Glial Gating: Biologically Inspired Channel Modulation and Pruning for Convolutional Networks

Jakub Nowak[1 [0000−0002−1572−3426]], Marcin Korytkowski[1,2 [0000−0002−6002−2733]], Rafał Scherer[1,3 [0000−0001−9592−262X]], Tomasz Nowak[1 [0000−0002−7028−2352]], Aneta Maćkiewicz[1 [0009−0002−9259−9956]], and Slava Voloshynovskiy[4 [0000−0003−0416−9674]]

[1] Department of Artificial Intelligence, Faculty of Computer Science and Artificial Intelligence, Czestochowa University of Technology
al. Armii Krajowej 36, 42-202 Częstochowa, Poland {jakub.nowak, marcin.korytkowski, rafal.scherer, tomasz.nowak2, aneta.mackiewicz}@pcz.pl
[2] University of Warmia and Mazury in Olsztyn,
ul. Michała Oczapowskiego 2, 10-719 Olsztyn, Poland
marcin.korytkowski@uwm.edu.pl
[3] Faculty of Computer Science
and Center of Excellence in Artificial Intelligence, AGH University of Krakow
al. Mickiewicza 30, 30-059 Kraków, Poland
[4] Department of Computer Science, University of Geneva, Geneva, Switzerland
svolos@unige.ch

Abstract. In recent years, convolutional neural networks have become the dominant tool in image processing and data classification tasks. Despite their high effectiveness, classical CNN architectures are characterized by a large number of parameters and limited capacity for adaptive control of information flow between layers. In response to these limitations, this paper proposes a biologically inspired control and pruning mechanism based on glial cell modeling. The presented architecture extends a classical convolutional network with a glial controller consisting of modules functionally corresponding to astrocytes, oligodendrocytes, and microglia. This controller generates dynamic channel gates that modulate the activations of convolutional feature maps in a manner dependent on the current state of the network. The learning process is carried out alternately for the glial and convolutional parts, enabling stable adaptation of both model components. The effectiveness of the proposed approach was experimentally verified on the publicly available CIFAR-10 datasets and on an original, deterministic sequential dataset built from network traffic logs containing user URLs. In the case of image data, classical classification scenarios were applied, while for sequential data, training sequences representing computer network user behavior were used. The obtained results indicate an improvement in classification quality compared to the base model lacking glial mechanisms, while maintaining the interpretability of the controller's operation. Additionally, the analysis of masks generated by the microglia module allowed for estimating the

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

potential reduction in the number of convolutional parameters by identifying redundant channels. The results confirm the validity of using glial cell mechanisms as an adaptive control layer in deep neural networks, both in image processing and network data analysis tasks.

Keywords: convolutional neural networks, glial cells, biologically inspired neural networks, adaptive modulation, image classification, cybernetic, network security, URL log analysis, user behavior classification

1 Introduction

Convolutional Neural Networks (CNNs) remain a central architecture in deep learning for processing grid-structured data, and are widely used in image classification, signal analysis, and related tasks. Despite their high performance on standard benchmarks, conventional CNN architectures are characterized by excessive parameterization, significant computational overhead, and a static approach to information processing. Upon completion of the training phase, feature channel activity is determined solely by fixed learned weights, lacking adaptation to the current input context. Such limitations constrain the computational efficiency of these models, as well as their interpretability and potential for further structural optimization.

Concurrently, modern neuroscience provides compelling evidence that information processing in the brain is not exclusively mediated by neurons. Glial cells, particularly astrocytes, oligodendrocytes, and microglia, play an active role in regulating signal flow, stabilizing neuronal activity, and remodeling synaptic connections [8, 11]. Astrocytes participate in long-term modulation of synaptic transmission; oligodendrocytes influence the efficiency and synchronization of signal propagation; while microglia facilitate the selection and elimination of redundant connections via synaptic pruning. These mechanisms suggest the existence of an extrinsic, slower-adapting control system overseeing the neuronal network and optimizing its structure over time.

Inspired by these observations, this paper proposes the *Glia-Controlled Convolutional Network* architecture, wherein a classical convolutional network is augmented with a glial controller serving as an adaptive control layer. This controller generates per-channel gates modulating the activations of convolutional layers and comprises three specialized modules functionally emulating the roles of astrocytes (invariant components and stabilization), oligodendrocytes (signal reinforcement), and microglia (soft channel selection and regularization signal). The implementation of an alternating learning strategy for the glial and convolutional components stabilizes the training process and enables the generation of interpretable channel importance masks.

