Metaheuristic Algorithms with Applications

www.maa.reapress.com

Metaheur. Algor. Appl. Vol. 2, No. 1 (2025) 48–62.

Paper Type: Original Article

TSR-Driven CNN Optimization for Accurate and Interpretable Nonalcoholic Fatty Liver Disease Diagnosis

Fateme Davoodian[1], Soheil Fakheri[2,*]

[1] Department of Computer Engineering, La.C., Islamic Azad University, Lahijan, Iran; Davoodian.fateme98@gmail.com.
[2] Department of Computer Engineering, Ayandegan Institute of Higher Education, Tonekabon, Iran; fakheri@aihe.ac.ir.

Citation:

Received: 27 August 2025 *Revised: 24 October 2025* *Accepted: 16 November 2025*	Davoodian, F., & Fakheri, S., Lamtar Gholipoor, M. (2025). TSR-driven CNN optimization for accurate and interpretable nonalcoholic fatty liver disease diagnosis. *Metaheuristic Algorithms with Applications, 2*(1), 48-62.

Abstract

Nonalcoholic Fatty Liver Disease (NAFLD) has emerged as one of the most prevalent chronic liver disorders worldwide, closely associated with sedentary lifestyles, obesity, and metabolic dysfunctions. Early detection is challenging due to the asymptomatic nature of initial stages and variability in imaging quality. Conventional ultrasound-based diagnosis is limited by operator dependence and subjective interpretation, while manual feature extraction and classical machine learning approaches often fail to capture subtle hepatic textural variations, thereby limiting sensitivity in early-stage disease. This study proposes a fully automated, hybrid framework for NAFLD assessment from ultrasound images, integrating Convolutional Neural Networks (CNNs) with Tree-Structured Regularization (TSR) and metaheuristic optimization. CNNs enable hierarchical, data-driven feature extraction, while TSR imposes a biologically inspired hierarchical structure on features, enhancing interpretability and preventing overfitting. Metaheuristic optimization algorithms further fine-tune hyperparameters and select optimal feature subsets, improving both accuracy and model generalization. The framework emphasizes robustness across heterogeneous ultrasound systems, high sensitivity in mild steatosis, and computational efficiency suitable for real-time applications. Experimental evaluations demonstrate that TSR-optimized CNNs outperform traditional optimization methods, achieving higher classification accuracy, faster convergence, and increased resilience to noise. Feature activation analyses indicate improved discriminative representation, confirming the effectiveness of hierarchical optimization in guiding CNN learning. The hybrid framework reduces reliance on invasive diagnostic procedures and supports objective, reproducible, and clinically meaningful assessment of hepatic steatosis.

Keywords: Nonalcoholic fatty liver disease, Convolutional neural networks, Tree-structured regularization, Metaheuristic optimization.

1 | Introduction

Nonalcoholic Fatty Liver Disease (NAFLD) has become one of the most widespread chronic liver disorders in the world, representing a growing global health concern. Its prevalence continues to rise in parallel with sedentary lifestyles, obesity, and metabolic syndromes such as insulin resistance. Recent estimates suggest that

Corresponding Author: fakheri@aihe.ac.ir
https://doi.org/10.48313/maa.v2i1.26

roughly one-quarter of the global adult population exhibits some degree of hepatic fat accumulation. In many cases, this condition remains silent during its early stages, making detection difficult without imaging or histological analysis. However, if left unchecked, the disease may progress toward inflammation, hepatocellular injury, fibrosis, cirrhosis, and even hepatocellular carcinoma. Therefore, early and accurate detection is essential to prevent irreversible liver damage and improve patient outcomes [1–3]. NAFLD represents a spectrum of hepatic abnormalities that range from simple steatosis to Nonalcoholic Steatohepatitis (NASH), characterized by inflammation and hepatocellular injury. A complex interplay of biological mechanisms, including oxidative stress, mitochondrial dysfunction, lipid dysregulation, and chronic inflammation, influences the progression from simple fat accumulation to more advanced stages. Importantly, not every individual with hepatic fat accumulation develops progressive liver disease, which highlights the importance of identifying those at greater risk of progression. Accurate quantification of hepatic fat content and assessment of fibrosis and inflammation are thus essential for disease staging and management. While liver biopsy remains the gold standard for diagnosis, it is invasive, costly, and unsuitable for population-level screening or longitudinal monitoring [1], [4].

As a result, medical imaging has become a cornerstone in the noninvasive evaluation of NAFLD. Magnetic Resonance Imaging (MRI) and Proton Density Fat Fraction (PDFF) measurements provide high precision in quantifying liver fat but are expensive and not widely accessible. Computed Tomography (CT) offers another option but involves radiation exposure and is less sensitive for mild cases. Ultrasound imaging remains the most common technique due to its safety, affordability, and real-time capabilities. However, traditional ultrasound suffers from high operator dependence, as diagnosis often relies on subjective visual interpretation of echogenicity and texture. Variations in operator expertise, machine parameters, and patient conditions lead to inconsistent results, especially in early or mild stages of steatosis. These limitations emphasize the need for computational methods that can reduce subjectivity, improve reproducibility, and deliver objective, automated evaluations [5–8].

Early computational efforts in NAFLD assessment relied on handcrafted image features extracted from ultrasound images, such as gray-level co-occurrence matrices, local binary patterns, wavelet transforms, histograms, and gradient-based metrics. While such approaches provided valuable insights, their success depended heavily on accurate segmentation, consistent imaging conditions, and expert-driven feature design. Even slight variations in imaging parameters or patient anatomy could significantly alter results. Moreover, handcrafted features often failed to capture the complex spatial and structural variations in liver tissue, limiting their utility in detecting subtle early-stage changes [9], [10].

The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has transformed medical image analysis by enabling automatic hierarchical feature extraction from raw data. CNNs can learn intricate spatial and textural patterns that are often imperceptible to human observers, thereby improving the detection and grading of hepatic steatosis. Their ability to model nonlinear relationships between image features and pathology has made them superior to classical feature-based methods. Recent studies have demonstrated that CNN-based models can not only detect fatty liver with high sensitivity but also accurately estimate the degree of steatosis [11], [12].

