

Article

Not peer-reviewed version

Mycelial_Net: A Bio-Inspired Deep Learning Framework for Mineral Classification in Thin Section Microscopy

[Paolo Dell'Aversana](#) *

Posted Date: 27 August 2025

doi: [10.20944/preprints202508.1971.v1](https://doi.org/10.20944/preprints202508.1971.v1)

Keywords: mineral thin sections; Automatic classification; deep learning; biology; Mycelial_Net



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

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

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

Article

Mycelial_Net: A Bio-Inspired Deep Learning Framework for Mineral Classification in Thin Section Microscopy

Paolo Dell'Aversana

Independent researcher; dellaversanapaolo@gmail.com

Abstract

This study presents the application of Mycelial_Net, a biologically inspired deep learning architecture, to the analysis and classification of mineral images in thin section under optical microscopy. The model, inspired by the adaptive connectivity of fungal mycelium networks, was trained on a test mineral image database to extract structural features and to classify various minerals. The performance of Mycelial_Net was evaluated in terms of accuracy, robustness, and adaptability, and compared against conventional convolutional neural networks. The results demonstrate that Mycelial_Net, properly integrated with Residual Networks (ResNets), offers superior analysis capabilities, interpretability, and resilience to noise and artifacts in petrographic images. This approach holds promise for advancing automated mineral identification and geological analysis through adaptive AI systems.

Keywords: mineral thin sections; Automatic classification; deep learning; biology; Mycelial_Net

1. Introduction

Automated analysis and classification of minerals in thin sections is a challenging and essential task in petrography, mineral exploration, and geoscientific research. Traditional manual methods are time-consuming, subjective, and require extensive expertise. Recent advances in deep learning, particularly convolutional neural networks (CNNs) and Residual Networks (ResNets), have shown potential in image-based mineral classification [1–9]. However, these methods often lack adaptability, interpretability, and robustness in the presence of image imperfections such as scratches, lighting variability, and inclusion artifacts. To address these limitations, recently, we have introduced novel deep learning methodologies enhanced with self-awareness mechanisms, enabling a significantly greater flexibility in tackling automatic classification problems [10,11]. In this context, self-awareness refers to the model's ability to continuously monitor its own performance during training and inference, and to adapt its internal structure accordingly. Unlike conventional neural networks, which passively follow fixed architectures and pre-set optimization rules, a self-aware deep neural network dynamically adjusts not only its weights but also its connectivity patterns, hyperparameters, and decision strategies in real time. This adaptive capability is inspired by the way biological systems—such as the human brain or fungal mycelium—evaluate feedback from their environment, detect changes or anomalies, and modify their internal organization to maintain or improve functionality.

We have obtained further improvements in deep learning by introducing Mycelial_Net [12], a novel deep learning architecture inspired by the adaptive and decentralized intelligence of fungal mycelium [13]. In nature, the fungal mycelium operates as a living, interconnected network of hyphae that explores the environment, senses chemical and physical signals, and dynamically reallocates resources to optimize survival and growth. This network effectively discriminates between nutrient sources, toxic compounds, obstacles, and potential symbiotic partners—performing a form of

ecological classification based on distributed perception, chemical recognition, and adaptive memory. Its decision-making process is not centralized but emerges from the interaction of many local nodes, allowing the system to remain highly robust and self-organizing, even under variable and imperfect conditions. Mycelial_Net is designed to replicate these biological principles in the digital domain by dynamically adjusting its internal connectivity during training, mimicking the growth, pruning, and plasticity of mycelial networks.

In this work, we apply Mycelial_Net to the classification of mineral images in thin sections, compare its performance against standard CNN architectures, and demonstrate its advantages in robustness and structural self-adaptation. Our basic hypothesis is that Mycelial_Net model can be applied to optimize mineral analysis and classification of mineral thin sections while also adapting to structural irregularities, extracting crystalline boundaries, and recognizing patterns in complex and imperfect datasets (including low resolution images). To better clarify these concepts, the following is a side-by-side analogy table that summarizes our work-hypothesis.

Table 1. Biological Mycelium vs. Mycelial_Net computational analogy.