The effectiveness of the proposed approach is verified across two complementary application domains. In the first stage, image classification performance is analyzed on the CIFAR-10 datasets [12], comparing baseline models with glia-controlled variants and versions subjected to structural reduction based on microglia-derived masks. This allows for an assessment of glial control's impact

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

on both classification quality and the potential for parameter reduction without significant performance degradation.

The second stage of the study focuses on a practical cybersecurity task: user identification within computer networks based on URL activity sessions. The selection of this problem is motivated by the fact that visited URL sequences constitute a stable and difficult-to-forge behavioral fingerprint, enabling passive and continuous authentication. Prior research has demonstrated that URL log analysis effectively profiles users and facilitates behavior-based continuous authentication using both machine learning methods and sequential models [18, 19]. This article demonstrates that glial control mechanisms can be successfully adapted to one-dimensional convolutional networks processing textual data, facilitating structural reduction and enhancing model interpretability.

The paper is structured as follows. First, we review the existing literature in Section 2. Next, we introduce our approach in Section 3 and datasets in Section 4. We present and discuss the results of the experiment in Sections 5 and 6.

2 Related Work

The global literature extensively documents the application of convolutional neural networks in image processing and data classification tasks. Classical CNN architectures, such as VGG or ResNet, achieve superior performance across numerous pattern recognition benchmarks; however, their operation relies upon a static set of convolutional filters whose activity remains unmodulated during the inference process [1, 2].

In the conventional approach, the neural network structure remains static upon completion of the training process, with learned convolutional filter parameters applied identically across all input samples. Consequently, the network lacks mechanisms for adaptive information flow regulation between layers based on the current data context. A notable departure from this paradigm is found in transfer learning methods, where specific network segments, typically convolutional layers, serve as general-purpose feature extractors, while modifications are restricted to the classification component [3]. However, such solutions do not introduce dynamic channel activity control, focusing instead on limiting the scope of parameter retraining or stabilizing the learning process.

To enhance neural network efficiency, numerous regularization and information flow control mechanisms have been proposed, including dropout, batch normalization, and channel attention mechanisms. Dropout involves the stochastic deactivation of neurons during training to mitigate model overfitting [4]. Batch normalization stabilizes the distribution of activations between layers, thereby accelerating training convergence [17]. Conversely, channel attention mechanisms, such as Squeeze-and-Excitation blocks, enable adaptive weighting of feature channels based on global activation statistics [5]. Despite their high effectiveness, these methods remain purely mathematical constructs lacking direct biological inspiration.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

Concurrently, biologically inspired neural network models have emerged, focusing primarily on neuronal behavior and plasticity mechanisms. In recent years, increasing attention has been directed toward neuron glia interactions, specifically the role of astrocytes in modulating synaptic activity. However, most extant research in this domain is confined to simulation models or small scale architectures, lacking direct integration with deep convolutional networks. Recent studies have attempted to incorporate mechanisms inspired by neuron glia communication into deep learning frameworks. A prominent example is the GliaNet model, proposing adaptive network structure learning through signals generated by modules emulating glial cells [6]. Within this framework, glial cells participate in synaptic connection modulation, facilitating dynamic architectural restructuring during training. Similarly, the MA-Net model proposes a multicell astrocyte neuron framework enabling bidirectional, temporal, and collective modulation of inter-neuronal connections [7]. The authors demonstrated that employing multiple astrocytes facilitates significant parameter reduction while maintaining or improving classification accuracy in vision tasks.

In contrast to these approaches, the architecture proposed herein focuses on the channel-wise modulation of convolutional filter activity, executed by specialized glial cell modules functioning as an external network controller. This approach enables simultaneous information flow control, identification of redundant parameters, and subsequent structural model reduction without necessitating costly architecture search procedures. In the field of network security, machine learning methods are widely deployed for traffic analysis, user classification, and anomaly detection based on URL logs. Current solutions primarily emphasize feature engineering or architectural selection, seldom utilizing mechanisms for adaptive model structure control.

