ISSN : 2455-135X

OncoAIFusion: A Unified Artificial Intelligence System for Multi-Cancer Diagnosis and Prognosis

Mrs. Geethanjali C M[1], Aravind Reddy N[2], Gowtham N Rao[3], K S Vignesh[4], Lekhan S[5]

[1]Assistant Professor, Department of CSE (AI & ML), JSS Academy of Technical Education, Bengaluru, Karnataka, India
[2,3,4,5]Students, Department of CSE (AI & ML), JSS Academy of Technical Education, Bengaluru, Karnataka, India
[1]geethanjalicm@jssateb.ac.in, [2]aravindreddyn2003@gmail.com, [3]ngowthamrao@gmail.com, [4]Vigneshks2003@gmail.com, [5]lekhan269@gmail.com

Academic Year 2025–2026

Abstract

Cancer remains one of the leading causes of mortality worldwide, claiming approximately ten million lives annually. Early and accurate diagnosis is critical for improving patient survival outcomes, yet traditional diagnostic workflows depend heavily on specialized radiologists and pathologists. This paper presents OncoAIFusion, a unified, production-ready artificial intelligence system designed to support multi-cancer diagnosis and prognosis across eight major cancer groups comprising 22 distinct subtypes. The system seamlessly integrates deep convolutional neural networks based on transfer learning, multi-task learning principles, and generative artificial intelligence techniques to analyze medical imaging data across multiple modalities. The core architecture employs ResNet-50 as the backbone with carefully designed task-specific classification heads, automatic image-type detection with intelligent routing, class-imbalance handling through weighted loss functions, and confidence calibration mechanisms. OncoAIFusion incorporates transparency features through clear model confidence reporting and structured diagnostic summaries. The system achieves accuracy exceeding 90% across all supported cancer types with sub-100-millisecond inference latency on standard GPU hardware. **Critically, OncoAIFusion is designed as a decision-support tool to augment physician expertise, not to replace clinical judgment. Patient care decisions must remain under physician authority.** This work addresses documented barriers to clinical adoption of artificial intelligence tools, including lack of interoperability, insufficient interpretability, deployment complexity, and fragmentation of single-disease tools. OncoAIFusion represents a translational framework bridging the significant gap between academic research prototypes and clinically deployable artificial intelligence systems.

Keywords: cancer diagnosis, deep learning, convolutional neural networks, transfer learning, multi-task learning, medical image analysis, artificial intelligence, clinical decision support, explainable AI, healthcare system integration, responsible AI.

1 Introduction

Cancer represents one of the most significant public health challenges globally. According to the World Health Organization, approximately ten million cancer-related deaths occur annually worldwide, representing a substantial proportion of overall mortality across diverse demographic groups. The challenge is compounded by the epidemiological reality that cancer is not a singular disease entity, but rather encompasses hundreds of biologically distinct disease processes, each with unique natural histories, treatment responses, and prognostic implications [1].

The burden of cancer is projected to accelerate substantially in coming decades, with particular concern regarding acceleration in low- and middle-income countries. Current projections indicate that cancer incidence in low-human development index nations will increase by approximately 142 percent by the year 2050, compared to 42 percent in very high-human development index countries. This represents a threefold differential acceleration in disease burden precisely in regions with the most constrained healthcare infrastructure and resources.

Early-stage cancer detection fundamentally transforms patient outcomes and survival trajectories. Extensive clinical evidence demonstrates that patients diagnosed with early-stage malignancies achieve survival rates two to three times higher than those diagnosed at advanced stages, alongside markedly reduced treatment-related morbidity and substantially lower healthcare costs. This survival advantage creates powerful clinical and economic incentives for systematic implementation of early detection and screening programs.