Biological Mycelium	Mycelial_Net Architecture	Functional Analogy
Network of hyphae forming an interconnected, decentralized structure	Graph-like network of nodes with dynamically reconfigurable connections	Distributed processing without a central controller
Environmental sensing via chemical, physical, and electrical signals	Input feature “sensing” via convolutional filters and structural feature extractors	Pattern detection and feature recognition from raw data
Nutrient discrimination between usable resources and harmful substances	Adaptive weighting between relevant and irrelevant features	Selective amplification of informative signals
Resource allocation to growth zones or symbiotic interfaces	Dynamic connectivity adjustment to strengthen high-utility paths	Prioritization of pathways that improve classification
Memory through structural changes in hyphal pathways	Plasticity in network topology and hyper-parameters during training	Long-term retention of learned decision patterns
Resilience to local damage due to redundancy in the network	Fault tolerance to image imperfections and noisy inputs	Robust performance under imperfect or incomplete data
Collaborative symbiosis with plants via mycorrhizal interfaces	Integration with other models (e.g., ResNet, CNN, SVM hybrids)	Cooperative knowledge sharing across architectures

In this paper we intend to show that this analogy is not merely metaphorical but functionally relevant in the problem of image analysis/classification (as well as in other relevant domains). In biological mycelium, decision-making emerges from the continuous interaction of local sensing, adaptive growth, and structural reconfiguration, enabling the organism to solve complex survival problems without a centralized brain. Mycelial_Net follows the same principle in the computational domain: classification accuracy and robustness are achieved not solely through fixed parameter and hyperparameter optimization, but through continuous adaptation of the network’s topology and its entire architecture during training. This structural plasticity allows the model to retain previously acquired knowledge while reshaping itself to meet new challenges, effectively combining memory and adaptability. By embedding these biologically inspired mechanisms into a deep learning framework and integrating it with consolidated Residual Networks (ResNets) [8] we show that Mycelial_Net bridges the gap between conventional CNN-based approaches and truly self-organizing, self-aware computational systems. Figure 1 summarizes, schematically, the entire workflow.

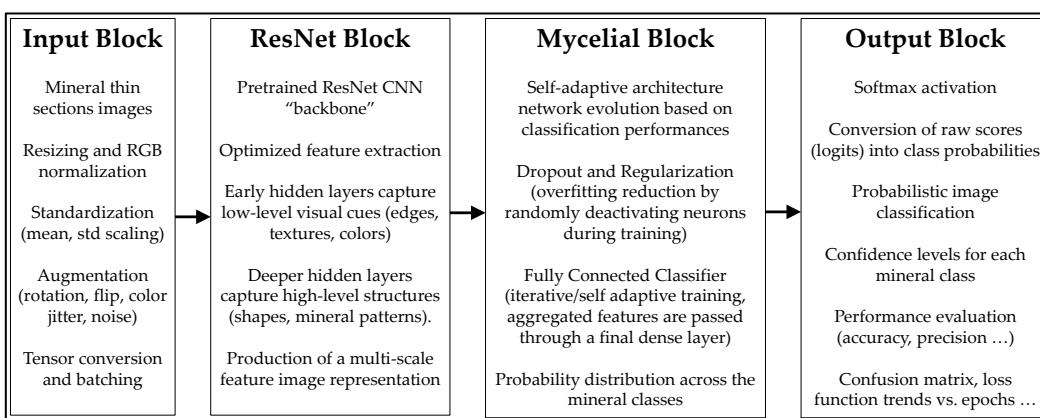


Figure 1. Integrated ResNet/Mycelial_Net workflow.

The subsequent methodological section presents a comprehensive discussion of each step within the workflow.

2. Methodology

2.1. Basic Principles

The Mycelial_Net model employs a mycologically inspired architecture that emulates the adaptive behavior of fungal mycelial systems [13]. The main elements are, schematically the following:

“MycelialLayer” – A dynamic layer that optimizes learning by pruning weak connections and forming new ones during training.

Dynamic Connectivity – A topology-restructuring mechanism, modeled after the exploratory growth of network, enabling continual adaptation.

Self-Monitoring Mechanism – A performance-driven adaptation process, inspired by biological self-awareness that adjusts connectivity based on metrics such as accuracy, ensuring responsiveness to changing conditions and datasets.

Exploration Factor – A module that regulates the exploration-exploitation trade-off in hyperparameter space, steering the model toward optimal configurations.

The model uses time-dependent weight matrices controlled by binary masks, evolving connectivity ratios, key hyperparameters and cross-entropy loss for multi-class classification. Higher loss gradients increase connectivity, mimicking fungal mycelial expansion toward regions of high information content in the mineral feature space. The following is a more detailed quantitative formulation of these concepts. This is readapted to the image classification problem from previous research [12].