3 Glia-Controlled Convolutional Network

The proposed architecture is inspired by contemporary neuroscience, which emphasizes the pivotal role of glial cells in information processing within the human brain. For a significant period, glia were regarded merely as supportive elements for neurons; however, neurobiological research has conclusively demonstrated their active participation in the regulation of synaptic activity, neuronal plasticity, and long-term organization of neural connections [8]. As Fields underscores, glia can be viewed as a "second brain" operating in parallel with the neuronal network, modulating signal flow and optimizing the performance of brain regions according to functional context [8]. These mechanisms encompass both short-term synaptic regulation and long-term neural network restructuring. Consequently, this study proposes a convolutional neural network architecture augmented with an auxiliary controller inspired by glial cell functionality. This approach introduces adaptive control over convolutional filter activity during both training and inference while maintaining compatibility with classical CNN frameworks.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

Astrocytes play a critical role in regulating synaptic activity and maintaining the homeostasis of the neuronal environment. Research indicates that astrocytes participate in the modulation of synaptic transmission by regulating neurotransmitter and ion concentrations in the extracellular space, thereby influencing neuronal excitability [9]. In the proposed model, the astrocyte function is emulated through a slow-varying modulation mechanism of feature channel activity, providing stabilization of long-term activation patterns within the convolutional network.

Oligodendrocytes are responsible for axonal myelination, enabling efficient and rapid nerve impulse conduction over long distances. This process not only increases signal propagation speed but also influences the synchronization of neural network activity [10]. Within the proposed architecture, the oligodendrocyte function is mapped to a mechanism for the stable reinforcement of signals in selected feature channels, facilitating the preferential processing of salient information in deep convolutional networks.

Microglia serve as the brain's immune system, participating in monitoring, removing redundant synaptic connections, and remodeling neural networks. This phenomenon, termed *synaptic pruning*, plays an essential role in optimizing neural network structure during development and learning [11]. In the proposed model, microglia-inspired mechanisms are utilized to identify low-salience channels, enabling the detection of potentially redundant parameters and supporting subsequent structural model reduction.

3.1 Overall Architecture

Figure 1 illustrates the schematic of the proposed glia-controlled convolutional neural network architecture. The primary data processing stream consists of a conventional CNN framework, comprising sequential convolutional layers interspersed with a channel activity modulation mechanism. For each convolutional layer, global feature map statistics are computed to serve as input for the corresponding glial controller. This controller models the synergistic interaction of three glial cell types: astrocytes, oligodendrocytes, and microglia. The outputs of these individual modules are integrated into an adaptive gating vector, which subsequently scales the activity of the respective feature channels. This modulation is executed via element-wise multiplication of the feature maps by the gating coefficients, enabling selective reinforcement or suppression of filters based on the input data context. Furthermore, the microglia module generates activity masks, which can be interpreted as a measure of channel salience and utilized to assess the network's susceptibility to parameter reduction.

3.2 Glial Controller Design

The glial controller represents a specialized control module designed to generate adaptive channel gating vectors that govern the information flow within the convolutional network. In contrast to conventional channel attention mechanisms,

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

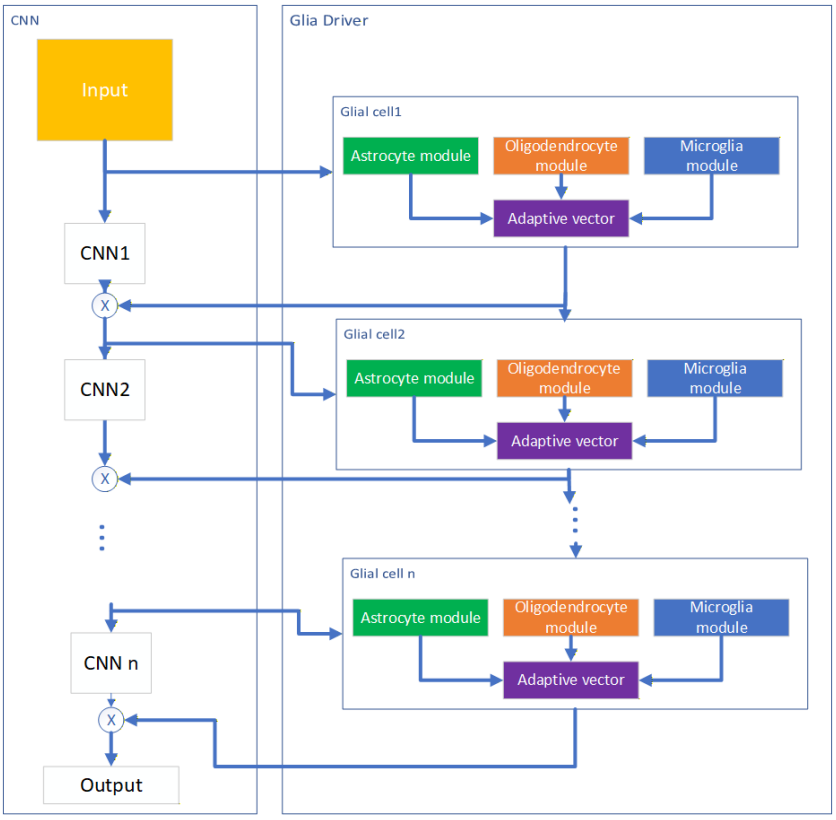

Fig. 1. Schematic of the glia-controlled convolutional neural network architecture.

the proposed controller is based on the synergistic interaction of several biologically inspired submodules, which model the functions of various glial cell types.

