BABEŞ-BOLYAI UNIVERSITY
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Blockchain-based neural network training

PhD thesis summary

PhD student: Gyulai-Nagy Zoltán Valentin
Scientific supervisor: Prof. dr. Czibula Gabriela

2026

Contents

List of publications

The ranking of publications was performed according to the CNATDCU (National Council for the Recognition of University Degrees, Diplomas and Certificates) standards applicable for doctoral students enrolled after October 1, 2018. All rankings are listed according to the classification of journals[1] and conferences[2] in Computer Science.

Publications in international journals

[GN25a] **Gyulai-Nagy Zoltán-Valentin**, *Accelerated Neural Network Initialization via Training of Symmetrically Partitioned Models*, Acta Universitatis Sapientiae, Informatica (AUS 2025), Volume 17, Article no. 12 (**indexed Web of Science**, Q4)

Rank C - 2 points.

Publications in conferences proceedings

[GN23] **Gyulai-Nagy Zoltán-Valentin**, *Acceleration of Neural Network training algorithm via FPGA Devices*, The 29th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems (KES), Procedia Computer Science, Volume 225, 2023, Pages 2674-2683 (**indexed Web of Science**)

Rank B, 4 points.

[GN24] **Gyulai-Nagy Zoltán-Valentin**. *Combining clustering algorithms to extract symmetric clusters from noisy data, applied to parking lots*. The 19th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2024), SciTePress, Volume 1: ENASE, 2024, Pages 362-369 (**indexed Web of Science**)

Rank B, 4 points.

[GN25c] **Gyulai-Nagy Zoltán-Valentin**, *Proof of Learning applied to Binary Neural Networks*, The 17th International Conference on Agents and Artificial Intelligence (ICAART), 2025, pp. 411-417 (**indexed Web of Science**)

Rank B, 4 points.

[1] https://uefiscdi.ro/premierea-rezultatelor-cercetarii-articole
[2] http://portal.core.edu.au/conf-ranks/

Publications in international books

[GN25b] **Gyulai-Nagy Zoltán-Valentin**, *Exploring the Usability and Applicability of Proof of Learning in Binary Neural Networks*, Lecture Notes in Artificial Intelligence (LNAI), 2025, in press (**indexed Web of Science**)

Rank C - 2 points.

Publications score: 16 points.

Introduction

This thesis investigates integrating blockchain technology with neural networks, focusing on optimizing computational efficiency and redefining validation mechanisms in decentralized systems. Blockchain networks, especially those using Proof-of-Work (PoW) consensus algorithms, face challenges because of computationally intensive hashing processes that consume significant energy [Sed20] without producing meaningful outputs beyond network security. At the same time, neural network training is resource-intensive and often limited by traditional hardware like CPUs and GPUs [FL20]. This thesis proposes addressing these challenges by using Field-Programmable Gate Arrays (FPGAs) to accelerate neural network training and repurposing blockchain validation functions to generate useful computational outcomes via BNNs (Binary Neural Networks).

The core problem addressed is the inefficiency of PoW mechanisms, which require participants to solve complex cryptographic puzzles that yield no tangible byproducts [Pil16]. This thesis suggests replacing these validation tasks with algorithms that perform meaningful computations, such as training neural networks, transforming blockchain networks into distributed platforms for high-performance computing. A key contribution is developing an FPGA-based architecture for neural network training, showing the potential for hardware acceleration to overcome bottlenecks in conventional processing units. The thesis also introduces a symmetrically partitioned model initialization technique to further reduce training time and resource consumption, support the distribution of models across a large pool of computers, and enable the use of these resources to evolve existing models in the network.

The research first establishes theoretical foundations for linking neural networks and blockchain systems. Once these aspects are demonstrated to be viable for a foundational architecture, the focus shifts to scaling the system, which also involves leveraging existing models and partitioning them to improve efficiency. To better illustrate a real-world scenario where this system provides tangible benefits, we present a simple project involving continuous training of its neural networks. Although the algorithms can be complex, their effectiveness mainly depends on the accuracy of neural networks, which can be enhanced by accumulating data over time and retraining frequently. This makes it an excellent use case for our project.

