Software Engineering and Artificial Intelligence Systems

Volume: 02 Issue: 01 (2026)

www.seais.org/

Robust Rigid Bone Registration Using Depth-Constrained Surface Optimization for Head-Mounted AR Surgical Systems

Pierre Dubois1, Antoine Moreau2, Camille Laurent3*
Université Paris-Saclay, Faculty of Medicine, Le Kremlin-Bicêtre 94270, France
***Corresponding author:** camille.laurent@univ-paris-saclay.fr

Key word: Rigid registration; orthopedic AR navigation; depth-based tracking; curvature feature extraction; surface optimization; markerless guidance

A B S T R A C T

Rigid anatomical structures such as vertebrae and long bones require sub-millimeter registration accuracy in AR-guided orthopedic procedures. This work introduces a depth-constrained surface optimization method that enhances markerless rigid alignment using curvature-aware feature matching. A two-stage pipeline is designed. First, curvature-weighted keypoints are extracted from depth maps using multi-scale Laplacian surface analysis. Second, a robust RANSAC-enhanced rigid transformation model is optimized with a depth-consistency constraint to reject outliers caused by occlusion. Evaluation was performed on 18 cadaveric spine datasets and 10 synthetic CT-to-depth alignment tasks. Mean rotational error decreased from 1.92° to 0.84°,and translational error reduced from 2.16 mm to 0.93 mm compared with baseline depth-ICP. The approach achieved stable tracking at 45 FPS on an AR headset, supporting intraoperative bone navigation without fiducial markers.

INTRODUCTION

Rigid registration of bony anatomy is a fundamental requirement in augmented reality (AR)–assisted orthopedic surgery, where even small spatial deviations may lead to inaccurate implant positioning, suboptimal trajectory guidance, or reduced procedural safety. With the rapid development of head-mounted AR devices, surgeons are now able to visualize preoperative plans, three-dimensional anatomical models, and guidance cues directly within the operative field [1]. This capability has created new possibilities for intraoperative navigation by improving spatial perception and reducing attention shifts between the patient and external display systems. Recent studies have also shown that depth-enabled wearable AR platforms can support markerless registration and tracking during surgical guidance, highlighting their potential for streamlined intraoperative workflows without external fiducials [2]. Despite this progress, accurate and robust alignment between patient anatomy and virtual models remains one of the main technical barriers to reliable clinical deployment [3]. Compared with soft-tissue navigation, orthopedic AR guidance places stricter demands on registration precision because the target anatomy is rigid and the acceptable error tolerance is narrow. In procedures involving the spine, pelvis, or long bones, clinical usability often requires sub-millimeter translational accuracy and sub-degree rotational precision to ensure safe screw placement, osteotomy execution, and implant alignment [4]. Even minor angular or positional mismatch may be amplified during tool guidance, especially in minimally invasive workflows where direct

