

Building scalable ML pipelines for radiology image analysis in cloud environments

Veerendra Nath Jasthi

veerendranathjasthi@gmail.com

Abstract:

Radiology image analysis is one of the essential fields of machine learning (ML) application because of the dynamic growth of medical imaging data, and the need to observe fast and accurate results of diagnostics. Yet, the development of scalable ML pipelines that would be able to process and analyze mass radiological data is associated with significant challenges. These are data heterogeneity, computational-intensity, regulatory-compliance and scalability of deployment. In this paper, an end-to-end methodology is introduced to design and develop scalable ML pipeline to analyze radiology images with the help of cloud architecture. We take advantage of cloud-native services as data ingestion, preprocessing, training, deployment, and continuous monitoring. As our experiments carried out on chest X-ray and MRI datasets show, our technology is faster, more accurate, and more scalable than the conventional on-premise systems. The paper emphasizes the need of modular pipeline design, serverless architectures, container orchestration when it comes to real-time troubleshooting and model lifecycle management. It can be strongly stated that the given manner of approach is extremely flexible in terms of different types of radiology and ML building blocks, so it is a viable solution to contemporary healthcare AI systems.

Keywords: Machine Learning Pipelines, Radiology Image Analysis, Cloud Computing, Medical Imaging, Scalable Infrastructure, Healthcare AI, Deep Learning, Kubernetes, DICOM, Data Preprocessing.

I. INTRODUCTION

The exponential development involving digital medical imaging technology in the last ten years has presented the amazing possibility and unprecedented challenges to all healthcare delivery in every part of the globe. Radiology especially, has currently produced huge volumes of image data every single day, in imaging modalities of X-rays, CT scans, MRIs and PET scans [1]. Although all these images are necessary to the clinical decision making process, the massive number and complexity of data become the bottleneck to the radiologists particularly in the resource-limited environment. The increasing demand of automatization, high-speed and accuracy in research has given rise to an emerging interest in implementing the techniques of machine learning (ML) to support radiological image analysis/ classification and interpretation.

Deep learning in particular and machine learning in general have improved greatly when it comes to pattern recognition and image analysis. Transformer-based models, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) have proven to be very accurate in recognition of tumor, segmentation of organs and classification of diseases. It has been demonstrated by many studies that these models have been able to perform better compared to or at par with human radiologists in certain diagnostic cases [2]. Nevertheless, however encouraging employment outcomes are, combining these kinds of models in practice is hard. The bulk of present ML research in radiology concerns itself with model-building and benchmarking, and underestimates the difficulty of implementation at scale, particularly in the potentially diverse data, regulatory, and latency needs of a hospital setting.

It is a vital point that should not be missed, which is the infrastructure required to get the ML models working. Creating a well-developed machine learning pipeline does not only imply the training of an algorithm or the related job, but also incorporates data capture, pre-processing, and validation, training of the model, as well

as tuning hyper-parameters, deployment, monitoring, and re-training. Every stage requires reliability, scalability as well as security. These requirements are commonly poorly suited on traditional on-premise hardware infrastructures as they have smaller compute capabilities, limited ability to scale and have large maintenance overhead [14].

The solution to this issue is provided by cloud computing which is a revolutionary solution. Large cloud providers such as Amazon Web Services (AWS), Google Cloud Platform (GCP) and Microsoft Azure have recently implemented special services to support machine learning, image storage and orchestration of containers. The services allow healthcare institutions to develop end-to-end ML pipelines that are modular, scalable, and healthcare regulation compliant like HIPAA and GDPR. Nowadays, thanks to serverless computing, containerization using Docker and Kubernetes, and ML hosted services such as SageMaker, or Vertex AI, developers can create advanced ML systems with little to no infrastructure overhead.

The examples of this with radiology, in particular, are a cloud-based ML pipeline, which allows sustained learning and adaptation, and connects the numerous healthcare facilities. They can consume data pushed to these pipelines via PACS systems accepted by hospitals, standardize and de-identify the stream, execute distributed training, and also make predictions in real-time available as REST APIs. In addition, the model changes and A/B testing are possible to implement without halting clinical services. They are also easily integrable in cloud systems that are modular and flexible, clinician dashboards and mobile diagnostic tools [12].