A single glial controller cell is illustrated in Figure 2. This cell incorporates two distinct information sources: the current signal from the CNN in the form of a channel statistics vector, and a historical signal originating from the preceding glial cell. This dual-input configuration enables the controller to incorporate operational context across successive network layers.

Formally, for each convolutional layer l, the glial controller generates a channel gating vector $\mathbf{g}^{(l)} \in \mathbb{R}^{C_l}$, which is subsequently utilized to modulate the feature maps within the CNN via per-channel scaling.

Controller Input Streams. A channel-wise statistics vector is computed for the layer l $\mathbf{s}^{(l)} \in \mathbb{R}^{C_l}$ via global average pooling of the feature maps:

$$\mathbf{s}^{(l)} = \mathrm{GAP}\left(\mathbf{X}^{(l)}\right), \qquad \mathbf{X}^{(l)} \in \mathbb{R}^{B \times C_l \times H \times W}. \tag{1}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

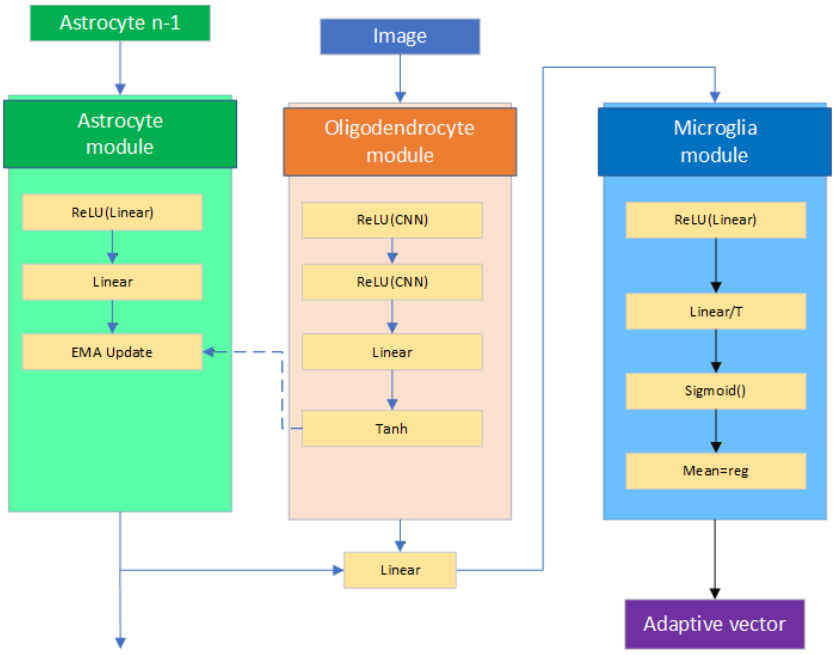

Fig. 2. Schematic of a single glial controller cell. The cell incorporates two distinct information sources: the current signal from the CNN (represented as a channel statistics vector) and a historical signal from the preceding glial cell. The outputs from the astrocyte, oligodendrocyte, and microglia modules are integrated into an adaptive gating vector, which modulates the channel activity of the convolutional network.

Vector $\mathbf{s}^{(l)}$ describes the current activation state of the layer filters.

Furthermore, the controller can incorporate a recurrent signal from the previous glial unit, designated as $\mathbf{h}^{(l-1)} \in \mathbb{R}^{C_{l-1}}$. This signal can be interpreted as either the preceding gating vector or the controller's state vector. In the event of a change in the number of channels between successive layers, a linear projection is applied:

$$\hat{\mathbf{h}}^{(l-1)} = \mathbf{P}^{(l)}\mathbf{h}^{(l-1)}, \qquad \hat{\mathbf{h}}^{(l-1)} \in \mathbb{R}^{C_l}. \tag{2}$$

Astrocyte module (slow-varying component). The astrocyte module generates a slow-varying modulation signal. $\mathbf{a}^{(l)}$ based on the historical signal. The processing is implemented using a multi-layer perceptron (MLP) and an exponential moving average (EMA) mechanism, which stabilizes temporal variations:

$$\mathbf{y}^{(l)} = f_{\text{astro}}\left(\hat{\mathbf{h}}^{(l-1)}\right), \tag{3}$$

$$\mathbf{u}_t^{(l)} = \lambda \mathbf{u}_{t-1}^{(l)} + (1-\lambda)\,\text{mean}_B\left(\mathbf{y}_t^{(l)}\right), \tag{4}$$

$$\mathbf{a}_t^{(l)} = \mathbf{y}_t^{(l)} + \mathbf{u}_t^{(l)}, \tag{5}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

where $\lambda \in (0,1)$ is the EMA smoothing factor, while $\text{mean}_B(\cdot)$ denotes averaging across the batch dimension.