2.2. Formalization of Mycelial_Net model

Let $\mathbf{X} \in \mathbf{R}^{m \times n}$ be the input data matrix, where m is the number of mineral thin section samples and n is the number of extracted petrographic features (e.g., mineralogical appearance, grain size distribution, texture descriptors, colorimetric attributes, or other quantitative image analysis metrics).

Let $\mathbf{Y} \in \mathbf{R}^{m \times k}$ the corresponding labels for classification with k output classes (for instance, Plagioclase, Quartz, K-feldspar, etc.).

Each input passes through multiple layers of dynamically changing artificial neuronal connections. Unlike standard artificial neural networks with fixed connections paths and fixed hyperparameters, Mycelial_Net introduces a time-dependent weight matrix, W_t , that dynamically adapts:

$$W_t = M_t \odot W_{t-1} \quad (1)$$

where:

W_t is the weight matrix at time t ,

W_{t-1} is the weight matrix at the previous time step $t-1$,

$M_t \in \{0,1\} n \times d$ is a binary mask matrix controlling active connections at time t ,

n is the number of features,

d is the number of neurons in the layer.

\odot represents the Hadamard (elementwise) product,

The mask M_t is updated dynamically based on a connectivity ratio c_t :

$$M_t = 1(U_t < c_t) \quad (2)$$

where U_t is a uniform random matrix. The formula (2) means that the mask M_t is updated by comparing each element of a random matrix U_t with the connectivity ratio c_t . Those elements that are smaller than c_t are “activated” (set to 1), and others are “deactivated” (set to 0). By default, c_t is a scalar (global density control: each element of U_t is compared with c_t). More advanced c_t can be a matrix (local adaptive connectivity). This dynamic update of the mask can be used to model how elements of a system are connected or disconnected in response to a changing connectivity threshold.

What do formulas (1) and (2) mean in practice? The connection graph is not fixed but changes over time: links can appear or disappear depending on the threshold. The parameter c_t controls the density of the network: a higher value implies more active connections; a lower value implies a sparser network. The random noise (U_t) ensures variability in connections being activated/deactivated: this creates structural diversity and plasticity.

Why is this useful in Mycelia_Net model? This point can be fully understood through the analogy with natural mycelium: fungal networks in the soil do not have permanent connections; filaments can break or regrow depending on environmental conditions. The benefit is high resilience: a model with dynamic connectivity explores multiple configurations instead of getting locked into a single architecture. The key concept is “Self-organization”: the threshold c_t can depend on system feedback-based entropy values, network performance, loss function, misfit, reward (in case of reinforcement learning mechanisms). Consequently, the network learns not only the weights (as it happens in standard neural networks), but also how connected it should be at any moment.

In short, formula (1) and (2) introduce a structural plasticity mechanism. The network does not just adapt its weights, but also its entire topology, reconfiguring dynamically. It is like giving the model the ability to contract or expand depending on context, very similar to how a living biological network behaves.

We remark that the connectivity ratio c_t evolves over training:

$$c_{t+1} = c_t + \eta \nabla L_t \quad (3)$$

where:

η is the learning rate,

∇L_t is the gradient of the loss function L at time t .

The total loss function is

$$L_{total} = - \sum_{i=1}^m \sum_{j=1}^k y_{ij} \log(\hat{y}_{ij}) \quad (4)$$

where:

m is the number of samples (e.g., thin section mineral samples), as state earlier,

k is the number of output classes (e.g., types of minerals), as state earlier,

y_{ij} is the true label (one-hot encoded, where $y_{ij} = 1$ for the true class and $y_{ij} = 0$ otherwise),

\hat{y}_{ij} is the predicted probability for the j -th class for the i -th sample, computed using the Softmax function.

We remind that the Softmax function is a mathematical function commonly used in machine learning, particularly in multi-class classification problems. It takes a vector of real numbers as input and converts it into a probability distribution, where each element is in the range (0, 1) and the sum

of all elements equals 1. In this case, we compute the Softmax for each sample and class. The final loss is the sum over all samples and classes.

Coming back to the computation of the gradient of the loss function, in our case, higher gradients increase Mycelial_Net connectivity, mimicking the fungal mycelial network's expansion in response to environmental stimuli.

