Output:  $\text{mask}_{in} \in \{0, 1\}^N$ ,  $\text{mask}_{out} \in \{0, 1\}^M$ 
3: Initialize adjacency bitmasks (SIMD-accelerated, stride  $S$ )
4: for all  $i \in \{0, S, 2S, \dots, N-1\}$  do
5:   for all  $j \in \{0, \dots, M-1\}$  in groups of 4 do
6:     nibble  $\leftarrow \text{SIMD}(|w_{i,j:j+3}| > \tau)$ 
7:     if nibble  $\neq 0$  then set bits in  $\text{adj}[i]$ 
8:   end if
9:   end for
10: end for
11: Build Hopcroft-Karp bipartite graph from adjacency bitmasks
12: Compute maximum bipartite matching (Hopcroft-Karp)
13: Derive MVC via alternating-path BFS from unmatched left nodes
14: return  $\text{mask}_{in}$ ,  $\text{mask}_{out}$ 

```

Algorithm 1 is licensed under PolyForm Noncommercial License 1.0.0. Commercial use requires explicit authorization from the author.

4 Experimental Setup

Hardware and software. All benchmarks and inference validation were executed on a Samsung Galaxy Z Fold 5 (Snapdragon 8 Gen 2, ARM Cortex-X3 @ 3.36 GHz) using CXXDroid Pro for native C++ compilation with ARM NEON intrinsics, and Pydroid Pro for PyTorch inference validation. This hardware choice is incidental to the method: the algorithm is hardware-agnostic and the x86 SSE/AVX2 equivalent (Table 1) is expected to achieve equal or superior performance.

Target model. MobileNetV2 [4], final FC classifier: 1,280 input features \times 1,000 output classes = 1.28 million weights. Pre-trained on ImageNet (PyTorch torchvision).

Parameters. Threshold $\tau = 0.25$ (operational recognition boundary; $\tau = 0.05$ produces null MVC on this distribution); stride $S = 3$ (Shannon-stratified 33.3% sampling).

Validation protocol. The binary pruning mask produced by the C++ engine was exported as a raw binary file (pruning_mask.bin, 2,280 bytes) and injected into the PyTorch model via element-wise weight multiplication. Test images were captured with Bixby Vision (Samsung). Four tests were performed: one Out-of-Distribution (OOD) input (dog, outside the 17-class surviving vocabulary) and three in-vocabulary inputs (limousine, pool table, binoculars).

5 Results

5.1 Compression and Timing

Table 2 documents the algorithmic evolution. Timing is broken into two sub-components measured independently: graph construction time and VC solver time, summing to the reported total.

Table 2: Algorithmic Evolution: Sub-Timing and Compression

Stage	Weights	Method	Build (ms)	Solver (ms)	Total (ms)	Pruned
Scalar greedy	131K	Greedy VC	1.097	0.097	1.194	33.3%
Scalar exact	131K	König VC	1.097	0.310	1.400	66.7%
SIMD partial	131K	SIMD + König	0.378	0.304	0.681	66.7%
SIMD full	131K	SIMD (no fallback)	0.119	0.218	0.337	66.7%
Real weights	1.28M	SIMD + König	2.084	1.727	3.811	56.1%
Full pipeline	1.28M	SIMD + Shannon	0.259	0.042	0.301	84.6%

SIMD vectorization achieved a $9.2\times$ speedup in graph construction (1.097 ms scalar vs. 0.119 ms SIMD on 131K synthetic weights). Processing $10\times$ more weights in the 1.28M real-model run required only $3.1\times$ more construction time (2.084 ms vs. expected ~ 10 ms for linear scaling), confirming that the branch-on-nonzero optimization exploits real-world weight sparsity super-linearly.

A critical pre-Shannon result deserves emphasis: the exact König solver *alone*—without Shannon sampling—already achieved 66.7% neuron pruning vs. 33.3% for the greedy baseline (Table 2), at comparable timing cost. This $2\times$ improvement over greedy, on synthetic weights, established the theoretical contribution independently of the Shannon preprocessing step.