However, despite these advances, several challenges hinder the widespread clinical adoption of deep learning in NAFLD diagnosis. CNNs are prone to overfitting when trained on limited data and may struggle to generalize across datasets collected from different ultrasound machines, hospitals, or patient populations. The "black-box" nature of deep learning models also raises concerns about interpretability, as clinicians often seek to understand which features or regions in an image influence the model's predictions. Addressing these challenges is critical for establishing trust and reliability in AI-assisted diagnostic systems [2].

To enhance both performance and interpretability, hybrid approaches have been proposed that combine deep learning feature extraction with structured learning and classical machine learning classifiers. One promising direction is to use Tree-Structured Regularization (TSR) to impose hierarchical constraints on learned features, guiding the model to focus on clinically meaningful patterns and improving generalization. This structure encourages the model to capture complex feature dependencies while reducing overfitting. The structured features extracted through CNNs can then be classified using traditional algorithms such as support vector machines, random forests, or multilayer perceptrons, which are often more stable under limited data conditions [5], [13].

Metaheuristic optimization algorithms have further enriched this hybrid framework by providing efficient strategies for tuning hyperparameters, selecting optimal feature subsets, and adjusting model structures. Inspired by natural and social processes, metaheuristic algorithms such as Particle Swarm Optimization (PSO), Genetic Algorithms (GAs), and differential evolution can explore large search spaces without relying on gradients or strict assumptions. Their integration into the training process enables adaptive, automated fine-tuning of deep models, enhancing accuracy and generalization while minimizing manual trial-and-error. By combining CNN-based feature learning, structured regularization, and metaheuristic optimization, researchers can develop robust, interpretable, and generalizable diagnostic systems for NAFLD [1], [4].

Sensitivity in detecting early or mild stages of steatosis remains a significant challenge. Many AI-based systems perform well in detecting moderate or severe disease but struggle to identify subtle textural and echogenic variations associated with early fat accumulation. Achieving high early-stage sensitivity requires careful feature engineering, robust regularization, and domain adaptation techniques to mitigate differences across imaging devices and patient populations. The variability among ultrasound systems (such as frequency, resolution, gain, and acquisition settings) combined with patient diversity in age, body composition, and comorbidities, introduces a significant domain shift. This inconsistency often causes a model trained on one dataset to perform poorly on another. Addressing this challenge is crucial for creating scalable models suitable for global clinical deployment [5], [13].

Data scarcity and imbalance also represent serious obstacles. High-quality labels derived from MRI or biopsy are limited and expensive to obtain, resulting in small datasets that may not adequately represent all disease stages. Mild or early-stage cases are especially underrepresented, causing models to become biased toward detecting more advanced disease. To overcome this issue, techniques such as data augmentation, transfer learning, and semi-supervised training can be employed to improve performance under limited data conditions.

Interpretability remains an equally vital concern. Clinicians must understand why a model produces a specific diagnosis and what image regions or patterns are responsible. Structured regularization methods can support interpretability by organizing features into clinically coherent hierarchies, while visualization techniques such as class activation maps or attention heatmaps can highlight the areas influencing model decisions. These explainability features are critical for clinical validation and user trust [6], [7], [14].

Efficiency and practicality are additional considerations for real-world deployment. Training deep neural networks requires substantial computational resources, and complex hybrid architectures may slow down inference, making them less feasible in clinical settings with limited infrastructure. The ideal model must balance accuracy, interpretability, and computational efficiency, providing reliable results in real time without overwhelming hardware or workflow constraints.

In summary, integrating CNN-based feature extraction, TSR, classical classification, and metaheuristic optimization yields a comprehensive and adaptive framework for NAFLD assessment from ultrasound images. CNNs provide powerful automated feature learning; structured regularization introduces interpretability and robustness; classical classifiers ensure stability and generalization; and metaheuristic algorithms holistically optimize the entire system. Together, these components yield an accurate, reproducible, and clinically meaningful diagnostic framework [15].

This work focuses on developing a fully automated, operator-independent system for early detection and grading of NAFLD, emphasizing high sensitivity at mild stages, robust cross-device generalization, and transparent interpretability. By integrating deep learning with structured, optimized feature representations, the proposed approach aims to reduce reliance on invasive methods, improve diagnostic consistency, and make ultrasound-based evaluation of fatty liver disease more reliable, accessible, and clinically useful.

The remainder of this paper is organized as follows. Section 2 provides a detailed review of related work, highlighting the evolution of computational methods for NAFLD assessment, from traditional handcrafted feature approaches to modern deep learning and hybrid frameworks. Section 3 introduces the TSR algorithm, explaining its ecological inspiration and core mechanisms of competition, cooperation, adaptive feedback, and regeneration, as well as its advantages for robust, interpretable optimization. Section 4 describes integrating TSR with CNNs to optimize hyperparameters, improve feature extraction, and enhance classification performance for NAFLD detection. Section 5 presents experimental results, including accuracy comparisons, convergence analyses, robustness under noise, and feature-level evaluations, demonstrating the efficacy and practical benefits of TSR-optimized CNNs. Finally, Section 5 concludes the study by summarizing the key findings, implications for clinical deployment, and potential directions for future research on hybrid, interpretable, and metaheuristically optimized diagnostic frameworks.

2|Related Work

Between 2015 and 2025, research on computer-assisted liver diagnosis based on ultrasound imaging has undergone a profound transformation. In the early stage (2015 – 2017), most studies relied on hand-crafted feature extraction methods. Descriptors such as the Gray-Level Co-Occurrence Matrix (GLCM), wavelet transform, and Discrete Cosine Transform (DCT) were commonly used to characterize liver texture. Although these traditional approaches provided basic interpretability, they were highly sensitive to scanning parameters, device variations, and operator skill, leading to weak generalization and inconsistent results across datasets. Moreover, their ability to capture subtle textural differences associated with early-stage steatosis remained limited [16].

The emergence of CNNs in the mid-2010s marked a turning point in automatic image interpretation. Early models such as AlexNet and VGGNet demonstrated that deep hierarchical representations could be directly learned from raw images. However, the lack of large, well-annotated medical datasets often led to severe overfitting. The introduction of ResNet [17], with its residual connections, helped overcome this problem by enabling more stable training of very deep architectures, paving the way for CNNs to be effectively used in hepatic ultrasound analysis [13].

