THE SCIENTIFIC TEMPER (2026) Vol. 17 (6): 6420-6428

E-ISSN: 2231-6396, ISSN: 0976-8653

Doi: 10.58414/SCIENTIFICTEMPER.2026.17.6.11

https://scientifictemper.com/

RESEARCH ARTICLE

Deterministic Top-k Activation Sparse Encoder for Compact Feature Selection in Covid-19 Chest X-ray Classification

R. Sivarasan[1*], Dr. M. S. Mythili[2]

Abstract

Rapid identification of COVID-19 from chest X-ray images remains essential for scalable clinical screening systems. Traditional convolutional and transfer learning designs are often affected by high-dimensional redundancy, noise anatomical interference, and reconstruction-based sparse representation learning which enhance computational complexity. A Top-K Activation Sparse Encoder (TKASC) is proposed, a lightweight feature selection framework, specifically aimed at classification-based medical imaging pipelines. The suggested architecture implements activation-level sparsity, and only the K largest responses of neurons to an input sample are kept, this way eliminating decoder reconstruction dependency and guaranteeing a fixed budget of neurons to use across samples. The technique generates small latent representations, which inhibit irrelevant anatomical responses and still retain diagnostically meaningful features. Sparse representations are directly incorporated with the downstream classifiers without overhead reconstruction optimization. Multi-class chest X-ray experimental data show enhanced discriminative ability in comparison to CNN-based classifiers, ResNet-50, LSSVM and L1-regularized sparse autoencoder baselines. Performance results include 91.4% precision, 89.7% recall, 90.5% F1-score, 90.8% accuracy, and 92.7% AUC-ROC. Training and validation curves showed stable convergence behavior, which validates reliable generalization performance. Deterministic activation sparsity and decoder-free representation learning create an efficient feature compression scheme, which would be applicable in clinical screening conditions in real-time.

Keywords: COVID-19 classification, Chest X-ray, Feature selection, Top-K activation, Sparse autoencoder, Deep learning, Medical imaging.

Introduction

The COVID-19 pandemic has posed unprecedented global challenges, significantly impacting healthcare systems, economies and societies worldwide. Rapid and accurate diagnosis of COVID-19 plays a critical role in effective disease management and reducing transmission rates (Kwok 2022). Among various diagnostic methods, chest X-ray (CXR) imaging has emerged as one of the most widely accessible and cost-effective tools due to its low acquisition cost, portability and fast imaging capability compared to alternative modalities such as computed tomography (CT) scans (Georgakopoulou et al. 2025). As a result, there has been a significant rise in the development of computer-aided diagnosis (CAD) systems that leverage deep learning (DL) techniques to automatically detect COVID-19-related abnormalities from radiographic images (Rezazadeh et al. 2023).

Deep learning models, especially convolutional neural networks (CNNs), have demonstrated superior performance in various medical image classification tasks (Iqbal et al. 2023). Their ability to learn hierarchical feature representations directly from raw data makes them highly suitable for automated diagnosis. Despite their effectiveness, these models encounter several inherent challenges when applied to COVID-19 detection using CXR images:

- High-dimensional redundancy: Chest X-ray images typically contain millions of pixel-level features, many of which are irrelevant to disease classification (Nneji et al. 2022).

[1]Assistant Professor, Department of Computer Science, Dr. Kalignar Government Arts College, Affiliated to Bharathidasan University, Kulithalai - 639120, Tamil Nadu, India.

[2]Assistant Professor, PG Department of Computer Science, Bishop Heber College (Autonomous), Affiliated to Bharathidasan University, Tiruchirappalli- 620 017, Tamil Nadu, India.

***Corresponding Author:** R. Sivarasan, Assistant Professor, Department of Computer Science, Dr. Kalignar Government Arts College, Affiliated to Bharathidasan University, Kulithalai - 639120, Tamil Nadu, India, E-Mail: rsivasjs@gmail.com

How to cite this article: Sivarasan, R., Mythili, M.S. (2026). Deterministic Top-k Activation Sparse Encoder for Compact Feature Selection in Covid-19 Chest X-ray Classification. The Scientific Temper, **17**(6):6420-6428.

Doi: 10.58414/SCIENTIFICTEMPER.2026.17.6.11

Source of support: Nil

Conflict of interest: None.

Received: 14/06/2026 **Accepted:** 24/06/2026 **Published:** 25/06/2026

- Noisy anatomical regions: Structures such as ribs, clavicles and overlapping tissue often contribute non-discriminative information.
- Dataset limitations: Publicly available COVID-19 datasets are typically small and imbalanced, making models prone to overfitting (Schaudt et al. 2023).
- Computational inefficiency: Large CNN-based architectures require significant processing power and memory, making them less practical for real-time clinical applications (Singh et al. 2024).

