AI AND MACHINE LEARNING

RenalCare AI for Kidney Stone Detection and Clinical Guidance

Steevan Raj[1], Sumant Murarji Hosmane[1], T. K. Varun Kumar[1], C. N. Varun Muthappa[1], J. K. Guruprasanna[2]

[1]Student, Department of Computer Science and Engineering, Vivekananda College of Engineering and Technology, Puttur, Karnataka, India
[2]Assistant Professor, Department of Computer Science and Engineering, Vivekananda College of Engineering and Technology, Puttur, Karnataka, India

Published November 2026

Abstract

Kidney stone assessment often depends on imaging interpretation and follow-up guidance that may not be readily available to every patient. This paper presents RenalCare AI, a web-based clinical guidance system that combines kidney-stone image analysis, explainable visual feedback, health tracking, risk insights, and report generation. The proposed approach supports CT and ultrasound as a modality-extensible workflow; the current implementation centers on CT image upload. A vision model based on EfficientNet-B4 or MobileNetV2 is paired with Grad-CAM heatmaps to make prediction regions visible. The application is implemented with a React interface, FastAPI services, SQLite persistence, and token-based access control. It also records hydration, meal, medication, appointment, and health-goal information, while an AI-guidance service is backed by a rule-based fallback. The uploaded project evaluation reports perfect scores on its held-out split; these preliminary findings require independent, multi-institutional validation before clinical deployment. The system is designed as decision support and is not a substitute for professional medical diagnosis.

Keywords: Clinical decision support; deep learning; Grad-CAM; kidney stones; medical image analysis; multimodal learning

1 Introduction

1.1 Background and motivation

Kidney stones are a recurrent urological problem for which imaging plays a central role in assessment. Computed tomography provides detailed diagnostic information, while ultrasound is radiation-free and accessible but can be less reliable for small stones. These trade-offs motivate tools that can assist with image review, present interpretable evidence, and connect imaging results with practical follow-up information. Recent work demonstrates the value of deep learning for image analysis [1] and the feasibility of automated kidney-stone detection and segmentation on noncontrast CT scans [5].

RenalCare AI is proposed as an end-to-end patient-facing and clinician-supporting platform. It brings together image-based stone assessment, Grad-CAM visual explanations, lifestyle tracking, risk insights, appointment support, and downloadable reports in one protected web application. The design intentionally separates measured information, model outputs, and generated guidance so that the interface remains auditable and appropriate for supervised use.

1.2 Contributions

The main contributions are: an extensible CT and ultrasound workflow; a vision-analysis component with Grad-CAM evidence overlays; a secure web architecture for tracking hydration, meals, medication, appointments, and reports; and a guidance layer that combines language-model responses with a rule-based fallback. The work is presented as an engineering prototype intended to support, rather than replace, clinical judgement.

2 Related work and requirement analysis

2.1 Related work

Deep neural networks provide the foundation for contemporary medical-image classifiers [1]. EfficientNet uses compound scaling to balance model capacity and computational cost [2], and attention mechanisms can support the fusion of complementary features [4]. Grad-CAM supplies a practical way to visualize image regions that contribute to a model decision [3]. For kidney-stone tasks, CT-oriented datasets and automated detection studies provide useful baselines for data preparation and validation [5,6].

2.2 Functional and nonfunctional requirements

The application accepts a patient profile and an image upload, returns stone-presence information with a confidence score and a heatmap, and retains scan history. It also supports daily water logging against a three-litre goal, meal entries for sodium and oxalate awareness, medication records, appointment requests, health goals, and report download. Nonfunctional requirements include secure authenticated access, responsiveness, maintainability, reliable fallback behavior, and inference performance suitable for interactive use. The project implementation uses Python 3.10, PyTorch, FastAPI, React, Tailwind CSS, SQLite, and GPU-capable hardware for model execution.

3 Proposed method

3.1 Image processing and prediction workflow

The proposed workflow begins with consented image upload and patient-context collection. CT images are prepared with clinically relevant intensity-windowing, while ultrasound images can be denoised with bilateral filtering. Separate visual encoders extract modality-specific features. In the multimodal configuration, cross-attention combines these feature representations before classification; an implementation with MobileNetV2 or EfficientNet-B4 can be selected according to the available compute budget. Training follows a transfer-learning approach in which encoders are first frozen and then fine-tuned.

