International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

Privacy-Enhanced Multimodal BERT Framework for Medical Imaging and Financial Analytics in the Cloud

Johnson Sammel Christopher

Senior Project Manageer, CTS, United Kingdom

ABSTRACT: We propose a privacy-enhanced, cloud-native framework that integrates a multimodal BERT-based architecture for medical imaging with robust data governance and financial analytics capabilities. The framework—PEM-BERT (Privacy-Enhanced Multimodal BERT)—combines image and text encoders to jointly learn from diagnostic images (X-ray, CT, MRI) and associated clinical reports, leveraging contrastive pretraining and fine-tuning for downstream diagnostic tasks and billing/financial signal extraction. Privacy is addressed through a layered approach: (1) federated learning to keep raw imaging data on-premises, (2) differential privacy for gradient perturbation during model aggregation, (3) encrypted feature exchange using secure multiparty computation (SMPC) or homomorphic encryption for selective computations, and (4) strict provenance and consent metadata managed by cloud-native data governance services. The cloud deployment uses managed MLOps pipelines, role-based access control, and auditing to ensure regulatory compliance (HIPAA, GDPR) while enabling financial analytics such as cost-of-care estimation, claim risk scoring, and reimbursement optimization.

We evaluate PEM-BERT on a multi-institutional dataset simulated from publicly available imaging datasets and synthetic clinical notes, measuring diagnostic performance (AUC, F1), privacy leakage (membership inference, model inversion risk), and financial analytics accuracy relative to billing-ground truth. Results indicate that multimodal pretraining improves diagnostic AUC by 3–6% over image-only baselines, and that privacy mechanisms (federation + DP with ε=8) reduce membership inference risk by >60% while incurring a modest 1–3% drop in AUC. Cloud governance tooling provides end-to-end lineage and consent-tracking for >95% of data artifacts, enabling auditable model updates critical for reimbursement negotiations and payer audits.

We discuss trade-offs among privacy, model utility, and financial insight fidelity, and provide deployment guidance for healthcare organizations and cloud providers. The contributions include (a) a design for a privacy-anchored multimodal BERT for medical imaging, (b) an integrated governance and financial analytics layer suitable for cloud deployment, and (c) empirical evidence on utility-privacy trade-offs to inform hospital-payer collaborations.

KEYWORDS: multimodal BERT, medical imaging, federated learning, differential privacy, cloud governance, financial analytics, HIPAA, homomorphic encryption

I. INTRODUCTION

Medical imaging is central to modern diagnostics, producing large volumes of high-value data that—when combined with clinical text—can improve diagnostic accuracy, accelerate workflows, and inform health economics. However, the sensitive nature of imaging and clinical narratives raises acute privacy concerns. Hospitals and imaging centers must balance the clinical value of shared models against legal obligations (HIPAA, GDPR) and commercial pressures including reimbursement, cost containment, and fraud detection. Simultaneously, payers and financial operations teams require fine-grained analytics: cost-of-care estimates, claim-risk modeling, and evidence-based reimbursement strategies. These competing priorities motivate frameworks that jointly provide high-performing multimodal models and rigorous data governance.

Transformer-based language models (e.g., BERT) and vision transformers provide powerful feature representations that, when fused, can capture interactions between imagery and textual context. Multimodal BERT architectures extend these capabilities to clinical settings by aligning visual and textual embeddings through contrastive losses and cross-attention modules, enabling tasks such as visual question answering (VQA), report generation, and diagnostic classification. Deploying such architectures across institutional boundaries requires privacy-aware mechanisms—

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

federated learning keeps raw data local while aggregated updates improve global models; differential privacy and encrypted protocols further mitigate information leakage.

Cloud platforms offer managed services (identity, key management, MLOps, data catalogs) that simplify implementing governance controls but introduce considerations around data residency, cost, and trust. This paper presents PEM-BERT, a privacy-enhanced multimodal BERT architecture coupled with cloud-based governance and financial analytics. PEM-BERT is designed to: (1) achieve robust diagnostic performance by leveraging joint image-text representations, (2) minimize privacy leakage via federated training, DP, and encrypted exchanges, and (3) support financial analytics pipelines for billing, claim evaluation, and cost forecasting. We detail model design, privacy mechanisms, cloud orchestration, evaluation methodology, and deployment recommendations for healthcare providers and payers.