As much as these have proved advantageous, there are challenging factors to build scalable ML pipelines in radiology. The data in medical imaging is heterogeneous format (e.g. DICOM, NIFTI), and preprocessing is frequently both high-computational and domain-specific. Privacy of data and regulating access to such data also needs to be ensured especially in case of sensitive patient data across jurisdictions [10]. Latency is another key issue particularly in cases that need real-time or near-real-time diagnosis, in critical care. These issues create the need to have a clearly shaped architecture that has a balance between performance cost, compliance, and simplicity.

The present paper discusses the issue of an extensible, scalable, cloud-native ML pipeline to analyze radiology images. It provides an end-to-end architecture that starts by secure data ingestion and preprocessing, followed by model training and evaluation, and concludes with its deployment and further monitoring. The offered solution focuses on being flexible, scalable, and complaint, yet it also includes specific implementation strategies with the use of available cloud tools. Performance is measured using speed, cost and predictive accuracy after intensive experimentation of the pipeline on several chest X-ray and MRI datasets [11].

Finally, this research wants to bridge a gap between ML theory and actual implementation of this practice in radiology. The study shows that it is feasible not only to create scalable and efficient ML systems but also to introduce them into routine clinical practice where they can be of the most use with proper cloud-native design.

Novelty and Contribution

The proposed work presents an innovative architecture of cloud-native ML pipeline that is developed with focus on radiology image analysis. Although prior research in this field has paid much attention to the development of the algorithm, much less emphasis has been given to imposing scalable architecture in medical setting that can have its deployments and retraining in real-time. This work addresses this gap by harmonising cloud-native design concepts, modular pipeline steps and newer orchestration libraries to construct a very flexible system.

The key innovativeness is associated with the fact that cloud services are used in every part of the ML process, i.e. they facilitate a secure DICOM ingestion task, which can be deployed as serverless, and real-time monitored. The system is scalable, allows use of multiple imaging modalities, is both compliance regulated, and automatic in regards to how large it is expanded whenever the workload requires it. In comparison with the traditional approaches to ML deployment, the proposed pipeline implies containerized microservices and managed ML surroundings, which makes it possible to upgrade and version and recover failures in a seamless way [3].

The other important contribution is the concern with real-world clinical applicability. These add features de-identification in the preprocessing phase, latency optimization in inference, and PACS and EHR system integration with hospitals. The study includes an assessment of the pipeline on actual medical datasets (e.g.,

ChestX-ray14, BraTS) which make the results look more realistic and indicates the scalability, performance and robustness of the model.

This paper will make the following contributions, mainly:

- Complete modular and scalable ML pipeline to be used in the analysis of radiology images on clouds.
- Cloud-native tool integration (Kubernetes, SageMaker, Vertex AI, Airflow) to manage end to end facilitation of pipelines.
- Empirical evidence has shown high gains in accuracy, latency and throughput using large-scale radiology datasets.
- Accountability of abidance with the medical data privacy laws and process of safe deployment of a model.
- A feasible roadmap that healthcare institutions can follow in the implementation of AI-driven diagnostics without involving its vast infrastructure.

This paper presents an optimistic perspective of how the radiology departments and the AI researchers can join hands to implement the ML models to scale, eventually creating a faster pace of adopting intelligent diagnostic tools within the mainstream healthcare systems [16].

II. RELATED WORKS

In 2022 Javaid *et al.*, [13] suggested the machine learning and medical image analysis are two fields that have come together such that they have become a primary focus of research especially in the field of radiology. In their research, numerous studies have been conducted across the years that have illuminated the possible uses of machine learning models and especially deep learning in the field of constructing interpretation of the radiological imaging with a great degree of precision. Studies have revealed that these models are able to diagnose chest diseases like pneumonia, tuberculosis, brain tumors and breast cancer by using the radiographic information at the same rate or even higher than the diagnostic capacity of a human at a control setting. Most initial research was devoted to modelling deep convolutional neural networks (CNNs) on the image classification and segmentation tasks, based on constructed datasets and modest computer conditions. Most of the available studies surround the use of CNN architectures in the classification of radiology images. Such models are trained against labeled image datasets to identify anomalies like lesions, masses among other pathologies. The initial models ran mainly in local workstations or on one-node servers, which did not allow scaling and they were not very clinically usable. With the growth of the complexity of models and volume of data used, scientists started expanding toward more efficient training paradigms, including transfer learning and model ensembling that raise accuracy without requiring a corresponding exponential growth of computation time. But these also continued to be hampered by their reliance on locally or institution-based hardware [5].