By closing the divide between hardware optimization and blockchain innovation, this study adds to the expanded dialogue on sustainable and utility-driven decentralized systems, offering a framework for harnessing neural networks as a dual-purpose tool for both machine learning and distributed consensus.

Approached Problems

Our thesis first addresses the challenge of performing calculations more efficiently compared to traditional computing units like CPUs or GPUs. In blockchain technology, ASICs are widely used due to their lower energy consumption and faster calculation speeds, as they are designed for such tasks at the

hardware level [BT13]. General-purpose hardware architectures, though powerful for diverse tasks, can have their speed enhanced through hardware-level optimizations. However, these devices typically perform only specific types of operations, which align with the needs of specialized blockchain hardware. Before manufacturing ASIC chips and deploying them in real-world applications, FPGA devices are used for prototyping. Thus, this research focuses on developing FPGA architectures that serve as effective solutions for this problem.

The second direction of our research involves repurposing the validation algorithm to generate usable byproducts. In the current context, many AI products demand significant power consumption [ZLL+24]. Meanwhile, blockchains have an untapped potential where anyone can contribute their computing power to support a broad ecosystem of training algorithms. This is particularly valuable because training models is one of the most expensive operations involved in developing new models for the consumer market. This entails replacing the existing validation process with a new one capable of validating blocks through the work performed by computers training neural network models. The fundamental idea of hashing blocks into a specific shape to demonstrate computational effort has inherent inefficiencies [G+16], as it depends on the processing power of the units due to the competition among participants. We seek to address this problem by shifting the focus towards supporting the proof-of-work mechanism through the evaluation of real-world models, which we suggest could be binary neural networks, rather than relying solely on hashes that must conform to a particular shape. This also includes the ability to distribute such workloads across multiple computers on the blockchain. As seen in existing blockchain systems, participation is not limited by a predefined amount of computing power that computers must have. This would limit the network's scalability and prevent smaller potential participants from contributing their computing power to sustain the blockchain. Instead, the concept of mining pools is common among participants with consumer-grade computing units who still want to contribute. We seek to explore scenarios where we can enable such individuals to expand the capacity of the blockchain.

Next, we aim to integrate and showcase a project that can be implemented in real-life scenarios. Even though our proposal in this thesis might be scientifically sound and the algorithms efficiently solve the problems mentioned above, we want to demonstrate a real-life scenario where its use would be beneficial. We believe that supporting the research with applications that provide examples is helpful and keeps the context more practical than purely empirical, where the reader is left to assume where the system can be used.

All of the problems mentioned above aim to create a more efficient blockchain system that optimizes computing power usage while preserving the fundamental principles of blockchain technology.

Original Contributions

Our research focused on utilizing FPGA devices and BNNs to develop an enhanced blockchain algorithm. The close integration of these technologies offers an innovative approach with practical applications in blockchain repurposing and large-scale model training, which in the end are applied to solving the problem of recognizing cars in a parking lot. The methods discussed in Chapters 2, 3 and 4 are summarized below:

Hardware perspective of the blockchain

Given the nature of the problem, the first aspect we need to analyze is the ability to transpose calculations required for training neural networks on a hardware level. This facilitates integrating the entire ecosystem into specialized hardware, ensuring a competitive advantage that makes blockchains

more efficient in specific tasks. In this context, we contributed by developing a comprehensive hardware solution on FPGA devices, which can be easily adapted into full ASIC processors. Our key contributions include implementing and researching hardware architectures that can accelerate the calculations necessary for our proposed blockchain.

The base operations that lie at the foundation of neural networks are multiplication and sum operations. The challenge comes in when there are plenty of them, and they also have dependencies that need to be satisfied when doing calculations between layers.