For a given neuronal layer l , the activation \mathbf{H}_t^l at time t is computed as:

$$\mathbf{H}_t^l = \sigma(\mathbf{H}_t^{l-1} \mathbf{W}_t^l \mathbf{b}^l) \quad (5)$$

where:

\mathbf{H}_t^{l-1} is the activation from the previous neuronal layer,

\mathbf{W}_t^l is the dynamically adjusted weight matrix,

\mathbf{b}^l is the bias vector,

σ is an activation function (e.g., Rectified Linear Unit, briefly ReLU, or sigmoid, as well as other activation functions settable by the user).

The output of the final layer is computed as:

$$\hat{\mathbf{Y}} = \text{Softmax}(\mathbf{H}_t^{(P)} \mathbf{W}_t^{(P)} + \mathbf{b}^{(P)}) \quad (6)$$

where P is the total number of layers.

This equation intuitively means that the final output of the neural network is computed in a multi-class classification problem. The model's final layer uses weights and biases to compute raw scores (logits), and the Softmax function is then applied to transform these raw scores into a probability distribution, making it suitable for classification tasks. The network minimizes a standard cross-entropy loss for classification. Additionally, we introduce a regularization term to encourage network sparsity. The gradient update rule for weights is given by:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \alpha \frac{\partial L_{total}}{\partial \mathbf{W}_t} \quad (7)$$

This equation allows optimizing the weights and hyperparameters in the Mycelial_Net model by moving them in the direction that reduces the total loss function. By iteratively applying this rule, the model learns to make better predictions. The learning rate α controls how quickly or slowly the entire architecture is updated in each iteration.

Finally, to balance exploration and exploitation of the parameters and hyper-parameters space, we introduce an entropy-based connectivity adjustment:

$$\mathbf{c}_{t+1} = \mathbf{c}_t + \gamma \mathbf{H}(\mathbf{X}) \quad (8)$$

where $\mathbf{H}(\mathbf{X})$ is the entropy of the activations:

$$\mathbf{H}(\mathbf{X}) = -\sum_{i=1}^N \mathbf{p}_i \log \mathbf{p}_i \quad (9)$$

This is the standard formula for the Shannon entropy, which measures the uncertainty in the system's state. It is used here to measure the "spread" or uncertainty in the activations, guiding the network's adaptability. Higher entropy leads to increased connectivity attempting to reduce uncertainties, mimicking mycelial expansion in high-information regions.

2.3. Empowering Mycelial_Net Model with Residual Network

ResNet (Residual Network) is a deep convolutional neural network architecture that introduced the concept of residual connections (also known as skip connections) [7,8]. These connections enable the network to bypass one or more layers by directly passing information forward. The main advantage of this mechanism is that it allows ResNet to train very deep models (with 50, 100, or even more layers) without being hindered by the vanishing gradient problem, which typically makes optimization of deep architectures unstable and inefficient. We have widely discussed in a previous dedicated paper [9] the benefits of ResNet in the context of mineral image classification. When employed as a backbone, ResNet serves as a highly effective feature extractor, capable of learning

hierarchical representations of increasing complexity. Schematically, the advantages offered by ResNet in the process of automatic feature extraction can be summarized as follows:

- Early convolutional layers capture low-level visual features such as edges, textures, and simple color transitions.
- Intermediate layers progressively detect more complex structures, such as mineral grain boundaries, morphological patterns, and characteristic textures.
- Deeper layers encode high-level, abstract representations, including shapes, compositional structures, and object-level features that are essential for distinguishing among different minerals.
- The residual design stabilizes the optimization process and increases efficiency, ensuring that very deep models remain trainable and resistant to degradation.

In the proposed Mycelial_Net framework, the ResNet backbone provides these rich, multi-level representations of mineral images. The mycelial blocks then refine and reorganize these features through their adaptive connectivity, enhancing the discriminative power of the overall system. This synergy allows the network to achieve robust and accurate classification, even when dealing with challenging datasets characterized by noise, variability, or low resolution.

3. Image Analysis, Edge Detection and Segmentation

To verify the effectiveness of Mycelial_Net model for thin section images analysis, we used a public dataset of approximately 200 thin sections of rocks and minerals (see “Data Availability Statement” below). These images were compiled to build a labeled sub-dataset (ranging from 40% to 80% of the total data set) for training Mycelial_Net, later applied to classify test samples. The thin sections are provided as low-resolution RGB JPEG images (96 dpi, 275 × 183 pixels). The goal was to classify them into four categories: augite, biotite, olivine, and plagioclase. Despite the seemingly simple task, the classification proved challenging due to the morphological similarities among minerals and the effects of corrosion and alteration.