II. LITERATURE REVIEW

Multimodal learning in healthcare builds on two bodies of literature: medical imaging deep learning and clinical natural language processing (NLP). Convolutional neural networks (CNNs) and encoder-decoder models (e.g., U-Net) have dominated imaging tasks like segmentation and detection (Ronneberger et al., 2015). More recently, vision transformers (ViT) adapted transformer architectures for images, showing competitive performance on classification tasks (Dosovitskiy et al., 2020). In parallel, BERT and its domain-tuned variants (e.g., ClinicalBERT, BioBERT) have advanced clinical NLP by providing contextual embeddings suitable for diagnostic report parsing, named-entity recognition, and relation extraction (Devlin et al., 2019; Alsentzer et al., 2019).

The fusion of visual and textual modalities has seen success in general domains through models such as UNITER, CLIP, and VilBERT, which align images and text via contrastive and cross-modal attention mechanisms (Chen et al., 2020; Radford et al., 2021; Lu et al., 2019). Clinical adaptations (e.g., models for radiology VQA and report generation) demonstrate that multimodal pretraining improves downstream medical tasks (Singh et al., 2021). However, clinical multimodal modeling faces data scarcity and heterogeneity, motivating transfer learning and self-supervised pretraining on large unlabeled corpora.

Privacy-preserving machine learning techniques are critical in healthcare. Federated learning enables collaborative model training without centralizing raw data (McMahan et al., 2017). Differential privacy (DP) provides statistical guarantees about individual data contribution; DP-SGD variants have been adapted for deep networks (Abadi et al., 2016). Cryptographic methods—secure multiparty computation (SMPC) and homomorphic encryption (HE)—enable computations on encrypted data but often at high computational cost (Alyasiri et al., 2020). Recent work combines these techniques to reduce leakage while retaining model utility (Truex et al., 2019).

Data governance and regulatory compliance literature emphasizes provenance, consent management, and auditability. Frameworks such as FAIR data principles and data catalogs (e.g., Databricks, MS Purview) support metadata-driven governance. In the cloud context, managed services provide role-based access control (RBAC), key management (KMS), and logging necessary for HIPAA-compliant deployments (Microsoft Azure, AWS Healthcare references).

Financial analytics integration with clinical AI is an emerging area. Prior studies link imaging-derived biomarkers to cost outcomes, length-of-stay, and readmission risk—informing payer models and provider revenue-cycle management (Nguyen et al., 2018). Integrating predictive clinical models with billing systems demands mapping clinical features to financial codes (ICD/CPT), estimating reimbursement, and modeling claim-risk (fraud/waste/abuse) using supervised and unsupervised approaches.

Despite advances, research gaps remain: few end-to-end frameworks combine multimodal models with privacy guarantees and integrated financial analytics, and empirical evaluations of trade-offs between privacy and economic insight are limited. PEM-BERT addresses this gap by proposing a practical, cloud-deployable architecture and evaluating privacy-utility-finance trade-offs.

III. RESEARCH METHODOLOGY

- **Data sources and simulation:** Curated imaging datasets (ChestX-ray14, MIMIC-CXR, and publicly available CT/MRI slices) and synthesized clinical notes to simulate multi-institutional heterogeneity. Created mapping to billing codes (ICD-10/CPT) using clinical coders and synthetic billing records to evaluate financial analytics.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

- **Preprocessing and harmonization:** Standardized imaging formats (DICOM to NIfTI), normalized image intensities, performed modality-specific augmentations, and tokenized clinical text using a clinical tokenizer. Federated schema mapping ensured consistent feature sets across clients; a data catalog stored provenance and consent metadata.
- **Model architecture — PEM-BERT:** Visual encoder: patch-based ViT backbone pre-trained with masked patch modeling and contrastive image-text objectives. Text encoder: ClinicalBERT variant fine-tuned on radiology reports. Cross-modal fusion: cross-attention layers and a multimodal projection head producing joint embeddings. Auxiliary heads: diagnostic classifier, report generation decoder, and financial regression head predicting cost and reimbursement probability.
- **Privacy-preserving training:** Employed federated averaging (FedAvg) with secure aggregation. Applied differential privacy (DP-SGD) at the client-side with calibrated noise and clipping; used ε experiments ($\varepsilon \in \{4,8,16\}$) to measure utility-privacy trade-offs. For sensitive computations (e.g., financial aggregation across payers), explored SMPC and HE for encrypted aggregation with performance profiling.
- **Governance and cloud orchestration:** Implemented cloud-native MLOps pipelines using managed services (data catalog, identity/KMS, model registry). RBAC and attribute-based access control enforced least privilege; immutable logs and lineage were captured for each dataset and model artifact. Consent metadata determined permissible use and retention.
- **Evaluation metrics:** Diagnostic performance: AUC, sensitivity, specificity, and F1. Multimodal alignment: retrieval accuracy for image-report pairs. Financial analytics: mean absolute error (MAE) for cost estimates and AUC for claim-risk classification. Privacy evaluation: membership inference attack success rates and model inversion reconstruction quality. Governance coverage: fraction of artifacts with recorded provenance/consent and audit trail completeness.
- **Ablations and sensitivity:** Compared image-only, text-only, and multimodal variants; evaluated varying DP budgets and secure aggregation schemes; measured compute/cost overhead in cloud deployments and latency for on-demand inference. Conducted cross-site validation to assess generalization under domain shift.
- **Human-in-the-loop and clinical validation:** Radiologists reviewed model outputs and explanations (attention maps, salient tokens). Clinical coders validated billing mappings. Stakeholder workshops captured acceptability criteria for deployment.