Oligodendrocyte Module (Signal Reinforcement) The oligodendrocyte module processes the current channel statistics $\mathbf{s}^{(l)}$ and generates reinforcement signal $\mathbf{o}^{(l)}$ constrained to the interval $(-1,1)$:

$$\mathbf{o}^{(l)} = \tanh\Big(f_{\text{oligo}}\Big(\mathbf{s}^{(l)}\Big)\Big). \tag{6}$$

Microglia Module (Selection Mask and Regularization) The microglia module is responsible for the soft selection of channels through the generation of a mask $\mathbf{m}^{(l)} \in (0,1)^{C_l}$:

$$\mathbf{m}^{(l)} = \sigma\left(\frac{f_{\text{micro}}\big(\mathbf{s}^{(l)}\big)}{T}\right), \tag{7}$$

where $T > 0$ is a temperature parameter that scales the output of the logistic function. Additionally, a regularization term is defined:

$$\mathcal{L}^{(l)}_{\text{micro}} = \frac{1}{C_l}\sum_{c=1}^{C_l} m_c^{(l)}. \tag{8}$$

This value describes the average "openness" of the gates and can be utilized to control the sparsity of the channel masks.

Signal Integration and Gating Generation The outputs of the glial submodules are integrated into a single channel gating vector. In the proposed implementation, a reinforcement component is first determined based on the signals from the astrocyte and oligodendrocyte modules:

$$\mathbf{z}^{(l)} = \tanh\Big(w_a\mathbf{a}^{(l)} + w_o\mathbf{o}^{(l)}\Big), \tag{9}$$

$$\mathbf{gain}^{(l)} = \exp\Big(\tfrac{1}{2}\mathbf{z}^{(l)}\Big), \tag{10}$$

which is subsequently modulated by the microglia mask:

$$\mathbf{g}^{(l)} = b \cdot \mathbf{gain}^{(l)} \odot \Big(\mathbf{m}^{(l)}\Big)^{\text{clamp}(w_m,0,3)}, \tag{11}$$

where w_a, w_o and w_m b are learnable parameters, b is a (gating) scale constant, $\odot$ denotes the element-wise product, whereas clamp$(\cdot)$ constrains the value of the exponent.

Signal Propagation Between Cells and stop-gradient To ensure training stability, a *stop-gradient* mechanism is introduced when passing information between successive glial cells. This means that the historical signal $\mathbf{h}^{(l)}$ passed to the subsequent cell is computed as:

$$\mathbf{h}^{(l)} = \text{stopgrad!}\left(\mathbf{g}^{(l)}\right), \tag{12}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

which corresponds to the implementation call `h = gate.detach()`. The application of this operation restricts gradient propagation between successive glial cells, preventing unstable recurrence while simultaneously preserving the ability to utilize historical information within the network control process.

3.3 Training Strategy

The training process of the proposed architecture was designed to mitigate instability resulting from the simultaneous adaptation of two interdependent components: the convolutional neural network (CNN) and the glial controller (GLIA), which dynamically modulates channel activations. Rather than updating all parameters concurrently, an *alternating training* strategy was employed, in which each epoch consists of two phases. In Phase A, only the glial component parameters are updated while the CNN weights are frozen. Conversely, in Phase B, only the convolutional part is trained while the glial controller weights remain frozen. In practice, this was implemented by toggling the *required_ grad* flags for subsets of parameters and performing two separate optimization steps within the same epoch. Such a scheme reduces the effect of *co-adaptation*, where rapid changes in gating could destabilize filter updates, and conversely, fast adaptation of the CNN could lead to controller oscillations.

