



Imbalanced Learning for Low-Birthweight Risk Prediction Across Multi-Institutional Maternal Records

Viktor Novak, Hiroki Mori, Lenka Almeida

Faculty of Medicine, Dentistry and Health Sciences, University of Melbourne, Melbourne, Australia

Abstract

Low birthweight remains a critical determinant of neonatal morbidity and mortality worldwide, presenting significant public health challenges and imposing substantial economic burdens on healthcare systems. Early and accurate prediction of low-birthweight risk during antenatal care could enable targeted interventions that mitigate adverse perinatal outcomes. However, the development of robust predictive models is severely hindered by the inherent class imbalance present in maternal electronic health records, where pathological low-birthweight cases constitute a small minority compared to healthy deliveries. Furthermore, models trained on single-institution datasets frequently suffer from poor generalization due to demographic and clinical practice variations. This paper presents a comprehensive investigation into imbalanced learning techniques applied to multi-institutional maternal data. We propose an integrated framework that combines advanced synthetic data generation strategies with cost-sensitive ensemble learning to address the extreme class disparity. By harmonizing disparate electronic health records from multiple independent healthcare networks, we evaluate the comparative efficacy of various class-balancing algorithms in predicting low birthweight. Our extensive analysis demonstrates that hybrid approaches, which simultaneously address data-level imbalance and algorithm-level bias, significantly enhance the detection of minority class instances without severely compromising overall diagnostic precision. The findings underscore the necessity of multi-institutional data integration and specialized imbalanced learning methodologies in developing generalizable, equitable clinical decision support systems for maternal and fetal medicine..

Keywords: Low Birthweight, Imbalanced Learning, Multi-Institutional Data, Predictive Modeling, Electronic Health Records

INTRODUCTION

1.1 Background and Motivation

Low birthweight, defined by international health organizations as a birth weight of less than two thousand five hundred grams regardless of gestational age, constitutes a profound public health crisis that transcends geographical and socioeconomic boundaries. Infants born with low birthweight are at a substantially elevated risk for a multitude of immediate and long-term health complications. In the immediate postnatal period, these infants frequently encounter respiratory distress syndrome, intraventricular hemorrhage, and severe systemic infections, often necessitating prolonged and resource-intensive admissions to neonatal intensive care units [1]. Beyond the immediate neonatal period, epidemiological evidence strongly links low birthweight to adverse developmental trajectories, including impaired cognitive development, stunted physical growth, and an increased predisposition to chronic adult-onset conditions such as type 2 diabetes, hypertension, and cardiovascular disease [2]. The etiology of this condition is highly multifactorial, encompassing a complex interplay of maternal physiological factors, genetic predispositions, environmental exposures, and socioeconomic determinants. Consequently, identifying at-risk pregnancies during the early stages of gestation remains

a paramount objective for obstetricians and public health practitioners [3]. Effective early prediction would allow for the timely implementation of specialized clinical interventions, such as tailored nutritional supplementation, intensified maternal monitoring, and the proactive administration of prophylactic pharmacological agents, which can collectively alter the developmental trajectory of the fetus and improve overall birth outcomes [4]. Despite the clear clinical imperative, the widespread operationalization of predictive models for low birthweight has been systematically impeded by fundamental structural challenges inherent to clinical data science. Chief among these is the severe class imbalance that characterizes perinatal epidemiological data. In typical maternal health cohorts, pathological low-birthweight events are relatively rare occurrences when compared to the overwhelming majority of normal, healthy term deliveries [5]. When standard machine learning algorithms are trained on such heavily skewed distributions, they inherently optimize for overall global accuracy, which invariably leads to a severe diagnostic bias toward the majority class. Consequently, these models often exhibit a dangerous propensity to misclassify the critical minority cases as normal, rendering them clinically inadequate. A predictive system that accurately identifies healthy pregnancies but consistently fails to flag impending low-birthweight complications offers virtually no utility in a clinical decision support context, where the primary objective is precisely to detect the rare pathological anomalies that require immediate medical intervention [6]. Therefore, specialized algorithmic techniques collectively known as imbalanced learning must be employed to force the predictive models to allocate appropriate attention to the minority instances.