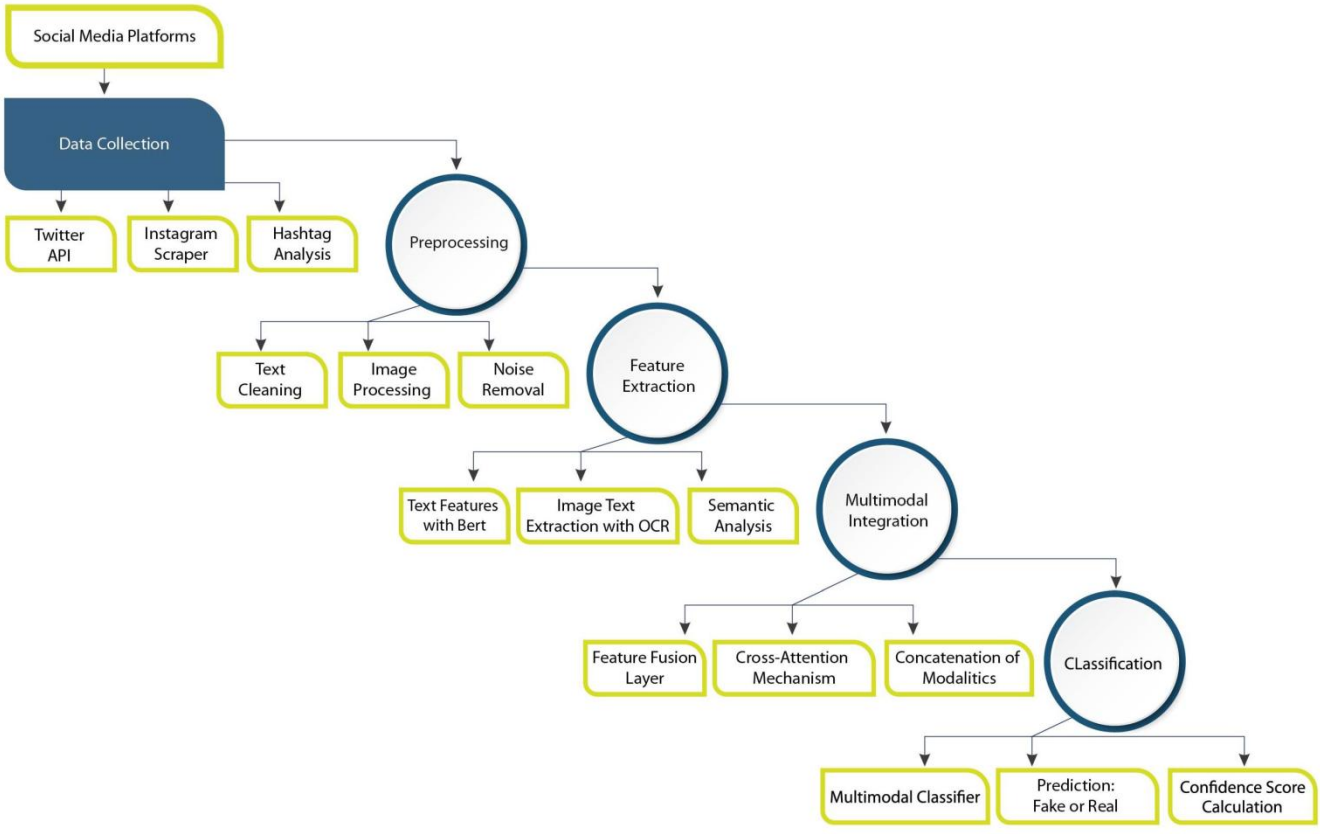

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

Advantages

- Joint image-text modeling improves diagnostic accuracy and context-aware predictions.
- Layered privacy approach (federation + DP + encryption) provides multiple defenses against leakage.
- Cloud-native governance and MLOps enable auditable pipelines and regulatory compliance.
- Integrated financial analytics aligns clinical insights with reimbursement and operational decision-making.