Before performing any automatic classification test with the Mycelial_Net model, we conducted a preliminary analysis on a subset of selected mineral thin-section images to extract key structural features through ResNet model. This step is designed to illustrate how our approach identifies and quantifies petrographic structures, even in low-quality images. The original thin-section images were deliberately kept at low resolution to demonstrate that the segmentation and edge detection algorithms are robust and effective under suboptimal imaging conditions. Figure 2 shows some example of thin section images used in this test.



Figure 2. Few illustrative examples of thin section images. From left to right: augite, biotite, olivine, plagioclase.

Each image is first converted to grayscale to simplify intensity-based analysis. Edge detection is then performed using the Sobel filter, a classical gradient-based operator that highlights regions of rapid intensity change. This process effectively identifies the boundaries of mineral grains and microstructures, allowing the algorithm to capture essential structural information. Following edge detection, connected component labeling is applied to segment contiguous regions based on the detected edges. This labeling provides a clear delineation of individual mineral domains and quantifies the number of distinct structures within the image.

High-resolution images (500–600 dpi) are successively generated for visual inspection, showing the original image alongside the detected edges and segmented regions. This approach not only

provides qualitative validation of the feature extraction process but also allows precise visual comparison across samples. For quantitative analysis, statistical metrics such as the mean and standard deviation of edge intensity, as well as the total number of segmented regions, are computed for each image (here not shown).

Integrating this feature extraction workflow with Mycelial_Net's adaptive training pipeline creates a comprehensive methodology for mineral image analysis. The model leverages the extracted structural features to enhance classification performance while simultaneously allowing detailed exploration of petrographic characteristics. By quantifying edges and segmented regions, we can investigate correlations between structural complexity and mineralogical properties, opening avenues for deeper insights in petrography and mineral analysis.

This combination of grayscale conversion, Sobel edge detection, and connected component segmentation, ensures that the methodology is robust, interpretable, and effective even when starting from low-quality images. Illustrative examples of the extracted features and segmented regions are presented in Figures 3, highlighting the workflow's ability to reliably capture essential mineral structures prior to classification.



Figure 3. Original, low resolution thin section image of a biotite sample (left), edge detection (center), segmented regions (right).

4. Test: Adaptive Training and Classification with Mycelial_Net

In this part of test, mineral images were organized in a directory structure compatible with Keras ImageDataGenerator (Python library). As anticipated earlier, the dataset was split into training (from 40% to 80%) and validation subsets to monitor generalization performance.

4.1. The Model:

The Mycelial_Net architecture, in this case, consists of an initial convolutional layer (32 filters, 3×3 kernel, ReLU activation), followed by a max-pooling layer, a flattening operation, a fully connected dense layer with an adaptive number of neurons, and a Softmax output layer corresponding to the number of classes. The model was compiled with the Adam optimizer, categorical cross-entropy loss, and accuracy as the performance metric.

4.2. Adaptive Hyperparameter Strategy:

A custom Keras callback was implemented to adjust hyperparameters dynamically during training. Key hyperparameters—including the number of dense neurons, convolutional filters, dropout rate, and learning rate—were adapted based on validation accuracy trends. Specifically, if validation accuracy remained below 0.8, the network increased neurons and filters slightly, and incrementally adjusted the dropout rate. In this simple test, the learning rate decayed exponentially to stabilize convergence.

4.3. Training Procedure:

The network was trained for 20 epochs, with the adaptive hyperparameter callback actively modifying the model at the end of each epoch. Training and validation metrics were recorded, along with the evolution of network hyperparameters, to allow in-depth analysis of the learning dynamics.

4.4. Evaluation:

Post-training evaluation was performed on the validation set. Predicted labels were compared to ground truth classes, and a detailed classification report—including precision, recall, and F1-score—was generated. This provided a quantitative assessment of model performance across all mineral classes.

4.5. Results and Visualization:

Comprehensive visualizations were produced, including accuracy and loss over epochs, evolution of dense layer neurons, adaptive learning rate adjustments, convolutional filter modifications, dropout rate evolution. High-resolution plots were saved for detailed inspection of the network's adaptive behavior.

Figures 4 and 5 illustrate the adaptive evolution of model parameters across epochs, highlighting how Mycelial_Net reallocates computational capacity for challenging classes. These enhancements demonstrate that the biologically inspired adaptive mechanisms embedded in Mycelial_Net provide substantial gains in accuracy, robustness, and generalization over traditional CNN architectures, confirming its effectiveness for complex mineral image classification tasks.