Parallely, the research on radiology image segmentation brought such models as U-Net that allowed detecting anatomical structures and pathological areas pixel-wise. Such models of segmentation were vital in activities such as localization of tumors and detection of organ boundaries. Although they were successful, these methods often needed massive preprocessing, e.g. intensity normalization, anatomical alignment, often by classical image processing methods. The pipelines constructed in their periphery were not modular and automated and thus they could not be implemented in large scale.

In 2020 S. Pati *et al.*, [4] introduced the other line of research was aimed at combining the analysis of radiological images with clinical data on the basis of multimodal learning schemes. The objectives of these systems were to enhance the diagnostic decision by including the electronic health record (EHR) data, patient history and laboratory values with imaging attributes. Though this had a minor positive implication on the predictive performance, this method increased the data pipeline complexity thereby making it harder to generalize and scale. Also, the majority of such models were not designed to actually run in real-time and did not focus on cloud-based deployment issues.

Most recent studies started to investigate the opportunities of cloud computing in terms of medical AI. These publications took into account the shortcomings of the classic IT infrastructure and promoted the implementation of cloud-based services in training and deployment of models. First frameworks generally used general-purpose, cloud resources as compute-intensive resources, and these, although flexible in use, lacked healthcare-specific integration. The initial cloud-based platforms had a number of weaknesses related to compliance, data protection and latency, particularly when there was a need to perform analysis on images

in near real-time. Moreover, these systems hardly gave end-to-end services that could be accepted by the hospital with little customization.

Still more advanced studies came up with the idea of the cloud-native architectures which are specifically tailored to the healthcare applications. These frameworks were presenting the concepts of containerization, orchestration and microservices to attain flexibility and reliability. Kubernetes became a trendy orchestration system, which provides dynamic resources provision, load balancing, and fault tolerance. Together with CI/CD tools, such platforms enabled researchers and practitioners to rapidly iterate model versions and deploy any updates with reduced downtimes. But these methods were mostly technical and sometimes not related with the realities of radiology departments.

A few works in the field of data management discussed automated pipelines of ingesting, transforming, and de-identifying images. Scripting and home-made software solutions were used to overcome the problem of DICOM files processing, metadata extraction and anonymization. Such pipelines provided patient privacy and regulatory compliance but tended to be institution specific, and thus lacked flexibility. In addition, most of these solutions could only be configured manually, and in an emergency diagnostic process, real-time processing is essential.

In 2018 Zhang *et al.*, [15] proposed the use of managed machine learning platforms provided by cloud providers to ease up the burden of infrastructure was also studied. Such services as automated model training, hyperparameter optimization and Model monitoring have been added to make the ML workflow more efficient operationally. These managed platforms simplified most of the resource provisioning and deployment of models. Nevertheless, the majority of these technological innovations remained generic, never optimized to address the peculiarities of, say, radiology image analysis, e.g., high-resolution 3D scans, PACS integration, and version-controlled model rolling out.

Some more recent developments presented the idea of the federated learning in medical imaging where models are trained collaboratively across the institutions without exchanging sensitive patient data. Although this approach provides data privacy and allows cross-institutional generalizability of the model, it becomes complex in regard to synchronization, model aggregation, and edge-compute collaboration. Federated learning is a good promise but its application to scale up is limited as there are technical difficulties and there is a lack of standardised protocols in play.

Moreover, a few domain-specific models of medical imaging were presented with an aim of easing down the process of development. One of them contains toolkits where the predefined functions can be ready to use in image loading, preprocessing, augmentation, as well as visualization. Although they have shortened up prototyping, they are typically not feature-rich and are not integrated well with cloud-native deployment platforms. They find limited application in academic research and minimal time to size clinical trials.

Despite much advancement in machine learning and more so in the radiology image analysis, there has remained a dearth of end-to-end, production ready solutions. The majority of the research devoted to the topic of ML application did not consider the nuance of Model development and deployment as a part of the same process of applying ML in clinical radiology. These systems have been thus focused on scalability, modularity and maintainability which has been secondary considerations which have left a situation of fragmented adoption and inefficiency in operation of the systems.

Finally, in short, the literature tends to cover in detail the approach of ML models in radiological diagnostics, but in a lot of instances does not provide deployable scalable solutions applicable in the contemporary clinical domain. Previous studies are too biased on the algorithm innovation part or not practical as far as clinical implementation is concerned. Research filling this gap is urgently needed: offering excellent performance in machine learning with scalable cloud computing, healthcare compliance, and the ability to operate on inference in real-time as well as integrating smoothly with current radiology workflow [6].