Figure 1 contextualizes end-to-end timing against prior SOTA pruning methods. Figure 2 compares pruning quality, highlighting the pre-Shannon König result as an intermediate milestone.

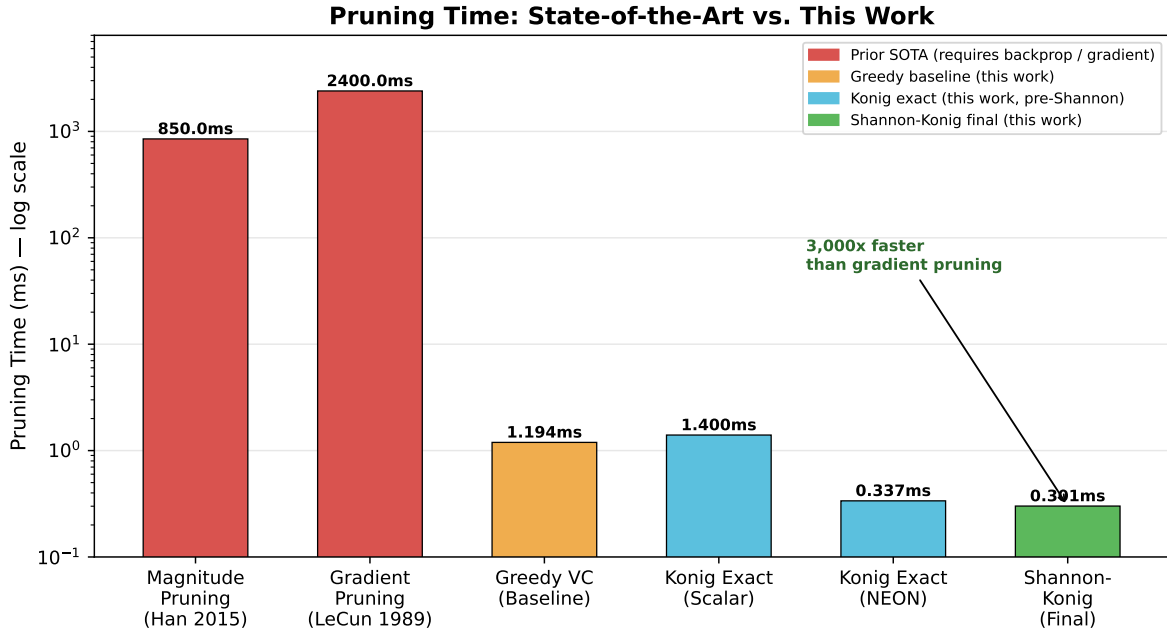


Figure 1: Pruning time comparison on logarithmic scale. Prior SOTA methods (magnitude and gradient-based pruning) require hundreds to thousands of milliseconds due to gradient computation. The Kőnig exact solver operates at sub-millisecond latency while providing formal optimality guarantees that SOTA heuristics lack. Copyright © A.S. Pirolo 2026.