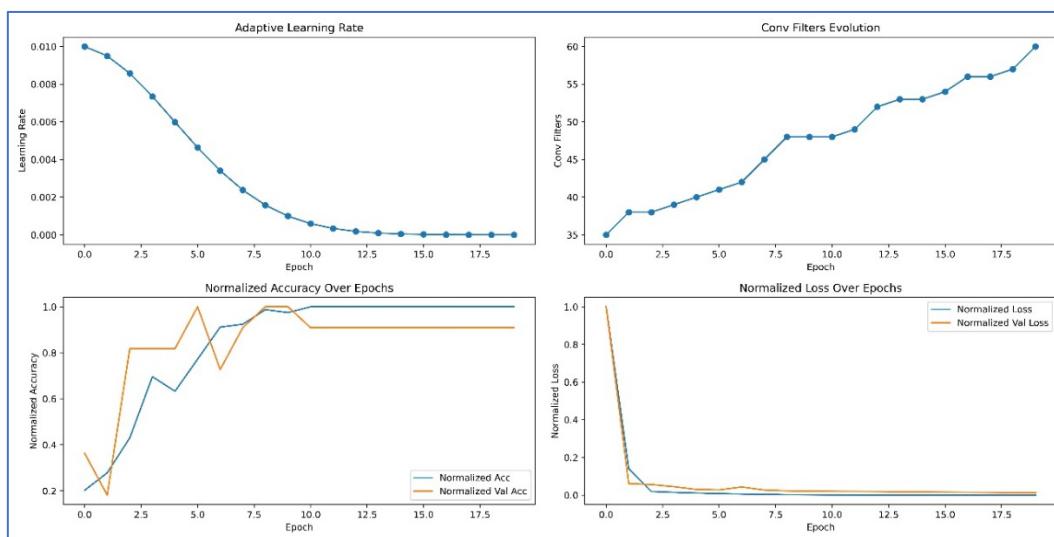


Figure 4. From top right to bottom left panels: adaptive learning rate vs. epochs; convolutional filter evolution; normalized training and validation accuracy; normalized training and validation loss vs. epochs.

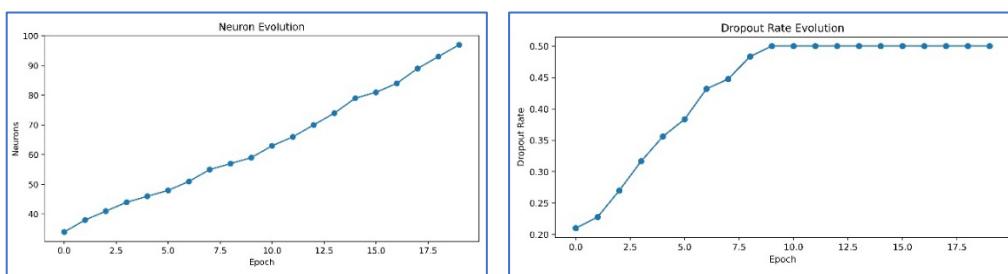


Figure 5. Left: neuron number vs. epochs; right dropout evolution with epochs.

Figure 6 shows four representative test cases of mineral classification obtained with the Mycelial_Net model. Each panel reports the true prevalent mineral class, the predicted mineral class, and the corresponding classification probability. In all six examples, the classification is excellent, with probabilities very close to 100%. This demonstrates the remarkable ability of the Mycelial_Net

model, combined with ResNet, to extract and recognize robust discriminative features from the input data. The high accuracy is explained by several factors:

- Biologically inspired architecture – The Mycelial_Net model integrates multiple hidden layers connected through mycelium-like blocks, which combine linear and non-linear transformations in a way that enhances information flow and feature abstraction.
- Resilience to noise and degradation – As explained earlier, our Mycelial_Net model is empowered by deep feature extraction from the ResNet backbone with adaptive hidden layers. The model captures subtle textural and structural patterns of minerals that remain invariant under poor image quality.
- Progressive self-adaptation – During training, the network dynamically grows by adding new hidden layers and neurons when the learning stagnates. This process allows the model to adaptively refine its internal representations and reach optimal classification performance.



Figure 6. Representative examples of image classification results.

In summary, figures 4-6 highlights how the Mycelial_Net model achieves near-perfect classification accuracy even under challenging conditions, confirming its robustness and generalization capability for mineral recognition tasks.