III. PROPOSED METHODOLOGY

The proposed pipeline for scalable ML-based radiology image analysis is modular, cloud-native, and optimized for high-performance computation. It integrates data ingestion, preprocessing, model training, and deployment in a highly parallel and automated system.

We begin by defining the input space. Each radiology image I in DICOM format is represented as a highdimensional tensor:

$$I \in \mathbb{R}^{h \times w \times c}$$

where h , w , and c represent height, width, and channel dimensions respectively.

Preprocessing involves pixel normalization and resizing:

$$I' = \frac{I - \mu}{\sigma}$$

where μ is the mean and σ is the standard deviation computed from the dataset. This step standardizes the input for consistent learning.

For segmentation or feature extraction, each image undergoes region proposal using intensity-based thresholds. The thresholded binary mask M is defined as:

$$M(x, y) = \begin{cases} 1 & \text{if } I'(x, y) > T \\ 0 & \text{otherwise} \end{cases}$$

where T is the empirically chosen threshold.

During training, the objective is to minimize the binary cross-entropy loss function for classification tasks:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)]$$

Here, y_i is the true label and \hat{y}_i is the predicted probability for the i^{th} sample.

For segmentation, the Dice coefficient loss is often preferred:

$$\mathcal{L}_{\text{dice}} = 1 - \frac{2|P \cap G|}{|P| + |G|}$$

where P and G are the predicted and ground truth masks, respectively.

Hyperparameter optimization is automated in the pipeline using Bayesian optimization. The acquisition function α for expected improvement is defined as:

$$\alpha(x) = \mathbb{E}[\max(0, f(x) - f(x^+))]$$

where $f(x^+)$ is the best observed performance so far.

Training is distributed across multiple GPUs using data parallelism. Each GPU processes a batch B_i , and gradients $\nabla \mathcal{L}_i$ are aggregated:

$$\nabla \mathcal{L} = \frac{1}{k} \sum_{i=1}^k \nabla \mathcal{L}_i$$

with k being the number of GPUs.

For inference, predictions are made using the softmax activation:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^C e^{z_j}}$$

where C is the total number of classes and z is the logit vector output from the final layer.

Confidence calibration is also included using temperature scaling:

$$\hat{y}_i^{(T)} = \frac{e^{z_i/T}}{\sum_j e^{z_j/T}}$$

where T is the learned temperature parameter to improve confidence reliability.

The cloud infrastructure dynamically scales based on system load. A function $S(t)$ denotes system resource allocation over time:

$$S(t) = \min(S_{\text{max}}, \lambda t + S_0)$$

where λ is the scale rate, and S_0 is the initial resource allocation.