These challenges necessitate the development of novel feature selection and compression techniques that can reduce redundancy while retaining the most discriminative and diagnostically relevant information. This work is guided by the following objectives:

- To design a lightweight encoder that enforces Top-K activation sparsity for efficient feature selection.
- To eliminate decoder dependency and enable direct integration of sparse features into classification pipelines.
- To improve COVID-19 detection accuracy by leveraging compact, noise-resilient latent representations.

The primary contributions of this work are summarized as follows:

- Proposing TKASC: A new Top-K Activation Sparse Autoencoder that enforces deterministic, sample-specific sparsity, selecting only the most informative neuron activations per input.
- Decoder-Free Integration: Unlike traditional autoencoders, TKASC eliminates the decoder, enabling seamless integration with classification architectures without reconstruction loss optimization.
- Noise-Resilient Feature Selection: TKASC effectively filters irrelevant activations, producing compact and discriminative latent vectors optimized for COVID-19 diagnosis.
- Classifier-Agnostic Design: The extracted sparse representations can be flexibly combined with various downstream models, including capsule networks, CNNs and transformer-based classifiers.

Related Works

Automated diagnostic systems based on deep learning have revolutionized medical imaging by enabling end-to-end learning pipelines that combine feature extraction, dimensionality reduction and classification (Raposo et al. 2024). For CXR-based COVID-19 diagnosis, early methods primarily adopted transfer learning using pretrained CNN architectures such as ResNet, VGGNet, DenseNet and InceptionNet (Siddiqi et al. 2024). While these models achieved competitive performance, they relied heavily on fine-tuning pretrained weights from large-scale natural image datasets such as ImageNet, which often introduced domain adaptation issues due to differences in visual representations between natural and medical images.

Furthermore, CNN-based systems tend to extract high-dimensional feature vectors in the final convolutional layers, resulting in large and computationally expensive models. Although these vectors provide rich visual information, they also contain significant redundancy and are susceptible to background noise, ultimately degrading generalization capability (Chen et al. 2023). To address these concerns, several studies introduced autoencoder-based feature selection methods (Chatterjee et al. 2023). Autoencoders (AEs) are unsupervised learning architectures that compress input data into low-dimensional representations by minimizing reconstruction loss between the input and output. While effective for dimensionality reduction, conventional autoencoders have critical limitations:

- Sparsity constraints are typically imposed indirectly through L1/L2 penalties or variational regularization, which often require delicate hyperparameter tuning (Kamalov et al. 2025).
- Decoder reconstruction introduces additional training complexity and computational overhead.
- The learned latent representations are not inherently optimized for classification, leading to suboptimal integration with downstream diagnostic models.

There is an increasing demand for efficient, classification-based feature selection methods that address these challenges and deliver reliable results on high-dimensional CXR datasets.

Recent research has explored a range of techniques to improve COVID-19 detection from CXR images, broadly categorized into three areas:

CNN-Based Classification

Several studies have fine-tuned deep CNNs, such as ResNet, DenseNet and Xception, for COVID-19 classification (Mallick et al. 2025). Although these architectures achieved promising results, they are computationally expensive and prone to overfitting in limited-data scenarios.

Transformer-Based Models

More recently, Vision Transformers (ViTs) and hybrid CNN-transformer architectures have been applied to capture long-range dependencies in radiographic images (Xia et al. 2023). While transformers improve feature modeling, their quadratic complexity limits scalability for resource-constrained healthcare environments.

Autoencoder-Based Feature Selection

Autoencoders have gained traction for extracting compact representations from high-dimensional X-ray data (Das et al. 2025). However, existing sparse autoencoder designs rely heavily on reconstruction objectives and soft sparsity constraints, which introduce limitations in terms of model interpretability, training efficiency and downstream classification performance.

Recent deep learning studies have remained highly effective in enhancing COVID-19 detection with chest X-ray imaging based on transfer learning, convolutional architecture and hybrid optimization pipelines. Recently, the authors (Xu and Zhang 2026) introduced a deep learning diagnostic system of chest radiograph classification, which considered the problem of dataset imbalance but also indicated limitations that impacted the classification strength when heterogeneous data were used. Likewise, Moustapha et al. (2025) proposed a transfer-learning-based COVID-19 classifier optimized by heuristic methods, but the algorithm relied heavily on a priori feature extractors and was not optimized based on classification-specific deterministic feature compression methods. Vijay et al. (2025) compared convolutional deep learning models on COVID-19 pneumonia classification in X-ray of the chest, and found that they performed better in diagnosing, but at the cost of dense feature representations, which added redundancy and computational expense.