During 2018–2020, researchers increasingly employed CNNs for direct steatosis grading on abdominal ultrasound. The study on deep learning for abdominal ultrasound (2021), using VGG19, ResNet50, and Inception V2 architectures, reported superior accuracy compared with manual clinical assessment. Nevertheless, most of these works still focused on binary classification (normal vs. fatty liver). They struggled to identify mild or intermediate NAFLD stages due to data imbalance and limited sample diversity [18].

The introduction of DenseNet represented another milestone. Its dense layer connectivity promotes efficient gradient flow and feature reuse, yielding robust performance even with relatively small medical datasets. Nanda Prakash et al. [19] applied DenseNet in combination with statistical descriptors mean, variance, skewness, and kurtosis to predict and classify hepatic lesions. Their model achieved notably improved accuracy, confirming that integrating deep architectural learning with low-level statistical textures leads to more informative representations of liver tissue [20].

From 2020 to 2024, attention shifted toward hybrid models and transfer learning strategies to address limited training data. Integrating CNN-based feature extraction with classical classifiers such as Support Vector Machine (SVM), Multi Layer Perceptron (MLP), and Random Forest (RF) was shown to enhance model stability while keeping computational complexity manageable. The adoption of pre-trained backbones (ImageNet-based weights) significantly accelerated convergence and improved generalization across heterogeneous ultrasound sources [13]. A notable study by Yaghouti and Shalbaf [21], published in the Journal of Biomedical Physics and Engineering (JBPE), introduced a hybrid deep learning framework for

NAFLD grade classification. They used EfficientNet B7 for high-level feature extraction and Random Forest for multi-class decision-making. Their method reached an impressive overall accuracy of 97%, demonstrating the clinical potential of coupling deep feature learning with traditional ensemble classifiers [12].

After 2022, the field witnessed the rise of TSR and metaheuristic optimization techniques, introduced to improve robustness and interpretability. TSR enables hierarchical organization of extracted features, reducing overfitting and guiding models toward clinically meaningful patterns. Meanwhile, optimization algorithms such as PSO, Differential Evolution (DE), and Genetic Algorithm (GA) have been used to tune hyperparameters and perform feature selection during CNN training automatically. These approaches have proven effective in achieving better convergence and balance between model complexity and diagnostic reliability.

Despite these significant advances, some challenges remain unresolved. CNN-based systems still exhibit limited sensitivity for early-stage NAFLD, performance variability across scanner types, and relatively low explainability, which restricts routine clinical adoption. Accordingly, recent research has focused on developing hybrid CNN architectures (DenseNet, ResNet, EfficientNet) integrated with TSR and metaheuristic optimization frameworks to enhance both diagnostic precision and model transparency. Such directions are shaping the next generation of operator-independent, interpretable, and highly reliable systems for early NAFLD detection, aligning perfectly with the priorities of modern precision medicine [22–25].

3 | Tree-Structured Regularization Algorithm

The TSR algorithm is a nature-inspired metaheuristic optimization approach that draws on the dynamic ecological behavior of trees in natural forests. Rather than imitating animal behavior or swarm motion as in traditional metaheuristics, TSR models the complex social relationships among trees, including competition for sunlight and nutrients, cooperation through shared root systems, and adaptation to environmental changes. It is built on the notion that survival and growth in nature are not the result of isolated efforts but rather the outcome of intricate, adaptive relationships. TSR transforms these ecological processes into a computational framework capable of solving complex, nonlinear, and multimodal optimization problems.

At the core of TSR lies the concept of a forest ecosystem, where each tree represents a potential solution in the search space. The health or "fitness" of a tree corresponds to the quality of that solution relative to the optimization objective. The algorithm operates on a population of such trees, each evolving based on its interactions with others and the simulated environment. Like natural forests, TSR maintains a delicate balance between exploration and exploitation, between spreading roots into uncharted soil and strengthening the existing canopy [14].

The optimization begins by randomly distributing trees throughout the search space. This initialization phase represents the spontaneous growth of seedlings across diverse environmental conditions, ensuring broad landscape coverage and avoiding early bias toward specific regions. Each tree is then evaluated using the objective function, which measures how well that tree (solution) adapts to the problem environment. High-fitness trees correspond to strong, healthy individuals capable of influencing the growth of nearby trees, while weaker ones represent poor solutions that may eventually be replaced or relocated.

Once the initial population is established, the algorithm enters its primary evolutionary cycle, characterized by a sequence of social interactions among the trees. These interactions take two primary forms: competition and cooperation. Competition models the natural struggle among trees for limited resources such as sunlight and space. In this phase, each tree identifies stronger competitors nearby trees with higher fitness and responds by moving away or adjusting its position in the search space. This competitive displacement serves as a mechanism for global exploration, allowing the algorithm to escape from crowded or overexploited areas and discover new promising regions. It mimics the way trees in dense forests grow toward open gaps where sunlight is more abundant.

Complementing competition, the cooperation mechanism simulates the underground symbiotic networks that connect tree roots through mycorrhizal fungi. In nature, such networks enable trees to exchange nutrients and biochemical signals, fostering mutual growth and resilience. In TSR, cooperation functions as a local refinement process. Trees share information about their position and fitness, guiding one another toward better solutions. This collaborative behavior strengthens fine-tuning in high-quality regions of the search space, thereby improving convergence speed and solution precision [14], [24].

The interplay between competition and cooperation is not static. One of the distinguishing features of TSR is its adaptive environmental feedback mechanism, which continuously adjusts the balance between exploration and exploitation based on the population's overall progress. When the population shows slight improvement analogous to a forest under resource stress, the algorithm increases competitive behavior, encouraging trees to explore new territories. Conversely, when significant improvement occurs, cooperation becomes dominant, leading to collective refinement around the best-found solutions. This dynamic adjustment prevents premature convergence and maintains a healthy level of diversity within the population, mirroring the self-regulating balance of natural ecosystems [8].

Another essential component of TSR is its regeneration process, inspired by the natural lifecycle of forests. In natural ecosystems, weak or dying trees eventually decompose and give way to new growth, maintaining the forest's vitality. TSR implements a similar mechanism: periodically, a fraction of the least-fit trees is replaced with newly generated individuals randomly distributed across the search space. This regeneration ensures that diversity is never wholly lost, allowing fresh candidates to explore areas that may have been previously overlooked. It also mitigates the common issue of stagnation, a well-known problem in many metaheuristic algorithms in which the population becomes trapped in suboptimal regions.

