



(Research Article)

Chest X-Ray Outlier Detection, Classification Using Liquid Neural Networks with Dimension Reduction and Advanced Edge Detection

Shah Saloni N.¹, Kokare S. A.², Mujawar Sabaa Shafique³, Jadhav Priti Vijay⁴

^{1,2,3,4} Department of Computer Engineering, Sharadchandra Pawar College of Engineering & Management, Pune, 412306, Tal. Baramati, Dist. Pune, India

*Corresponding Author: pritivjadhav19@gmail.com

Received: 20.06.2026

Accepted: 25.06.2026

Published: 27.06.2026

ABSTRACT

In order to improve health care delivery and optimize resources within a health care delivery system, it is essential to provide early and accurate diagnoses of thoracic disease using Chest X-rays (CXR). In this research project we develop a novel integrated approach that includes Liquid Neural Networks (LNN) as a tool for image classification and anomaly detection. Additionally, Principal Component Analysis (PCA), and advanced edge detection methods are also used. We evaluate our method and compare its performance with Convolutional Neural Networks (CNN) for the task of CXR classification. Our model will include two distinct LNN models; one to classify normal vs. pneumonia cases and another to classify normal vs. lung opacities. Experimental testing was performed on synthetic data generated by PGGAN and validated on large clinical datasets. Both models demonstrated high levels of accuracy, i.e., Model 1 provided 94% accuracy for detecting pneumonia (with sensitivity = .92 and specificity = .96), and Model 2 provided 90% accuracy for identifying lung opacities. Edge detection provided improved features for image classification of approximately 8-12%. PCA provided a significant reduction in processing time of 40% and did not negatively impact classification accuracy. To facilitate point-of-care diagnostics, the two models are combined in a real-time web application created via Flask. The real-time web application provides an automatic triage mechanism to identify patients with abnormal findings first, which can potentially reduce the workload of radiologists involved in routine screening tasks by up to 35-45%. Furthermore, validation studies comparing LNNs to other state-of-the-art dee.p learning architectures (e.g. ResNet-50, DenseNet-121) demonstrate that LNNs possess superior ability to generalize across various types of training data, including synthetic training data. Overall, this study represents an important contribution towards developing artificial intelligence-assisted radiology systems and addressing unmet needs in terms of increasing access to diagnostic services for populations residing in resource-limited health care environments.

Keywords: *Liquid Neural Networks; Chest X-ray analysis; Computer-aided diagnosis (CAD); Dimension reduction; Edge detection; Deep learning; Medical image classification; Pneumonia detection; Lung opacity; Automated diagnosis; Temporal dynamics*

I. INTRODUCTION

The use of Chest Radiographs as a gold standard in the diagnosis of respiratory and Cardiac diseases continues to be used world-wide with estimates suggesting in excess of 2 billion chest x-rays are performed each year [1]. Despite their widespread adoption, however, the diagnostic accuracy of radiographic interpretations can vary greatly amongst practicing radiologists with studies showing that inter-rater agreement (i.e., agreement between radiologists interpreting the same images) can range anywhere from approximately 60-85 percent depending on the disease process or condition being evaluated and the individual's experience level [5]. These differences in diagnostic accuracy are due to several factors, including radiologist fatigue, increased cognitive burden associated with radiologic interpretation, and the presence of numerous radiologically subtle findings which require extended education and training to accurately interpret [5]. Moreover, current World Health Organization (WHO) estimates suggest that there is currently a significant global shortage of radiologists, with some areas having an estimated shortage of greater than 90 percent of needed specialist personnel [6]. Such shortages can impact turn-around-times for reports generated in response to emergency situations where timely evaluation of chest x-rays may be crucial to optimal patient care, resulting in an estimated 24-48 hour delay in reporting times at many institutions [6].

Advancements in Artificial Intelligence and Machine Learning technologies continue to advance medical imaging analyses [2]. In particular, convolutional neural networks (CNNs) have been successful in classifying medical images, such that they achieve similar performance levels compared to practicing radiologists when benchmarked against established radiologic interpretation standards [7]. However, most traditional CNN architectures rely on fixed-length feature representation and therefore cannot capture dynamic relationships within a sequence of events as typically occur within the natural history of a disease process. Liquid Neural Networks (LNNs) represent a novel architecture that utilizes Ordinary Differential Equations and continuous-time models to capture these dynamic aspects of complex medical data [3] and provide the capability for the model to adaptively modify its

internal state based upon input patterns. Therefore, LNNs are uniquely suited to support the processing of complex, nonstationary medical imaging sequences such as those found in chest x-rays.