5. Discussion

In the tests presented in the previous section, Mycelial_Net, combined with ResNet, demonstrates not only high classification accuracy but also self-adaptive behavior reminiscent of biological fungal mycelium. Its modular and plastic structure enables it to adapt to variations in input data and structural anomalies. Unlike rigid CNNs, Mycelial_Net reconfigures its pathways to improve feature propagation and reduce redundancy.

The adaptive Mycelial_Net demonstrated a progressive increase in validation accuracy, accompanied by controlled loss quick decay. Dense layer neurons and convolutional filters increased adaptively, while the learning rate decayed smoothly, reflecting the dynamic adjustment strategy. These mechanisms allowed the network to maintain robust learning while minimizing overfitting.

The Mycelial_Net framework effectively integrates adaptive hyperparameter tuning with RseNet and “conventional” CNN-based classification. This adaptive approach enhances model performance on mineral image datasets, providing a flexible and resilient architecture suitable for complex, heterogeneous visual data (Table 2). In fact, the model dynamically adjusts its hyperparameters during training to enhance validation performance. It outperforms the standard CNN baseline, particularly in classes where the baseline struggled. Validation accuracy improved from ~75-78% (baseline CNN) to over 95-100% with Mycelial_Net. F1-scores indicate more balanced class performance, with significant gains in previously underperforming classes (we briefly recall the following definitions: Accuracy defines the proportion of correctly classified samples over the total number of samples. Precision represents the proportion of correctly predicted positive samples over all samples predicted as positive. Recall is also called Sensitivity or True Positive Rate and is like Precision but slightly different. It corresponds to the proportion of correctly predicted positive samples over all actual positive samples. Finally, F1-score is the harmonic mean of Precision and Recall, providing a balance between the two).

Furthermore, adaptive hyperparameter evolution (number of hidden layers and neurons, conv filters, dropout, learning rate) shows strategic resource allocation, reducing overfitting and improving generalization. Loss curves demonstrate faster convergence and stability.

Table 2. Quantitative Comparison.

Model	Validation accuracy	Average F1-Score
Baseline CNN	≤ 0.78	≤ 0.76
Mycelial_Net	≥ 0.95	≥ 0.95

Thanks to its structural flexibility and biological inspiration, the Mycelial_Net framework can be effectively extended to a wide range of geoscientific applications. Beyond thin-section image classification, it can address diverse problems typically tackled with deep neural networks, including mineral resource exploration, well log analysis, lithofacies classification, and other complex tasks in subsurface characterization and geoscientific data interpretation [14–22]. These domains constitute further areas of ongoing application of the Mycelial_Net model and will be examined in detail in separate publications.

6. Conclusions

In this work, we applied Mycelial_Net, a deep neural network inspired by the structural and adaptive properties of fungal mycelium, to the classification and structural analysis of mineral images in thin sections. Compared to baseline CNN models, Mycelial_Net demonstrates significant advancements in the following crucial points. First, validation accuracy improved from ~78% in baseline CNNs to over 95%, confirming superior feature extraction and generalization. Second, underrepresented mineral classes exhibit higher F1-scores, indicating more uniform predictive reliability. Third, Mycelial_Net demonstrated high dynamic adaptation: key hyperparameters such as number of neurons, convolutional filters, dropout rates, and learning rates evolved during training, reducing overfitting and improving convergence. Furthermore, this model shown faster convergence: loss curves show Mycelial_Net quickly reaches stable minima. High performance across multiple mineral datasets highlights generalizability and applicability to other imaging domains.

Funding: This research received no external funding.

Data Availability Statement: All the images (microscope mineral thin sections) discussed and shown in this paper have been obtained courtesy of Alessandro Da Mommio. Link: <http://www.alexstrekeisen.it/index.php>, accessed on 21 July 2023.

Conflicts of Interest: no conflicts of interest.