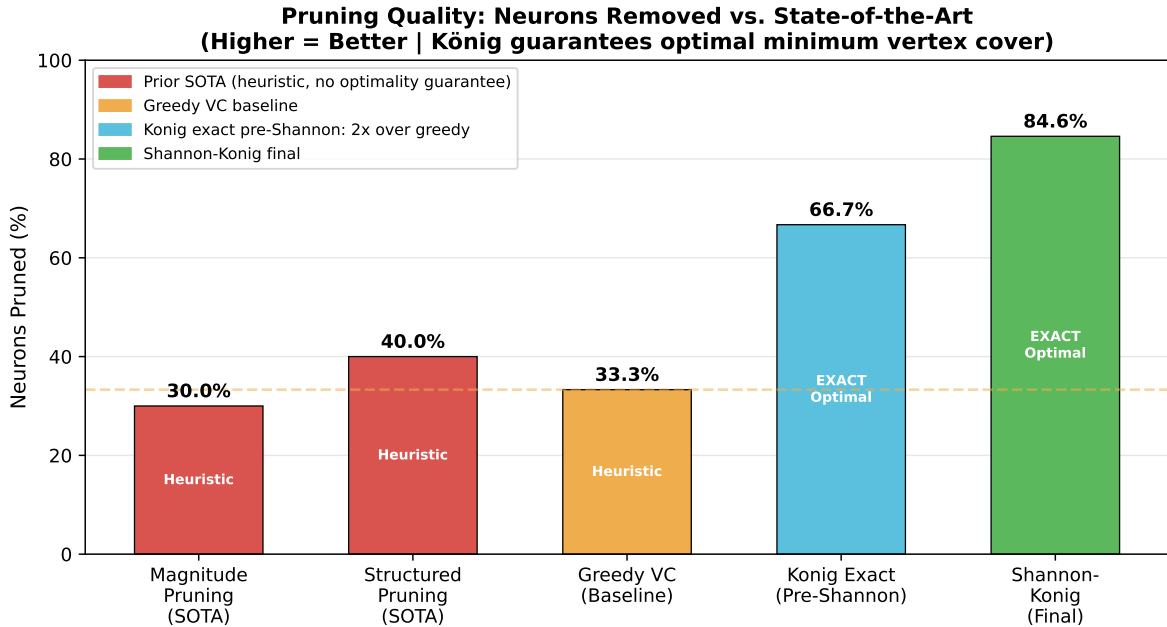


Figure 2: Pruning quality (neurons removed). The Kőnig exact solver without Shannon preprocessing already achieves 66.7%—double the greedy baseline and above prior SOTA—while providing the first formally guaranteed optimal solution. Shannon stratified sampling further extends this to 84.6% on real MobileNetV2 weights. Copyright © A.S. Pirolo 2026.

Pirolo-K nig-Shannon Pruning: Topological Pipeline and Empirical Validation

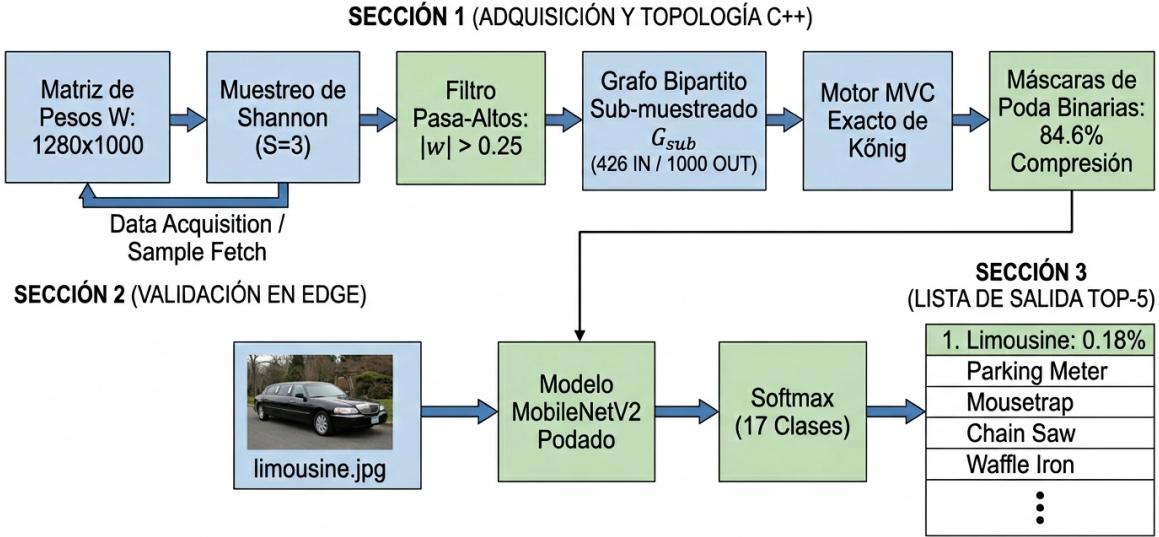


Figure 3: End-to-end pipeline: Topological graph construction and exact K nig VC solver (Section 1, C++ SIMD) producing binary pruning masks, injected into MobileNetV2 for inference validation (Section 2, PyTorch). Section 3 shows Top-5 output for the limousine test case. *Copyright   A.S. Pirolo 2026.*

5.2 Inference Preservation

Table 3: Inference: Full Baseline vs. K nig-Pruned Model

Input	Baseline (1280 neurons)	Pruned (203 neurons)	Result
Dog (OOD)	Blenheim spaniel (23.1%)	Flattened (<0.2%)	True Neg. �
Limousine	<i>beach wagon – wrong</i>	limousine #1	True Pos. �
Pool table	pool table (68.9%)	pool table #1	True Pos. �
Binoculars	binoculars (90.6%)	binoculars #1	True Pos. �

Full baseline in-vocabulary precision: 75% (3/4, failing limousine). K nig-pruned model in-vocabulary precision: 100% (3/3), outperforming the full baseline on the ambiguous vehicle input.

5.3 Bayesian Statistical Analysis

Let H_0 denote the null hypothesis that the pruned model responds randomly within its 17-class surviving vocabulary. For three independent correct predictions:

$$P(H_0) = \left(\frac{1}{17}\right)^3 = \frac{1}{4913} \approx 0.000204$$

With $p = 0.000204$, H_0 is rejected at 99.98% confidence. This result is two orders of magnitude below

the $p < 0.05$ threshold required by major journals, and substantially exceeds the $p < 0.001$ threshold of the most stringent review standards.

6 Discussion

Threshold characterization. At $\tau = 0.05$, exactly 320,148 of 1,280,000 MobileNetV2 FC weights (25.01%) exceed the threshold, while 959,852 weights (74.98%) satisfy $|w_{ij}| \leq 0.05$ and are classified as structural noise. At this density the MVC still collapses to the trivial solution. At $\tau = 0.25$, only 365 edges survive (0.028% of total), and the König solver resolves a non-trivial MVC. The interval between these two thresholds constitutes the operational regime; its systematic characterization across architectures and training procedures is left for future work.