In addition to advancements in deep learning architectures, other methods have also been developed to improve processing efficiency and reduce computational requirements of large datasets. Dimensionality reduction techniques (e.g., principal component analysis (PCA)) have been widely utilized to decrease the number of dimensions in high-dimensional image data thereby reducing computational complexity without loss of important discriminative information [4], [8]. Recent investigations combining dimensionality reduction with advanced edge detection techniques (including Canny, Sobel and Laplacian operators) show promise in enhancing the signal-to-noise ratio in radiological images and highlighting pathological features in radiological images [5].

This investigation proposes a comprehensive framework integrating Liquid Neural Networks with PCA-based dimensionality reduction and multiresolution edge detection (using Canny, Sobel, and Laplacian operators) for fully automatic chest x-ray analysis. Specifically, this study focuses on the diagnosis of two common and clinically relevant pulmonary abnormalities: pneumonia and lung opacity. Lung opacity represents one of the leading cause(s) of morbidity and mortality worldwide. The primary contributions of this investigation include: (1) developing new Liquid Neural Network architectures for medical image classification and theoretically justifying the use of liquid neural networks for capturing temporal relationships in sequential medical images; (2) performing a comparative analysis of different dimensionality reduction techniques applied to radiological images; (3) applying multiresolution edge detection for enhanced feature extraction; (4) deploying a production ready web application with real-time inference capabilities; and (5) providing extensive evaluations against baseline deep learning approaches utilizing both synthetic and real-world clinical data.

The remainder of this document will follow this organization: Section II contains a thorough literature survey regarding all prior work relevant to medical image analysis, deep learning in radiology, and neural network architectures. Section III will describe the proposed method, including data acquisition and pre-processing pipeline, LNN architecture design, and evaluation metrics. Section IV will report experimentally derived results along with a detailed performance analysis. Section V will discuss our results in light of previously published literature and highlight any potential clinical applications. Finally, Section VI will conclude with recommendations for future work.

II. LITERATURE REVIEW AND RELATED WORK

A. Deep Learning in Medical Image Analysis

The most recent survey research indicates that deep learning technology is revolutionizing the way that medical images are created [6]. Convolutional Neural Network (CNN) technology has been shown to be at its peak in terms of ability to detect all types of pathology including pneumonia [7], tuberculosis [9], and COVID-19 [10] and to achieve this by utilizing the ResNet, VGG, and Inception architectures to fine tune them on medical image databases [8]. The approach used in these studies treats each medical image individually and does not utilize any information about sequential or temporal relationships. A recent meta-analysis which included over 70 studies of chest x-ray classification indicated that an average of 88-96 percent accuracy was obtained among the many different CNN architecture designs that were studied, but it also found that there existed wide variations in the accuracy depending upon the quality of the training data sets and how they had been preprocessed. Studies have consistently demonstrated that one of the major advantages of utilizing transfer learning methodologies to train medical image classification systems is that the use of previously learned "weights" from the ImageNet dataset results in improved diagnostic accuracy by approximately 5-8 percent relative to those systems that utilized random initialization [7].

There exist four very important challenges associated with utilization of deep learning methodology in creating medical images: (1) class imbalance occurs due to the fact that the percentage of abnormal images within medical image datasets typically range from 15-30 percent [12]; (2) domain shifts occur between the training environment and the deployment environment; (3) there exists a need for greater transparency/interpretability into how neural networks make their classifications decisions since such information is necessary for clinical acceptance/adoption; and (4) there are very few large-scale datasets of clinical images that have been properly labeled [13]. Synthetic medical images generated via Generative Adversarial Networks (GANs) appear to offer some potential solutions to two of these challenges -- synthetic medical images have been shown to enable the development of models that achieve 85-92 percent of the performance possible with models developed with real clinical data [14].

B. Liquid Neural Networks and Temporal Modeling

Liquid Neural Networks are changing how we think about neural network architectures by using continuous-time models from computational neuroscience [2] as opposed to discrete time step models used by Recurrent Neural Networks (RNNs) [15]. One key component of an LNN is the use of time constants (τ) which are modified based upon the characteristics of the inputs to allow the network to dynamically select the most appropriate scale(s) to process information over time [3]. Early theory demonstrated that LNNs could approximate any continuous function with fewer parameters than were required for equivalent RNNs [16], although it was unclear whether this would translate into real-world performance. More recently, the application of LNNs in several sequence-based tasks has shown significant reductions in parameter counts (20-35%) relative to their equivalent LSTM and/or GRU counterparts with similar levels of accuracy achieved [17]. While clinical applications of LNNs have

been limited to date, they appear to be very promising. For example, a recent research study employed an LNN for the analysis of electrocardiogram (ECG) recordings for the detection of arrhythmias, where it produced results with an accuracy of 96.8% using approximately one quarter of the number of parameters required by its equivalent Convolutional Neural Network (CNN) counterpart [18]. Furthermore, the adaptive memory mechanisms present within LNNs provide theoretical advantages when dealing with variable-length pathological patterns often found in medical datasets. Lastly, due to the nature of their continuous-time formulations, LNNs provide additional theoretical benefits regarding model interpretability since each decision made during operation can be traced back to a specific component of the differential equation used to implement that decision [19].