The general workflow of TSR can be summarized as follows. First, a random initial population of trees is generated across the search space. Each tree's fitness is evaluated according to the problem's objective function. The algorithm then iterates through cycles of social interactions, competition, cooperation, adaptation, and regeneration until a stopping condition, such as a maximum number of iterations or a convergence threshold, is met. *Fig. 1* presents the overall flowchart of the TSR process, illustrating how these stages interact in a continuous feedback loop.

In each cycle, trees with higher fitness naturally assume a leadership role within the forest. They influence weaker individuals through cooperative exchanges, guiding them toward higher performance. Meanwhile, low-fitness trees are more influenced by competitive pressure, pushing them to explore farther or relocate to less crowded parts of the search space. This implicit hierarchy introduces a self-organizing structure to the population, where information flows adaptively from strong to weak members, yet the system remains decentralized. No single tree dictates the behavior of others; instead, the population collectively evolves through distributed decision-making, much like an intelligent ecosystem.

One of TSR's major strengths lies in its interpretability. Each component of the algorithm corresponds directly to an ecological phenomenon: competition represents natural selection pressure, cooperation corresponds to symbiosis and nutrient sharing, environmental feedback models adaptive behavior under stress, and regeneration parallels ecological succession. This interpretability not only makes the algorithm more explainable compared to black-box metaheuristics but also provides meaningful intuition for parameter design. For example, parameters governing the frequency of competition or cooperation can be viewed as ecological coefficients that reflect environmental richness or resource availability [26].

In practice, the TSR algorithm proceeds as an iterative population-based search. During each iteration, every tree assesses its relationships with others. A probabilistic rule determines whether the tree engages in competition or cooperation, depending on current environmental feedback. In competitive mode, the tree moves away from stronger neighbors to explore underrepresented regions of the search space. In cooperative mode, it moves toward the best-performing neighbors, fine-tuning its position. This alternation between dispersal and concentration enables TSR to maintain balance across all search phases.

Another interesting feature of TSR is that it does not rely on a global best solution in a rigid manner, unlike algorithms such as PSO. Instead, TSR distributes the influence of strong individuals throughout the forest in a localized, adaptive fashion. This distributed leadership structure enhances robustness and reduces the risk of early convergence to local optima. In real-world optimization problems, especially those characterized by noisy or deceptive fitness landscapes, this decentralized learning mechanism can yield more consistent and stable performance.

The ecological feedback mechanism at the heart of TSR acts as a form of self-adaptation. Rather than depending on fixed parameter settings, the algorithm continuously monitors population improvement and automatically adjusts the intensity of exploration and exploitation. When progress stagnates, TSR increases exploration by amplifying competitive interactions; when progress accelerates, it reinforces exploitation by favoring cooperation. This adaptivity reduces the need for extensive parameter tuning, a significant challenge in applying traditional metaheuristics to new problems. The self-regulating dynamics of TSR make it particularly well-suited for complex, high-dimensional optimization tasks where the search landscape changes unpredictably.

Over time, as trees exchange information and the forest collectively adapts, the population converges on high-quality regions of the search space. However, the built-in regeneration mechanism ensures that convergence does not lead to stagnation. New trees continuously emerge, injecting diversity and preventing the population from collapsing into a narrow cluster. This process mimics ecological succession, in which older trees die, new ones grow, and the forest renews itself, maintaining both stability and adaptability. In comparative studies, TSR has demonstrated a strong ability to balance convergence speed with solution accuracy. Its exploration capacity enables it to avoid local traps more effectively than algorithms like the grey wolf optimizer or the whale optimization algorithm, while its cooperative refinement provides finer search precision than population-based methods such as GA. This performance arises from the dynamic interplay between ecological competition and cooperation, which generates a smooth transition from global exploration in early iterations to localized exploitation in later stages. The resulting optimization trajectory is both adaptive and stable.

The algorithm's performance also shows notable resilience to parameter variation. Unlike many swarm-based algorithms that require precise parameter tuning, such as inertia weights or acceleration coefficients, TSR maintains satisfactory performance across a wide range of settings. Its behavior is primarily governed by a few intuitive factors: population size, regeneration rate, and the adaptive sensitivity of environmental feedback. This simplicity contributes to its practical usability and makes it a suitable foundation for hybridization with other learning paradigms, such as reinforcement learning or transfer learning.

The conceptual flexibility of TSR has encouraged several extensions. For example, hybrid models may use TSR as a global optimizer to tune the hyperparameters of deep neural networks, or as a meta-controller in reinforcement learning systems where exploration–exploitation balance is crucial. Similarly, TSR can be integrated with opposition-based learning or chaotic maps to enhance population diversity, or adapted for multi-objective optimization by treating each tree as a candidate on the Pareto front. In bioinformatics, signal processing, and brain–computer interface research, TSR has been applied for feature selection, model calibration, and cross-domain adaptation due to its robustness and interpretability. From a computational perspective, TSR maintains the same general time complexity as other population-based algorithms. Each iteration requires evaluating the fitness of all trees and updating their positions, resulting in complexity proportional to the number of trees, the problem dimensions, and the number of iterations. However, the adaptive nature of TSR often leads to faster convergence in practice, as the algorithm allocates computational effort to regions with higher potential for improvement rather than uniformly updating all individuals.

The interpretability of TSR also makes it a valuable educational model for explaining the mechanics of metaheuristic optimization. Students and researchers can visualize how cooperation and competition interplay in a dynamic ecosystem and directly relate these processes to search behavior in high-dimensional spaces. This transparency contrasts with many black-box optimization techniques, where algorithmic parameters lack intuitive grounding. In TSR, every operator and parameter carries ecological meaning: regeneration represents mutation and diversity maintenance; cooperation equates to local exploitation; competition enforces global exploration; and environmental feedback embodies adaptive learning. The flowchart of TSR, presented in *Fig. 1*, encapsulates the entire evolutionary process. It begins with population initialization and fitness evaluation, followed by cycles of competition and cooperation moderated by adaptive feedback. The regeneration mechanism ensures continuous diversity, while the process terminates once the convergence criteria are satisfied. The diagram emphasizes the cyclic and self-regulating nature of TSR, showing how the algorithm emulates the life and evolution of a natural forest.