anatomical exposure is limited. This accuracy requirement distinguishes orthopedic navigation from many other AR applications and makes registration robustness a decisive factor in system performance. Optical marker–based tracking remains a common solution for achieving acceptable registration accuracy in image-guided surgery [5]. These systems can provide stable pose estimation under controlled conditions and have been widely used in research and commercial navigation platforms. However, their dependence on fiducial placement, external tracking infrastructure, and uninterrupted line-of-sight increases workflow complexity and may interfere with surgical access or sterile field management [6,7]. In practical operating environments, marker displacement, occlusion by instruments or personnel, and setup burden may all reduce usability. For this reason, markerless depth-based registration has received increasing attention in orthopedic AR navigation as a more flexible alternative that may better fit routine surgical workflow [8]. Integrated depth sensors in head-mounted displays offer a promising technical basis for markerless registration because they can capture local surface geometry rapidly and repeatedly during surgery [9,10]. Structured-light and time-of-flight technologies enable point cloud acquisition without ionizing radiation and support real-time interaction between the user and the surgical scene [11]. In orthopedic settings, these advantages are particularly relevant because exposed bone surfaces can provide geometrically informative structures for alignment with preoperative CT-derived models. Repeated depth acquisition also allows intraoperative re-registration when the viewing pose changes or when initial tracking confidence decreases. Among existing registration methods, the Iterative Closest Point (ICP) algorithm and its variants remain the most widely used approaches for rigid alignment between intraoperative surface scans and preoperative models [12]. ICP is attractive because of its conceptual simplicity, efficient implementation, and compatibility with dense geometric data. Nevertheless, conventional ICP is strongly affected by initialization error, incomplete overlap, surface occlusion, and nonuniform point sampling. These problems become more pronounced in orthopedic surgery, where the visible field is often narrow and anatomical features may be partially blocked by blood, retractors, soft tissue, or surgical tools [13]. Under such conditions, nearest-neighbor correspondences may become unreliable, causing convergence to local minima or unstable rigid transformations. Robust trimming and probabilistic weighting strategies have been introduced to reduce the influence of outliers, but their performance often deteriorates when only a limited portion of the bone surface is observable [14]. To improve initial alignment quality, feature-based rigid registration methods have been explored in parallel with dense correspondence refinement. Surface curvature descriptors, spin images, and geometry-aware embeddings can enhance matching stability under noise by emphasizing structurally informative regions of the anatomy [15]. Multi-scale geometric analysis is particularly attractive in spinal and orthopedic navigation because ridges, spinous processes, cortical edges, and other shape discontinuities can serve as relatively distinctive landmarks for correspondence estimation [16]. By focusing on these salient regions, feature-based methods may reduce ambiguity during coarse alignment and improve the convergence basin of later optimization stages. Even so, several limitations remain. Many feature extraction strategies have been developed for offline reconstruction, shape retrieval, or post-processing pipelines rather than wearable real-time AR systems. Their computational overhead may be difficult to accommodate on head-mounted devices, where power, memory, and latency budgets are limited. In addition, local curvature descriptors alone do not always provide sufficient discrimination on repetitive bony surfaces. Anatomical regions such as laminae, long cortical segments, or smooth cylindrical structures may exhibit similar local geometry, making incorrect correspondence assignment more likely when only descriptor similarity is considered. This issue is particularly relevant in constrained intraoperative views, where only a subset of anatomical landmarks is available. To strengthen registration robustness under incomplete visibility, recent studies have investigated depth-consistency constraints during optimization [17]. By enforcing geometric consistency along the viewing direction, depth-aware models can reduce alignment drift and suppress implausible transformations