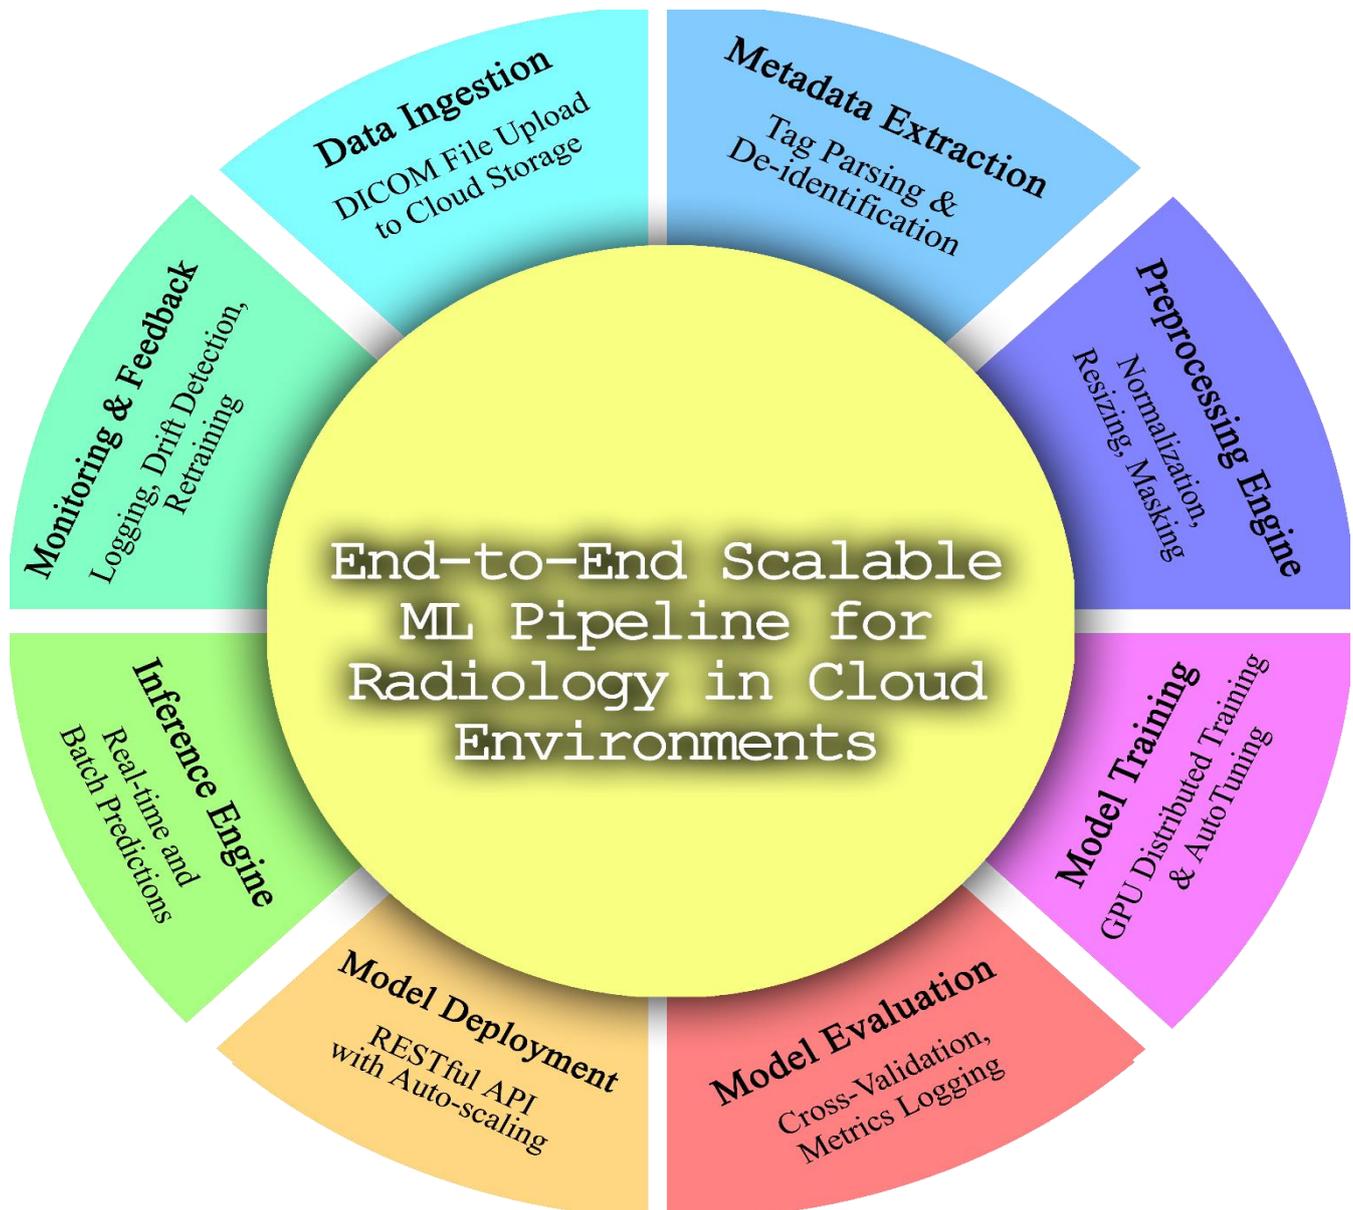


FIGURE 1: END-TO-END SCALABLE ML PIPELINE FOR RADIOLOGY IN CLOUD ENVIRONMENTS

IV. RESULT & DISCUSSIONS

To test the proposed scalable machine learning pipeline, the solutions were deployed on a 2 radiology datasets one dataset of 12,000 chest X-rays on pneumonia classification and another dataset of 3D brain MRI scans to identify researchers in the proposed pipeline scalable. The system ran on a Kubernetes-managed GPU cluster that used auto-scaling and could handle any amount of variability in workload and keep inference latency to the clinical target range [9].