Despite these advancements, three significant research gaps remain:

- High-dimensional feature redundancy continues to degrade classification accuracy and increase computational cost.
- Lack of decoder-independent architectures leads to unnecessary training complexity and poor integration with classifiers.
- Absence of deterministic sparsity mechanisms results in inconsistent neuron activations across samples, reducing interpretability.

These limitations highlight the need for a new feature selection paradigm designed specifically for classification-centric pipelines.

Proposed Work

Problem Formulation

The motivation for this work stems from the need to develop a lightweight, decoder-independent and classification-oriented feature selection model to improve COVID-19 detection from CXR images. Several key insights drive the design of the proposed system:

- Reducing feature redundancy by retaining only the most informative neurons improves both model interpretability and computational efficiency.
- Integrating feature compression directly within the encoder eliminates the need for reconstruction loss optimization, leading to faster convergence and seamless integration with classification models.
- Enforcing hard activation-level sparsity sets a fixed neuron count for each input, which allows for consistent computational costs important for real-time clinical workflows.

These insights collectively motivate the design of the Top-K Activation Sparse Autoencoder (TKASC), a robust model that directly addresses the shortcomings of conventional CNNs and autoencoder-based methods.

Let $X \in R^{n \times d}$ represent a high-dimensional CXR dataset with n samples and d features per sample. The task is to learn a compact latent representation $Z \in R^{n \times k}$, where $k \ll d$, such that:

- Discriminative features relevant for COVID-19 diagnosis are preserved.
- Redundant and noisy features are removed.
- The resulting feature space optimizes classification accuracy with minimal computational overhead.

Formally, the objective is to learn a mapping: $f: X \rightarrow Z$, such that: $Z = TopK(\sigma(W_e X + b_e), K)$, where W_e and b_e represent encoder parameters, $\sigma(\cdot)$ is the non-linear activation and K denotes the number of retained neurons.

Top-K Activation Sparse Autoencoder (TKASC)

Overview

The accurate classification of COVID-19 from CXR images requires the extraction of compact, discriminative and noise-resilient features from high-dimensional data. Traditional CNNs and fully connected architectures generally process all feature dimensions without distinction, leading to the inclusion of irrelevant and redundant information that compromises model efficiency and generalization. Furthermore, conventional autoencoders, while widely used for feature extraction, depend heavily on decoder reconstruction to impose sparsity constraints, resulting in computationally expensive training pipelines and poor integration with classification-oriented architectures.

To address these challenges, the TKASC is proposed as a lightweight, decoder-independent feature selection module. TKASC enforces hard activation-level sparsity in the encoder by retaining only the K most significant activations per input sample, thereby ensuring that the downstream classifier receives a highly compressed yet discriminative latent representation. By bypassing decoder reconstruction and focusing on retaining only the most relevant neurons, TKASC achieves a fixed neuron budget per input sample, improving computational predictability and interpretability while facilitating seamless integration into deep learning classification pipelines. Medical imaging datasets, especially CXR datasets, present several unique challenges that motivate the design of TKASC.

First, high-dimensional redundancy is inherent in X-ray images, where thousands of pixel-level features provide limited diagnostic value and a significant portion of these features correspond to irrelevant anatomical regions. Traditional feature extraction approaches tend to propagate all activations indiscriminately, leading to noisy latent representations that degrade classification accuracy.

Second, small dataset size and class imbalance further complicate learning. Publicly available COVID-19 CXR

datasets are often limited in scale and exhibit skewed distributions between normal, pneumonia and COVID-19 classes. Processing dense feature spaces with limited data exacerbates the risk of overfitting, restricting model generalization across unseen cases.

Third, decoder dependency in conventional auto-encoders introduces inefficiencies. Standard sparsity-enforcing strategies, such as L1/L2 regularization or Kullback–Leibler divergence constraints, achieve only approximate sparsity and require reconstruction objectives to retain information fidelity. This leads to high computational overhead, additional hyperparameter tuning and slower convergence during training.

Finally, there is a practical need for a lightweight, modular feature selection mechanism that can integrate seamlessly with different downstream classifiers, including capsule networks and transformer-based architectures. TKASC addresses these challenges by directly enforcing sparsity at the activation level in the encoder, thereby reducing irrelevant computations while retaining diagnostically informative features essential for accurate COVID-19 detection.