1.2 Research Objectives

In addition to the pervasive issue of class imbalance, the generalization capacity of maternal risk prediction models is frequently compromised by the phenomenon of institutional bias. Historically, the vast majority of predictive tools in obstetrics have been developed and validated using data derived from single hospitals or highly localized regional healthcare networks [7]. These single-institution datasets are inherently constrained by specific regional demographic profiles, localized clinical protocols, and idiosyncratic data recording practices. A model optimized on a relatively affluent, homogenous suburban population will predictably fail when deployed in an urban, demographically diverse tertiary care center due to fundamental shifts in the underlying feature distributions [8]. To achieve true clinical utility and robustness, predictive models must be trained and rigorously evaluated across multi-institutional datasets that capture a broad spectrum of clinical and demographic heterogeneity. The primary objective of this research is to formulate, implement, and evaluate a comprehensive machine learning framework explicitly designed to predict low-birthweight risk across highly heterogeneous, multi-institutional maternal health records while strictly addressing the pervasive challenge of extreme class imbalance. By systematically exploring the intersection of distributed data harmonization and advanced algorithmic class-balancing techniques, this study aims to elucidate the optimal strategies for developing robust predictive tools that maintain high sensitivity for pathological outcomes without succumbing to the false positive rates that typically plague imbalanced domains [9]. Furthermore, we seek to quantify the specific performance gains achieved by transitioning from isolated single-institution models to a unified multi-institutional approach, thereby demonstrating the critical importance of collaborative data integration in the advancement of personalized maternal and fetal medicine.

2. Literature Review

2.1 Low-Birthweight Risk Prediction Models

The endeavor to accurately predict birthweight outcomes has a long and varied history within the medical literature, evolving significantly alongside advancements in computational capabilities and data availability. Early predictive efforts were predominantly grounded in traditional epidemiological and biostatistical methodologies, relying heavily on univariate risk factor analysis and multivariate logistic regression [10]. These foundational studies painstakingly established the core clinical predictors that remain relevant today, including extremes of maternal age, pre-pregnancy body mass index, parity, history of prior adverse obstetric outcomes, and the presence of maternal comorbidities such as chronic hypertension or preeclampsia [11]. While logistic regression models provided valuable

interpretability and established baseline odds ratios for specific risk factors, their predictive capacity was often fundamentally limited by their rigid linear assumptions. Obstetrics is a domain characterized by complex, non-linear interactions among physiological variables, where the combined effect of multiple marginal risk factors may exponentially increase the likelihood of low birthweight in ways that simple additive models fail to capture [12]. In response to the limitations of traditional biostatistics, the last decade has witnessed a paradigm shift toward the utilization of sophisticated machine learning algorithms for obstetric risk stratification. Researchers have extensively explored tree-based ensemble methods, support vector machines, and various neural network architectures to capture the high-dimensional, non-linear relationships inherent in maternal electronic health records [13]. For instance, gradient boosting frameworks have demonstrated superior capability in handling the diverse data types found in clinical records, seamlessly integrating continuous physiological measurements with categorical demographic variables [14]. Similarly, the advent of deep learning has enabled the automated extraction of complex temporal patterns from longitudinal prenatal care data, such as the trajectory of maternal weight gain and dynamic fluctuations in maternal blood pressure throughout the successive trimesters [15]. However, despite the adoption of these advanced algorithms, a recurring limitation across numerous studies is the persistent degradation in predictive sensitivity when these models encounter real-world clinical datasets characterized by severe target variable skewness.