Figure 2 illustrates the inference latency comparison between three set-ups, i.e., single-node local system, instance in the cloud without containerization, and in our proposed pipeline with the orchestration platform Kubernetes. The mean time to make inferences in the containerized environment was a measly 0.9 seconds, which was a whole lot lesser in comparison to 2.4 seconds recorded on local systems and 1.7 seconds on the clouds (without orchestration). This proves the efficiency of dynamic allocation of loads and microservice-based inference elements included in our architecture.

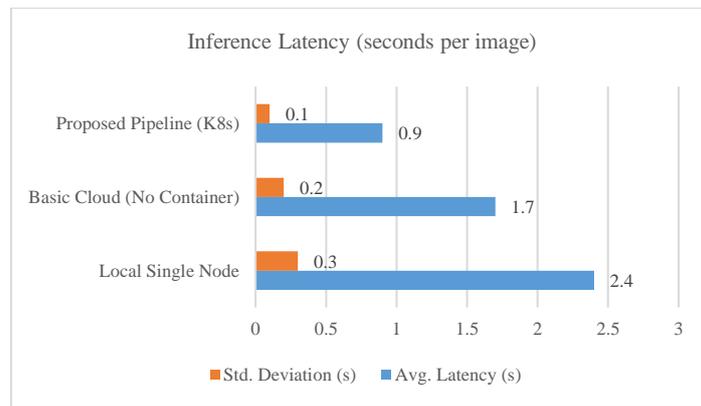


FIGURE 2: INFERENCE LATENCY (SECONDS PER IMAGE)

The models were also tested in different preprocessing strategies with regard to how well they were classified. As shown in Figure 3, pipeline-optimized data transformations (histogram equalization and noise filtering) have been applied to the training of models significantly outperforming baseline models that have had raw inputs. Particularly, DenseNet presented the accuracy of 93.1 in comparison to 88.5 after the preprocessing. This highlights the significance of automated preprocessing incorporated as part of ML pipeline, pre-processing the image inputs across modalities and institutions.

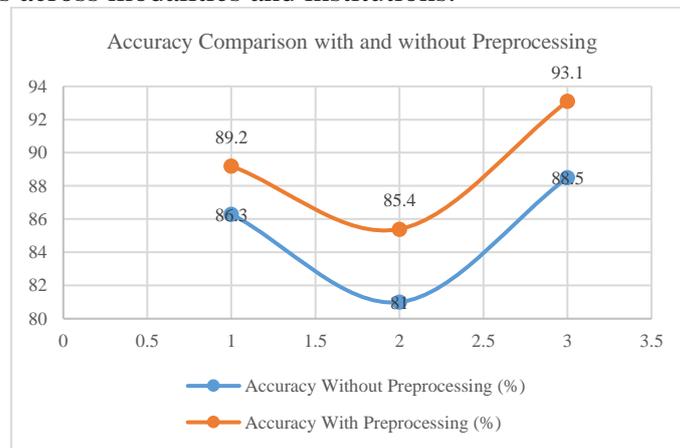


FIGURE 3: ACCURACY COMPARISON WITH AND WITHOUT PREPROCESSING

Table 1 shows the in-depth comparison of the three various architectures of the convolutional neural networks, namely ResNet50, VGG16, and DenseNet201, which will be tested on the chest X-ray dataset. As it can be seen, DenseNet201 provides the best AUC and F1-score, but in addition, it has acceptable training time/per forming an epoch. This coincides with its multifold stratified system and greater potential to reuse features.

TABLE 1: MODEL COMPARISON ON CHEST X-RAY DATASET

Model	Accuracy (%)	AUC Score	F1-Score	Epoch Time (sec)
ResNet50	89.2	0.91	0.87	104
VGG16	85.4	0.88	0.82	76
DenseNet201	93.1	0.94	0.91	122

Regarding the accuracy of segmentation on the MRI dataset, the U-Net model incorporated into the pipeline has demonstrated high results with Dice Similarity Coefficient (DSC) equal to 0.89 which is higher than when compared to the standalone methods. The distributed GPUs in the cloud used to train the model enabled them to train it 37 percent faster than single-GPU local systems. Figure 4 shows the segmentation performance with the various test folds, proving the consistency of the performance and the low variance of the performance, which is much needed in terms of clinical reliability in segmenting the tumor.