- The hardware acceleration approach outlined in the original paper [GN23], introduced in chapter 2, aims to address these challenges. Our research centers on developing multiplication and addition cores capable of meeting these requirements. These cores are modular and configurable, allowing scalability based on the FPGA hardware used. They can be arranged on hardware to mirror the structure and layers of neural networks, enabling streamlined data flow between cores to boost processing speeds. Connected by specially designed memory blocks, one core always writes to a block while a single core reads from it. This setup eliminates the need for custom prioritization logic for scheduling calculations. The system's bus width is adjustable, supporting different data types from 2-bit to 64-bit values. This flexibility allows training models with quantizations suited to specific application requirements. Calculations are performed in parallel whenever possible. Essentially, all numbers on the bus can be multiplied or summed in a single operation, increasing speed. The cores are designed for easy scalability, limited primarily by available space on FPGA or ASIC hardware to widen the bus. Using specialized hardware also reduces energy consumption, as focusing on a limited set of operations removes unnecessary peripheral functions found in general-purpose processors, leading to more efficient energy use.

Overall, our contribution through the computing cores aims to offer a solution for enabling blockchain extensibility on custom hardware designed to accelerate operations. They highlight the extensive possibilities for further specialization or optimization within the hardware domain.

Algorithmic trust through repurposed validation

The purpose of blockchains is to establish a robust foundation that supports multiple critical functions while ensuring availability and consistency. It possesses inherent properties that uphold its core principles. A crucial element is the validation algorithm, which verifies whether a block is part of the blockchain and whether it meets the necessary criteria to be added, ensuring all requirements are satisfied. These requirements are designed to ensure that, within large distributed computing systems, participants follow the established rules. Our research explores the possibility of maintaining such a framework while modifying how PoW blockchains operate. This aims to harness computing capacity more effectively and give it a new purpose, providing additional benefits to blockchain users by creating byproducts that benefit from the computing power of a blockchain and are usable in real-world applications.

- In our original paper [GN25c], introduced in Chapter 3, describing an altered validation algorithm, we present a validation design that links blocks by verifying that participants who submitted them actually trained a neural network model. The proof submitted by contributors takes the form of a block containing all the standard fields, like transactions and embedded information. Additionally, it includes a binary neural network. It also considers scenarios where limited datasets cause network models to have a higher likelihood of converging to a

specific state. To prevent this, evolutionary algorithms are used to evolve the models and create new shapes and configurations. The same method applies to models that appear perfect based on backpropagation training algorithms. Since they are binary neural networks, they can be further evolved by adjusting weights to previously unknown states, aiming to enhance model performance. This process requires testing the model with data to identify key performance indicators, which consumes processing power. Our research indicates that alternative validation methods, compared to traditional PoW, can effectively track contributor efforts in training and validating neural networks.

- To complete the algorithm, the original article [GN25a], introduced in Chapter 3, explores ways to enable the distribution of workload among all network contributors, reducing the need for high-performance hardware. Typically, mining pools in blockchains work together to solve the computational challenges presented by the blockchain. In our case, to enable the same mechanism, we aim to split training iterations across nodes to achieve faster results. However, we need to consider that the models must fit within hardware constraints, which may not always be feasible. Therefore, training the same large neural network model in batches might not be suitable. Ideally, the model could be divided either by layers or by cutting through layers, with smaller portions distributed to contributors, who then combine these parts to form a trained model. Such training methods necessitate frequent synchronization between nodes, so we proposed a different approach: pre-training sub-sections of the model weights on a neural network, with the goal of further training them after concatenation. Our research successfully shows that this method can initialize model weights in a distributed manner, thereby reducing the overall training time for the complete model. This contribution opens the door for mining pools that distribute rewards among participants to be more inclusive, especially for those lacking the processing power to train larger neural networks.

- Building on our previous article, we expanded our research in [GN25b], introduced in Chapter 3, by incorporating additional use cases and conducting further testing to verify its feasibility. This involved exploring the benefits of using blockchain-generated models, aiming to enhance security and ensure they support blockchain architecture constraints. Key aspects examined include the ability to adjust network difficulty, which influences how rewards are scaled and regulated. We also assessed the diversity of models, which helps strengthen security against spoofing and potential attacks that could compromise the blockchain. These findings demonstrate the approach's viability, showing it can be effective not only in controlled settings but also in large-scale distributed systems where external actors might attempt to manipulate results and undermine the validation mechanism.