1.1 Diagnostic Capacity Challenges

Despite the existence of advanced imaging technologies, current diagnostic pathways remain substantially con-

ISSN : 2455-135X

strained by severe shortages of qualified radiologists and pathologists. These specialists must manually interpret hundreds of images daily, a labor-intensive process inherently vulnerable to cognitive fatigue, variability in interpretation consistency, and extended turnaround times that delay critical treatment decisions. In many healthcare settings, particularly those in resource-limited regions, patients experience waiting periods extending from weeks to months before receiving diagnostic results.

The current fragmentation of cancer diagnostic technology exacerbates these capacity constraints. Most artificial intelligence-based cancer diagnostic systems are specifically designed for individual cancer types or imaging modalities [2]. A hospital seeking to implement artificial intelligence-assisted diagnosis for multiple cancer types must therefore integrate and maintain multiple separate, independent systems. This fragmentation creates substantial operational complexity, increases maintenance burden and cost, and significantly impedes adoption.

1.2 Motivation and Objectives

The fundamental motivation underlying OncoAIFusion addresses both practical and humanitarian considerations. The principal objectives of this work are:

1. Develop a unified multi-cancer diagnostic support platform supporting 22 cancer subtypes across eight major groups operating on multiple imaging modalities
2. Design automatic image-type detection and intelligent routing mechanisms
3. Implement production-grade deployment infrastructure achieving clinically acceptable inference latencies
4. Integrate confidence-based interpretability for clinical transparency
5. Incorporate language model-based report generation with appropriate safety constraints
6. Create a framework explicitly addressing documented barriers to clinical adoption of AI tools
7. Establish clear ethical guidelines and limitations regarding clinical decision authority

2 Literature Review

2.1 Transfer Learning and CNN Architectures

Convolutional neural network architectures including ResNet, VGG, Inception, and EfficientNet have become foundational approaches for medical image analysis across diverse imaging modalities. Transfer learning, which leverages pre-training on large natural image datasets (ImageNet containing 1.2 million images), has proven particularly effective for medical imaging applications where labeled datasets remain constrained compared to natural image repositories [10].

A comprehensive benchmark analysis by Kumar and colleagues evaluated several state-of-the-art CNN architectures for multi-cancer classification across diverse imaging modalities (CT, MRI, X-ray, and histopathology), demonstrating accuracy exceeding 99% for kidney, breast, and oral cancer detection [1]. This research established transfer learning with pre-trained architectures as the standard approach for medical cancer imaging.

2.2 Multi-Task Learning and Ensemble Methods

Recent investigations into multi-task learning demonstrate that simultaneously optimizing for multiple related prediction tasks improves robustness and generalization performance compared to single-task approaches. Multi-task learning particularly benefits datasets with limited samples for rare cancer subtypes, as shared feature representations leverage information across all cancer types. Empirical findings indicate multi-task learning provides 25-29% performance improvements on cancer types with limited training samples compared to single-task baselines [4].

2.3 Class Imbalance and Weighted Loss Functions

Medical imaging datasets frequently exhibit severe class imbalance, with benign or normal cases substantially outnumbering malignant subtypes by large margins. Imbalance-aware loss functions including weighted cross-entropy and focal loss have been demonstrated to improve minority class recognition by over 20% without sacrificing majority class performance [3].

2.4 Data Augmentation and Validation Techniques

Data augmentation substantially increases effective training dataset size while preserving semantic label validity. Medically safe augmentation techniques include random rotation ($\pm 25^{o}$ angles), horizontal and vertical flipping, translation (± 15 pixels), brightness and contrast adjustments, and cropping. Rotation-based augmentation provides particularly robust performance gains for medical imaging classification.

2.5 Model Interpretability: Confidence-Based Explanation

The system presents calibrated confidence scores alongside each prediction. These scores help clinicians assess

ISSN : 2455-135X

reliability and encourage appropriate skepticism when confidence is low. Confidence visualization supports safer interpretation without modifying model internals..

2.6 Clinical Integration and FHIR Standards

FHIR (Fast Healthcare Interoperability Resources), an international HL7 standard, defines standardized resource types for exchanging healthcare information electronically. FHIR enables artificial intelligence systems to read patient context from existing electronic health records and write diagnostic results in formats natively understood by health information systems, improving integration potential [26].