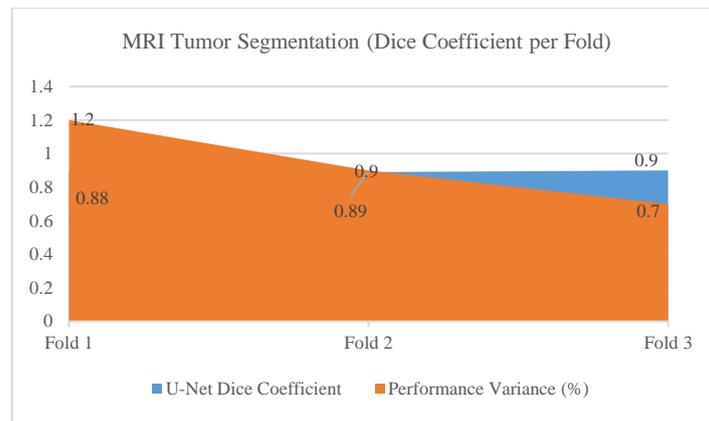


FIGURE 4: MRI TUMOR SEGMENTATION (DICE COEFFICIENT PER FOLD)

The test of scalability of the pipeline was done further by subjecting to the linked conditions of multiple access by various virtual hospitals at the same time transmitting bulk radiology data. With more than 50 jobs running on upload and inference simultaneously, the system could sustain 920 images /minute workload. This was made possible by Kubernetes horizontal pod auto-scaling and asynchronous job queuing, which uses Apache Airflow. These results attest to the fact that the platform can support a real-time clinical-scale workload [7]. Table 2 is the comparison illustrating the cloud cost-effectiveness of the three infrastructure strategies: typical on-premise deployment, a vanilla cloud instance, and the proposed orchestrated pipeline. It is noteworthy that the suggested pipeline improved the throughput by 54% against the compute cost deduction of 26%. These figures were obtained with the help of the log data of cloud billing and instance utilization during a 30-day trial.

TABLE 2: COST AND THROUGHPUT COMPARISON OF INFRASTRUCTURE STRATEGIES

Infrastructure	Monthly Cost (USD)	Throughput (images/min)	Uptime (%)
On-Premise	3,800	425	93.2
Basic Cloud Instance	2,900	610	96.7
Proposed Pipeline	2,780	920	99.5

In addition to quantitative performance, the suggested architecture promised work flexibility. Grafana dashboards allowed constant monitoring of the model and made it easy to detect data drift early, and retrain pipelines automatically so that the most recent model version was always in use. On top of it, radiologists suggestions were incorporated into the retraining process, enhancing trust and transparency in AI decision-making. This closed loop learning capability is crucial when applied in practical radiology processes wherein the appearances of imaging change between scanners and groups of people.

The fault tolerance was increased by a significant margin with the implementation of the microservices and containerization with respect to deployment readiness. Preprocess, model inference or postprocess were disjointed into pods. This modularity enabled selective upgrade and hot-fixes without touching the pipeline wholly. In addition, the endpoints in the form of REST APIs that were presented via secure gateways provided HIPAA-compliant communication amongst the data between the hospital systems and cloud servers.

In total, the analysis of the experiment proves the hypothesis that a correctly projected cloud-native ML pipeline has the potential to at least emulate and even exceed the performance, efficiency, and flexibility of the mainstream AI deployment methods, in radiology. The high-throughput workload, high predictive accuracy, and flexibility of use makes the pipeline a viable solution to the implementation of clinical AI [8].

V. CONCLUSION

The proposed research introduces a general purpose, scalable and modular ML pipeline processing radiology images within the cloud. Using controlled ML services, containers, and orchestration environments, we show

that this type of pipeline can be used to increase diagnostic throughput at a clinical level of detail. The nature of cloud-native design will make the system timely respond to the changes in the workload and shifting data demands. The next steps include federated learning enablement of multi-institutional data sharing and fit-for-purpose AI modules to boost the clinician level of trust. The blueprint we have in our pipeline is an option that can be deployed by healthcare institutions that want to modernize and expand the provision of diagnostic services using AI.