caused by occlusion or partial overlap [18]. These constraints are especially useful in AR-guided surgery because the observed surface is inherently viewpoint-dependent, and reliable registration should remain compatible with the sensor acquisition geometry. However, many existing depth-consistent methods have been validated mainly on limited phantom models, controlled benchtop settings, or simulated data, while evidence from cadaveric or clinically realistic scenarios remains comparatively scarce [19,20]. As a result, their practical robustness under realistic anatomical complexity and intraoperative visual disturbance is still not fully established. Another gap in the current literature is the limited integration of geometric feature selection and depth-constrained optimization within a single wearable AR-oriented framework. In many markerless pipelines, dense ICP refinement remains the dominant component, while robust estimation is introduced only after correspondence generation rather than being embedded into the initialization stage itself. This design weakens the system's ability to reject poor matches early in the registration process. At the same time, studies that emphasize local features do not always incorporate explicit depth consistency during rigid transformation estimation. Consequently, the benefits of feature-guided initialization and viewpoint-aware geometric validation are often not jointly exploited. This separation reduces robustness when surface visibility is incomplete, correspondence quality is uneven, or the initial pose uncertainty is large. Current validation practice also leaves room for improvement. Experimental assessment is frequently performed on small datasets, limited phantoms, or simplified alignment tasks, with insufficient comparison against clinically meaningful baselines [21]. In orthopedic AR navigation, this is a significant limitation because real-world performance depends not only on average registration error but also on stability across viewpoints, tolerance to partial occlusion, and computational efficiency under real-time constraints. Wearable AR systems impose strict latency requirements, and a method that achieves good alignment but cannot maintain interactive performance may have limited translational value. Therefore, practical registration frameworks should be evaluated in terms of rotational accuracy, translational precision, robustness, and runtime in datasets that better reflect realistic surgical scenarios. Against this background, this study presents a depth-constrained surface optimization framework for rigid bone registration in head-mounted AR systems. The proposed method combines multi-scale Laplacian surface analysis with curvature-weighted keypoint extraction to identify geometrically informative regions for coarse correspondence establishment. A RANSAC-enhanced rigid transformation model with explicit depth-consistency constraints is then introduced to reject outliers associated with occlusion, incomplete exposure, and uneven surface sampling. By coupling discriminative geometric initialization with depth-aware transformation estimation, the framework is designed to improve alignment robustness without sacrificing the computational efficiency required for wearable deployment. The purpose of this work is to provide a practical markerless registration solution for AR-guided orthopedic navigation that is both accurate and compatible with real-time intraoperative use. Rather than relying primarily on dense refinement alone, the proposed framework emphasizes robust initialization, viewpoint-consistent optimization, and efficient implementation tailored to head-mounted devices. Its significance lies in three aspects: improving rigid registration reliability under limited surface visibility, strengthening the feasibility of markerless bone navigation in clinically relevant AR workflows, and offering an effective technical pathway toward more streamlined intraoperative guidance without external fiducials. Validation on cadaveric spine datasets and synthetic CT-to-depth alignment tasks is conducted to assess rotational accuracy, translational precision, stability, and computational efficiency. Through this evaluation, the study aims to demonstrate that depth-constrained and geometry-aware rigid registration can better support accurate, markerless, and workflow-friendly orthopedic AR navigation.

2.Materials and Methods

2.1 Sample and Study Site Description

The evaluation used 18 human cadaveric spine specimens and 10 synthetic CT-to-depth alignment datasets. The cadaveric samples included thoracic and lumbar segments without severe deformity or prior implants. Each specimen underwent high-resolution CT scanning with isotropic voxel spacing between 0.6 and 0.8 mm. During testing, the specimens were placed on a surgical table under standard operating room lighting. Surface depth data were collected using a head-mounted AR device equipped with a time-of-flight sensor. The synthetic datasets were created by rendering CT-based bone models into depth images under different viewpoints and controlled noise conditions. This dataset design allowed assessment under both realistic anatomical settings and controlled simulation environments.

2.2 Experimental Design and Control Setup

The depth-constrained surface optimization method was compared with a baseline depth-based Iterative Closest Point (ICP) algorithm. The experimental group applied curvature-weighted keypoint selection followed by RANSAC-based rigid alignment with a depth-consistency constraint. The control group used a standard point-to-plane ICP method without curvature weighting or depth filtering. Both approaches were initialized using the same coarse alignment based on anatomical bounding boxes. Multiple camera viewpoints were tested for each dataset to simulate intraoperative motion and partial occlusion. Evaluation metrics included rotational error, translational error, repeatability across trials, and processing speed. This design enabled direct comparison of accuracy and robustness under identical conditions.

2.3 Measurement Methods and Quality Control

Registration accuracy was assessed by comparing the estimated rigid transformation with a reference alignment obtained through fiducial-based calibration. Rotational error was measured in degrees using axis-angle representation. Translational error was calculated in millimeters using Euclidean distance. Tracking speed was recorded directly from the AR device during continuous operation. Each registration task was repeated five times with randomized initial offsets to evaluate stability. The depth sensor was calibrated before each session using a planar reference board to reduce systematic bias. Point clouds were preprocessed using statistical neighborhood filtering to remove noise and isolated outliers. All experiments were performed on the same hardware platform to ensure consistent timing measurement.