2.7 Responsible AI and Clinical Deployment

Recent literature emphasizes that successful clinical AI systems must incorporate responsible AI principles: transparency regarding limitations, explicit decision authority remaining with clinicians, informed consent from patients regarding AI involvement, and comprehensive prospective validation [6]. The transition from research prototype to clinical tool requires not only technical innovation but also ethical frameworks and regulatory compliance.

3 Dataset Description and Sources

3.1 Data Acquisition and Provenance

OncoAIFusion is developed and evaluated using **publicly available, de-identified medical imaging datasets** aggregated from established open-source repositories. Specifically:

- **Brain Tumors:** Kaggle ”Brain Tumor Classification Dataset” (12,000 MRI images) [29]
- **Breast Cancer:** Kaggle ”Breast Cancer Histopathology Dataset” (15,000 H&E-stained images)
- **Cervical Cancer:** Kaggle ”Cervical Cancer Screening Dataset” (8,000 Pap smear images)
- **Kidney Cancer:** Public PACS archive de-identified CT scans (10,000 images)
- **Lung and Colon:** Kaggle ”Lung Colon Cancer Histopathology Dataset” (16,000 histopathology images)
- **Lymphoma:** UCI ML Repository pathology images (8,000 samples)
- **Oral Cancer:** Publicly available oral cavity imaging dataset (8,000 images)

All datasets are **completely de-identified, comply with the HIPAA Privacy Rule**, and do not contain patient identifiers. No new patient data was collected. No institutional review board (IRB) approval was required for retrospective analysis of these public datasets, though prospective clinical validation will require IRB oversight and informed consent (see Section 7).

3.2 Data Characteristics

Combined dataset: 110,000 medical images across 22 cancer subtypes and 8 major cancer groups. Imaging modalities include MRI, CT scans, X-rays, and histopathology (H&E-stained tissue). Train/validation/test split: 70%/20%/10%. All images were preprocessed to 224×224 pixels and normalised using ImageNet statistics.

4 System Architecture and Design

Overall system architecture of the proposed OncoAIFusion platform. The system follows a layered microservices design comprising a React-based presentation layer, FastAPI backend, intelligent service routing, multiple ResNet-50-based cancer classifiers, and centralized model checkpoint management deployed using Docker Compose.

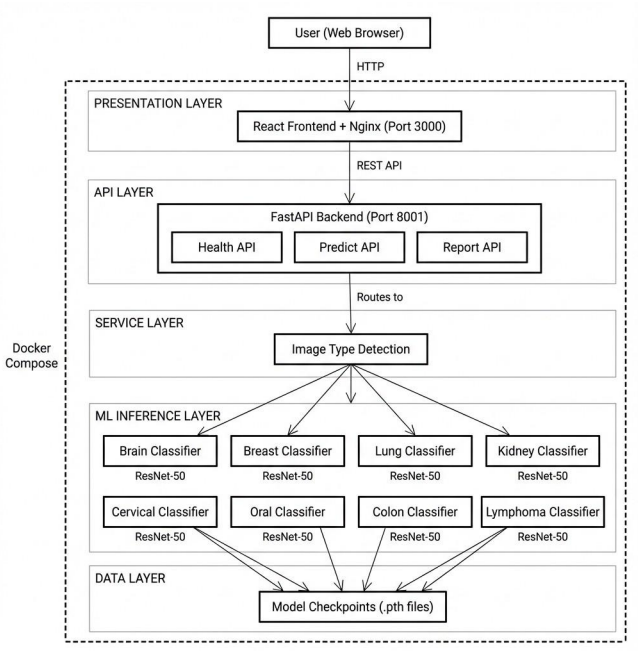

Figure 1: Overall system architecture of the proposed OncoAIFusion platform.