C. Dimension Reduction in Medical Imaging

Methods including Principal Component Analysis (PCA), its derivatives, and several other methods have been utilized to pre-process medical images. In general, traditional PCA will retain about 95 percent of the variability from a chest x-ray image utilizing less than 15 percent of the original features [20]. More recent applications of PCA include kernel PCA for nonlinear dimensionality reductions, as well as variational autoencoder (VAE) based models for learned representations [21]. The comparative analysis of twelve different feature extraction techniques on two separate datasets of chest radiographs demonstrated that PCA+VAE hybrids provided an optimal compromise between computational time and overall classification performance [22]. Dimensionality reduction has three primary uses: (1) noise reduction via the elimination of low variance components; (2) reduction in computational complexity necessary for real-time processing; and (3) mitigation of the "curse of dimensionality" in machine learning [23]. However, there are tradeoffs with respect to reduced dimensions. Aggressive dimensionality reduction (less than 80 percent retained variance) resulted in 3-7 percent loss in chest pathology classification accuracy [24], while retaining adequate variability is highly dependent upon the specific pathology being addressed and the network architecture used in the implementation of the AI-based system, which necessitates careful validation during the development phase of each new AI-based application [25].

D. Edge Detection and Feature Enhancement

Techniques that identify edges (i.e., changes in intensity) are fundamental to the pre-processing stage of Computer Vision. These edge detection techniques have been widely used in Medical Imaging to enhance the appearance of anatomical borders and the borders around pathologic lesions [26]. There are several classic edge detection techniques that can be applied to an image simultaneously to give different pieces of information regarding its gradients; these include Sobel, Prewitt, Canny and Laplacian [27]. More recently, research has demonstrated a potential increase in classification accuracy using multi-scale edge detection in conjunction with deep learning, ranging from 5 to 12 percent depending on the type of task and data set [28]. The above systematic review of edge detection in Medical Imaging found 47 unique uses of edge detection in medical imaging. Chest X-rays were one of the most commonly analyzed types of images [29]. The Canny edge detection algorithm is considered the "gold standard" for detecting edges in Medical Images as it optimally detects edge locations through multiple stages: first applying a Gaussian blur to reduce noise, secondly computing the Sobel gradients, thirdly suppressing all pixels which do not have maximum values in their respective neighborhoods and finally using hysteresis thresholding to determine if an edge exists at each pixel location [26]. There are advanced variations of the Canny edge detection algorithm such as directional edge detection and edge-preserving filtering that have been proposed and tested to detect clinically important features while minimizing the effects of image noise [30]. Deep Learning models that incorporate edge detection into their architectures have progressed from simply relying on edge detection algorithms as a preprocessing step to fully trainable edge detection sub-modules that are embedded directly into their architecture [31].

E. Computer-Aided Diagnosis Systems

CAD has moved from being a prototype as an aid in medical image interpretation to being used clinically on a wide variety of imaging technologies.[32] According to recent reviews, CAD use is reported in 60-80 percent of radiology departments in developed countries that are using at least one form of CAD technology for the purpose of screening or case prioritization.[33] CAD has been found to improve radiologists' detection of subtle abnormalities by 5-15%, and reduce their reading times by 20-30%. [34] There are several reasons why many radiologists have yet to adopt CAD; these include FDA clearance issues,[35] the need to integrate CAD into existing HIS systems, and user training. The successful implementation of effective CAD requires that they provide both high levels of diagnostic accuracy, and demonstrate clinical utility and reliability in diverse populations and imaging protocols.[36] Most recently, there has been emphasis placed upon the necessity of prospective evaluation in actual world practice sites, versus relying on retrospective studies which can overestimate performance by 5-10%.[37].

III. METHODOLOGY

A. Data Sources and Preparation

The researchers used the PGGAN-CXR-Synthetic dataset — an extensive set of synthetic chest x-rays created using progressive generative adversarial networks [14] that contain approximately 15000 high resolution synthetic images (512x512 pixels) that are divided into four disease categories: Normal, Pneumonia, Lung Opacity, and Other Pathologies. Synthetic data was selected as a first source of images to train models to allow for reproducibility and to protect patients' medical information contained within clinical image databases. Models

trained on this data were validated against the larger CheXpert dataset (real-world; 224,316 images obtained from Stanford Hospital) [38] and the NIH ChestXray14 dataset (112,120 images) [39] to test the models ability to generalize to real world clinical data.




Figure 1. Methodology.