In Phase A (GLIA training), the objective function to be minimized was the sum of the classification loss and the regularization term derived from the microglia module. For an input sample $\mathbf{x}$ and label y, the cross-entropy (*cross-entropy*) loss was employed for multi-class classification, while the microglia regularization was defined as the mean value of the mask $\mathbf{m}$ across the channels:

$$\mathcal{L}_{\text{GLIA}} = \mathcal{L}_{\text{CE}}(\hat{\mathbf{y}}, y) + \lambda_{\text{micro}} \cdot \frac{1}{C} \sum_{c=1}^{C} m_c, \tag{13}$$

where λ_{micro} is a weighting coefficient. Since $m_c \in (0, 1)$ (following the sigmoid activation), minimizing the mean mask value exerts pressure toward sparser channel "openings" and facilitates subsequentfilter excludability analysis (in the sense of soft *pruning*).

In Phase B (CNN training), only $\mathcal{L}_{\text{CE}}$ was optimized, utilizing gates generated by the frozen controller. In both phases, the AdamW optimizer (Adam with decoupled *weight decay*) was employed, which is a standard choice for deep network training due to its stability and robust convergence [14, 15]. Separate optimizers with independent learning rates were used for the CNN and GLIA components. Additionally, a Cosine Annealing scheduler was implemented for both optimizers, reducing the learning rate according to a cosine curve throughout the epochs [16].

In the convolutional backbone, Batch Normalization was integrated into the convolutional layers, along with Dropout applied at several points in the architecture, to enhance training stability and mitigate overfitting [4, 17]. For the CIFAR dataset, standard input augmentations were employed, including (*random crop-*

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

ping) with padding and (*horizontal flipping*), followed by channel normalization; a typical training configuration for image classification on this benchmark.

Model quality assessment was performed after each epoch on the validation/test set, reporting the loss and classification metrics, specifically Accuracy and the *macro* and *weighted* versions of the F1-score. The *macro* variant allows for evaluating model behavior in the presence of class imbalance, while the *weighted* variant accounts for class frequencies. Additionally, based on the microglia masks, the proportion of potentially excludable convolutional parameters for a given threshold (e.g. $m < 0.2$) was estimated by counting channels with a low mean mask activation and calculating the corresponding number of weights within the convolutional filters.

4 Datasets

To comprehensively evaluate the proposed glial network structures, the study was conducted in two stages, utilizing datasets of distinct characteristics. In the first stage, experiments were performed on the publicly available CIFAR-10 image datasets, proposed in [12]. These datasets serve as a standard benchmark in convolutional neural network research and are widely used to analyze the effectiveness of new architectures, regularization methods, and model optimization techniques. CIFAR-10 consists of 60 000 RGB color images with a resolution of 32×32 pixels, divided into 10 classes.

In the second stage of the study, the proposed solutions were verified using real-world data derived from firewall logs, describing user activity within a computer network. This stage aimed to evaluate the utility of glial networks in the practical task of profiling and identifying users based on their network behavior, as well as to verify the feasibility of transferring concepts developed on image data to cybersecurity problems. At a comparable model complexity following the structural pruning stage, the proposed architecture achieves higher classification effectiveness than the GliaNet model [6], reaching an accuracy of 97.40% compared to the 96.90% reported in the comparative study.

4.1 Network User URL Log Dataset and Encoding

The experiments regarding user classification based on network activity utilized a dataset derived from real-world firewall logs. The data originated from a Palo Alto device cluster operating in *active-active* mode and contained HTTP/HTTPS traffic events recorded as URL requests. The logged traffic resulted from both user activity (e.g. web browsing) and background application activity (e.g. system updates, antivirus software). A total of 1,408,000 records were collected for 200 unique user accounts, and the dataset was partitioned into training and test sets in an 80:20 ratio. To evaluate the temporal robustness of the solution, additional logs were collected from the same environment approximately three months after the system construction, forming a separate validation set. Prior to analysis, the data were anonymized and verified for sensitive information;

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

strings corresponding to email addresses and phone numbers were replaced with `<email>` and `<phone>` tokens.

The stream of URL requests was transformed into training examples by grouping events into user activity sessions, where each session contained between 8 and 300 URLs. An additional segmentation criterion was time-based: if the interval between two events for a given user exceeded 30 minutes, the requests were assigned to separate sessions. To enable the application of convolutional neural networks to textual data, each session was encoded using a (*one-hot* method at the character level, defining an alphabet of size $n = 70$ encompassing letters, digits, and special characters typical for URL syntax (e.g. `/`, `:`, `.`, `-`, `_`, `?`, `#`, `&`, `%`). Each character was represented by a binary vector $\mathbf{e} \in \{0,1\}^{70}$, and the entire session as a binary matrix $\mathbf{X} \in \{0,1\}^{L\times 70}$, where L denotes the length of the character sequence in the session. These encoded data served as input to a baseline 1D-CNN architecture, comprising three convolutional layers coupled with *max pooling* layers, followed by three *fully connected* layers for user classification. Due to the high-dimensional nature of the character encoding, a constraint on the maximum number of URLs per session was introduced to reduce memory costs and facilitate GPU-accelerated training.