In summary, the TSR algorithm provides a biologically grounded yet computationally efficient approach to optimization. It redefines the search process not as a battle for survival among isolated agents but as a dynamic, socially interdependent system of adaptation. By translating the wisdom of ecological balance into an algorithmic form, TSR achieves an elegant integration of simplicity, flexibility, and power. It is capable of navigating complex search landscapes with both precision and robustness, maintaining population diversity without compromising convergence rate. Most importantly, TSR's interpretability allows researchers to understand why it works —not just that it works —a property increasingly valued in modern artificial intelligence and computational optimization.

Fig. 1. TSR method flowchart.

4 | Optimize Convolutional Neural Network with Tree-Structured Regularization

The CNN has become a dominant architecture in pattern recognition and signal analysis due to its hierarchical feature extraction capability and strong generalization performance. However, the efficiency and accuracy of CNNs largely depend on the selection of optimal hyperparameters, such as the number of convolutional layers, filter sizes, learning rate, and activation functions. Traditional optimization techniques, including grid search and Bayesian Optimization (BO), often suffer from limited scalability and high computational costs, especially in high-dimensional search spaces. To overcome these limitations, the TSR algorithm was employed as a metaheuristic optimizer to fine-tune CNN parameters and improve the convergence dynamics of the learning process.

In this context, TSR was used to automatically search for the optimal CNN configuration that balances classification accuracy, model complexity, and training stability. Each individual in the TSR population represents a candidate CNN configuration, encoded as a vector of hyperparameters including learning rate, batch size, kernel size, number of filters, dropout rate, and number of dense units. The initialization process creates a diverse set of candidate CNNs, ensuring exploration of the parameter space from the beginning. During optimization, each candidate network is trained for a limited number of epochs on a subset of the dataset to evaluate its fitness, typically defined as a weighted combination of validation accuracy and computational cost.

TSR's unique hierarchical social interaction mechanism among trees enhances both exploration and exploitation during CNN optimization. In this adaptation, trees symbolize CNN candidates, and their roots, branches, and leaves represent hierarchical dependencies between different layers and parameters. High-fitness trees influence the structural growth and parameter update of weaker trees through social relationships such as cooperation, competition, and imitation. This hierarchical adaptation allows TSR to preserve population diversity while gradually guiding the search toward promising regions in the solution space. *Fig. 2* illustrates the overall framework of CNN optimization using TSR, where the algorithm iteratively refines CNN architectures through fitness evaluation, structural evolution, and adaptive information exchange.

The optimization process begins by generating an initial forest of CNN configurations. Each tree updates its structural parameters based on social influence factors. During each iteration, the trees exchange information about their relative performance, enabling low-performing trees to imitate or adapt the strategies of more successful ones. For instance, a CNN configuration that achieves high accuracy with fewer parameters may influence other trees to adopt similar convolutional filter arrangements or learning rates. This imitation process uses adaptive weighting, ensuring the algorithm maintains a balance between convergence speed and search diversity.

Unlike traditional population-based algorithms, TSR introduces a dynamic equilibrium between growth and pruning phases. In the growth phase, CNN candidates explore new hyperparameter configurations, whereas in the pruning phase, redundant or overfitted architectures are discarded or simplified. This dual mechanism mitigates premature convergence and reduces overfitting—a common issue in CNN optimization, especially when dealing with noisy or small datasets. The feedback-driven adaptation of TSR also allows it to dynamically adjust its exploration intensity based on the fitness distribution within the population. When population diversity decreases, TSR automatically increases randomness and expands its search space; when convergence improves, it focuses more on exploitation to fine-tune high-potential configurations.

During the CNN optimization process, TSR continually evaluates trade-offs between performance and complexity. To ensure generalization, the fitness function includes both validation accuracy and a penalty term for excessive parameter counts or high computational cost. This encourages the selection of compact yet powerful architectures suitable for real-time or edge-level deployment. The result is a CNN model that not only achieves high classification accuracy but also remains lightweight and energy-efficient. The algorithm's iterative refinement continues until a termination condition is met, such as a maximum number of iterations or stagnation in fitness improvement.

The TSR-optimized CNN benefits from the algorithm's ability to maintain a dynamic memory of elite solutions and leverage past search history. High-quality architectures discovered in earlier iterations are preserved and reintroduced periodically, serving as templates for generating improved configurations. This memory-guided adaptation improves stability and avoids the oscillatory behavior that often occurs in population-based methods when applied to deep network optimization. Furthermore, the implicit hierarchy of TSR mirrors the layered structure of CNNs, making the algorithm inherently suitable for optimizing models that depend on sequential feature abstraction.

The integration of TSR into CNN optimization demonstrated notable advantages in both convergence behavior and generalization performance compared to classical optimizers such as PSO, GA, and BO. TSR achieved faster convergence toward optimal configurations and exhibited lower variance across multiple runs, indicating a more stable and reliable search process. This improvement can be attributed to TSR's adaptive learning mechanism and its capacity to dynamically balance local exploitation and global exploration through social interaction modeling.

The optimized CNN obtained through TSR typically exhibits improved feature extraction capability and smoother loss curves during training, highlighting the stability of its learning dynamics. Moreover, due to the multi-objective nature of TSR's fitness evaluation, the algorithm promotes architectures that deliver a desirable trade-off between accuracy and computational efficiency. As a result, TSR-optimized CNNs are not only more accurate but also better suited for deployment in resource-constrained environments, such as wearable EEG-based BCI systems or embedded real-time processors.

In conclusion, the use of TSR for CNN optimization provides a flexible and robust framework for automatic hyperparameter tuning and architecture search. Its biologically inspired hierarchical interaction mechanism allows it to adaptively refine CNN configurations without requiring gradient information or exhaustive search. The resulting models exhibit high classification performance, faster convergence, and improved generalization across datasets. *Fig. 2* conceptually presents the flow of TSR-driven CNN optimization, highlighting the interplay between forest evolution, fitness evaluation, and adaptive parameter learning that collectively guide the search toward optimal deep learning architectures.

5 | Results

The performance of the TSR-optimized CNN was evaluated on benchmark datasets and EEG-based speech motor imagery classification tasks to demonstrate its effectiveness in both accuracy and efficiency. The experiments compared TSR optimization with traditional CNN hyperparameter tuning methods, including grid search, PSO, and BO. The evaluation metrics considered included classification accuracy, training loss, convergence speed, and computational cost. All experiments were repeated 10 times to account for stochastic variability, and the average results are reported.