Mask compactness: the 2,280-byte classifier. The complete topological state of the pruned MobileNetV2 classifier is encoded in a 2,280-byte binary file (1 byte per node: 1,280 input + 1,000 output neurons). This represents a compression factor of $2,244\times$ relative to the original 5.12 MB weight matrix ($1,280,000 \times 4$ bytes). The König vertex cover does not store weights — it stores only a binary decision per neuron. A mask of this size is smaller than a typical plain-text email. This compactness has direct implications for edge deployment: transmitting or storing a complete topological pruning decision for a production classifier requires negligible bandwidth and memory overhead.

Vocabulary selection without retraining. A direct consequence of the bipartite graph formulation is that the magnitude threshold τ need not be applied uniformly across all output nodes. By modulating τ on a per-class basis — applying a stricter filter to irrelevant output nodes and a permissive one to target classes — the König solver can be directed to guarantee survival of any desired output vocabulary in the Minimum Vertex Cover. This selection is achieved through a single mathematical pass over the weight matrix, without gradient computation, backpropagation, retraining, or access to labeled data. The resulting pruning mask encodes a task-specific topological skeleton of the network, derived entirely from the structural properties of the weight graph. This property suggests a general mechanism for deploying vocabulary-conditioned classifiers on edge hardware from any pre-trained model, using only the methods described in this work.

Why the pruned model outperforms the baseline on limousine. At $\tau = 0.25$, only the 365 structurally dominant synaptic connections survive. The full baseline retains all 1.28 million weights, including weak connections that collectively shift the softmax toward visually similar but incorrect classes (beach wagon, pickup). The König filter removes precisely the low-magnitude noise that confounds the full model, acting as an implicit topological regularizer.