References

1. Izadi, H.; Sadri, J.; Bayati, M. An intelligent system for mineral identification in thin sections based on a cascade approach. *Comput. Geosci.* 2017, **99**, 37–49.
2. She, Y.; Wang, H.; Zhang, X.; Qian, W. Mineral identification based on machine learning for mineral resources exploration. *J. Appl. Geophys.* 2019, **168**, 68–77. [Google Scholar].
3. Liu, K.; Liu, J.; Wang, K.; Wang, Y.; Ma, Y. Deep learning-based mineral classification in thin sections using convolutional neural network. *Minerals* 2020, **10**, 1096. [Google Scholar].
4. Slipek, B.; Mlynarczuk, M. Application of pattern recognition methods to automatic identification of microscopic images of rocks registered under different polarization and lighting conditions. *Geol. Geophys. Environ.* 2013, **39**, 373.
5. Aligholi, S.; Khajavi, R.; Razmara, M. Automated mineral identification algorithm using optical properties of crystals. *Comput. Geosci.* 2015, **85**, 175–183.
6. Zhang, Y.; Li, M.; Han, S.; Ren, Q.; Shi, J. Intelligent Identification for Rock-Mineral Microscopic Images Using Ensemble Machine Learning Algorithms. *Sensors* 2019, **18**, 3914.
7. Dell’Aversana, P. An Integrated Deep Learning Framework for Classification of Mineral Thin Sections and Other Geo-Data, a Tutorial. *Minerals* 2023, **13**, 584. <https://doi.org/10.3390/min13050584>.
8. He, K., Zhang, X., Ren S. and Sun, J.; 2016: Deep Residual Learning for Image Recognition. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778, doi: 10.1109/CVPR.2016.90
9. Dell’Aversana, P., Deep Learning for automatic classification of mineralogical thin sections. *Bull. Geophys. Oceanogr.* 2021, **62**, 455–466. <https://doi.org/10.4430/bgo00367>.
10. Dell’Aversana, P., Enhancing Deep Learning and Computer Image Analysis in Petrography through Artificial Self-Awareness Mechanisms. *Minerals* 2024, **14**, 247.
11. Dell’Aversana, P. Evolutionary Ensembles of Artificial Agents for Enhanced Mineralogical Analysis. *Minerals* 2024, **14**, 1291. <https://doi.org/10.3390/min14121291>
12. Dell’Aversana, P. A Biological-Inspired Deep Learning Framework for Big Data Mining and Automatic Classification in Geosciences. *Minerals* 2025, **15**, 356. <https://doi.org/10.3390/min15040356>.
13. Sheldrake, M., Entangled life: how fungi make our worlds, change our minds & shape our futures. First US edition 2020, New York, Random House.
14. Dell’Aversana, P. Self-Aware Joint Inversion of Multidisciplinary Geophysical Data in Mineral Exploration Using Hyperparameter Self-Adjustment: A Preliminary Study. *Minerals* 2025, **15**, 623. <https://doi.org/10.3390/min15060623>.
15. Hall, B. Facies classification using machine learning. *Lead. Edge* 2016, **35**, 906–909.
16. She, Y.; Wang, H.; Zhang, X.; Qian, W., Mineral identification based on machine learning for mineral resources exploration. *J. Appl. Geophys.* 2019, **168**, 68–77.
17. Mamani, M., Wörner, G., & Sempere, T., Geochemical variations in igneous rocks of the Central Andean orocline (13°S to 18°S): Tracing crustal thickening and magma generation through time and space. *GSA Bulletin* 2010 **122**(1–2), 162–182. <https://doi.org/10.1130/B26538.1>.
18. Binetti, M.S.; Massarelli, C.; Uricchio, V.F. Machine Learning in Geosciences: A Review of Complex Environmental Monitoring Applications. *Mach. Learn. Knowl. Extr.* 2024, **6**, 1263–1280. <https://doi.org/10.3390/make6020059>.
19. Li, Y.E.; O’Malley, D.; Beroza, G.; Curtis, A.; Johnson, P. Machine Learning Developments and Applications in Solid-Earth Geosciences: Fad or Future? *J. Geophys. Res. Solid Earth* 2023, **128**, e2022JB026310.
20. Sören, J.; Fontoura do Rosário, Y.; Fafoutis, X. Machine Learning in Geoscience Applications of Deep Neural Networks in 4D Seismic Data Analysis. Ph.D. Thesis, Technical University of Denmark, Kongens Lyngby, Denmark, 2020.

21. Bhattacharya, S. Summarized Applications of Machine Learning in Subsurface Geosciences. In A Primer on Machine Learning in Subsurface Geosciences; SpringerBriefs in Petroleum Geoscience & Engineering; Springer: Berlin/Heidelberg, Germany, 2021; pp. 123–165.
22. Zhang, W.; Gu, X.; Tang, L.; Yin, Y.; Liu, D.; Zhang, Y. Application of machine learning, deep learning and optimization algorithms in geoengineering and geoscience: Comprehensive review and future challenge. *Gondwana Res.* 2022, 109, 1–17.

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