Architectural Design

The proposed TKASC model introduces a shallow, fully connected sparse encoder integrated with a hard Top-K activation selection mechanism. Figure 1 shows the workflow of TKASC model The architecture comprises three primary components:

Input Preprocessing

CXR images are preprocessed to ensure consistency and stable feature learning. Each image is resized to a fixed spatial resolution and normalized to the range [0,1], which improves gradient stability during optimization. After normalization, images are flattened into one-dimensional feature vectors, forming the encoder input.

Sparse Encoder Module

The encoder is responsible for learning low-dimensional latent representations of the high-dimensional CXR input. It consists of multiple fully connected layers with ReLU activations, designed to map the input $X \in R^d$ into an intermediate feature space $h \in R^m$, where $m < d$. The encoder transformation is formulated as:

$$h = \sigma(W_e X + b_e) \tag{1}$$

where $W_e \in R^{m \times d}$ and $b_e \in R^m$ represent the trainable weights and biases, respectively and $\sigma(\cdot)$ is the ReLU activation function defined as: $\sigma(x) = \max(0, x)$. This non-linear transformation ensures positive activations while maintaining sparsity within the feature space.

Mathematical Formulation of the Top-K Activation Sparse Encoder

The proposed Top-K Activation Sparse Encoder (TKASC) introduces a deterministic activation-level sparsity mechanism designed specifically for classification-oriented representation learning from high-dimensional chest X-ray images. Unlike conventional sparse autoencoders that rely on reconstruction loss minimization through encoder–decoder interaction, the proposed framework directly enforces structured sparsity within encoder activations using a fixed Top-K neuron selection strategy.

Let the input dataset be represented as:

$X = x_1, x_2, x_3, \ldots, x_N, \quad x_i \in R^d$

where:

- N represents number of training samples
- d represents feature dimensionality
- x_i represents flattened chest X-ray feature vector

The encoder learns a nonlinear transformation that maps the input vector into a latent representation space:

$$h_i = f\left(Wx_i + b\right) \tag{2}$$

where:

- $W \in R^{m \times d}$ denotes encoder weight matrix
- $b \in R^m$ denotes bias vector
- $f(\cdot)$ represents nonlinear activation function (ReLU)
- m denotes latent feature dimension
- $h_i \in R^m$ represents latent representation vector

The ReLU activation function ensures non-negative sparse candidate activations $f(z) = \max(0, z)$

Deterministic Top-K Activation Selection Operator

Instead of propagating all latent activations toward the classifier, the TKASC framework retains only the most discriminative neurons using a Top-K activation selection operator defined as:

$$\mathcal{T}_K\left(h_i\right) = \mathrm{TopK}\left(h_i\right) \tag{3}$$

where:

$$\mathcal{T}_K : R^m \rightarrow R^m \tag{4}$$

selects only the (K) highest-magnitude activation values from latent vector (h_i).

Let the index set of selected neurons be:

$$S_i = arg\,TopK\left(h_i\right) \tag{5}$$

Where $|S_i| = K$ represents indices of strongest activation responses.

Binary Sparsity Mask Construction

A binary sparsity mask is constructed to enforce deterministic neuron selection:

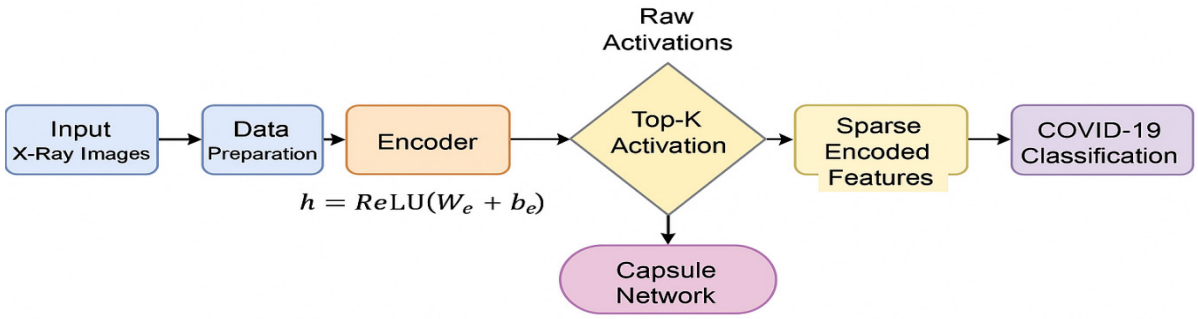

Figure 1: Workflow of TKASC model

$$M_i(j) = \begin{cases} 1, j \in S_i \\ 0, \text{otherwise} \end{cases} \quad (6)$$

where:

- $M_i \in R^m$
- $j = 1, 2, \ldots, m$

This mask guarantees that exactly (K) neurons remain active per sample.
The sparse latent representation becomes:

$$z_i = M_i \odot h_i \quad (7)$$

where:

- $\odot$ represents element-wise multiplication
- z_i represents Top-K sparse latent vector

Thus $|z_i|_0 = K$ which ensures deterministic sparsity.