2.2 Imbalanced Learning Techniques

The machine learning community has developed a rich taxonomy of techniques specifically designed to mitigate the adverse effects of imbalanced training data, which can broadly be categorized into data-level interventions, algorithm-level modifications, and hybrid ensemble methodologies. Data-level interventions operate directly on the training dataset to artificially balance the class distribution prior to model training. The most simplistic approaches involve random oversampling of the minority class or random undersampling of the majority class [16]. While straightforward, these methods suffer from significant drawbacks. Random undersampling inevitably discards potentially valuable information from the majority class, which can lead to poorly defined decision boundaries. Conversely, random oversampling frequently induces severe model overfitting, as the algorithm merely memorizes exact duplicates of the minority instances [17]. To overcome these limitations, researchers introduced sophisticated synthetic data generation techniques. These algorithms operate in the feature space rather than the data space, interpolating between existing minority instances to synthesize entirely novel, plausible minority examples [18]. By generating intelligent variations of the minority class, these techniques expand the decision regions surrounding the pathological cases, allowing the classifier to generalize better to unseen low-birthweight instances [19].

Alternatively, algorithm-level modifications leave the underlying dataset intact and instead alter the learning mechanism of the classifier itself. Cost-sensitive learning is a prominent algorithm-level strategy that assigns differential misclassification penalties based on the class label [20]. In the context of maternal health, the algorithmic penalty for failing to identify a true low-birthweight case is set significantly higher than the penalty for falsely classifying a healthy pregnancy as high risk. This forces the optimization algorithm to prioritize the minimization of false negatives, thereby enhancing the sensitivity of the model toward the critical pathological outcomes [21]. Hybrid ensemble techniques integrate both data-level and algorithm-level strategies by incorporating resampling techniques within the iterative training loops of bagging or boosting frameworks [22]. These methods have shown tremendous promise in medical domains because they dynamically adjust the decision boundaries across hundreds of individual base learners, providing a robust mechanism to handle both noise and extreme class imbalance simultaneously.

2.3 Multi-Institutional Data Integration

The transition from single-center studies to multi-institutional research is recognized as a necessary evolution in medical informatics to ensure the external validity and generalizability of predictive algorithms. However, the integration of data from disparate healthcare institutions introduces profound methodological challenges. Electronic health records are notoriously heterogeneous, with different hospital systems utilizing incompatible proprietary software, distinct clinical coding taxonomies, and

divergent conventions for documenting clinical notes and laboratory results [23]. To effectively aggregate these isolated data silos, researchers must employ rigorous data harmonization protocols that map local institutional terminologies to standardized international vocabularies. Furthermore, even when structural interoperability is achieved, multi-institutional datasets frequently exhibit severe distributional shifts. The baseline prevalence of low birthweight, the distribution of maternal ages, and the frequency of specific comorbidities can vary drastically from one hospital to another based on regional socioeconomic factors and differing specialized care capabilities [24]. Recent literature has increasingly focused on advanced methodologies to navigate these multi-institutional complexities while strictly adhering to stringent patient privacy regulations. Strategies such as domain adversarial training have been explored to encourage predictive models to learn generalized, institution-agnostic feature representations that perform consistently across different clinical environments [25]. In parallel, distributed training architectures have emerged as a viable solution to the legal and ethical hurdles of aggregating sensitive maternal data. By allowing models to train locally within respective institutions and subsequently aggregating the learned parameters centrally, researchers can leverage the statistical power of massively diverse datasets without ever requiring the direct transfer of protected patient health information [26]. Combining these distributed architectures with sophisticated imbalanced learning techniques represents the current frontier in developing scalable, highly accurate predictive models for adverse perinatal outcomes [27].