As shown in the research directions above, establishing an algorithmic trust that accurately represents the distributed nature of blockchains is achievable, even while modifying its fundamental validation mechanism. These changes preserve the core principles of blockchains and incorporate all existing functionality, ensuring the same capabilities are maintained. Utilizing the processing power of such networks can enable further development scenarios for both new and existing blockchains.

A practical use-case

Empirical research often concludes with a methodology that appears suitable for real-world applications, but sometimes it becomes disconnected from practical reality. This can lead to notable studies that remain unused until an ideal scenario arises. For example, the transformer architecture was

known and studied by many scientific groups, but it only became practically useful after the development of the attention mechanism, which extended its capabilities. Given this context, we aim to demonstrate a real-world application that leverages our repurposed validation mechanism to validate its viability. Aside from potential consumer-grade training linked to a blockchain transaction pricing-based model, we want to emphasize a large-scale application. If implemented at the city or national level, such a system could harness the vast computational power of a blockchain to enhance its effectiveness.

- We demonstrated a complex application involving multiple algorithms that aim to deterministically identify parking lot availability. As shown in our paper [GN24], introduced in Chapter 4, our contribution is a well-defined algorithm that does not depend on external variables to specify parking spaces and automatically detect occupancy. We employed clustering and segmentation techniques to estimate the likelihood of a car being parked there. Despite the algorithm's deterministic nature and good performance, we found that ongoing retraining based on positive detections is essential for it to improve its accuracy to a more dependable value. Cars are identified in camera images using convolutional neural networks, which, since writing, have evolved even more, and remain a standard method for object detection. Environmental, lighting, and weather conditions significantly influence how cars appear in images. Regular retraining using relevant data collected by on-site cameras about cars enhances the system's accuracy. This approach illustrates how blockchain validation can be used to ensure the system remains accurate despite constantly changing external factors like reflections, lightning, and not limited to, car models. Our research highlights the importance of continuous model retraining to adapt to numerous external and evolving factors affecting parking lot detection.

Using the above example, we aim to clarify for the reader the types of applications that can benefit from our system. The market's demand for training models for various specialized tasks is likely to persist in the future. A centralized repository of models, accessible for review and verification by everyone, increases the likelihood of ensuring that models that fit existing scenarios are not redundantly retrained. Although we do not delve deeply into economic or pricing strategies for our validation algorithm, the application serves as an example of how untapped blockchain computing power can be repurposed toward a more sustainable model, thereby easing the load on data centers responsible for providing computing resources.

Thesis structure

The remainder of this thesis is structured as follows. Chapter 1 offers a general introduction to the proposed problem, explaining the mechanics and purpose of blockchains, along with their inefficiencies. It also lays the foundational knowledge necessary to understand neural network research and offers insights into how these technologies complement each other. It highlights areas where they can benefit from one another and extend their functionalities. Additionally, it provides a comprehensive review of existing literature, focusing on the current state of research on these topics.

Chapter 2 focuses on the hardware aspects involved in implementing blockchain accelerator components. It emphasizes the significance of hardware acceleration and its advantages in terms of speed and energy efficiency. The chapter highlights key building blocks that provide a foundation for designing customizable hardware architectures tailored to a full blockchain. Several architectures are discussed as potential solutions for performing mathematical operations on FPGA devices, referred to as processing cores. These cores form the basic components for developing hardware-accelerated

devices for blockchain applications. Although the chapter does not specifically focus on ASICs, it outlines the benefits of using FPGA devices, which can outperform traditional computing units. Transitioning FPGA architectures into ASIC chips enhances performance by removing limitations inherent to FPGA devices, resulting in chips with higher processing power. This is also a widespread practice for blockchains. The chapter also stresses the importance of scaling and integrating these cores into configurable hardware architectures that meet the evolving needs of blockchain systems.