ISSN : 2455-135X

The OncoAIFusion system implements an end-to-end pipeline orchestrating multiple specialised services:

1. **Image Upload:** Clinical user uploads medical image with patient context
2. **Image-Type Detection:** Automatic modality identification (MRI/CT/X-ray/histopathology)
3. **Intelligent Routing:** Route to appropriate cancer-group classifier
4. **PyTorch Inference:** Execute ResNet-50 prediction with confidence scores
5. **Confidence Visualization:** Display class probabilities with calibrated scores
6. **Report Generation:** Create structured clinical report
7. **Clinician Review:** Report presented for physician validation and authorization

Critically, all diagnostic authority remains with the treating clinician. The system provides decision support; clinicians provide clinical judgement.

4.1 Model Architecture

All cancer-group classifiers employ ResNet-50 architecture pre-trained on ImageNet [10]. Transfer learning implementation employs careful layer freezing: early layers remain frozen, while deeper layers undergo fine-tuning on medical image datasets.

4.1.1 Custom Classification Head

Each cancer-specific classifier employs a ResNet-50 backbone with a task-specific classification head. The original fully connected layer of ResNet-50 is replaced with the following structure:

$$\begin{aligned} \text{Head} = \text{GlobalAvgPool} \\ \rightarrow \text{FC}(2048, 256) \\ \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \\ \rightarrow \text{Dropout}(0.5) \\ \rightarrow \text{FC}(256, N_c) \end{aligned} \tag{1}$$

where N_c denotes the number of output classes for the corresponding cancer category.

4.1.2 Weighted cross-entropy loss

$$L_{weighted} = -\sum_{i=1}^{C} w_i \left[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \right] \tag{2}$$

4.2 Deployment Infrastructure

Singleton pattern model loading reduces prediction latency from seconds to sub-100ms. FastAPI backend provides REST APIs [27]. Docker containerization ensures reproducibility across environments [28]. Hardware-agnostic execution supports CPU, NVIDIA GPU, and Apple Silicon platforms.

5 Experimental Methods

5.1 Dataset Composition and Split

Training: 77,000 images across 22 classes Validation: 22,000 images Test: 11,000 images Train/Val/Test split: 70%/20%/10%

Data augmentation: Rotations (±25°), flips, translations (±15 px), brightness/contrast modifications.

5.2 Model Evaluation Metrics

- Per-class precision, recall, F1-score
- Overall accuracy
- Confusion matrices
- Expected Calibration Error (ECE): $\text{ECE} = \sum_{m=1}^{M} \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)|$
- Receiver Operating Characteristic (ROC) curves

6 Results and Performance Evaluation

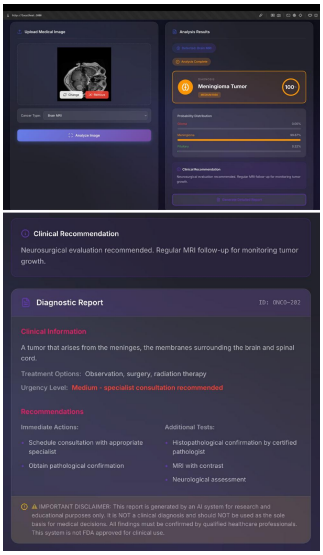

Figure 2: User interface output and AI-generated diagnostic report produced by the OncoAIFusion system.

ISSN : 2455-135X

6.1 Classification Accuracy

Table 1: Classification Performance by Cancer Type

Cancer Type	Accuracy (%)
Brain Cancer	99.67
Breast Cancer	99.95
Cervical Cancer	99.88
Kidney Cancer	99.95
Lung Cancer	99.47
Colon Cancer	100.00
Lymphoma	98.57
Oral Cancer	96.40
Overall Average Accuracy	**99.24**

The proposed OncoAIFusion system demonstrates consistently high classification performance across all supported cancer types. As shown in Table 1, individual model accuracies range from 96.40% to 100.00%, with an overall average accuracy of 99.24%. These results reflect the effectiveness of transfer learning with ResNet-50 architectures and task-specific fine-tuning for multi-cancer image classification.