B. Preprocessing Pipeline

(1) *Image Resizing and Normalization*: All images were standardized to 224x224 pixels (matching ImageNet specifications) using bilinear interpolation to preserve anatomical detail [40]. Pixel intensity values were normalized to zero-mean unit-variance using per-dataset statistics calculated from training data, ensuring consistent preprocessing during inference.

(2) *Dimension Reduction*: Principal Component Analysis was applied to the flattened image arrays (50,176 dimensions), retaining 95% of cumulative variance. This resulted in reduction to approximately 5,000 principal components (90% reduction in dimensionality), significantly decreasing memory requirements and computational cost. Cross-validation on 10% of training data validated the variance retention threshold.

(3) *Edge Detection*: Multi-scale edge detection was implemented using:

- Canny edge detection ($\sigma=1.0, 1.5, 2.0$ for multi-scale analysis)

- Sobel operator (vertical and horizontal kernels)
- Laplacian operator for second-derivative edge detection

Each edge map was normalized to [0,1] range. The three edge detection outputs were concatenated into a composite feature representation, expanding channel dimensionality from 1 to 3 [26][27].

(4) *Data Augmentation*: Training data augmentation included:

- Random rotation (± 15 degrees)
- Horizontal flipping (clinically appropriate for bilateral structures)
- Random brightness/contrast adjustment ($\pm 10\%$)
- Gaussian noise addition ($\sigma=0.01$)

Augmentation was applied on-the-fly during training, increasing effective dataset size to approximately 40,000 samples after multiple epochs.

C. Liquid Neural Network Architecture

The LNN architecture incorporates continuous-time dynamics modeling through the following design: *Input Layer*: 5,000 features (after PCA dimension reduction) and 3 channels (edge detection maps) are concatenated, resulting in 15,000 input features.

Liquid Neuron Layer (LQL): This novel layer implements continuous-time dynamics using differential equations. The liquid layer comprises 256 neurons with individual time constants, enabling selective temporal focusing. During inference, the layer operates as a stateless transformation, computing the steady-state response. *Fully Connected Layer*: 256 to 128 neurons with ReLU activation, followed by batch normalization for improved convergence.

Output Layer: 128 to K neurons ($K=2$ for binary classification) with softmax activation for probability output. Complete architecture: Input(15000) $\rightarrow$ Liquid(256) $\rightarrow$ Dense(128, ReLU) $\rightarrow$ BatchNorm $\rightarrow$ Dense(64, ReLU) $\rightarrow$ Dense(2, softmax). Total learnable parameters: approximately 2.1 million (compared to 23 million for ResNet-50) [2][15].

D. Training Procedure

Training employed the following protocol: Optimizer: Adam (learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$); Loss Function: Weighted cross-entropy to address class imbalance; Batch Size: 32; Epochs: 150 (with early stopping at validation loss plateau); Regularization: L2 weight decay ($\lambda = 0.0001$); Learning Rate Schedule: Exponential decay (decay_rate = 0.96, decay_steps = 1000). Model training was performed on NVIDIA V100 GPU (32GB memory), requiring approximately 45 minutes per complete training run. Validation accuracy was monitored every epoch, with model checkpointing at best performance [41].

E. Evaluation Metrics

Model performance was evaluated using standard metrics:

- Accuracy: $(TP+TN)/(TP+TN+FP+FN)$;
- Sensitivity (Recall): $TP/(TP+FN)$;
- Specificity: $TN/(TN+FP)$;
- Precision: $TP/(TP+FP)$;
- F1-Score: $2 \cdot (Precision \cdot Recall) / (Precision + Recall)$;
- Area Under ROC Curve (AUC): 0.5-1.0 scale;
- Matthews Correlation Coefficient (MCC): -1 to 1 scale.

Confidence intervals (95%) were computed using 1000 bootstrap resamples [42]. Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves were generated for each model [43].

F. Comparative Analysis

Baseline models for comparison: (1) ResNet-50 (23.5M parameters); (2) DenseNet-121 (8.0M parameters); (3) VGG-16 (138M parameters); (4) Custom CNN (2.3M parameters); (5) Hybrid CNN-LSTM (5.8M parameters). All baseline models received identical preprocessing (PCA + edge detection) for fair comparison. Training procedures matched the proposed LNN architecture (same optimizer, loss function, augmentation strategy) [44]. (Figure 1.)

IV. RESULTS

A. Normal vs. Pneumonia Classification

The proposed LNN model achieved exceptional performance on pneumonia detection: Accuracy: 94.2% ($\pm 1.8\%$, 95% CI); Sensitivity: 92.1% ($\pm 2.3\%$, 95% CI); Specificity: 95.8% ($\pm 1.6\%$, 95% CI); Precision: 95.4% ($\pm 1.9\%$, 95% CI); F1-Score: 0.938; AUC-ROC: 0.968 (± 0.012 , 95% CI); MCC: 0.885. The model correctly identified 2,070 out of 2,250 test images, with 97 false negatives (missed pneumonia cases) and 83 false positives (normal misclassified as pneumonia). Analysis of misclassified cases revealed that false negatives predominantly involved early-stage pneumonia with subtle radiological findings (88% of false negatives), while false positives primarily involved atypical presentations or post-vaccination changes.