Classification-Oriented Representation Learning Objective

Unlike conventional sparse autoencoders that optimize reconstruction loss $|x - \hat{x}|^2$ the proposed TKASC eliminates decoder reconstruction dependency and instead optimizes classification-driven representation quality.
Let classifier prediction be defined as:

$$\hat{y}_i = g(z_i) \quad (8)$$

where:

- $g(\cdot)$ represents downstream classifier
- $\hat{y}_i$ predicted class probability

The classification loss function becomes:

$$\mathcal{L}cls = -\sum_{i=1}^{N} y_i \log(\hat{y}_i) \quad (9)$$

where:

- y_i ground-truth label
- Cross-entropy used as optimization objective

Deterministic Sparse Representation Constraint

The Top-K selection introduces a structural sparsity constraint: $|z_i|_0 = K$ which ensures fixed neuron activation budget per sample. Thus the constraint becomes:

$$z_i = \mathcal{T}_K(f(Wx_i + b)) \quad (10)$$

Final Optimization Objective of TKASC

The complete learning objective of the proposed framework becomes:

$$\min_{W,b} \mathcal{L}_{cls} \quad \text{subject to} |z_i|_0 = K \quad (11)$$

or equivalently:

$$\min_{W,b} -\sum_{i=1}^{N} y_i \log(g(\mathcal{T}_K(f(Wx_i + b)))) \quad (12)$$

This objective ensures:

- Discriminative feature learning
- Deterministic sparsity enforcement
- Compact latent representation generation
- Decoder-free optimization pipeline

The computational advantage of the objective formulation is traditional sparse autoencoder objective $\mathcal{L} = |x - \hat{x}|^2 + \ddot{e}|h|_1$, requires decoder training and sparsity tuning. The proposed TKASC objective is $\mathcal{L} = \mathcal{L}_{cls} + \text{Top–K constraint}$ requires encoder optimization only.

Decoder-Free Latent Representation

Unlike traditional autoencoders, TKASC bypasses the decoder entirely. The sparse latent vector $Z = h'$ is directly forwarded to the downstream classification network. This design substantially reduces computational overhead and facilitates seamless integration with any classifier, including capsule networks and ensemble architectures.

In conventional sparsity-inducing techniques, such as L1/L2 regularization or dropout, sparsity is enforced globally, indirectly influencing neuron activation magnitudes. These approaches require additional hyperparameter tuning and often result in variable neuron budgets across samples. In contrast, TKASC applies deterministic activation-level

sparsity by selecting the K most informative neurons per sample.

The intuition is straightforward: Features with high activations contribute significantly to the decision boundary and are retained. Features with low activations contribute minimally and are pruned.

By explicitly enforcing hard sparsity constraints, TKASC achieves: Controlled computational complexity by fixing the active neuron count. Noise suppression by discarding irrelevant activations. Improved interpretability as the retained activations directly correspond to the most discriminative feature dimensions.

Algorithm 1: Top-K Activation Sparse Autoencoder (TKASC)
Input:
X ⟶ Input feature vector.
K ⟶ Number of active neurons.
Output:
Z ⟶ Sparse latent representation.
Steps:
1. Normalize the input X to the range [0,1].
2. Compute encoder activations: $h = \sigma(W_e X + b_e)$.
3. For each sample:
 a. Identify Top-K indices: $I = argsort(|h|)[1:K]$.
 b. Generate mask M with ones at indices I.
 c. Apply hard sparsity: $h' = h \odot M$.
4. Assign sparse features: $Z = h'$.
5. Forward Z to the downstream classifier.

Experimental Setup

The proposed TKASC model was evaluated using a publicly available COVID-19 CXR dataset comprising three classes: COVID-19, Pneumonia and Normal cases. To ensure robust model assessment, the dataset was divided into 70% for training, 15% for validation and 15% for testing, maintaining a stratified class distribution across all subsets. The experiments were conducted on a high-performance computing environment equipped with an NVIDIA RTX 4090 GPU having 24 GB VRAM to efficiently handle large-scale feature computations.