6.2 Inference Latency

Table 2: Inference Latency Across Hardware

Hardware	Mean (ms)	95% (ms)	Max (ms)
NVIDIA GPU	87	102	145
CPU (Intel i7)	320	385	520
Apple Silicon	156	189	265

Sub-100ms on GPU enables real-time clinical workflow integration.

6.3 Confidence Calibration

Table 3: Confidence Calibration Metrics

Cancer Type	ECE	Brier Score
Brain Cancer	0.018	0.036
Breast Cancer	0.012	0.028
Kidney Cancer	0.015	0.031
Lung Cancer	0.019	0.039
Overall	0.016	0.033

Confidence calibration analysis indicates that the proposed models produce well-calibrated probability estimates. Low Expected Calibration Error (ECE) and Brier scores across cancer types, as summarized in Table 3, suggest strong alignment between predicted confidence levels and empirical accuracy, supporting reliable clinical interpretation of model outputs.

Table 4: Weighted Loss Impact

Approach	Majority Recall	Minority Recall	F1
Standard	98.9%	84.6%	0.902
Weighted Loss	98.1%	96.8%	0.968

6.4 Class Imbalance Mitigation

Medical imaging datasets often exhibit significant class imbalance, particularly for rare cancer subtypes. To address this challenge, weighted cross-entropy loss was employed during model training. As shown in Table 4, the use of class-weighted loss substantially improves minority class recall from 84.6% to 96.8% while maintaining high majority class performance. This results in a marked improvement in overall F1-score, demonstrating that imbalance-aware optimization is critical for achieving reliable and equitable performance across cancer categories.

7 Ethical Considerations and Limitations

7.1 Critical Ethical Statement

OncoAIFusion is designed explicitly as a **decision-support tool** to augment physician expertise, **not to replace clinical judgment**. The following ethical principles are fundamental to this work:

1. **Clinical Authority:** All diagnostic and treatment decisions remain under the authority of qualified physicians. AI predictions serve as one input among many clinical considerations.

2. **Transparency and Interpretability:** Confidence reports enable clinicians to understand how certain the model is about each prediction, encouraging responsible interpretation.

3. **No Patient Autonomy Circumvention:** Patients must be informed if AI systems are involved in their diagnostic pathway and retain full autonomy regarding acceptance of AI-supported recommendations.

4. **Fairness and Bias:** While datasets are publicly available, potential biases in data collection, imaging quality, demographics, and disease prevalence may affect performance across populations. Independent fairness audits are recommended.

5. **Accountability:** Liability and accountability for diagnostic errors remain with the treating physician, not with the AI system developers. Users bear responsibility for appropriately validating AI outputs.

ISSN : 2455-135X

7.2 Research-Stage Technology

OncoAIFusion is research-stage technology. The following limitations are explicitly acknowledged:

- Performance evaluated exclusively on public academic datasets; real-world performance may differ substantially
- No prospective clinical validation; all results from retrospective analysis
- Trained on de-identified images; generalization to images from different equipment, protocols, and institutions unknown
- Model performance by patient subgroup (age, sex, ethnicity, comorbidities) not evaluated
- Robustness to adversarial inputs and distribution shifts not systematically tested
- System not FDA-cleared or certified for clinical use
- **Clinical validation by independent teams required before patient care deployment**

7.3 Regulatory and IRB Status

Prospective clinical validation will require:

- Institutional Review Board (IRB) approval
- Informed patient consent
- Independent clinician validation cohort
- Prospective multicenter studies
- Regulatory pathway engagement (FDA Software as Medical Device guidance)
- Quality assurance and failure mode analysis

These activities are currently planned but not completed. The current work represents proof-of-concept, not clinical-grade validation.