5 Results and Discussion

This section presents the experimental results aimed at evaluating the effectiveness of the proposed glial control mechanism in convolutional neural networks. All vision-based experiments were conducted on the same CIFAR-10 dataset, enabling a direct comparison of the impact of glial control across architectures of varying complexity. Two network variants were considered: a classical CNN architecture and a deep residual network, ResNet-18. In both cases, an identical experimental procedure was applied: (i) a comparison between the baseline model and the Glia-controlled variant, (ii) estimation of channel excludability based on microglia masks, and (iii) structural pruning, followed by the removal of the glial controller and the final evaluation of the streamlined architecture.

5.1 CIFAR-10: CNN and ResNet-18 comparison

In the first stage, a classical CNN architecture was analyzed, followed by the application of the same glial control mechanism to a deep ResNet-18 architecture. The results, summarized in Table 1, demonstrate that in both cases, the Glia-controlled variant achieves higher classification quality than its corresponding baseline model. Simultaneously, the microglia masks enable the identification of channels with a low impact on the classification outcome.

For both architectures, structural pruning was performed on channels satisfying the condition $m < 0.2$, after which the glial controller was completely removed. In both instances, it was observed that the reduction in the number of parameters does not lead to a significant degradation in classification quality. This effect is particularly evident for ResNet-18, where it was possible to

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

Table 1. CIFAR-10: Comparison of baseline models, Glia-controlled variants, and networks after channel pruning ($m < 0.2$) and glial controller removal.

Architecture	Variant	Parameters [mln]	Accuracy	F1 (macro)
CNN	Baseline	4.63	87.22%	0.8711
CNN	Glia-controlled	5.93	89.31%	0.8923
CNN	Pruned (bez GLIA)	4.11	88.98%	0.8891
ResNet-18	Baseline	11.69	96.58%	0.9656
ResNet-18	Glia-controlled	16.18	97.46%	0.9744
ResNet-18	Pruned (bez GLIA)	4.75	97.40%	0.9738

reduce the parameter count by 6.94 million while maintaining an accuracy level of 97.40%.

5.2 User identification based on URLs

The final experiment addressed the practical task of identifying computer network users based on character-level encoded URL sessions. The application of glial control within a 1D convolutional neural network (1D-CNN) yielded a minor yet stable improvement in classification quality compared to the baseline model. An analysis of the microglia masks enabled the identification of feature maps with a low impact on the classification decision. Following their removal and the detachment of the glial controller, an optimized CNN was obtained. Its performance was comparable to the Glia-controlled variant, while simultaneously reducing model complexity. These results confirm the utility of the glial mechanism also in tasks involving text and network data analysis.

The experiment was conducted on a dataset encompassing 200 network users, where each input sample represented a user activity session consisting of 8 to 300 URLs. The input data were encoded as binary matrices, and a one-dimensional convolutional neural network (1D-CNN) with three convolutional layers and three fully connected layers was employed for classification. Similarly to the previous experiments, a glial controller was introduced after each convolutional layer to generate per-channel saliency masks.

In the first stage, the baseline model (a 1D-CNN without the glial component) was compared with the variant extended by glial control. Subsequently, based on the mean values of the microglia masks, feature maps with a low impact on the classification decision were identified. Feature maps for which the sigmoid output values were below the threshold of $m < 0.3$ across all classes were treated as candidates for removal. Following the selection stage, the architecture was reduced by removing the designated feature maps and their corresponding filters. In the final step, the glial controller was detached, and the streamlined CNN was subjected to evaluation.

The final results are summarized in Tables 2 and 3. The application of the glial mechanism enabled both an improvement in classification quality and a significant reduction in the number of feature maps, without causing a substantial degradation in the system's performance.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

Table 2. User Identification (URL): Comparison of Final Metrics for the Baseline Model and the Glia-Controlled Variant

Metrics	Baseline 1D CNN	Glia-controlled 1D CNN
Loss	0.4261	0.4017
Accuracy	81.00%	81.40%
Precision (macro)	0.8085	0.8132
Recall (macro)	0.8100	0.8140
F1 (macro)	0.8089	0.8130

Table 3. User Identification (URL): Pruning of feature maps identified by microglia and removal of the glial controller.