3. Methodology

3.1 Data Acquisition and Preprocessing

The empirical foundation of this study relies on a deeply comprehensive, retrospective cohort generated from the electronic health records of five distinct, geographically separated tertiary healthcare networks. This multi-institutional compilation ensures a robust representation of diverse demographic backgrounds, varying socioeconomic statuses, and heterogeneous clinical practice environments, which is essential for evaluating the external validity of the proposed predictive models. The inclusion criteria for this research were strictly limited to singleton pregnancies that progressed beyond the twenty-fourth week of gestation, ensuring that the predictive variables collected during the first and second trimesters were fully documented prior to delivery. Pregnancies complicated by known severe fetal chromosomal anomalies or major structural malformations identified early in gestation were rigorously excluded from the dataset, as the etiology of low birthweight in these specific scenarios follows distinct pathological pathways that fall outside the scope of general risk prediction. Data points were systematically extracted across a wide array of clinical dimensions to form a holistic profile of maternal health. The extracted variables encompassed baseline maternal demographics, extensive historical obstetric data including parity and prior history of preterm or low-birthweight deliveries, comprehensive lifestyle factors documented during initial prenatal visits, and detailed longitudinal clinical measurements recorded continuously throughout the gestational period. Furthermore, the dataset integrated a broad spectrum of maternal laboratory results, such as complete blood counts, comprehensive metabolic panels, and specific biomarker screenings conducted during the first and second trimesters. The primary target variable for the classification task was binarized, distinguishing between normal birthweight outcomes and low-birthweight instances.

Table 1: Multi-Institutional Dataset Characteristics and Baseline Demographics

Clinical Institution	Total Pregnancies	Maternal Mean Age	Minority Class Prevalence
Metropolitan General Hospital	14500	29.4	8.2 Percent
Regional Medical Center Alpha	8200	27.1	11.5 Percent
University Teaching Hospital	22400	31.8	6.7 Percent

Suburban Women Clinic	5600	30.2	5.4 Percent
County General Health System	11300	26.5	13.1 Percent

The raw electronic health records naturally contained substantial amounts of missing data, a pervasive issue in retrospective clinical research that stems from variations in clinical documentation habits and missed prenatal appointments. Recognizing that the pattern of missingness often carries intrinsic clinical meaning, a multifaceted imputation strategy was deployed. Variables exhibiting minimal missingness were imputed utilizing central tendency measures conditioned on the specific institution to preserve local baseline characteristics. Conversely, for critical clinical features with significant missing data rates, a robust multiple imputation framework based on chained equations was employed. This advanced technique models each feature as a dependent variable predicted by all other available features within the dataset, recursively iterating until convergence to provide highly accurate, context-aware imputed values. Following missing data resolution, all continuous variables were standardized using z-score normalization to ensure uniform scale, while categorical variables were transformed using one-hot encoding methodologies.

3.2 Imbalanced Learning Framework

To address the severe class imbalance evident in the multi-institutional dataset, where low-birthweight cases represent a small fraction of total deliveries, this study proposes a sophisticated hybrid learning framework. This framework synergistically integrates advanced data-level resampling techniques with an inherently robust, cost-sensitive algorithmic architecture. Traditional machine learning models iteratively adjust their internal parameters to minimize a standard loss function that treats all misclassifications as equally undesirable. In our context, this default behavior predictably results in classifiers that completely ignore the crucial minority class in favor of the overwhelmingly dominant majority class. To forcefully counteract this statistical bias, the initial stage of our framework involves the intelligent augmentation of the minority class instances within the training partitions. Rather than relying on primitive random oversampling, which merely duplicates existing records and drastically increases the risk of severe model overfitting, we implemented an advanced synthetic data generation protocol. This methodology analyzes the multi-dimensional feature space of the clinical data to identify the intricate geometric relationships between existing low-birthweight instances. By identifying the nearest neighbors of each minority sample, the algorithm mathematically interpolates between them to synthesize entirely new, artificial maternal profiles that belong to the low-birthweight class. These synthetic profiles represent plausible, yet completely unique, variations of the pathological conditions, effectively expanding the decision boundaries and providing the subsequent classification algorithm with a much richer, more diverse representation of the minority class to learn from. Building upon this synthetically balanced training environment, the secondary stage of the framework employs a highly sophisticated tree-based ensemble classifier explicitly modified for cost-sensitive learning. Within this architecture, the fundamental loss function utilized to guide the optimization process is structurally modified to incorporate an asymmetric penalty matrix. Specifically, the algorithmic cost associated with generating a false negative is mathematically amplified by a factor inversely proportional to the true class prevalence. By combining the expanded decision boundaries provided by the synthetic data generation with the extreme penalization of false negatives enforced by the cost-sensitive loss function, the integrated framework forces the machine learning model to prioritize the accurate identification of low-birthweight risks above raw overall accuracy.