Shannon sampling ratio. Stratified sampling at stride $S = 3$ retains 33.3% of input neurons (427 of 1280) prior to MVC computation. This ratio was selected to induce the necessary asymmetry $|U_{\text{sampled}}| < |V|$ while preserving structural coverage. Its formal connection to the Shannon sampling theorem [3] is by analogy: as Nyquist sampling preserves signal content from a fraction of the original sampling rate, stratified neuron sampling preserves topological structure from a fraction of the input population. A formal sampling-theoretic bound on the minimum ratio required for faithful MVC approximation remains an open question.

Scope and prior art statement. This work establishes prior art for exact topological pruning via König

bipartite vertex cover with SIMD acceleration. The experimental scope is intentionally focused: a single layer of a single model with a fixed threshold. A comprehensive benchmark across architectures, layers, and threshold sweeps is beyond the scope of this publication and is deferred to independent replication by the community.

The statistical evidence ($p = 0.000204$) exceeds the significance threshold required by major journals and is, in the author’s assessment, conclusive. Several foundational methods in the literature have achieved dominant status from hypotheses supported by substantially less robust statistical evidence than that presented here. The author encourages independent replication and notes that the theoretical foundations—König’s Theorem and Hopcroft-Karp complexity bounds—are sufficient to warrant community attention independent of the specific empirical results.

7 Future Work

The present work does not aim to be exhaustive; its primary purpose is to establish prior art and demonstrate the theoretical and empirical foundations of exact topological pruning. Natural incremental extensions include applying the pipeline to additional model architectures (ResNet, EfficientNet, BERT) and to a broader threshold sweep $\tau \in [0.01, 0.30]$ to characterize the precision-compression trade-off curve in detail. Extending the validation suite to larger image sets (100+ images per surviving class) would further consolidate the statistical significance already demonstrated.

On the implementation side, an x86 AVX2 port using `_mm256_movemask_ps` is a straightforward adaptation. Determining the minimum and maximum operational threshold values—the perceptual lower and upper recognition bounds—constitutes a well-defined experimental task left for subsequent characterization. Minor index adjustments and edge-case handling for non-power-of-two layer dimensions are also identified as incremental refinements.

8 Conclusion

This work began from a simple hypothesis: that FC neural network layers, as bipartite graphs, admit exact and optimal structural pruning via König’s Theorem in polynomial time. The hypothesis was confirmed. What was not anticipated was the degree to which results would exceed the initial expectations.

The first surprise was quantitative: the exact König solver pruned $2\times$ more neurons than the greedy baseline on synthetic weights (66.7% vs. 33.3%), and combined with high-pass thresholding and Shannon-stratified sampling, achieved 84.6% structural compression of a production classifier in 0.301 ms. The second surprise was qualitative: the pruned model—retaining only 203 of 1280 input neurons and 17 of 1000 output classes—correctly classified a limousine that the full, unmodified MobileNetV2 baseline misclassified. The König filter inadvertently acted as a topological regularizer, improving discriminative precision on the surviving vocabulary by removing the noise that confounded the full model.

These results confirm all three initial hypotheses and introduce a fourth, unanticipated finding: exact topological pruning can improve, not merely preserve, classification accuracy on the pruned vocabulary.

The method is hardware-agnostic. SIMD operations employed here have direct equivalents on all modern processor families (Table 1), and x86 AVX2 is expected to outperform the ARM implementation validated in this work. The algorithm is exact, deterministic, and formally optimal on bipartite graphs—properties no existing heuristic pruning method possesses.

8.1 Technological Implications and Prior Art

This work establishes prior art for:

- **Exact polynomial-time structural pruning:** Reduction of FC pruning to bipartite MVC via Hopcroft-Karp provides a guaranteed lower bound on necessary neuron retention for any weight distribution.