Chapter 3 covers the software perspective of the application. It discusses how repurposing the validation algorithm presents key challenges that must be addressed to create a working blockchain. The chapter reviews the core functionalities and describes necessary algorithmic adjustments as solutions for integrating the new validation mechanism. It provides insights into the algorithm's logic and discusses the importance of the blockchain's distributed nature, which is essential for its operation. One potential implementation involves distributing workload by training symmetrical partitions of neural network models. Overall, this chapter outlines the main proposal of the thesis, enabling a functional validation algorithm that can be integrated into blockchains, thereby allowing more purposeful use of available processing power.

Chapter 4 adopts a different approach, illustrating a use case for blockchain technology. It diverges from the style of earlier chapters to explore and describe a complex problem that requires extensive algorithmic integration to function effectively. The problem tackled in this chapter involves detecting free and occupied parking spaces. Although it describes a purpose-built approach with good accuracy, it concludes that higher accuracy depends on frequent retraining sessions with newly acquired data from the system. By scaling this application to a larger integration level with multiple parking lots, it's evident that the blockchain can generate new models continuously based on real-time data, enabling ongoing neural network improvements without specialized hardware, all while maintaining security.

Finally, the last chapter presents the conclusions of this thesis and outlines future directions for exploring the advantages of the proposed approach.

Chapter 1

Background

In this chapter, we are going to provide insight into the current state of the main technologies used in the project. The proposed solution is an enhancement; therefore, in order to understand the problem we are trying to solve, we are going to present the functionality and architecture of blockchains and neural networks. Communities that use them have already noticed a few areas where improvements can be applied; therefore, to understand the unique behavior of our project, we are also going to present some of the research that tries to provide a solution.

Chapter 2

Hardware perspective of the blockchain

This chapter explores the potential of FPGAs to revolutionize blockchain hardware ecosystems. It evaluates their benefits over traditional processing units (CPUs/GPUs) and quantifies productivity gains for developers, particularly in scenarios requiring rapid prototyping or parallel processing.

Blockchains, foundational to decentralized systems, rely heavily on Application-Specific Integrated Circuits (ASICs) for executing computationally intensive tasks such as mining and smart contract validation. While ASICs offer unmatched efficiency for predefined operations, their rigidity and high energy consumption pose significant challenges. As the demand for sustainable and adaptable blockchain solutions grows, the limitations of ASICs—particularly their inflexibility to evolving algorithms and environmental costs—necessitate innovative alternatives.

Field-Programmable Gate Arrays (FPGAs) emerge as a promising bridge between general-purpose processors and ASICs. Unlike ASICs, which are hardwired for specific tasks, FPGAs offer reconfigurable hardware architectures, enabling dynamic optimization for diverse blockchain workloads. This adaptability aligns with the evolving nature of blockchain protocols, such as proof-of-work (PoW) algorithms or machine learning-driven consensus mechanisms, while maintaining energy efficiency. By balancing flexibility with performance, FPGAs address the critical trade-off between computational power and energy consumption, positioning themselves as a viable intermediary step toward future ASIC designs.

The results presented in this chapter are included in one of the original papers [GN23]. Our contribution is the following:

- Two custom computing cores are designed to perform the essential computations for training neural networks. They handle sum and addition operations on buses with variable bit widths. Each bus encodes values using a configurable number of bits, allowing operations on different model quantizations. The cores are configured to process all values in the buses simultaneously, increasing parallelism. These cores can be integrated into hardware designs that mirror the structure of neural network training operations, enabling efficient, streamed calculations between them. They serve as a foundation for developing hardware that is optimized for specific neural network models.

By demonstrating FPGAs' capacity to reduce energy overhead and improve computational efficiency, this chapter contributes to the discourse on sustainable blockchain infrastructure. It underscores the potential of reconfigurable hardware to address current limitations while paving the way for scalable, energy-conscious solutions in decentralized systems.

The rest of the current chapter is organized into six key sections, each addressing specific aspects of FPGA-based hardware solutions for blockchain applications. The structure is designed to systematically explore the feasibility, performance, and productivity benefits of FPGAs and to address the

research questions. Section 2.1 introduces an acceleration approach for neural network training using FPGA devices. The problem statement is presented in Section 2.2, and the literature review of the state of the scientific literature is presented in Section 2.3. The methodology employed in our research is presented in Section 2.4, while Section 2.5 discusses the results and tests conducted to reaffirm the approach's viability. Finally, Section 2.6 concludes the research and identifies future work areas for exploration.