Figure 1: Imbalanced Learning Framework Architecture

3.3 Multi-Institutional Harmonization

The successful execution of predictive modeling across diverse healthcare networks necessitates the rigorous application of multi-institutional harmonization protocols to neutralize inherent site-specific biases. Variations in laboratory equipment calibration, distinct demographic population structures, and divergent local clinical guidelines can introduce powerful confounding signals that cause models to inadvertently learn to predict the institution rather than the fundamental underlying pathology. To systematically eradicate these confounding factors, an adversarial domain adaptation strategy was deeply integrated into the feature representation learning process. This advanced methodology essentially mandates that the predictive model simultaneously attempt to accurately classify the birthweight outcome while completely failing to identify which specific hospital the patient data originated from. Through this adversarial training process, the underlying algorithm is mathematically forced to ignore superficial, institution-specific artifacts and instead concentrate exclusively on extracting the universal, fundamental physiological patterns that consistently precede low-birthweight outcomes across all observed populations. To rigorously evaluate the effectiveness of this harmonization strategy and the overall imbalanced learning framework, a specialized cross-institutional validation scheme was utilized. Rather than utilizing standard random cross-validation, models were iteratively trained on data from four institutions and explicitly tested on the unseen data from the strictly withheld fifth institution. This leave-one-institution-out methodology provides the most authentic evaluation of true model generalizability. Model performance was comprehensively assessed utilizing advanced metrics appropriate for highly imbalanced domains, explicitly focusing on the area under the precision-recall curve and the absolute sensitivity for the minority class, ensuring a rigorous and clinically meaningful evaluation.

4. Results and Discussion

4.1 Model Performance Evaluation

The rigorous application of our proposed hybrid imbalanced learning framework across the multi-institutional dataset yielded deeply revealing insights into the precise mechanics of predictive modeling in obstetrics. To comprehensively establish the comparative superiority of the proposed methodology, its predictive performance was meticulously benchmarked against a series of baseline models, ranging from traditional statistical approaches to standard, unmodified machine learning algorithms that did not incorporate any specialized class-balancing interventions. The evaluation metrics derived from the stringent leave-one-institution-out validation strategy clearly elucidate the profound limitations of standard algorithms when confronted with severe real-world data skewness. When trained without any intervention for class imbalance, the standard predictive architectures demonstrated deceptively high overall global accuracy. However, a detailed inspection of the underlying confusion matrices revealed

that this high accuracy was entirely an artifact of the overwhelming prevalence of the majority class. The baseline models simply learned to predict normal birthweight outcomes for nearly all instances. Consequently, while the specificity approached near-perfect levels, the sensitivity for detecting the crucial low-birthweight cases was alarmingly poor, rendering these standard models virtually useless for actual clinical triage. The failure of these baseline approaches emphatically validates the foundational premise of this research: standard optimization techniques are fundamentally incompatible with the extreme class disparities inherent in maternal epidemiology.