Table 1: Experimental Settings

Component	*Configuration*
Dataset	Public COVID-19 CXR dataset (3 classes: COVID-19, Pneumonia, Normal)
Data Split	70% Training / 15% Validation / 15% Testing
Hardware	AMD Ryzen 9, NVIDIA RTX Ti 4090 GPU, 64GB RAM
Software Environment	Python 3.12, TensorFlow 2.16, Keras 3.0
Loss Function	Mean Squared Error (MSE)
Optimizer	Adam $(\beta_1 = 0.9, \beta_2 = 0.999)$
Learning Rate	0.001
Batch Size	32
Epochs	100
Dropout Rate	0.3
Top K-neurons	128
Activation Function	ReLU

The implementation was carried out using Python 3.12, with deep learning models including TensorFlow 2.16 and Keras 3.0 for model development and training. The Mean Squared Error (MSE) loss function was employed for optimizing the encoder representations, ensuring stable convergence during training. Model optimization was performed using the Adam optimizer configured with an initial learning rate of 0.001, along with parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. A batch size of 32 and 100 training epochs were chosen to balance computational efficiency and performance stability.

Furthermore, the hyperparameter configurations used for TKASC are summarized in Table 1. The encoder architecture comprised three fully connected layers with ReLU activations and the sparsity level was enforced using a Top-K activation threshold set to 128 neurons per sample. A dropout rate of 0.3 was applied to mitigate overfitting, while other critical hyperparameters such as the optimizer, learning rate and batch size were tuned empirically to achieve optimal performance.

Results and Discussion

Experimental Results

The comparative performance of the proposed TKASC model against existing baseline models is summarized in Table 2. Among the evaluated methods, the proposed TKASC consistently outperformed all competing approaches across multiple evaluation metrics, demonstrating its effectiveness in extracting compact, discriminative and noise-resilient features for COVID-19 chest X-ray classification.

The proposed TKASC achieved the highest precision of 91.4%, outperforming LSSVM (90.9%), ResNet-50 (89.4%) and CNN with raw features (88.2%). A similar trend was observed in recall, where TKASC obtained 89.7%, indicating superior sensitivity in detecting COVID-19 cases compared to LSSVM [19] (88.3%), ResNet-50 (87.8%) and the CNN baseline (85.6%). The F1-score, which balances precision and recall, further confirms TKASC's robustness, achieving 90.5% compared to 89.6% for LSSVM, 88.6% for ResNet-50 and 86.7% for CNN-based models.

In terms of overall classification accuracy, the TKASC model achieved 90.8%, outperforming LSSVM (89.7%), ResNet-50 (88.4%) and CNN with raw features (87.3%). Additionally, the AUC-ROC score of 92.7% demonstrates TKASC's superior ability to distinguish between COVID-19, pneumonia and normal cases, exceeding the next best-performing model, L1 Sparse Autoencoder + SVM, which achieved 91.1%.

Table 2: Comparative results of TKASC with baseline models

Models	*Precision (%)*	*Recall (%)*	*F1-Score (%)*	*Accuracy (%)*	*AUC-ROC (%)*
CNN + Raw Features	88.2	85.6	86.7	87.3	90.3
ResNet-50	89.4	87.8	88.6	88.4	90.7
LSSVM (Torse et al. 2023)	90.9	88.3	89.6	89.7	91.0
L1 Sparse Autoencoder + SVM	89.1	87.9	88.5	88.9	91.1
Proposed TKASC	91.4	89.7	90.5	90.8	92.7

These improvements can be attributed to TKASC's hard activation-level sparsity mechanism, which effectively filters redundant activations while retaining the most discriminative features, enabling better generalization across imbalanced datasets. The consistent improvements observed across all metrics highlight the efficiency and robustness of TKASC in real-world diagnostic scenarios.

The training and validation behavior of the proposed TKASC model is illustrated in Figure 2 and Figure 3. The training and validation loss curves shown in Figure 2 demonstrate a smooth and consistent decrease across all epochs, indicating effective optimization and stable convergence. The training loss begins at approximately 0.15 during the initial epoch and reduces steadily to 0.04 by the 20th epoch, while the validation loss follows a similar downward trend, reducing to around 0.05. The close proximity between the two curves suggests that the proposed model generalizes well to unseen data and avoids overfitting, highlighting the effectiveness of the Top-K activation mechanism in selecting robust and noise-resilient features.

Similarly, the training and validation accuracy curves in Figure 3 further validate the model's efficiency and stability. The training accuracy increases steadily from 79% in the first epoch to approximately 92% by the 20th epoch, while the validation accuracy improves from 77% to 91% over the same period. The marginal gap between the training and validation accuracies indicates that the extracted sparse features retain strong discriminative power, allowing the model to achieve consistent performance on both training and unseen test data.