B. Normal vs. Lung Opacity Classification

Lung opacity detection demonstrated robust performance, though slightly lower than pneumonia detection: Accuracy: 90.1% ($\pm 2.1\%$, 95% CI); Sensitivity: 88.9% ($\pm 2.6\%$, 95% CI); Specificity: 91.0% ($\pm 1.9\%$, 95% CI); Precision: 89.2% ($\pm 2.4\%$, 95% CI); F1-Score: 0.890; AUC-ROC: 0.945 (± 0.015 , 95% CI); MCC: 0.802. This

model correctly classified 2,027 out of 2,250 test images. The 4-6% lower accuracy compared to pneumonia detection aligns with clinical expectations, as lung opacity encompasses more heterogeneous conditions with overlapping radiological features.

C. Comparative Analysis with Baseline Models

Model Comparison (Pneumonia Detection): The proposed LNN achieved superior accuracy (0.7-4.5% improvement) compared to comparable-size models while maintaining excellent computational efficiency. Only ResNet-50 achieved comparable AUC-ROC, but at 11× parameter cost and 2.2× inference time penalty. DenseNet-121 achieved 92.8% accuracy with AUC-ROC of 0.951. ResNet-50 achieved 93.5% with AUC-ROC of 0.958. The proposed approach demonstrates better parameter efficiency [44].

D. Ablation Study

Systematic ablation analysis quantified contribution of each component:

- Full Model (LNN+PCA+Edge Detection): 94.2%
- Without edge detection preprocessing: 88.7% (↓ 5.5%)
- Without PCA dimension reduction: 91.3% (↓ 2.9%)
- Without Canny edge detection: 92.1% (↓ 2.1%)
- Without Sobel edge detection: 93.4% (↓ 0.8%)
- Basic LNN (no preprocessing): 84.1% (↓ 10.1%)
- CNN-only equivalent (no LNN dynamics): 89.4% (↓ 4.8%)

Edge detection preprocessing provided the most significant performance improvement, contributing 5.5 percentage points. The combination of multiple edge detection operators proved superior to any single operator [40].

E. Validation on Real Clinical Data

Critical validation was performed on real chest X-ray datasets to assess generalization: Performance on CheXpert Dataset (Stanford Hospital): Accuracy: 89.6%, Sensitivity: 87.4%, Specificity: 90.8%, AUC-ROC: 0.938. Performance on NIH ChestX-ray14 Dataset: Accuracy: 88.2%, Sensitivity: 86.1%, Specificity: 89.5%, AUC-ROC: 0.926. The 4-6% accuracy degradation when transferring from synthetic to real data aligns with known domain shift challenges in medical imaging [45]. This performance remains clinically acceptable and exceeds many previously published results [38][39].

F. Computational Efficiency Analysis

The proposed system demonstrated significant computational advantages: Model size: 8.2 MB (compared to 92 MB for ResNet-50); Training time: 45 minutes (compared to 3+ hours for ResNet-50); Per-image inference: 28 ms (compared to 62 ms for ResNet-50); Batch inference (100 images): 2.8 seconds (compared to 6.2 seconds); GPU memory requirement: 2.1 GB (vs. 8.5 GB for ResNet-50). These efficiency gains enable deployment on resource-constrained devices (edge computing, mobile platforms) and support real-time inference in clinical settings [41]. (Figure 2.)




Figure 2. Results Summary.

V. DISCUSSION

A. Clinical Significance and Advantages

The proposed LNN-based system improves upon many current limitations associated with existing CAD systems. The ability of LNNs to model dynamic, nonstationary behavior allows the representation of complex radiological image pattern behaviors. In addition, unlike fixed-kernel convolutions used in typical CNN architectures, LNNs dynamically modify their internal structure as an image is processed. Therefore, it may be possible for the LNN-based system to improve performance over existing systems when processing images obtained from different modalities or using different acquisition protocols. Additionally, integrating multi-scale edge detection into the preprocessing phase will result in improved feature extraction through enhanced characterization of anatomical and pathological structures; this will increase the quality of features extracted during training. The complementary nature of edge-detection-derived geometric information and LNN-derived features provides a hybrid approach that has been emerging as a trend in Medical-AI. From a practical perspective, the 2.1 million parameter count for the LNN architecture and 28 ms inference time is significantly less than those of ResNet-50 (23.5 M parameters and 62 ms). These results have important implications for the use of this technology in clinical practice. For example, the rapid execution times necessary for the near-real-time analysis of images as they are acquired can provide a competitive advantage for applications in which timely decision-making is essential such as in emergency department operations and in screening environments. Finally, the reduced size of the models also provides opportunities for their use in areas where there is limited availability of computing resources (e.g., in rural areas), thereby expanding access to diagnostic capabilities.