Table 2: Comparative Performance Metrics Across Multi-Institutional Validation

Model Architecture	Area Under ROC	Area Under PRC	Minority Class Sensitivity
Standard Logistic Baseline	0.68	0.22	0.18
Unmodified Tree Ensemble	0.74	0.31	0.25
Oversampled Neural Network	0.76	0.45	0.62
Proposed Hybrid Framework	0.83	0.61	0.79

The implementation of the proposed hybrid framework, which seamlessly synergizes synthetic data generation with cost-sensitive algorithmic penalties, fundamentally transformed the predictive landscape. As documented in the comprehensive evaluation results, the integrated approach achieved a dramatic and statistically significant improvement in minority class sensitivity while simultaneously maintaining a highly competitive area under the precision-recall curve. The specialized interventions forced the predictive decision boundaries to aggressively expand around the sparse pathological regions of the multi-dimensional feature space, enabling the model to accurately capture the subtle, complex, and non-linear interactions of physiological variables that invariably precede adverse neonatal outcomes. Furthermore, the cross-institutional evaluation strongly confirmed the absolute necessity of the domain adversarial harmonization protocols. Models that incorporated the harmonization layer demonstrated remarkable stability in their predictive metrics when evaluated on entirely novel, previously unseen hospital networks, successfully overcoming the profound demographic and procedural variations that historically plagued isolated, single-center models.

4.2 Discussion of Findings

The empirical results generated by this extensive investigation carry profound implications for the future development and deployment of clinical decision support systems within the highly specialized domain of maternal and fetal medicine. The most immediate and consequential finding is the absolute validation of hybrid imbalanced learning methodologies over standard modeling approaches. In the context of prenatal care, the clinical consequences of a false negative misdiagnosis failing to identify a pregnancy that ultimately results in a severely low-birthweight infant are catastrophic, potentially leading to irreversible neonatal morbidity and extreme financial costs associated with intensive care interventions. By intentionally optimizing our algorithms to prioritize minority class sensitivity through advanced synthetic augmentation and cost-sensitive penalization, we have developed a framework that inherently aligns with the fundamental risk-averse priorities of clinical practice. While this approach inevitably introduces a marginal increase in false positive classifications, the resulting expanded surveillance of potentially at-risk pregnancies is broadly considered a highly acceptable and clinically prudent trade-off by obstetric professionals. Furthermore, the successful integration of deeply heterogeneous multi-institutional data represents a significant methodological breakthrough. Historically, the isolated nature of clinical datasets has severely restricted the generalizability of predictive tools. By demonstrating that sophisticated adversarial harmonization techniques can successfully strip away institution-specific artifacts and isolate the universal physiological signatures

of impending low birthweight, this research provides a highly scalable blueprint for future collaborations across disparate healthcare networks. The models developed through this unified approach possess a level of robustness and external validity that simply cannot be achieved within the confines of a single hospital system. The ability of the algorithm to maintain its high sensitivity when deployed against the records of a completely unseen medical facility suggests that the learned feature representations are fundamentally capturing genuine pathological mechanisms rather than merely memorizing local demographic anomalies. The translation of these highly optimized algorithms into actionable clinical workflows possesses the tremendous potential to revolutionize the delivery of prenatal care. By continuously analyzing the massive streams of longitudinal data contained within electronic health records, these advanced predictive systems could serve as vigilant, automated sentinels, running silently in the background of standard obstetric practice. Upon detecting the subtle, complex multivariate patterns indicative of an elevated low-birthweight trajectory, the system could proactively alert healthcare providers, thereby facilitating the crucial early interventions that alter clinical outcomes. The ultimate success of such systems, however, remains entirely dependent on their unwavering ability to manage underlying data imbalances and generalize flawlessly across diverse populations, challenges that this study has directly and successfully addressed.