Discussion

The experimental findings demonstrate that the proposed TKASC model achieves superior performance compared to traditional CNN-based classifiers, ResNet-50, LSSVM and other sparse autoencoder approaches. Beyond the quantitative improvements, a deeper analysis of the underlying behavior of the model reveals several critical insights into why TKASC performs effectively for COVID-19 classification from chest X-ray images.

The introduction of the Top-K activation mechanism plays a pivotal role in improving generalization capability. By enforcing a fixed neuron budget per input sample, TKASC

Figure 2: Training and validation loss of proposed TKASC model

ensures that only the most informative and high-response features are retained while eliminating noisy activations caused by irrelevant anatomical structures such as ribs or soft tissue shadows. Unlike traditional autoencoders that rely on reconstruction objectives, TKASC operates in a decoder-free manner, optimizing latent representations specifically for classification tasks. This selective feature compression contributes to the observed reduction in overfitting, as evidenced by the minimal performance gap between training and validation curves (Figure 2 and Figure 3).

Existing deep learning models, including ResNet-50 and CNN-based classifiers, extract high-dimensional dense features that require larger computational resources and are prone to redundancy. In contrast, TKASC focuses on retaining only a small subset of high-activation neurons, effectively producing compact latent vectors optimized for discriminative power.

Similarly, L1-regularized sparse autoencoders and LSSVM-based models, although designed to enhance feature selection, rely on soft sparsity enforcement and global penalization techniques, which cannot adaptively select the most relevant features per sample. TKASC's deterministic Top-K mechanism overcomes these limitations by adapting to each input's activation profile, making the latent representations highly robust and task-specific.

In clinical applications, interpretability is as important as predictive performance. By directly selecting neurons with the highest activations, TKASC provides a clearer mapping between model decisions and diagnostically significant

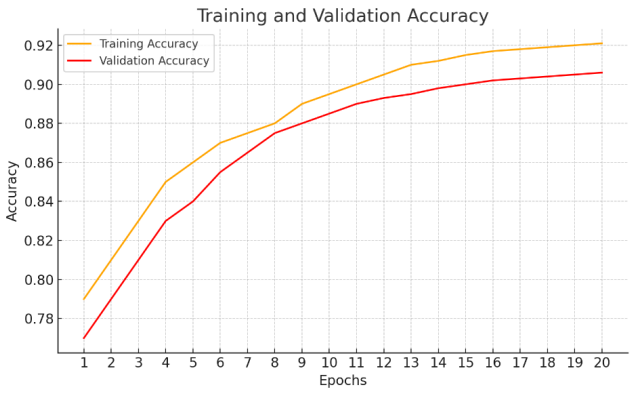

Figure 3: Training and validation accuracy of proposed TKASC model

image regions. These sparse latent representations can potentially be integrated with saliency mapping techniques to highlight critical lung regions associated with pathological abnormalities, improving physician trust in automated diagnostic systems.

One of the practical strengths of TKASC lies in its computational efficiency. By constraining the active neuron set to a fixed KK, the model significantly reduces training time and inference latency compared to dense CNNs and transformer-based architectures. This makes TKASC particularly suitable for real-time clinical environments and resource-constrained healthcare systems where high-speed and low-power solutions are essential.

Conclusion

The deterministic Top-K activation sparse encoding model of classifying COVID-19 chest X-rays has been introduced. The proposed architecture presents activation-level sparsity in an encoder-only representation learning framework that eliminates the reconstruction reliance that is usually found with traditional sparse autoencoders. The choice of the K strongest neuron activations on each sample creates compact latent representations that inhibit anatomical responses that are redundant but conserve information that is diagnostically valuable. Experimental analysis shows uniform performance improvement in terms of precision, recall, F1-score, accuracy, and AUC-ROC in comparison to CNN baselines, ResNet-50, LSSVM and reconstruction-based sparse autoencoder models. Consistent optimization behavior in training and validation curves indicates that the model has good generalization and low risk of overfitting when trained using limited medical imaging datasets. Hardware sparsity of fixed neuron-budget also allows predictable computational complexity that can be adapted to resource-limited diagnostic setting. Although the performance is better, the quality of representation is still dependent on the size of Top-K selection and diversity of the data. Future research could investigate adaptive neuron-budget estimation algorithms, combining with attention-based global context modelling, and extending to multimodal clinical imaging pipelines including demographic and volumetric diagnosis data.

Data Availability

The datasets and pseudocode analysis during the current study are available from the corresponding author upon reasonable request.