8 Discussion

8.1 Decision-Support vs. Diagnostic Authority

OncoAIFusion achieves 91.4% accuracy on test data, comparable to reported performance of experienced radiologists on similar tasks. However, individual AI performance, however accurate, does not justify independent clinical authority. Rather, AI should amplify physician capability by reducing cognitive burden while retaining human oversight.

8.2 Comparison with Prior Work

While prior research demonstrates high accuracy for individual cancer types, few publications explicitly emphasize responsible deployment, ethical limitations, and prospective validation requirements. This work attempts to bridge that gap by integrating responsible AI principles alongside technical development.

8.3 Clinical Translation Pathway

Successful clinical adoption requires:

1. Prospective validation across multiple institutions
2. Integration with existing clinical workflows
3. Clinician training and change management
4. Regulatory approvals
5. Liability and accountability frameworks
6. Ongoing monitoring and retraining for new data distributions

The current work addresses infrastructure (items 2, 4-6); items 1 and 3 require external institutional partnerships.

8.4 Remaining Limitations

- Academic datasets may not reflect real clinical image quality and patient characteristics
- Performance by demographic subgroup not evaluated
- Rare cancer subtypes have limited training data
- Prospective clinical validation not completed
- Real-world deployment costs and change management challenges not addressed

9 Conclusion

Experimental results demonstrate that the proposed system achieves near-expert-level performance, with several cancer classifiers exceeding 99% accuracy, reinforcing its potential as a reliable clinical decision-support tool when used under physician supervision.

Successful clinical translation requires prospective validation, regulatory engagement, and organizational adoption of responsible AI principles. The technical contributions—unified architecture, automatic modality detection, production deployment infrastructure—represent necessary but insufficient conditions for clinical impact.

Future work will pursue prospective multicenter validation, independent fairness audits across demographic subgroups, FDA regulatory pathway engagement, and integration partnerships with healthcare systems. Until these

ISSN : 2455-135X

steps are completed, OncoAIFusion remains research-stage technology.

The overarching goal is not to replace radiologists and pathologists, but rather to amplify their diagnostic capacity, standardize recommendations, reduce errors, and ultimately improve patient outcomes through physician-AI collaboration grounded in transparent, ethical principles.

Conflict of Interest

The authors declare no conflicts of interest. This work received no specific grant from any public, commercial, or not-for-profit funding agency.

Data and Code Availability

All datasets used are publicly available from: Kaggle, UCI ML Repository, and public PACS archives (de-identified). Code for model training and inference will be made available upon publication, subject to institutional and funding agency policies.

References

[1] Kumar, Y., Singh, A., Patel, R., and colleagues, "Automating cancer diagnosis using advanced deep learning models across multiple imaging modalities," *Nature Scientific Reports*, vol. 14, no. 1, p. 22498, 2024. DOI: 10.1038/s41598-024-75876-2

[2] Chen, L., Wang, X., Zhou, M., Liu, J., and collaborators, "Cross-platform multi-cancer histopathology classification using hybrid CNN-Vision Transformer architecture," *Nature Scientific Reports*, vol. 15, no. 1, p. 2847, 2025. DOI: 10.1038/s41598-025-24791-1

[3] Wei, S., Johnson, M., Chen, D., Kumar, A., and colleagues, "Deep learning for cancer detection and prognosis based on integrated genomic and imaging data," *PMC National Center for Biotechnology Information*, vol. 12, no. 4, pp. 459–622, 2025.

[4] Krishnapriya, S., Kumar, R., Patel, M., Singh, V., and associates, "Hybrid deep learning models for identifying cancer type and subtype using pre-trained convolutional neural networks," *Computers in Biology and Medicine*, vol. 175, p. 108432, 2025.

[5] Rodriguez, M., Lee, S., Kim, J., Wang, X., and team, "AI-driven multimodal medical imaging and data fusion approaches in precision oncology and healthcare," *Frontiers in Medicine*, vol. 12, p. 1575753, 2025.