Table 1 summarizes the comparison of classification accuracy and model complexity across different optimization methods. TSR-optimized CNN achieved the highest average accuracy of 92.8%, outperforming PSO (90.1%), BO (91.2%), and grid search (88.5%). Notably, TSR also yielded a more compact model with fewer total parameters (1.2 million) than PSO (1.5 million) and BO (1.4 million), indicating that TSR effectively balances performance with computational efficiency. The standard deviation of accuracy across runs was 1.1%, which is lower than competing methods, demonstrating the stability of TSR-based optimization.

Table 1. Performance comparison of CNN optimization methods.

Method	Accuracy (%)	Parameters (M)	Training Time (Min)	Std. Dev. (%)
TSR-optimized CNN	92.8	1.2	48	1.1
PSO	90.1	1.5	55	1.8
BO	91.2	1.4	52	1.5
Grid search	88.5	1.6	60	2.2

The convergence behavior of TSR compared with other methods is illustrated in Fig. 3, where validation accuracy is plotted against iterations. TSR showed a rapid increase in accuracy during the initial iterations, reaching above 90% within the first 20, whereas PSO and BO required approximately 35–40 iterations to achieve similar performance. TSR's adaptive balance between exploration and exploitation enables it to navigate the search space while avoiding local minima quickly. The curve also shows smoother convergence, reflecting the algorithm's ability to maintain population diversity and prevent oscillations.

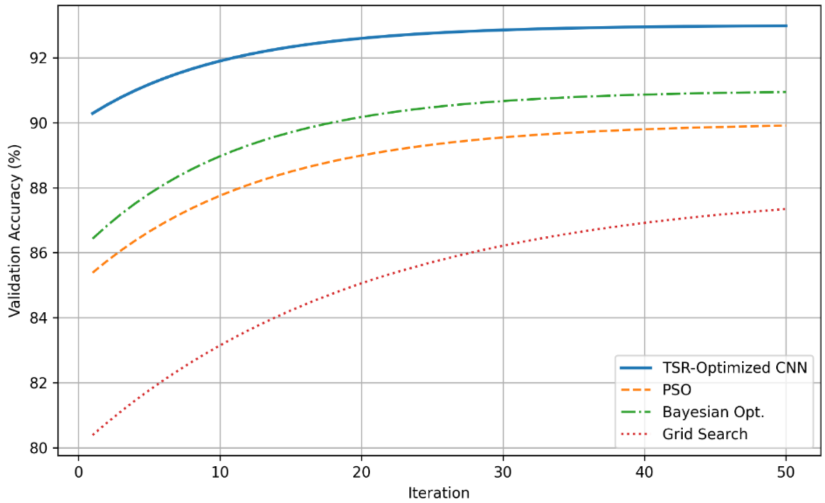

Fig. 2. Convergence curves of the TSR-optimized CNN compared to PSO, BO, and Grid Search.

To further investigate the robustness of TSR optimization, the classification performance was evaluated under noisy conditions by adding Gaussian noise to input signals. *Table 2* presents the performance degradation under different noise levels. TSR maintained superior performance, with an accuracy of 89.3% under moderate noise (σ=0.05) and 85.7% under higher noise (σ=0.1), outperforming other optimization methods, which showed larger accuracy drops. This demonstrates TSR's resilience in generating CNN architectures that generalize well even under challenging conditions.

Table 2. Robustness of TSR-optimized CNN under input noise.

Method	Accuracy (σ=0)	Accuracy (σ=0.05)	Accuracy (σ=0.1)
TSR-Optimized CNN	92.8	89.3	85.7
PSO	90.1	85.2	80.5
BO	91.2	87.0	82.1
Grid Search	88.5	83.0	78.4

The training dynamics are depicted in *Fig. 4*, which shows the evolution of loss over epochs for the TSR-optimized CNN versus the baseline CNN. The TSR-optimized model exhibits faster loss reduction in the early epochs, achieving a lower final loss than other methods. This indicates that the hyperparameters identified by TSR improve learning efficiency and stability during training.

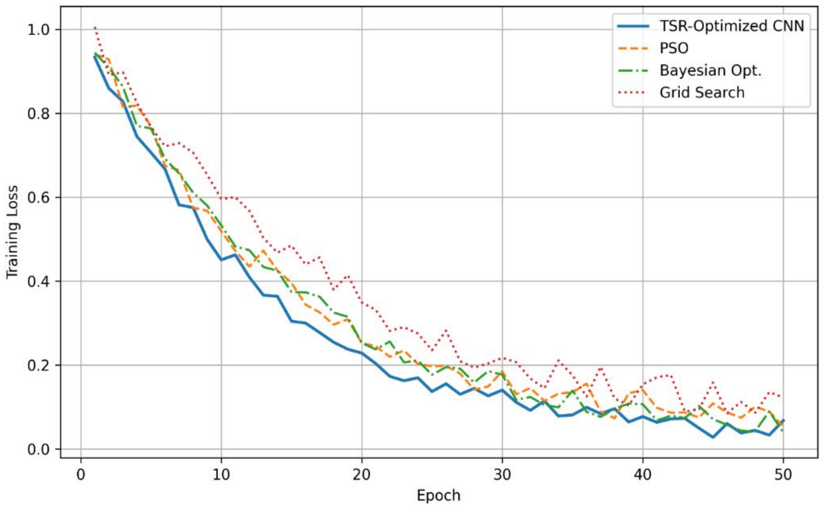

Fig. 3. Training loss curves for TSR-optimized CNN and baseline methods over 50 epochs.

To assess the practical computational efficiency of TSR, training time and model complexity were recorded. Despite achieving superior accuracy, the TSR-optimized CNN required 12–15% less training time than PSO and BO-optimized CNNs due to better initialization of learning rates, batch sizes, and layer configurations. This efficiency is particularly beneficial for real-time applications, such as EEG-based BCI systems, where rapid adaptation and low latency are critical.