Disadvantages

- Privacy technologies (HE, SMPC) add computational overhead and latency.
- DP degrades model utility; choosing ε requires stakeholder negotiation and risk tolerance.
- Federated training requires robust client infrastructure and coordination across institutions.
- Mapping clinical features to billing codes can be error-prone and labor-intensive.

IV. RESULTS AND DISCUSSION

We performed experiments simulating 12 clinical sites with heterogeneous data distributions. The multimodal PEM-BERT outperformed single-modality baselines: mean AUC improved from 0.86 (image-only) to 0.90 (multimodal) across key diagnostic tasks. Financial cost estimation MAE decreased by ~12% when using multimodal features versus text-only features. Federated training with DP (ε=8) reduced membership inference attack success by 62% while lowering AUC by an average of 1.8 percentage points relative to non-private federated training. SMPC-based encrypted aggregation introduced a 2–5× runtime overhead but preserved accuracy comparable to secure aggregation baselines.

Governance metrics showed >95% lineage coverage for synthesized datasets and auditable model updates. Radiologist evaluations rated multimodal explanations as clinically useful in 78% of cases. Error analysis indicated that performance degradation under DP was concentrated in rare disease classes, suggesting targeted strategies (e.g., selective fine-tuning without DP or data augmentation) to mitigate utility loss. Financial analytics were sensitive to billing code mapping errors; introducing a coder-in-the-loop verification reduced MAE by 7%.

Cost profiling on a major cloud provider indicated that federated orchestration and encrypted aggregation dominated operational costs; however, amortized over model lifecycle and shared across institutions, total cost-per-update was acceptable for consortium settings. The results highlight practical trade-offs and point to hybrid deployment options: local on-prem inference with cloud-hosted governance and analytics.

IV. CONCLUSION

PEM-BERT demonstrates that multimodal BERT architectures can be adapted for medical imaging with privacy-preserving training and cloud-native governance, while enabling integrated financial analytics. The layered privacy approach substantially reduces leakage risk with modest utility trade-offs; cloud-managed governance simplifies compliance and auditability. For real-world adoption, institutions must weigh privacy budgets, computational overheads, and the investment required for billing integration. We provide guidelines for deployment, emphasizing human oversight, provenance, and controlled use of encryption for sensitive aggregates.

V. FUTURE WORK

- Evaluate PEM-BERT in real-world multi-hospital consortia with production EHR and PACS data.
- Optimize HE/SMPC protocols for lower-latency encrypted aggregation tailored to imaging features.
- Investigate adaptive DP techniques that protect individual privacy while preserving rare-class utility.
- Extend financial analytics to longitudinal cost modeling and outcome-based payments.
- Integrate federated model monitoring and drift detection with automated governance workflows.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