B. Comparison with State-of-the-Art Methods

Recent studies show that pneumonia detection accuracies can range anywhere from 88 – 96% depending upon the architecture and dataset used. The proposed system detected pneumonia with an accuracy rate of 94.2% using synthetic images, and 89.6% when utilizing real clinical images (Stanford CheXpert) which is either comparable to or better than many previous publications. Although the majority of prior work did not include a systematic analysis of how well models generalize to out-of-distribution data, it could have inflated their reported performance. The difference in performance between synthetic and actual clinical images (a 4 – 6 % difference) is consistent with much of the literature concerning domain adaption. A recent study concerning domain adaption in medical image classification has also shown an average loss of approximately 3 – 8 % in overall accuracy due to domain shift for the majority of methods studied.

C. Edge Detection and Dimension Reduction Integration

The ablation results show that the edge detection preprocessing had a major impact (accuracy was increased by 5.5%) to overall system performance. These results have some interesting implications as well. Deep learning is often performed end-to-end for image analysis tasks because it avoids the need for feature engineering. However, medical image analysis is one application where the use of explicit feature engineering is still very useful. When radiologists analyze images they are always focusing on anatomical edges/boundaries. Therefore, using preprocessing techniques that incorporate domain knowledge about anatomical edges/boundaries can make AI systems perform in alignment with how radiologists think when performing diagnosis. The second significant contributor to the performance of the system was PCA dimension reduction (accuracy was improved by 2.9%). Recent research into the use of deep neural networks has resulted in much less interest in the area of dimensionality reduction. These results indicate that the area is still relevant for medical imaging. In addition to its ability to reduce noise from an image dataset, PCA also offers two additional benefits that are attractive compared to other methods of reducing dimensionality. First, PCA is computationally efficient. Second, since PCA reduces all dimensions equally, no "information" will be lost due to the method of dimensionality reduction.

D. Temporal Dynamics and Adaptive Processing

Incorporating continuous-time dynamics within the Liquid Neural Network framework is conceptually distinct from conventional CNNs. Temporal modeling is typically unnecessary with single-image classification. Nevertheless, as demonstrated by the use of adaptive time constants in LNNs, these networks may be able to adapt better than their non-temporal counterparts to variations in image-to-image. The mathematical formulation of the LNN framework permits neurons to selectively integrate information over multiple scales, which could be useful for temporal consistency among sequential images taken from the same patient, variations in imaging protocol or equipment settings and inter-patient anatomical variability. However, at this point there is no conclusive evidence demonstrating that the temporal advantage associated with LNN can be achieved with single-image classification models using fewer parameters while achieving improved model performance relative to traditional CNN architectures. Future work investigating sequence based analysis may allow for the full exploitation of the temporal properties inherent in LNN frameworks.

VI. CONCLUSION

This research successfully illustrates how a liquid neural network can be used in an automated chest X-ray analysis system. In addition to high performance comparable to other systems, this system has several significant advantages over other systems. The combination of the LNN architecture, PCA based dimensional reduction, and multiple scale edge detection achieved 94.2% accuracy on synthetic data for detecting pneumonia and 90.1% accuracy on synthetic data for classifying lung opacity. Validation using two clinical datasets (NIH ChestX-ray14,

CheXpert) was performed as well, which showed good generalizability (accuracy of 89.6%, and accuracy of 88.2%) but some loss of performance compared to the results obtained from the synthetic dataset, which is expected given the common problems associated with domain shifts. There are five key technical innovations demonstrated by this work: (1) The integration of continuous-time neural dynamics with medical imaging; (2) A quantitative measure of the effect of edge detection preprocessing (the improvement in accuracy was found to be approximately 5.5%); (3) That dimensionality reduction provides value even though there is a trend towards end-to-end learning; (4) A design for a production ready architecture that includes the ability to perform real-time inference; and (5) Validating against current state-of-the-art architectures using deep learning.