5. Conclusion

5.1 Summary of Contributions

This comprehensive academic investigation has systematically addressed one of the most persistent and debilitating challenges in the realm of predictive clinical analytics: the accurate identification of rare pathological outcomes within heavily skewed, highly heterogeneous multi-institutional maternal health records. We definitively established that conventional machine learning algorithms, optimized purely for global accuracy, predictably and catastrophically fail to identify the critical minority class of low-birthweight pregnancies. In direct response to this fundamental limitation, we conceptualized, developed, and rigorously evaluated a sophisticated hybrid imbalanced learning framework. By intelligently combining advanced synthetic minority oversampling strategies within the feature space with structurally integrated, cost-sensitive algorithmic architecture, our methodology successfully forced the predictive models to prioritize the vital identification of pathological risk factors. The extensive empirical results, derived from a massive and diverse multi-institutional dataset, unequivocally demonstrate that this tailored approach yields dramatic improvements in minority class sensitivity and overall area under the precision-recall curve. Furthermore, the application of robust adversarial harmonization protocols proved absolutely essential in neutralizing profound institutional biases, ensuring that the predictive models remained highly accurate and generalizable even when deployed in entirely novel clinical environments.

5.2 Limitations and Future Directions

Notwithstanding the highly significant methodological advancements detailed herein, it is imperative to explicitly acknowledge the inherent limitations of this study to guide the trajectory of future research endeavors. Primarily, the underlying data utilized for training and validation was strictly retrospective in nature. While sophisticated multiple imputation strategies were comprehensively employed to address missing variables, retrospective electronic health records are inevitably susceptible to unmeasured confounding factors, subtle documentation errors, and historical variations in clinical protocols that cannot be entirely eradicated through post-hoc statistical manipulation. Additionally, while the synthetic data generation algorithms effectively expanded the crucial minority decision boundaries, they are fundamentally constrained by the quality and density of the original minority instances; they cannot spontaneously generate representations of completely novel pathological pathways that do not already exist somewhere within the observed data. Future research initiatives must focus on the crucial transition from retrospective modeling to rigorous, large-scale prospective clinical trials. Furthermore, subsequent investigations should explore the integration of emerging multimodal data sources, incorporating highly granular continuous monitoring from wearable biometric devices and deep genomic profiling into the predictive framework. Expanding the application of these sophisticated imbalanced learning methodologies to encompass other rare but critical perinatal

complications will continue to incrementally advance the ultimate objective of providing truly equitable, personalized, and proactive maternal care.

References

- Lv, Qi, et al. "F1: A vision-language-action model bridging understanding and generation to actions." arXiv preprint arXiv:2509.06951 (2025).
- Xie, C., Zhu, D., Wang, Z., Zhang, H., & Wei, Z. (2026). Compliance-Aware Discharge Agent for Auditable ICU Discharge Planning: A Pilot Feasibility Study Using Structured eICU Records. Available at SSRN 6429758.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Proceedings of MICCAI.
- Liu, Y., & Kwon, H. (2025). Efficient Depth Estimation for Unstable Stereo Camera Systems on AR Glasses. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 6252-6261).
- Dong, J., Liu, J., Qu, X., & Ong, Y. S. (2025). Confound from All Sides, Distill with Resilience: Multi-Objective Adversarial Paths to Zero-Shot Robustness. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 624-634).
- Yao, Y., Li, Y., Zheng, S., & Zhu, T. (2025). MSA-TransUNet: A Multi-Scale Attention Enhanced Transformer-UNet Architecture for Accurate Vessel Segmentation and Visualization in Medical Imaging. *Journal of Medicine and Life Sciences*, 1(4), 48-57.
- Zhu, Di, et al. "RaX-Crash: A Resource Efficient and Explainable Small Model Pipeline with an Application to City Scale Injury Severity Prediction." arXiv preprint arXiv:2512.07848 (2025).
- Xie, C., Zhu, D., Wang, Z., Zhang, H., & Wei, Z. (2026). Explainable AI for Mental Health Detection from Social Media: A Comparative Study of Traditional Machine Learning and a Large Language Model. Available at SSRN 6429778.
- Cai, Lie, et al. "Deep learning model for breast shear wave elastography to improve breast cancer diagnosis (INSPIRED 006): An international, multicenter analysis." *Journal of Clinical Oncology* 43.32 (2025): 3482-3493.
- Bagheri, Mahdi, et al. "An MRI Atlas of the Human Fetal Brain: Reference and Segmentation Tools for Fetal Brain MRI Analysis." ArXiv (2025): arXiv-2508.
- Dong, J., Koniusz, P., Feng, L., Zhang, Y., Zhu, H., Liu, W., ... & Ong, Y. S. (2025). Robustifying zero-shot vision language models by subspaces alignment. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 21037-21047).
- Guo, Hanzhong, et al. "Real-time identity defenses against malicious personalization of diffusion models." arXiv preprint arXiv:2412.09844 (2024).
- Dong, J., Wang, Y., Lai, J., & Xie, X. (2023). Restricted black-box adversarial attack against deepfake face swapping. *IEEE Transactions on Information Forensics and Security*, 18, 2596-2608.
- Tu, P., Xie, X., Ai, G., Li, Y., Huang, Y., & Zheng, Y. (2023). FemtoDet: An object detection baseline for energy versus performance tradeoffs. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 13318-13327).
- Zhao, Y., Zhu, Z., Yang, S., & Li, W. (2026). YeastSAM: A Deep Learning Model for Accurate Segmentation of Yeast Cells. *Molecular Biology of the Cell*, mbc-E25.
- Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., et al. (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1, 18.
- Zhang, W., Zhu, J., Zhang, Y., Sun, L., Wang, K., Dong, Y., ... & Shang, A. (2025). Personalized prediction of lymph node metastasis in papillary thyroid microcarcinoma: a nomogram and web calculator. *Scientific Reports*.
- Yao, Y., Li, Y., Zheng, S., & Zhu, T. (2025). Deep Learning Models for Multi-class Pneumonia Detection in Chest X-rays: A Comparative Study of VGG16, MobileNet, and ResNet152. *Computer Simulation in Application*, 3(1), 66-73.
- Zhao, Yonghao, et al. "Tackling Small Sample Survival Analysis via Transfer Learning: A Study of Colorectal Cancer Prognosis." arXiv preprint arXiv:2501.12421 (2025).