Chapter 3

Algorithmic trust through repurposed validation

Blockchain validation is a cornerstone of decentralized systems, ensuring consensus and trust through mechanisms like Proof of Work (PoW) or Proof of Stake (PoS). In traditional blockchains, transactions are verified by nodes through cryptographic hashing, with consensus algorithms guaranteeing immutability and security. However, these methods often face challenges such as high computational costs, scalability limitations, and energy inefficiency. This chapter explores an innovative approach to validation by repurposing algorithmic techniques from machine learning (ML) to address these limitations, enabling efficient, distributed, and scalable validation frameworks.

The proposed method shifts from blockchain-centric validation to algorithmic trust through an alternative model validation, partitioning, and dataset segmentation. Key concerns are also addressed, where the distributed nature of the blockchain still needs to be kept intact to enable participants with varying computers to contribute. The results of this chapter are presented in three original papers [GN25c], [GN25b], [GN25a]. They are as follows:

- Proof of Learning applied to Binary Neural Networks is introduced in paper [GN25c]. Our main contribution is defining an algorithm that enables the creation and validation of blocks based on proof that neural networks have been trained. The core idea is to store binary neural networks in each block, which is practical due to their small size. Participants, when accepting new blocks, verify whether key metrics are met. The approach also handles complex scenarios, such as using evolutionary algorithms to ensure neural network models do not converge to similar results. The models' shapes and weights are evolved to maintain enough diversity, allowing them to keep up with the potential constant demand for new blocks. The researched approach establishes a foundation for repurposing blockchain processing power by storing BNNs, the byproducts, on the blockchain. These are accessible to the public as a central repository of models, all while preserving the blockchain's transactional functions.

- Our next contribution is presented in the [GN25b] article, which extends the research directions of the previous one. Building on our prior work, we expanded our research by incorporating additional use cases and rigorous testing to validate the feasibility of integrating blockchain-generated models into decentralized systems. This study focused on leveraging blockchain to enhance security while adhering to its architectural constraints. Key areas of investigation included: analyzing how network difficulty settings influence reward scaling and system stability, and evaluating the role of architectural variation in mitigating adversarial attacks and spoofing. These aspects highlight the system's capacity to follow and uphold the established principles

of blockchains. Our findings confirm the viability of this approach in both controlled environments and in empirical, large-scale distributed systems, where external manipulation attempts are common.

- Another contribution from our paper [GN25a] introduces a novel approach to distributed neural network initialization, enabling resource-constrained nodes to participate in training large models. The framework addresses both initialization and distributed training challenges, offering benefits such as reduced compute times, improved scalability, and enhanced accuracy for large and complex neural networks. By leveraging blockchain-inspired collaboration, the method allows low-resource participants to contribute meaningfully to training processes. The approach includes a model partitioning algorithm that splits neural networks into smaller segments for parallel training, distributed initialization of model weights to accelerate convergence, and padding techniques for handling odd-dimensional data. Empirical validation on datasets like MNIST, CIFAR-10, and Imagenet demonstrates improvements in training efficiency and accuracy.

The rest of the chapter is organized as follows. Section 3.1 details the proof of learning applied to the binary neural network algorithm. The literature review supporting the current research is presented in Subsection 3.1.1, while the problem statement and research questions are outlined in Subsection 3.1.2. Subsection 3.1.3 provides a detailed description of the algorithm and the results, including test 3.1.4. The conclusions regarding the algorithm are given in Subsection 3.1.5. Section 3.2 then introduces the accelerated neural network initialization algorithm, while the literature review is shown in Subsection 3.2.1. Its details and methodology are explained in Subsection 3.2.2, with test case results shown in Subsection 3.2.3. Finally, the conclusion for the initialization algorithm is presented in Subsection 3.2.4.