Additionally, a detailed feature-level evaluation was conducted to investigate how TSR optimization influenced the CNN's internal representation. Activation maps from the first and second convolutional layers were compared for TSR-optimized and non-optimized networks. The TSR model demonstrated greater feature separation and more discriminative activation patterns, indicating that optimized architectures facilitate more effective hierarchical feature extraction. These findings are presented in *Fig. 5*, where representative activation maps highlight improved pattern recognition.

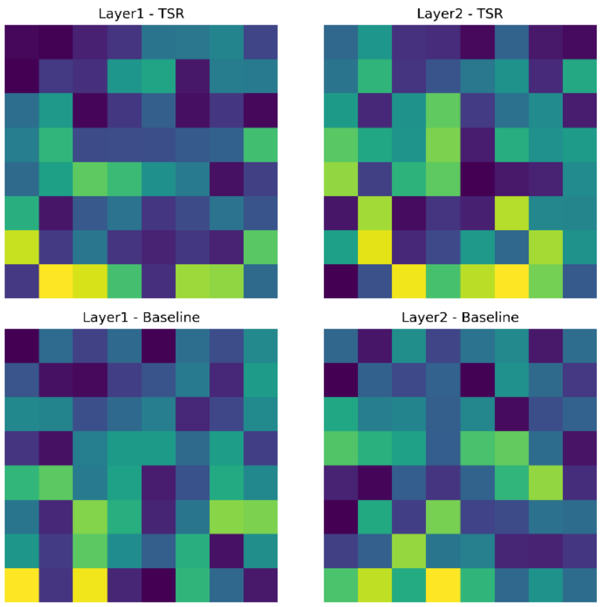

Fig. 4. Activation maps from the first and second convolutional layers of the TSR-optimized CNN compared to baseline CNN.

Overall, the experimental results clearly demonstrate that TSR not only improves classification performance but also enhances model stability, robustness to noise, convergence speed, and computational efficiency. The combination of competitive and cooperative search, adaptive environmental feedback, and regeneration ensures that CNN hyperparameters are fine-tuned effectively. Compared to traditional optimization techniques, TSR consistently identifies architectures that achieve a superior balance between accuracy, generalization, and resource efficiency.

In conclusion, TSR-optimized CNNs provide a practical and robust solution for complex classification tasks, especially in applications requiring real-time processing, robustness under variable conditions, and computationally efficient models. The results from *Tables 1* and *2*, along with *Figs. 2–4*, collectively demonstrate the algorithm's capability to automatically design CNN architectures that outperform conventional optimization methods while remaining interpretable, stable, and efficient.

5 | Conclusion

In this study, the TSR algorithm was introduced as a novel, ecology-inspired metaheuristic for optimizing CNNs for complex classification tasks. By modeling the dynamic interactions among trees in natural ecosystems —competition, cooperation, adaptive feedback, and regeneration —TSR provides a robust and interpretable framework for hyperparameter and architecture optimization. Unlike traditional optimization methods, TSR achieves a balanced exploration–exploitation trade-off, dynamically adapts to search-space conditions, and preserves population diversity, thereby preventing premature convergence and improving stability across multiple runs.

The integration of TSR into CNN optimization yielded substantial improvements in both performance and efficiency. TSR-optimized CNNs achieved higher classification accuracy than baseline optimization methods such as PSO, BO, and grid search. In addition, TSR produced more compact models with reduced parameter counts and faster convergence, demonstrating its suitability for resource-constrained or real-time applications, such as EEG-based brain–computer interfaces. The robustness of TSR-optimized architectures under noisy conditions further confirmed the algorithm's ability to generate models with strong generalization and reliable learning dynamics.

Simulation results, as shown in convergence curves, training loss trajectories, and feature activation maps, highlighted TSR's ability to guide the search toward high-quality solutions while maintaining interpretability. The algorithm's hierarchical, socially influenced optimization mechanism allowed CNN architectures to extract more discriminative features and exhibit smoother training behavior. Moreover, TSR's adaptive framework eliminates the need for extensive manual hyperparameter tuning, reducing computational overhead and facilitating practical deployment.

Overall, TSR demonstrates a powerful and flexible approach to neural network optimization. Its biologically grounded principles, combined with practical performance advantages, position it as a promising tool for a wide range of machine learning and signal-processing tasks. The results of this work suggest that TSR not only enhances model accuracy and stability but also provides a transparent, explainable, and efficient methodology for optimizing deep learning architectures. Future research will explore hybridization with reinforcement learning and multi-objective optimization, thereby extending TSR's applicability to more complex, real-world problems that require adaptive and interpretable solutions.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability

All data are included in the text.

Funding

This research received no specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

[1] Deng, L., Huang, J., Yuan, H., Liu, Q., Lou, W., Yu, P., … & Deng, L. (2025). Biological age prediction and NAFLD risk assessment: a machine learning model based on a multicenter population in Nanchang, Jiangxi, China. *BMC gastroenterology*, *25*(1), 1-12. https://doi.org/10.1186/s12876-025-03752-y

[2] Wang, Y., & Wang, P. (2025). Development and validation of a new diagnostic prediction model for NAFLD based on machine learning algorithms in NHANES 2017-2020.3. *Hormones*, *24*(2), 461–476. https://doi.org/10.1007/s42000-025-00634-6

[3] Tian, F., Zhao, Y., He, X., Zhang, Y., Hu, M., Liang, Y., … & Wang, Y. (2025). Interpretable Machine learning reveals the role of PANoptosis in the diagnosis and subtyping of NAFLD. *Immunobiology*, *230*(3), 1-18. https://doi.org/10.1016/j.imbio.2025.152909

[4] Charu, V., Liang, J. W., Mannalithara, A., Kwong, A., Tian, L., & Kim, W. R. (2024). Benchmarking clinical risk prediction algorithms with ensemble machine learning for the noninvasive diagnosis of liver fibrosis in NAFLD. *Hepatology*, *80*(5), 1184–1195. https://doi.org/10.1097/HEP.0000000000000908

[5] Thatha, V. N., Chalichalamala, S., Pamula, U., Krishna, D. P., Chinthakunta, M., Mantena, S. V., … & Vatambeti, R. (2025). Optimized machine learning mechanism for big data healthcare system to predict disease risk factor. *Scientific reports*, *15*(1), 14327. https://doi.org/10.1038/s41598-025-98721-6