2.4 Data Processing and Model Formulation

Depth images were converted into three-dimensional point clouds using intrinsic camera projection parameters. Surface normals were estimated through local plane fitting within a fixed-radius neighborhood. Curvature saliency values were computed using a multi-scale Laplacian operator and applied as correspondence weights. The rigid transformation between source and target points was defined as

$$p'_i = Rp_i + t,$$

where R is the rotation matrix and t is the translation vector. Parameter estimation minimized the following objective function:

$$E(R,t) = \sum_{i=1}^{N} w_i \| Rp_i + t - q_i \|^2 + \lambda \sum_{i=1}^{N} \| d_i^{est} - d_i^{obs} \|^2,$$

where w_i denotes the curvature-based weight, p_i and q_i represent corresponding points, and d_i^{est} and d_i^{obs} indicate estimated and observed depth values along the viewing direction. The parameter λ controls the contribution of the depth-consistency term. RANSAC sampling was applied to remove mismatched correspondences during optimization.

2.5 Statistical Analysis

Statistical comparison between the proposed method and the baseline ICP approach was performed using paired t-tests. Mean values and standard deviations were calculated for rotational and translational errors across all trials. A significance level of $p<0.05$ was used to determine statistical difference. Variance analysis was conducted to evaluate repeatability under different initial conditions. Processing speed was reported as the average frame rate over continuous tracking sequences longer than 60 seconds. These analyses provided objective assessment of accuracy, stability, and real-time performance.

3.Results and Discussion

3.1 Registration Accuracy on Cadaveric and Synthetic Data

On the 18 cadaveric spine datasets, the proposed method reduced the mean rotational error from 1.92° to 0.84° and the mean translational error from 2.16 mm to 0.93 mm compared with depth-ICP. The reduction was consistent across thoracic and lumbar segments and remained stable under moderate viewpoint variation. On the 10 synthetic CT-to-depth alignment tasks, similar improvements were observed under controlled noise and simulated occlusion. These results indicate that the improvement was not limited to a specific dataset but was related to better feature selection and stronger geometric constraints during optimization [22,23]. Compared with previously reported depth-based AR alignment systems that often show millimeter- to centimeter-level variability under intraoperative conditions, the present method demonstrates tighter error bounds and improved repeatability (Fig.1).