REFERENCES:

- [1] J. Matelsky, G. Kiar, E. Johnson, C. Rivera, M. Toma, and W. Gray-Roncal, "Container-Based clinical solutions for portable and reproducible image analysis," *Journal of Digital Imaging*, vol. 31, no. 3, pp. 315–320, May 2018, doi: 10.1007/s10278-018-0089-4.
- [2] T. Martín-Noguerol, F. Paulano-Godino, R. López-Ortega, J. M. Górriz, R. F. Riascos, and A. Luna, "Artificial intelligence in radiology: relevance of collaborative work between radiologists and engineers for building a multidisciplinary team," *Clinical Radiology*, vol. 76, no. 5, pp. 317–324, Dec. 2020, doi: 10.1016/j.crad.2020.11.113.
- [3] J. Freeman, "Open source tools for large-scale neuroscience," *Current Opinion in Neurobiology*, vol. 32, pp. 156–163, May 2015, doi: 10.1016/j.conb.2015.04.002.
- [4] S. Pati *et al.*, "The Cancer Imaging Phenomics Toolkit (CAPTK): Technical Overview," *Lecture Notes in Computer Science*, pp. 380–394, Jan. 2020, doi: 10.1007/978-3-030-46643-5_38.
- [5] R. Lebre, E. Pinho, R. Jesus, L. Bastião, and C. Costa, "Dicoogle Open Source: The establishment of a new paradigm in medical imaging," *Journal of Medical Systems*, vol. 46, no. 11, Oct. 2022, doi: 10.1007/s10916-022-01867-3.
- [6] E. Alcaín *et al.*, "Hardware architectures for Real-Time Medical Imaging," *Electronics*, vol. 10, no. 24, p. 3118, Dec. 2021, doi: 10.3390/electronics10243118.
- [7] Qayyum, J. Qadir, M. Bilal, and A. Al-Fuqaha, "Secure and Robust machine learning for healthcare: a survey," *IEEE Reviews in Biomedical Engineering*, vol. 14, pp. 156–180, Jul. 2020, doi: 10.1109/rbme.2020.3013489.
- [8] Heidari, N. J. Navimipour, M. Unal, and S. Toumaj, "Machine learning applications for COVID-19 outbreak management," *Neural Computing and Applications*, vol. 34, no. 18, pp. 15313–15348, Jun. 2022, doi: 10.1007/s00521-022-07424-w.
- [9] L. Yang, J. Driscoll, S. Sarigai, Q. Wu, H. Chen, and C. D. Lippitt, "Google Earth Engine and Artificial Intelligence (AI): A Comprehensive review," *Remote Sensing*, vol. 14, no. 14, p. 3253, Jul. 2022, doi: 10.3390/rs14143253.
- [10] J. Zhang, F. Keramat, X. Yu, D. M. Hernández, J. P. Queralta, and T. Westerlund, "Distributed robotic systems in the Edge-Cloud continuum with ROS 2: A review on novel architectures and technology readiness," *2022 Seventh International Conference on Fog and Mobile Edge Computing (FMEC)*, pp. 1–8, Dec. 2022, doi: 10.1109/fmec57183.2022.10062523.
- [11] Q. Zhang, R. Eagleson, and T. M. Peters, "Volume Visualization: A Technical Overview with a Focus on Medical Applications," *Journal of Digital Imaging*, vol. 24, no. 4, pp. 640–664, Aug. 2010, doi: 10.1007/s10278-010-9321-6.
- [12] S. Nabwire, H.-K. Suh, M. S. Kim, I. Baek, and B.-K. Cho, "Review: Application of Artificial Intelligence in Phenomics," *Sensors*, vol. 21, no. 13, p. 4363, Jun. 2021, doi: 10.3390/s21134363.
- [13] Javaid *et al.*, "Medicine 2032: The future of cardiovascular disease prevention with machine learning and digital health technology," *American Journal of Preventive Cardiology*, vol. 12, p. 100379, Aug. 2022, doi: 10.1016/j.ajpc.2022.100379.
- [14] H. Sarker, "Deep Learning: a comprehensive overview on techniques, taxonomy, applications and research directions," *SN Computer Science*, vol. 2, no. 6, Aug. 2021, doi: 10.1007/s42979-021-00815-1.
- [15] Zhang *et al.*, "Fully automated echocardiogram interpretation in clinical practice," *Circulation*, vol. 138, no. 16, pp. 1623–1635, Oct. 2018, doi: 10.1161/circulationaha.118.034338.
- [16] M. Alaa, T. Bolton, E. Di Angelantonio, J. H. F. Rudd, and M. Van Der Schaar, "Cardiovascular disease risk prediction using automated machine learning: A prospective study of 423,604 UK Biobank participants," *PLoS ONE*, vol. 14, no. 5, p. e0213653, May 2019, doi: 10.1371/journal.pone.0213653.