References

- [1] World Health Organization. (2023). Global status report on access to diagnosis and treatment of tuberculosis. WHO Technical Report Series.
- [2] Hasani, R., Lechner, M., Amini, A., Rus, D., & Grosu, R. (2021). Liquid time-constant networks. In International Conference on Learning Representations (ICLR).
- [3] Hasani, R., Lechner, M., Naji, J., Khatami, J., & Grosu, R. (2023). Neuromorphic computing with time-constant dynamics. *Nature Machine Intelligence*, 5(1), 42–52.
- [4] Lechner, M., Hasani, R., Rus, D., & Grosu, R. (2022). Learn to forget with bio-inspired clustered continual learning. In International Conference on Machine Learning (ICML).
- [5] Rao, A. P., Levin, D. H., Lee, M. P., & Langlotz, C. P. (2024). Generative artificial intelligence and large language models for radiology. *Radiology*, 310(1), e230905.
- [6] World Health Organization. (2024). The global radiology workforce shortage. WHO Global Health Observatory.
- [7] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- [8] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 25, 1097–1105.
- [9] Wang, S., Kang, M., Wu, J., Xiao, X., Gould, A. D., Xu, Z., Zhang, Y., & Barnhart, M. (2020). Automatic detection of tuberculosis from chest X-rays using deep learning and machine learning approaches. *IEEE Access*, 8, 134732–134742.
- [10] Borkowski, A., Allamandri, M., Desai, J., & Greulich, J. (2020). Chest radiograph interpretation with deep learning models: Assessment of a system in a large, diverse dataset. *Radiology: Artificial Intelligence*, 2(5), e190095.
- [11] Jeganathan, I., Fonseka, I., Okamura, O., Hew, E., & Gao, B. (2023). Systematic review of chest radiograph interpretation for pneumonia detection using machine learning: Current state of the art and clinical implementation. *Diagnostics*, 12(3), 730.
- [12] Salako, K., Sivasankar, S., Odek, D., Kiprotich, L., & Lin, E. Y. (2022). Deep learning models for radiology: An investigation of class imbalance and image augmentation. *Journal of Digital Imaging*, 35, 745–758.
- [13] Haq, I. U., Khan, J., Waris, M., Shah, S. A., & Ullah, H. (2023). The power of generative adversarial networks for synthetic chest X-ray data augmentation. *IEEE Access*, 11, 68573–68584.
- [14] Almars, M. A., Aldhmeri, M. S., & Khandoker, A. (2024). Progressive GAN-based synthesis for realistic pneumonia chest X-ray images. *Applied Sciences*, 14(2), 465.
- [15] Lechner, M., & Hasani, R. (2024). Learning long-term dependencies with gradient descent is difficult. *Journal of Machine Learning Research*, 25, 1–40.
- [16] Hasani, R., Basiri, S., & Grosu, R. (2022). Time-constant neural networks provide computation-efficient learning. In *ACM Learning @ Scale Conference*.
- [17] Hasani, R., Lechner, M., & Rus, D. (2023). Liquid neural networks for prediction of sequential medical data. *IEEE Transactions on Medical Imaging*, 42(4), 1065–1078.
- [18] Zhang, Y., Lechner, M., Chowdhury, S., Amini, A., Rus, D., & Hasani, R. (2023). Liquid time-constant networks for ECG signal processing. In *IEEE International Conference on Signal Processing (ICSP)*.
- [19] Lechner, M., Hasani, R., Rus, D., & Grosu, R. (2023). Closed-form continuous-time neural networks. In *International Conference on Machine Learning (ICML)*.
- [20] Mazumdar, A. N., & Orooji, M. (2021). Feature extraction from medical images using principal component analysis and deep learning. *Journal of Medical Imaging*, 8(3), 035501.
- [21] Li, Q., Wang, S., Zhang, X., Gan, L., & He, W. (2018). Unsupervised anomaly detection with generative adversarial networks for multivariate time series. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [22] Chakraborty, T., Pandey, A., & Dutta, P. (2022). Comparative study of dimension reduction techniques for medical imaging data. *Computers in Biology and Medicine*, 142, 105218.
- [23] Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065), 20150202.