Chapter 4

Practical usecase

Scientific research thrives on bridging theoretical frameworks with real-world applications. A practical use case is essential to validate the need for methodologies and demonstrate adaptability in dynamic environments. This chapter presents a practical use case: developing a robust, parameter-free approach to parking space detection in camera-monitored lots. Traditional clustering algorithms like OPTICS or DBSCAN rely on fixed parameters, which struggle with uneven data distributions and require manual tuning. This limitation is particularly acute in parking lots, where vehicle occupancy patterns vary spatially and temporally. To overcome this, we propose an iterative training framework that eliminates reliance on predefined thresholds, ensuring results remain consistent across resolutions and camera angles.

The most suitable use case for iterative training is parking space detection, where the detected images of cars are of varying shapes, and outliers are trimmed. Here, iterative refinement is indispensable: parking lots exhibit irregular usage (e.g., underutilized zones, image distortions) and require dynamic adaptation to environmental shifts. By iterating over cluster size thresholds and leveraging data-derived metrics (e.g., median cluster width), our approach automatically adapts to density variations, eliminating the need for human intervention. Although this approach, on a standalone basis, is able to detect parking spaces really well, cars sometimes are not identified due to environmental factors. To eliminate this, we can use the samples of the detected cars to retrain the model and improve the algorithm's real-time detection. This method not only improves accuracy for real-time users, but it also showcases the need for a system like the one presented in this thesis, where models can be continuously retrained based on new data to continuously adapt and scale the application.

- The original article [GN24] contributes a novel, flexible methodology for parking space detection, combining hierarchical clustering (HDBSCAN), self-organizing maps (SOM), and K-means. Key contributions include the following aspects. An iterative denoising process that systematically reduces noise and irregular clusters by tuning minimum cluster sizes and assessing symmetry with width generalized thresholds. Topological adaptation through Self-Organizing Maps (SOMs) maintains spatial relationships and uses short-term memory to identify topology changes, such as shifts in camera perspective. Real-time accuracy is enhanced by retrained YOLOv8 models that utilize detected bounding boxes, boosting detection precision and lowering false negatives. Dynamic metrics that generate confidence scores and analyze spatial overlaps provide occupancy insights, supporting effective lot management.

The application can operate autonomously but lacks environmental information about parking lots. While models improve over time, they always generalize from their training data and are somewhat limited by it. Factors like environmental conditions or different lighting can reduce system

performance, which is image-centric. Although there are other algorithms in the app aimed at adapting to these constraints, overall, they are bound by the quality of the image detections. Retraining the convolutional neural networks with new, use-case-specific data is crucial and can be achieved by a modified blockchain that supplies models as a byproduct. Therefore, this application serves as a good use case of where our modified blockchain algorithm can bring in its advantages.

The remainder of this chapter is organized as follows. Section 4.1 outlines the application and its purpose, while Section 4.2 reviews existing approaches from the literature. In Section 4.3, the problem statement is provided, accompanied by research questions that guide the scope of the study. Section 4.4 details the methodology and algorithms utilized in the application. Section 4.5 presents the test cases and results. Lastly, Section 4.6 summarizes the conclusions and discusses the potential for extending the application.

Conclusions

The provided thesis outlines a comprehensive study that integrates blockchain technology with neural networks to enhance validation mechanisms for real-world applications, particularly focusing on parking lot availability detection. The research aims to demonstrate how repurposed validation mechanisms can be effectively applied using blockchain technology, emphasizing large-scale applications.

The structure of the thesis is methodically laid out across several chapters:

Chapter 1 introduced the problem statement, mechanics of blockchains, their inefficiencies, and foundational knowledge of neural networks. It reviewed existing literature to establish how these technologies can complement each other in achieving desired results.

In Chapter 2, the focus shifts to hardware aspects, discussing blockchain accelerator components and emphasizing the importance of hardware acceleration for speed and energy efficiency. The chapter explored FPGA devices and ASICs as potential solutions for executing mathematical operations on blockchains, underlining their scalability and integration capabilities.

Chapter 3 provided a software perspective, addressing challenges in repurposing validation algorithms for blockchain applications. It described necessary algorithmic adjustments for integrating new validation mechanisms within distributed systems.