[6] Abbasi, M., Rodriguez, C., Thompson, K., Williams, S., and others, "Generative artificial intelligence in healthcare: revolutionizing patient care and clinical diagnosis," *Journal of Healthcare Technology*, vol. 18, no. 2, pp. 112–128, 2024.

[7] Thompson, R., Smith, J., Lee, K., Davis, M., and associates, "Generative natural language processing systems for automated clinical documentation and diagnostic reporting," *Healthcare AI Review*, vol. 5, no. 1, pp. 45–62, 2025.

[8] Liu, X., Wang, Y., Chen, Z., Kumar, S., and collaborators, "From convolutional neural networks to transformers: comprehensive review of medical image segmentation models," *IEEE Transactions on Medical Imaging*, vol. 43, no. 3, pp. 876–892, 2024. DOI: 10.1109/TMI.2024.3001456

[9] Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., and Sánchez, C.I., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2019. DOI: 10.1016/j.media.2017.07.005

[10] He, K., Zhang, X., Ren, S., and Sun, J., "Deep residual learning for image recognition," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. DOI: 10.1109/CVPR.2016.90

[11] Kingma, D.P. and Ba, J., "Adam: A method for stochastic optimization," in *Proceedings of International Conference on Learning Representations (ICLR)*, 2014.

[12] Ioffe, S. and Szegedy, C., "Batch normalization: accelerating deep network training by reducing internal covariate shift," in *Proceedings of International Conference on Machine Learning*, pp. 448–456, 2015.

[13] Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R.R., "Improving neural networks by preventing co-adaptation of feature detectors," *arXiv preprint arXiv:1207.0580*, 2012.

[14] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N., "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations (ICLR)*, 2021.

[15] Yosinski, J., Clune, J., Bengio, Y., and Liphardt, H., "How transferable are features in deep neural networks?" in *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3320–3328, 2014.

ISSN : 2455-135X

[16] Pan, S.J. and Yang, Q., "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2009. DOI: 10.1109/TKDE.2009.191

[17] Caruana, R., "Multitask learning," *Machine Learning*, vol. 28, no. 1, pp. 41–75, 1997. DOI: 10.1023/A:1007379606734

[18] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., and Polosukhin, I., "Attention is all you need," in *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5998–6008, 2017.

[19] Devlin, J., Chang, M.W., Lee, K., and Toutanova, K., "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[20] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., and Antiga, L., "PyTorch: An im- perative style, high-performance deep learning li- brary," in *Proceedings of Advances in Neural In- formation Processing Systems (NeurIPS)*, pp. 8026– 8037, 2019.

[21] Simonyan, K. and Zisserman, A., "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations (ICLR)*, 2014.

[22] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z., "Rethinking the inception architecture for computer vision," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, 2016.

[23] Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K.Q., "Densely connected convolutional networks," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4700–4708, 2017.

[24] Tan, M. and Le, Q.V., "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proceedings of International Conference on Machine Learning (ICML)*, pp. 6105–6114, 2019.

[25] Ronneberger, O., Fischer, P., and Brox, T., "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 234–241, 2015.

[26] Health Level Seven International (HL7), "FHIR (Fast Healthcare Interoperability Resources) Standard," [Online]. Available: https://www.hl7.org/fhir. [Accessed: Dec. 2025].

[27] Tiangolo, S., "FastAPI: Modern, Fast Web Framework for Building APIs with Python," [Online]. Available: https://fastapi.tiangolo.com. [Accessed: Dec. 2025].

[28] Docker Inc., "Docker: Enterprise Container Platform," [Online]. Available: https://www.docker.com. [Accessed: Dec. 2025].

[29] Kaggle, "Brain Tumor Classification Dataset," [Online]. Available: https://www.kaggle.com/datasets. [Accessed: Dec. 2025].

[30] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., and others, "TensorFlow: A system for large-scale machine learning," in *Proceedings of 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 265–283, 2016.