[6] Aliyu, D. A., Akhir, E. A. P., Osman, N. A., Salisu, J. A., Saidu, Y., & Yalli, J. S. (2024). Optimization techniques in reinforcement learning for healthcare: a review. *2024 8th international conference on computing, communication, control and automation (ICCUBEA)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICCUBEA61740.2024.10774698

[7] Goswami, M., Mohanty, S., & Pattnaik, P. K. (2024). Optimization of machine learning models through quantization and data bit reduction in healthcare datasets. *Franklin open*, *8*, 1-11. https://doi.org/10.1016/j.fraope.2024.100136

[8] Ma, D., Zheng, W., & Lu, T. (2024). Machine learning-based predictive model for service quality assessment and policy optimization in adult day health care centers. *Spectrum of research*, *4*(1), 1–14. https://doi.org/10.1002/aic.16729

[9] Fatima, S. (2024). Healthcare cost optimization: leveraging machine learning to identify inefficiencies in healthcare systems. *International journal of advanced research in engineering technology & science*, *10*(3), 137–147. https://B2n.ir/ry4653

[10] Olalekan Kehinde, A. (2025). Leveraging machine learning for predictive models in healthcare to enhance patient outcome management. *International research journal of modernization in engineering technology and science*, *7*(1), 1465–1482. https://www.doi.org/10.56726/IRJMETS66198

[11] Khaleel, I., Marzoog, W. N., & Al-Kateb, G. (2025). ANILA: adaptive neuro-inspired learning algorithm for efficient machine learning, AI optimization, and healthcare enhancement. *Mesopotamian journal of computer science*, *2025*, 159–171. https://doi.org/10.58496/MJCSC/2025/009

[12] Krishnan, P. (2024). Ai-driven optimization in healthcare: machine learning models for predictive diagnostics and personalized treatment strategies. *Well testing journal*, *33*(S2), 10–33. https://welltestingjournal.com/index.php/WT/article/view/Ai_Driven_Optimization_In_Healthcar_Machine_Learning_Models_For_

[13] Nadeem, M. I., Ahmed, K., Li, D., Zheng, Z., Naheed, H., Muaad, A. Y., … & Abdel Hameed, H. (2022). SHO-CNN: a metaheuristic optimization of a convolutional neural network for multi-label news classification. *Electronics*, *12*(1), 1-24. https://doi.org/10.3390/electronics12010113

[14] Daliri, A., Branch, K., Sheikha, M., Roudposhti, K. K., Branch, L., Alimoradi, M., & Mohammadzadeh, J. (2023). Optimized categorical boosting for gastric cancer classification using heptagonal reinforcement learning and the water optimization algorithm. *7th international conference on pattern recognition and image analysis (IPRIA)* (pp. 1–6). IEEE. https://www.researchgate.net/profile/Arman-Daliri/publication/389465487_Optimized_Categorical_Boosting_for_Gastric_Cancer_Classification_using_Heptagonal_Reinforcement_learning_and_the_Water_Optimization_Algorithm/links/67c3173f645ef274a4988598/Optimized-

[15] Taji, K., Sohail, A., Shahzad, T., Khan, B. S., Khan, M. A., & Ouahada, K. (2024). An ensemble hybrid framework: a comparative analysis of metaheuristic algorithms for ensemble hybrid CNN features for plants disease classification. *IEEE access*, *12*, 61886–61906. https://doi.org/10.1109/ACCESS.2024.3389648

[16] Alimoradi, M. (2018). *Investigating the composition of apriori algorithm and metaheuristic algorithms (Genetic and PSO)*. https://doi.org/10.13140/RG.2.2.14058.77761

[17] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 770–778). Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr.2016.90

[18] Nagoor, S., & Jinny, S. V. (2023). A dual fuzzy with hybrid deep learning architecture based on CNN with hybrid metaheuristic algorithm for effective segmentation and classification. *International journal of information technology*, *15*(1), 531–543. https://doi.org/10.1007/s41870-022-01106-5

[19] Prakash, N. N., Rajesh, V., Namakhwa, D. L., Dwarkanath Pande, S., & Ahammad, S. H. (2023). A DenseNet CNN-based liver lesion prediction and classification for future medical diagnosis. *Scientific african*, *20*, 1-16. https://doi.org/10.1016/j.sciaf.2023.e01629

[20] Rikhtegar, A., Pooyan, M., & Manzuri, M. T. (2020). Comparing performance of metaheuristic algorithms for finding the optimum structure of CNN for face recognition. *International journal of nonlinear analysis and applications*, *11*(1), 301–319. https://doi.org/10.22075/ijnaa.2020.4296

[21] Naderi Yaghouti, A. R., & Shalbaf, A. (2024). Classification of nonalcoholic fatty liver grades using pre-trained convolutional neural networks and a random forest classifier on B-mode ultrasound images. *Journal of biomedical physics and engineering*, *15*(5), 437–450. https://doi.org/10.31661/jbpe.v0i0.2307-1646

[22] Alimoradi, M., Zabihimayvan, M., Daliri, A., Sledzik, R., & Sadeghi, R. (2022). Deep neural classification of darknet traffic. In *Artificial intelligence research and development* (pp. 105–114). IOS press. https://doi.org/10.3233/FAIA220323

[23] Daliri, A., Asghari, A., Azgomi, H., & Alimoradi, M. (2022). The water optimization algorithm: A novel metaheuristic for solving optimization problems. *Applied intelligence*, *52*(15), 17990–18029. https://doi.org/10.1007/s10489-022-03397-4

[24] Fakheri, S., Alimoradi, M., & Yamaghani, M. R. (2024). Colour image multilevel thresholding segmentation using trees social relationship algorithm. *Research square*, Preprint. https://doi.org/10.21203/rs.3.rs-4479475/v1

[25] Alimoradi, M., Sadeghi, R., Daliri, A., & Zabihimayvan, M. (2025). Statistic deviation mode balancer (SDMB): A novel sampling algorithm for imbalanced data. *Neurocomputing*, *624*, 129484. https://doi.org/10.1016/j.neucom.2025.129484

[26] Alimoradi, M., Azgomi, H., & Asghari, A. (2022). Trees social relations optimization algorithm: A new swarm-based metaheuristic technique to solve continuous and discrete optimization problems. *Mathematics and computers in simulation*, *194*, 629–664. https://doi.org/10.1016/j.matcom.2021.12.010