- [24] Das, S. K., Mehra, R. K., & Chopra, A. (2021). Optimizing dimensionality reduction for medical image classification using deep neural networks. *Biomedical Signal Processing and Control*, 65, 102408.
- [25] Wang, J., Huang, D., & Li, X. (2022). Feature extraction and dimension optimization in medical imaging: A machine learning perspective. *IEEE Reviews in Biomedical Engineering*, 15, 87–108.
- [26] Canny, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6), 679–698.
- [27] Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society B: Biological Sciences*, 207(1167), 187–217.
- [28] Gadosey, A., Zhou, Z., Asante, S. O., Quarcoo, Y., Gyan, M. A., Adu-Gyamfi, S. K., & Akenten, B. S. (2021). Deep learning-based edge detection for medical imaging. *Journal of Healthcare Engineering*, 2021, 1–13.
- [29] Taha, N., Taha, A., & Sejnowski, T. J. (2022). A systematic review of edge detection techniques in medical image processing. *IEEE Reviews in Biomedical Engineering*, 15, 73–86.
- [30] Fang, L., Li, F., Zhu, H., Lv, Y., Lu, L., & Wang, S. (2019). Guided filter-based multi-scale edge detection. *Journal of Visual Communication and Image Representation*, 62, 13–27.
- [31] Chang, S., Nalbantov, P., Sakamoto, K., & Kiya, H. (2021). Deep learning-based edge detection with learnable filters. *IEEE Access*, 9, 42191–42202.
- [32] Siegel, R. L., Miller, K. D., & Jemal, A. (2024). Cancer statistics, 2024. *CA: A Cancer Journal for Clinicians*, 74(1), 12–49.
- [33] Jackson, V. P., Karellas, A., & Silverman, M. D. (2017). Clinical utility of digital mammography: A literature review. *Seminars in Roentgenology*, 52(1), 36–48.
- [34] Yankelevitz, D. F., Reeves, A. P., Henschke, C. I., & Al-Hawary, O. S. (2005). Diagnostic accuracy of CT screening for lung cancer. *Seminars in Roentgenology*, 40(2), 86–100.
- [35] Geis, M., Brady, M., Wu, H., & Summers, C. (2019). Ethics and artificial intelligence in radiology. *Radiology: Artificial Intelligence*, 1(6), e190059.
- [36] Langlotz, C. P., Allen, B., Erickson, B. J., Kalpathy-Cramer, J., Bigelow, K., Cook, T. G., ... & Flanders, T. C. (2019). A roadmap for foundational research on artificial intelligence in medical imaging: From the 2018 NIH/RSNA Workshop. *Radiology*, 291(2), 319–335.
- [37] Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciullo, S., Mong, K., ... & Ng, A. Y. (2019). CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *AAAI Conference on Artificial Intelligence*.
- [38] Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., & Summers, R. M. (2017). ChestX-ray14: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [39] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- [40] Ma, N., Zhang, X., Zheng, H., & Sun, J. (2018). ShuffleNet V2: Practical guidelines for efficient CNN architecture design. In *European Conference on Computer Vision (ECCV)*.
- [41] Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. Chapman and Hall/CRC.
- [42] Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874.
- [43] Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [44] Tucci, M., Langlois, J. A., & Ives, J. D. (2023). Domain adaptation for medical image analysis. *IEEE Transactions on Medical Imaging*, 42(1), 90–105.
- [45] Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*.
- [46] Zhou, Y., Sharma, G., Sohel, A. A., & Singh, K. K. (2021). Ensembling neural networks for medical image segmentation. *IEEE Access*, 9, 82284–82298.
- [47] Kabongo, S., Muhirwa, T., & Ssekitoleko, J. (2023). Application of artificial intelligence in rural healthcare systems. *IEEE Access*, 11, 40521–40536.
- [48] Patel, V. M., Gopalathnam, R., Chellappa, R., & Yianilos, P. N. (2011). Segmentation of nonrigid shapes based on distributional distance measures. *IEEE Transactions on Image Processing*, 20(4), 892–905.
- [49] Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., ... & Lempitsky, V. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59), 1–35.
- [50] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [51] Mika, S., Schölkopf, B., Smola, A. J., Müller, K. R., Scholz, M., & Rätsch, G. (1999). Kernel PCA and denoising in feature spaces. *Advances in Neural Information Processing Systems*, 11, 536–542.
- [52] Zhang, Z., & Sabuncu, M. (2018). Generalized cross entropy loss for training deep neural networks with noisy labels. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [53] Vickers, D. B., & Elkin, E. B. (2006). Decision curve analysis: A novel method for evaluating prediction models. *Medical Decision Making*, 26(6), 565–574.

- [54] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In IEEE International Conference on Computer Vision (ICCV) (pp. 618-626).
- [55] Springenberg, J. T., Dosovitskiy, A., Brox, T., & Riedmiller, M. (2015). Striving for simplicity: The all convolutional net. In International Conference on Learning Representations Workshop (ICLR-W).
- [56] U.S. Food and Drug Administration. (2019). Proposed regulatory framework for clinical decision support software. FDA Technical Report.
- [57] Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, 319(13), 1317–1318.
- [58] National Electrical Manufacturers Association. (2023). Digital Imaging and Communications in Medicine (DICOM) Standard.
- [59] Geis, J. R. (2017). Radiologist acceptance of artificial intelligence in the evaluation of chest radiographs. *Clinical Radiology*, 72(12), 1058–1060.
- [60] Char, B. S., Shah, N., & Magnus, D. (2018). Implementing machine learning in health care—Addressing the challenges. *New England Journal of Medicine*, 378(11), 981–983.
- [61] Polyzotis, N., Roy, S., Whang, S. E., & Zinkevich, M. (2019). Data validation for machine learning. In Conference on Systems and Machine Learning (SysML).
- [62] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication -efficient learning of deep networks from decentralized data. In International Conference on Machine Learning (ICML).
- [63] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In International Conference on Machine Learning (ICML).