REFERENCES

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
2. Poornima, G., & Anand, L. (2025). Medical image fusion model using CT and MRI images based on dual scale weighted fusion based residual attention network with encoder-decoder architecture. Biomedical Signal Processing and Control, 108, 107932.
3. Reddy, B. V. S., & Sugumar, R. (2025, June). COVID19 segmentation in lung CT with improved precision using seed region growing scheme compared with level set. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020154). AIP Publishing LLC.
4. Konda, S. K. (2023). Strategic planning for large-scale facility modernization using EBO and DCE. International Journal of Artificial Intelligence in Engineering, 1(1), 1–11. https://doi.org/10.34218/IJAIE_01_01_001
5. Gonepally, S., Amuda, K. K., Kumbum, P. K., Adari, V. K., & Chunduru, V. K. (2023). Addressing supply chain administration challenges in the construction industry: A TOPSIS-based evaluation approach. Data Analytics and Artificial Intelligence, 3(1), 152–164.
6. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). Fusion: Practice & Applications, 14(2).
7. Alsentzer, E., Murphy, J. R., Boag, W., Weng, W.-H., Jin, D., Naumann, T., & McDermott, M. (2019). Publicly available clinical BERT embeddings. *arXiv preprint arXiv:1904.03323*.
8. Alyasiri, S., Rana, O., & Krishna, C. (2020). Encrypted machine learning: A survey. *Journal of Cloud Computing*, 9(1), 1–29.
9. Sardana, A., Kotapati, V. B. R., & Ponnoju, S. C. (2025). Autonomous Audit Agents for PCI DSS 5.0: A Reinforcement Learning Approach. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 4(1), 130-136.
10. Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations (SimCLR). *Proceedings of the 37th International Conference on Machine Learning*, PMLR 119:1597–1607.
11. KESAVAN, E. (2025). THE FUTURE OF WORK: TRENDS AND IMPLICATIONS FOR MANAGEMENT. i-Manager's Journal on Management, 19(4).
12. Manda, P. (2023). A Comprehensive Guide to Migrating Oracle Databases to the Cloud: Ensuring Minimal Downtime, Maximizing Performance, and Overcoming Common Challenges. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(3), 8201-8209.
13. Thambireddy, S., Bussu, V. R. R., & Joyce, S. (2023). Strategic Frameworks for Migrating Sap S/4HANA To Azure: Addressing Hostname Constraints, Infrastructure Diversity, And Deployment Scenarios Across Hybrid and Multi-Architecture Landscapes. Journal ID, 9471, 1297. https://www.researchgate.net/publication/396446597_Strategic_Frameworks_for_Migrating_Sap_S4HANA_To_Azure_Addressing_Hostname_Constraints_Infrastructure_Diversity_And_Deployment_Scenarios_Across_Hybrid_and_Multi-Architecture_Landscapes
14. Shashank, P. S. R. B., Anand, L., & Pitchai, R. (2024, December). MobileViT: A Hybrid Deep Learning Model for Efficient Brain Tumor Detection and Segmentation. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 157-161). IEEE.
15. Kumar, R., Christadoss, J., & Soni, V. K. (2024). Generative AI for Synthetic Enterprise Data Lakes: Enhancing Governance and Data Privacy. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 7(01), 351-366.
16. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT*, 4171–4186.
17. Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.
18. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.
19. Kandula, N. (2024). Optimizing Power Efficient Computer Architecture With A PROMETHEE Based Analytical Framework. J Comp Sci Appl Inform Technol, 9(2), 1-9.
20. Nguyen, P. A., Nguyen, P., Hawkins, R., & Haug, P. J. (2018). Predicting hospital length of stay and billing from clinical notes and admissions data. *Journal of Biomedical Informatics*, 84, 52–62.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 7, Issue 6, November – December 2025 |

DOI:10.15662/IJEETR.2025.0706009

21. Sivaraju, P. S. (2024). PRIVATE CLOUD DATABASE CONSOLIDATION IN FINANCIAL SERVICES: A CASE STUDY OF DEUTSCHE BANK APAC MIGRATION. ITEGAM-Journal of Engineering and Technology for Industrial Applications (ITEGAM-JETIA).
22. Radford, A., Kim, J. W., Hallacy, C., et al. (2021). Learning transferable visual models from natural language supervision. *Proceedings of the 38th International Conference on Machine Learning*, PMLR 139:8748–8763.
23. Gosangi, S. R. (2023). Reimagining Government Financial Systems: A Scalable ERP Upgrade Strategy for Modern Public Sector Needs. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8001-8005.
24. Anbalagan, B., & Pasumarthi, A. (2022). Building Enterprise Resilience through Preventive Failover: A Real-World Case Study in Sustaining Critical Sap Workloads. International Journal of Computer Technology and Electronics Communication, 5(4), 5423-5441.
25. Adari, Vijay Kumar, "Interoperability and Data Modernization: Building a Connected Banking Ecosystem," International Journal of Computer Engineering and Technology (IJCET), vol. 15, no. 6, pp.653-662, Nov-Dec 2024. DOI:https://doi.org/10.5281/zenodo.14219429.
26. Kiran, A., Rubini, P., & Kumar, S. S. (2025). Comprehensive review of privacy, utility and fairness offered by synthetic data. IEEE Access.
27. Reddy, B. T. K., & Sugumar, R. (2025, June). Effective forest fire detection by UAV image using Resnet 50 compared over Google Net. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020274). AIP Publishing LLC.
28. Kakulavaram, S. R. (2023). Performance Measurement of Test Management Roles in 'A' Group through the TOPSIS Strategy. International Journal of Artificial intelligence and Machine Learning, 1(3), 276. https://doi.org/10.55124/jaim.v1i3.276
29. Archana, R., & Anand, L. (2025). Residual u-net with Self-Attention based deep convolutional adaptive capsule network for liver cancer segmentation and classification. Biomedical Signal Processing and Control, 105, 107665.
30. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 234–241.