Variant	Feature Maps	Accuracy	ΔAccuracy
Glia-controlled 1D CNN (BEST)	32000	81.40%	–
Pruned 1D CNN (without GLIA)	26273	81.00%	−0.0040

Interpretational remarks. A minor decrease in accuracy following the removal of the glial controller and the reduction of feature maps (below 0.5 percentage points) does not significantly impact system functionality and remains within the statistical variance of the training process. The obtained results confirm that the microglia mechanism provides a useful and interpretable signal for selecting user-specific feature maps. Consequently, it is possible to construct smaller, specialized models dedicated to specific user groups without the need to retrain the entire network.

6 Conclusions

This paper presented the Glia-Controlled Convolutional Network, a biologically inspired architecture that augments a standard CNN with a glial controller emulating the functions of astrocytes, oligodendrocytes, and microglia. The controller generates adaptive per-channel gating vectors that dynamically modulate feature map activations, while the alternating training strategy ensures stable co-adaptation of the convolutional and glial components. Experimental evaluation on CIFAR-10 demonstrated consistent improvements over baseline models for both a classical CNN and ResNet-18. Notably, after microglia-guided pruning and removal of the glial controller, ResNet-18 retained 97.40% accuracy while reducing its parameter count from 16.18M to 4.75M – a reduction of over 70% with negligible performance loss. The mechanism was further validated on a real-world cybersecurity task, where a 1D-CNN enhanced with glial control achieved improved user identification accuracy on character-level URL session data, and the microglia masks enabled meaningful structural reduction without retraining. These results confirm that biologically inspired glial mechanisms constitute an effective and interpretable adaptive control layer for deep convolutional networks, applicable across both image and sequential data domains. Future work may explore integration with transformer-based architectures, dy-

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

namic threshold selection for pruning, and extension to continual or federated learning settings.

References

1. Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. In: International Conference on Learning Representations (ICLR) (2015)
2. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. IEEE (2016)
3. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How Transferable Are Features in Deep Neural Networks? In: Advances in Neural Information Processing Systems (NeurIPS), pp. 3320–3328 (2014)
4. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research **15**(1), 1929–1958 (2014)
5. Hu, J., Shen, L., Sun, G.: Squeeze-and-Excitation Networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7132–7141. IEEE (2018)
6. Han, M., Pan, L., Liu, X.: GliaNet: Adaptive Neural Network Structure Learning with Glia-Driven Modulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1–13 (2025)
7. Han, M., Pan, L., Liu, X.: MA-Net: Rethinking Neural Unit in the Light of Astrocytes. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, pp. 2040–2048 (2024)
8. Fields, R.D.: The Other Brain: The Scientific and Medical Breakthroughs That Will Heal Our Brains and Revolutionize Our Health. Simon & Schuster, New York (2011)
9. Perea, G., Navarrete, M., Araque, A.: Tripartite synapses: astrocytes process and control synaptic information. Trends in Neurosciences **32**(8), 421–431 (2009)
10. Fields, R.D.: A new mechanism of nervous system plasticity: activity-dependent myelination. Nature Reviews Neuroscience **16**(12), 756–767 (2015)
11. Paolicelli, R.C., et al.: Synaptic pruning by microglia is necessary for normal brain development. Science **333**(6048), 1456–1458 (2011)
12. Krizhevsky, A.: Learning Multiple Layers of Features from Tiny Images. Technical Report, University of Toronto (2009)
13. Krizhevsky, A., Hinton, G.: Learning Multiple Layers of Features from Tiny Images. University of Toronto (2009), available at https://www.cs.toronto.edu/ kriz/cifar.html
14. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: International Conference on Learning Representations (ICLR) (2015)
15. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization. In: International Conference on Learning Representations (ICLR) (2019)
16. Loshchilov, I., Hutter, F.: SGDR: Stochastic Gradient Descent with Warm Restarts. In: International Conference on Learning Representations (ICLR) (2017)
17. Ioffe, S., Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In: Proceedings of the International Conference on Machine Learning (ICML), pp. 448–456 (2015)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29

18. Nowak, J., Holotyak, T., Korytkowski, M., Scherer, R., Voloshynovskiy, S.: Fingerprinting of URL logs: continuous user authentication from behavioural patterns. In: International Conference on Computational Science (ICCS), pp. 184–195 (2021)
19. Nowak, J., Korytkowski, M., Nowicki, R., Scherer, R., Siwocha, A.: Random forests for profiling computer network users. In: International Conference on Artificial Intelligence and Soft Computing (ICAISC), pp. 734–739 (2020)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_29