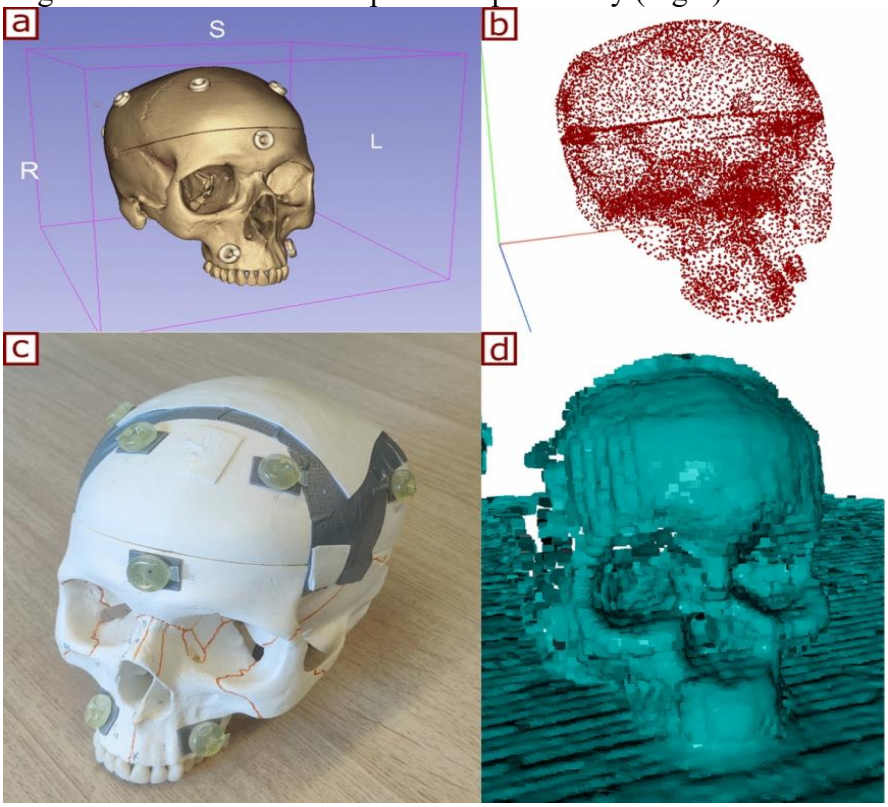


Fig.1.Processing pipeline for depth-based rigid registration, showing curvature-based keypoint selection and depth-constrained pose refinement for AR bone alignment.

3.2 Stability Under Occlusion and Partial Surface Visibility

The main difference between the proposed method and baseline ICP appeared in frames with limited overlap or local occlusion. In these cases, standard ICP frequently converged to locally consistent but globally incorrect solutions. The depth-consistency constraint reduced this effect by penalizing transformations that conflicted with measured depth along the camera viewing direction. As a result, large pose jumps were less frequent when surgical tools or soft tissue partially blocked the bone surface. Similar studies have shown that robust registration in AR

navigation requires explicit handling of partial visibility and mismatched correspondences rather than relying on dense point matching alone [24,25]. The curvature-weighted keypoints further improved robustness by emphasizing ridge and edge regions, which provide stronger geometric constraints than flat laminar areas. This design reduced the number of unstable RANSAC hypotheses and improved convergence reliability.

3.3 Real-Time Performance on Head-Mounted AR Hardware

The system maintained an average processing speed of 45 frames per second, which supports continuous navigation without noticeable delay. The computational load remained stable because the optimization was applied to a reduced set of keypoints instead of the full point cloud. The robust estimation stage was also limited to a fixed number of candidate transformations, which prevented excessive runtime growth. In comparison, some AR navigation pipelines require multi-frame fusion or external computing resources to achieve stable tracking, which may increase latency and reduce portability. The present framework demonstrates that geometry-based registration with lightweight constraints can meet real-time requirements on wearable hardware while maintaining high accuracy [26,27].

3.4 Comparison with Related AR Navigation Studies and Clinical Relevance

When compared with recent AR navigation studies, the reported sub-millimeter mean translation error is within the range considered suitable for rigid bone guidance. Many AR systems report mean accuracy values but also show wide error distributions during dynamic tracking. In contrast, the current method achieved both reduced mean error and narrower variance across repeated trials, indicating improved stability. Fig.2 illustrates a typical error distribution analysis reported in AR surgical navigation research. Although direct comparison is limited by differences in experimental setup, the present results show competitive accuracy without the need for fiducial markers or external tracking devices. Remaining limitations are mainly observed in low-curvature regions and cases of extreme surface loss. Future work may include adaptive re-initialization or view adjustment strategies to maintain overlap during extended tracking sequences [28,29].

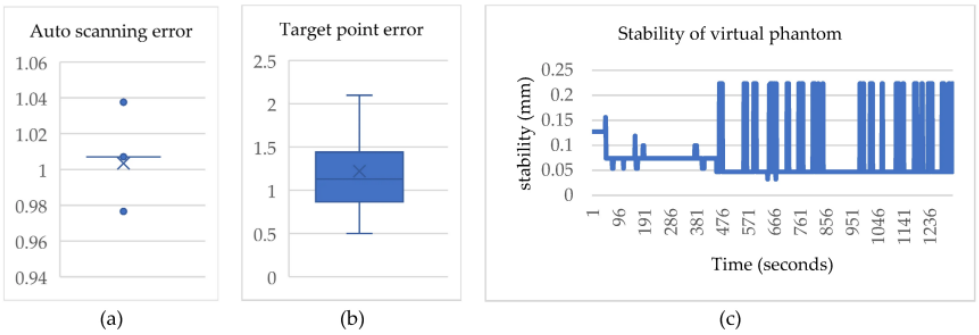


Fig.2.**Rotational and translational error distribution demonstrating registration accuracy and tracking stability during AR-guided surgery.**