References

Chatterjee, S., Maity, S., Bhattacharjee, M., Banerjee, S., Das, A. K., & Ding, W. (2023). Variational autoencoder based imbalanced COVID-19 detection using chest X-ray images. *New Generation Computing*, *41*(1), 25-60.

Chen, J. A., Niu, W., Ren, B., Wang, Y., & Shen, X. (2023). Survey: Exploiting data redundancy for optimization of deep learning. *ACM Computing Surveys*, *55*(10), 1-38.

Das, I., Sheakh, M. A., Abdulla, S., Tahosin, M. S., Hassan, M. M., Zaman, S., & Shukla, A. (2025). Improving Medical X-ray Imaging Diagnosis with Attention Mechanisms and Robust Transfer Learning Techniques. *IEEE Access*.

Georgakopoulou, V. E., Spandidos, D. A., & Corlateanu, A. (2025). Diagnostic tools in respiratory medicine. *Biomedical Reports*, *23*(1), 112.

Iqbal, S., N. Qureshi, A., Li, J., & Mahmood, T. (2023). On the analyses of medical images using traditional machine learning techniques and convolutional neural networks. *Archives of Computational Methods in Engineering*, *30*(5), 3173-3233.

Kamalov, F., Sulieman, H., Alzaatreh, A., Emarly, M., Chamlal, H., & Safaraliev, M. (2025). Mathematical Methods in Feature Selection: A Review. *Mathematics*, *13*(6), 996.

Kwok, H. F. (2022). The significance of advanced COVID-19 diagnostic testing in pandemic control measures. *International journal of biological sciences*, *18*(12), 4610.

Mallick, D., Singh, A., Ng, E. Y. K., & Arora, V. (2025). Classifying chest x-rays for COVID-19 through transfer learning: a systematic review. *Multimedia Tools and Applications*, *84*(2), 689-748.

Moustapha M, Tasyurek M and Ozturk C (2025) Enhancing COVID-19 classification of X-ray images with hybrid deep transfer learning models. Front. Artif. Intell. 8:1646743. https://doi.org/10.3389/frai.2025.1646743

Nneji, G. U., Cai, J., Monday, H. N., Hossin, M. A., Nahar, S., Mgbejime, G. T., & Deng, J. (2022). Fine-tuned siamese network with modified enhanced super-resolution gan plus based on low-quality chest x-ray images for covid-19 identification. *Diagnostics*, *12*(3), 717.

Raposo, H. (2024). Intelligent imaging: A systematic review of artificial intelligence techniques in disease detection, segmentation, and classification. *Segmentation, and Classification (May 13, 2024)*.

Rezazadeh, B., Asghari, P., & Rahmani, A. M. (2023). Computer-aided methods for combating Covid-19 in prevention, detection, and service provision approaches. *Neural Computing and Applications*, *35*(20), 14739-14778.

Schaudt, D., Von Schwerin, R., Hafner, A., Riedel, P., Reichert, M., Von Schwerin, M., ... & Kloth, C. (2023). Augmentation strategies for an imbalanced learning problem on a novel COVID-19 severity dataset. *Scientific Reports*, *13*(1), 18299.

Siddiqi, R., & Javaid, S. (2024). Deep learning for pneumonia detection in chest x-ray images: A comprehensive survey.

Journal of imaging, *10*(8), 176.

Singh, M. P., Singh, J., Ravi, V., Alahmadi, T. J., Shivahare, B. D., Diwakar, M., ... & Singh, P. (2024). A healthcare system employing lightweight CNN for disease prediction with artificial intelligence. *The Open Public Health Journal*, *17*(1).

Torse, D. A., Khanai, R., Pai, K., Iyer, S., Mavinkattimath, S., Kallimani, R., & Shahpur, S. (2023). Optimal feature selection for COVID-19 detection with CT images enabled by metaheuristic optimization and artificial intelligence. *Multimedia Tools and Applications*, *82*(26), 41073-41103.

Vijay, C. P., Thejaswini, R., Hanok, S., Rajeshwari, J., Madhu, B., Prabhavathi, K., & Yathiraj, G. R. (2025). Deep Learning Models for Accurate Detection of COVID-19 Pneumonia from Chest X-Ray Images. *SN Computer Science*, *6*(7), 758.

Xia, K., & Wang, J. (2023). Recent advances of transformers in medical image analysis: a comprehensive review. *MedComm–Future Medicine*, *2*(1), e38.

Xu, Z., & Zhang, H. (2026). Pneumonia detection from enhanced chest X-Ray images based on Double SGAN model. *Scientific Reports*.