The exploration continued in Chapter 4 with a use case involving parking space detection. This chapter emphasizes the need for frequent retraining with new data and suggests scaling the application to generate models continuously, thus improving neural network accuracy without specialized hardware.

Using an application like the one described in Chapter 4 on a larger scale requires training and centralizing numerous models and datasets. The steady stream of data and the need for frequent retraining demand substantial computational resources. Additionally, security and traceability are crucial; with centralized data, it must be protected from unauthorized modifications and access to sensitive personal information. Making these datasets publicly accessible and regularly updated offers the open source community the chance to leverage the data and develop innovative new applications.

Our thesis explored both the empirical and practical dimensions of the proposed approach. The foundational work supports full blockchains that incorporate our adapted validation algorithm. The testing examines various factors relevant to large-scale implementations of the chain. The chapters presented the core building blocks, covering hardware aspects essential for modern high-capacity computations and elaborating software algorithms capable of implementing a blockchain that offers additional value to a highly secure transactional ledger. Through our contributions across multiple domains related to this subject, we have scientifically reviewed the approach and demonstrated its viability.

In conclusion, the thesis summarized its findings and proposed future directions for leveraging the advantages of the proposed approach. It presented an innovative way by combining blockchain technology with neural networks to address real-world problems, such as parking lot availability detection, suggesting a sustainable model that could alleviate computational loads on data centers. The emphasis on continuous retraining and scalability highlights the potential for adaptive systems

capable of evolving with changing conditions.

Future explorations might focus on enhancing the efficiency and scalability of blockchain-integrated neural networks, exploring broader applications beyond parking lot detection, and addressing any identified limitations or challenges in the current approach.

Bibliography

[BT13] Michael Bedford Taylor. Bitcoin and the age of bespoke silicon. In *2013 International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES)*, pages 1–10, 2013.

[FL20] Zonghao Feng and Qiong Luo. Evaluating memory-hard proof-of-work algorithms on three processors. *Proc. VLDB Endow.*, 13(6):898–911, February 2020.

[G+16] Arthur Gervais et al. On the Security and Performance of Proof of Work Blockchains. In *CCS 2016*, page 3–16, 2016.

[GN23] Zoltán-Valentin Gyulai-Nagy. Acceleration of Neural Network training algorithm acceleration via FPGA Devices. In *The 27th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems (KES 2023)*, pages 2674–2683, 2023. Procedia Computer Science, Volume 225.

[GN24] Zoltán-Valentin Gyulai-Nagy. Combining clustering algorithms to extract symmetric clusters from noisy data, applied to parking lots. In *Proceedings of the 19th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2024)*, pages 1–8, 2024.

[GN25a] Zoltán-Valentin Gyulai-Nagy. Accelerated Neural Network Initialization via Training of Symmetrically Partitioned Models. *Acta Universitatis Sapientiae, Informatica*, 17:Article no. 12, 2025.

[GN25b] Zoltán-Valentin Gyulai-Nagy. Exploring the Usability and Applicability of Proof of Learning in Binary Neural Networks. *Lecture Notes in Artificial Intelligence (LNAI)*, page to be published, 2025.

[GN25c] Zoltán-Valentin Gyulai-Nagy. Proof of Learning applied to Binary Neural Networks. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025)*, pages 1–8, 2025.

[Pil16] Marc Pilkington. Blockchain technology: principles and applications. In *Research handbook on digital transformations*. Edward Elgar Publishing, 2016.

[Sed20] Buhl-H.U. Fridgen G. et al. Sedlmeir, J. The energy consumption of blockchain. *Business Information Systems Engineering*, 2020.

[ZLL+24] Haiyun Zhu, Jiaqi Liu, Xianxu Li, Zhiqin Huang, and Yong Zhang. A power consumption measurement method for large ai-based intelligent computing servers. In *Proceedings of the 2023 5th International Conference on Internet of Things, Automation and Artificial Intelligence*, IoTAAI '23, page 150–155, New York, NY, USA, 2024. Association for Computing Machinery.