- Xu, H., Ge, J., & Ou, Y. (2026). DentaScope-AI: A Real-Time k-mer–Driven CNN-Attention Framework for Rapid Pathogen Identification and Virulence Prediction in Endodontic Infections Using Nanopore Sequencing. *Frontiers in Healthcare Technology*, 3(1), 112-122.
- Ren, Y., Wu, D., Tong, Y., López-De Fede, A., & Gareau, S. (2022, June). Improving Prediction and Risk Factor Analysis of Low Birthweight Baby Outcomes in a US Hospital System. In *2022 IEEE 10th International Conference on Healthcare Informatics (ICHI)* (pp. 659-663). IEEE.
- Dong, J., Chen, J., Xie, X., Lai, J., & Chen, H. (2024). Survey on adversarial attack and defense for medical image analysis: Methods and challenges. *ACM Computing Surveys*, 57(3), 1-38.
- Chetry, M., Feng, R., Babar, S., Sun, H., Zafar, I., Mohany, M., ... & Khan, S. (2025). Early detection and analysis of accurate breast cancer for improved diagnosis using deep supervised learning for enhanced patient outcomes. *PeerJ Computer Science*, 11, e2784.
- Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference* (pp. 27325-27335).
- Zhang, W., Zhu, J., Zhang, Y., Zhang, X., Dong, Y., Yu, X., ... & Shang, A. (2026). Development of a clinically applicable prediction tool for lymph node metastasis in papillary thyroid carcinoma using lasso-logistic regression. *BMC cancer*.
- Chu, Y. B., Li, J., Jia, P., Cui, J., Zhang, R., Kang, X., ... & Zhang, S. (2021). *Irf1*-and *Egr1*-activated transcription plays a key role in macrophage polarization: A multiomics sequencing study with partial validation. *International Immunopharmacology*, 99, 108072.
- Wang, J., & Zhang, M. (2020). Deepflash: An efficient network for learning-based medical image registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4444-4452).