4.Conclusion

A depth-constrained surface optimization method for markerless rigid bone registration in head-mounted AR surgical systems was developed and evaluated. The framework combines curvature-based keypoint selection with a depth-consistency term within a robust rigid transformation model. This design improved alignment accuracy and stability when only partial bone surfaces were visible. Experiments on cadaveric spine datasets and synthetic CT-to-depth tasks showed clear reductions in both rotational and translational error compared with conventional depth-ICP, while maintaining real-time performance at 45 frames per second. These findings indicate that combining geometric feature weighting with depth-based constraints can improve rigid registration without the need for fiducial markers or external tracking systems. The method supports practical use in AR-guided orthopedic and spinal procedures that require precise bone alignment. Its lightweight computation makes it suitable for wearable devices in intraoperative environments. Limitations remain in areas with low

surface curvature and in cases of severe occlusion, where geometric information is limited. Future research may include adaptive re-initialization or the integration of additional sensing cues to further improve robustness during long surgical procedures.

References

Cekic, E., Cetin, G., Ozer, A., Baylarov, B., Tatar, I., & Hanalioglu, S. (2025). Digital 3D Models in Neurosurgical Training, Surgical Planning, and Intraoperative Guidance: Comprehensive Review of the Literature. Authorea Preprints.

Yang, Y., Leuze, C., Hargreaves, B., Daniel, B., & Baik, F. (2025). EasyREG: Easy Depth-Based Markerless Registration and Tracking using Augmented Reality Device for Surgical Guidance. arXiv preprint arXiv:2504.09498.

Wickramasinghe, N., Thompson, B. R., & Xiao, J. (2022). The opportunities and challenges of digital anatomy for medical sciences: Narrative review. JMIR Medical Education, 8(2), e34687.

Wang, Y., Wen, Y., Wu, X., Wang, L., & Cai, H. (2025). Assessing the Role of Adaptive Digital Platforms in Personalized Nutrition and Chronic Disease Management.

Olthof, K., Fusaglia, M., Kok, N., & Kuhlmann, K. (2025). Image-guided navigation in liver surgery. British Journal of Surgery, 112(11), znaf254.

Zhang, Y., & Wang, J. (2026). Design and Implementation of a Computer-Aided Full Lifecycle Quality Management System for Wind Farms in Upgrades, Renovations, and Subcontractor Supervision.

Frisk, H. (2025). Intraoperative imaging, navigation and monitoring (Doctoral dissertation, Karolinska Institutet).

Jiao, Y., Wang, A., Zhao, B., & Shi, T. (2026). The Impact of Visual Language Strategies in Public Art Creation on Community Spatial Perception and Public Behavior.

Doughty, M., Ghugre, N. R., & Wright, G. A. (2022). Augmenting performance: A systematic review of optical see-through head-mounted displays in surgery. Journal of Imaging, 8(7), 203.

Jiao, Y., Wang, A., Zhao, B., & Shi, T. (2026). Quantitative Study on the Construction and Application Effectiveness of Graffiti Wall Painting Teaching Models in Public Space Contexts.

Trujillo-Hernández, G., Rodríguez-Quiñonez, J. C., Flores-Fuentes, W., Sanchez-Castro, J. J., Real-Moreno, O., Castro-Toscano, M. J., & Hernández-Balbuena, D. (2024). Optical 3D Scanning System in Medical Applications. In Scanning Technologies for Autonomous Systems (pp. 159-196). Cham: Springer Nature Switzerland.