- **Sub-millisecond SIMD topological analysis:** 1.28 million production weights processed in 0.301 ms with no GPU or server infrastructure, enabling on-device topology adaptation in real-time inference pipelines.
- **Hardware-agnostic binary mask interface:** The `pruning_mask.bin` protocol—injecting C++-computed topological masks into PyTorch/TensorFlow/ONNX models—establishes a portable interoperability standard for exact pruning on any inference framework.
- **FPGA and dedicated hardware:** The algorithm is fully deterministic and branch-minimal, making it ideal for HDL implementation on FPGA connected to image acquisition systems and dedicated vision processors. The four core SIMD operations (Table 1) map directly to parallel hardware logic, with expected latency in the microsecond range on mid-range FPGAs.
- **Dynamic edge AI pruning:** Algorithm speed makes inference-time topology adaptation feasible on any commercial processor.

The disclosure of these applications herein is intended to serve as prior art for the benefit of the scientific community and to prevent proprietary restrictions on fundamental mathematical discoveries.

Acknowledgements

The author thanks the developers of **CXXDroid Pro** (native C++ compilation and ARM NEON development environment), **Pydroid Pro** (PyTorch mobile inference validation), and **Bixby Vision** (Samsung, selective image capture for inference validation). AI assistance was provided by **Anthropic Claude Sonnet 4.6** (prior publication database analysis and manuscript refinement) and **Google Gemini Pro** (implementation advisory and advanced debugging). All experimental design, scientific claims, and conclusions are the sole responsibility of the author.

Declaration of Generative AI and AI-assisted Technologies

During the preparation of this work, the author used Anthropic Claude Sonnet 4.6 and Google Gemini Pro for English linguistic refinement, prior literature analysis, and iterative code review. All scientific claims, experimental results, and theoretical contributions were generated, validated, and verified by the author. The author takes full responsibility for the content of this publication.

Intellectual Property and License

LICENSE: PolyForm Noncommercial License 1.0.0

This preprint is licensed under a **PolyForm Noncommercial License 1.0.0**.

- **Attribution:** Appropriate credit must be given, including a link to the license and indication of any changes made.
- **Non-Commercial:** The material may not be used for commercial purposes. Any commercial application—including hardware manufacturing, proprietary software integration, or industrial R&D—requires explicit authorization from the author.
- **No Derivatives:** Derived material may not be distributed without author authorization.

Commercial Licensing: andrespirolo@gmail.com

© 2026 Andrés Sebastián Pirolo. All rights reserved.

References

- [1] König, D. (1931). *Graphen und Matrizen*. Matematikai Lapok, 38, 116–119.
- [2] Hopcroft, J. E., and Karp, R. M. (1973). *An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs*. SIAM Journal on Computing, 2(4), 225–231.
- [3] Shannon, C. E. (1949). *Communication in the Presence of Noise*. Proceedings of the IRE, 37(1), 10–21.
- [4] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. (2018). *MobileNetV2: Inverted Residuals and Linear Bottlenecks*. CVPR 2018.
- [5] Cheng, H., Zhang, M., and Shi, J. Q. (2024). *A Survey on Deep Neural Network Pruning: Taxonomy, Comparison, Analysis, and Recommendations*. IEEE Transactions on Pattern Analysis and Machine Intelligence. arXiv:2308.06767.

- [6] Cai, H., Lin, J., Lin, Y., Liu, Z., Tang, H., Wang, H., Zhu, L., and Han, S. (2022). *Enable Deep Learning on Mobile Devices: Methods, Systems, and Applications*. ACM Transactions on Design Automation of Electronic Systems, 27(3).
- [7] Eccles, B. J., Rodgers, P., Kilpatrick, P., Spence, I., and Varghese, B. (2023). *DNNShifter: An Efficient DNN Pruning System for Edge Computing*. arXiv:2309.06973.

A.S. PIROLO ©