Liu, S., & Yim, J. (2025). Research on Generative AI Creation Systems Based on Visual Language Modeling: Human-Machine Collaboration and Cognitive Feedback Mechanisms. Available at SSRN 6139770.

Liehn, M., & Schlautmann, H. (2022). 101 of surgical instruments: naming, recognizing, handling and presenting. Springer Nature.

Yu, X. (2026). AI-Driven Personalization across Domains for Local Categorical Query Understanding and Context-Aware Retrieval.

Chen, J., Xiao, H., Zeng, Z., Du, X., & Xiao, R. (2026). Diffusion-step attention consistency for multivariate time series anomaly detection. Knowledge-Based Systems, 115716.

Chen, X., Xiao, H., Zeng, Z., Zhang, S., & Xiao, R. (2025). Fine-Grained Multivariate Time Series Anomaly Detection via Causal Inference. Knowledge-Based Systems, 114765.

Hossain, T., Shi, P., & Kovacs, L. (2026). Dual-Detector Vision and Depth-Aware Back-Projection for Accurate Apple Detection and 3D Localisation for Robotic Harvesting. Robotics, 15(2), 47.

Wang, S. (2025). W-PCA Based Gradient-Free Proxy for Efficient Search of Lightweight Language Models. arXiv preprint arXiv:2504.15983.

CUT, A. A., Georgiou, L., Christofi, A., & CUT, C. D. (2025). MRI evaluation of the thermal heating.

Wen, Y., Wu, X., Wang, L., Cai, H., & Wang, Y. (2025). Application of Nanocarrier-Based Targeted Drug Delivery in the Treatment of Liver Fibrosis and Vascular Diseases. Journal of Medicine and Life Sciences, 1(2), 63-69.

Amirian, M., Bach, M., Jimenez-del-Toro, O., Aberle, C., Schaer, R., Andrearczyk, V., ... & Depeursinge, A. (2025). A Multi-Centric Anthropomorphic 3D CT Phantom-Based Benchmark Dataset for Harmonization. arXiv preprint arXiv:2507.01539.

Wang, Y., Chen, J., Arias, R., Wang, Y., & Yin, X. (2026). Development and Validation of a Patient-Friendly Digital Assessment Platform for Precision Screening of Oral Anti-Obesity Medications (AOMs).

Alhussan, A. A., Abdelhamid, A. A., El-Kenawy, E. S. M., Ibrahim, A., Eid, M. M., Khafaga, D. S., & Ahmed, A. E. (2023). A binary waterwheel plant optimization algorithm for feature selection. IEEE Access, 11, 94227-94251.

Gao, G., Ma, X., Lu, C., & Gao, R. (2026). Reliability Analysis and Application Research of SMS Communication Systems in Medical Notification Scenarios.

Wang, Y., Chen, J., Wang, Y., & Yin, X. (2026). Application of Obtainable Biological Agent Characteristics in Efficacy Stratification of Oral Anti-Obesity Drugs.

Haroon, U., AlMughrabi, A., Zoumpekas, T., Marques, R., & Radeva, P. (2026). Vole: a point-cloud framework for food 3D reconstruction and volume estimation. Scientific Reports.

Zhang, Y., Gu, W., & Wang, J. (2026). Construction of Wind Farm Asset Health Index Based on Multi-Dimensional Indicators and Analytic Hierarchy Process and Its Correlation with Operational Performance. Authorea Preprints.

Zebarjadi, M., Organ, A., Grafstrom, L. A., Zachs, D., & Lim, H. H. (2026). Adaptive fusion of accelerometer and ultrasound signals for real-time motion tracking in therapeutic ultrasound stimulation. Physics in Medicine and Biology.

Rocha, R. D. L., & de Figueiredo, F. A. (2023). Beyond land: A review of benchmarking datasets, algorithms, and metrics for visual-based ship tracking. Electronics, 12(13), 2789.