For each completed inference, Grad-CAM is generated from the selected vision model. The heatmap is displayed as explanatory evidence rather than as an independent diagnosis. This distinction is important because highlighted regions can support review but do not establish a clinical finding on their own. Figure 1 summarises the use cases and key application features.

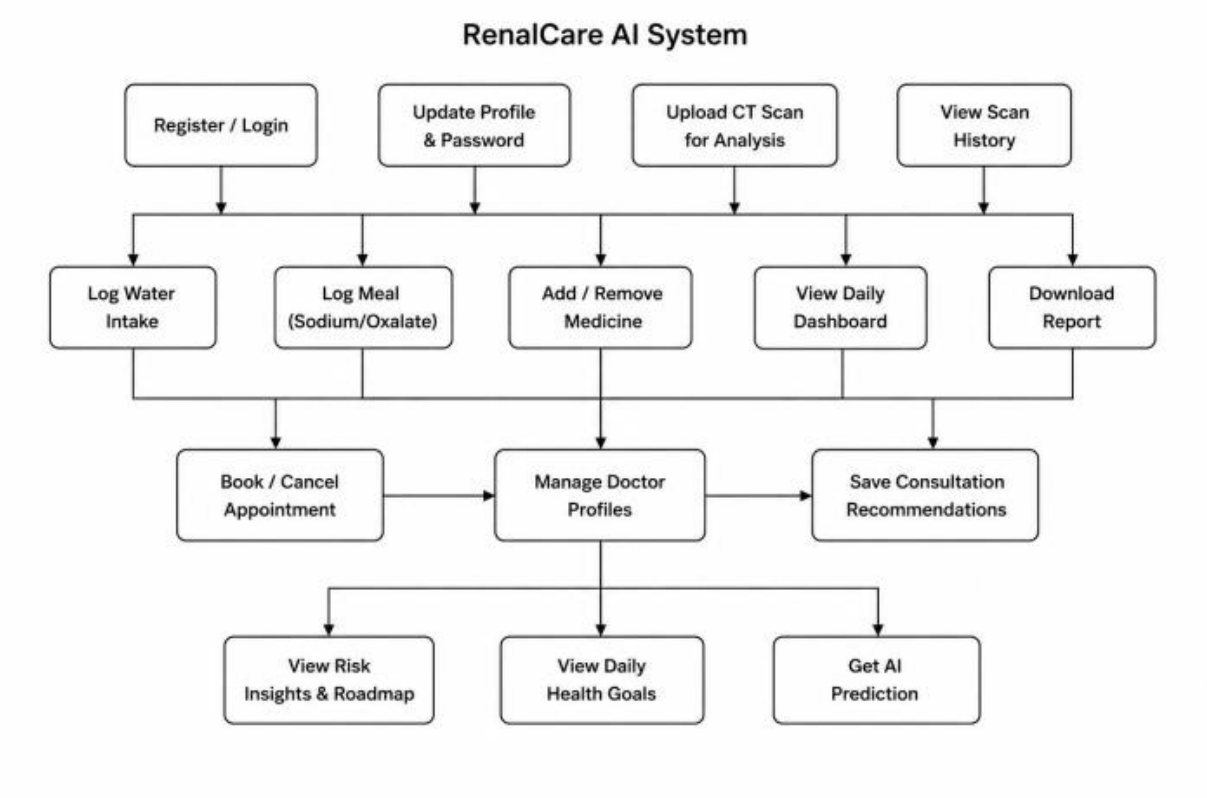


Figure 1. RenalCare AI use cases and key application features.

3.2 Guidance, monitoring and risk insights

After analysis, the system records patient-entered hydration, meal, medication, and recurrence information. A transparent risk calculation combines scan severity weights, hydration-compliance penalties, and recurrence history to provide a short-term risk indicator. The guidance component can use a language-model API to draft personalized diet, hydration, and medical-consultation prompts; a rule-based fallback preserves availability when the external service is unavailable. Guidance is framed as educational decision support and must be reviewed by qualified healthcare professionals when clinical action is considered.

4 System design and implementation

4.1 Layered architecture

The system uses a layered client-server design. The React presentation layer provides registration, login, dashboard, scan upload, health logs, reports, and profile settings. FastAPI services manage authentication, user data, scan analysis, risk assessment, recommendations, and report generation. The data layer stores patient and application records, file references, and model outputs; the infrastructure layer hosts API services and AI inference. This separation allows modules to be maintained and tested independently. The layered architecture is shown in Figure 2.

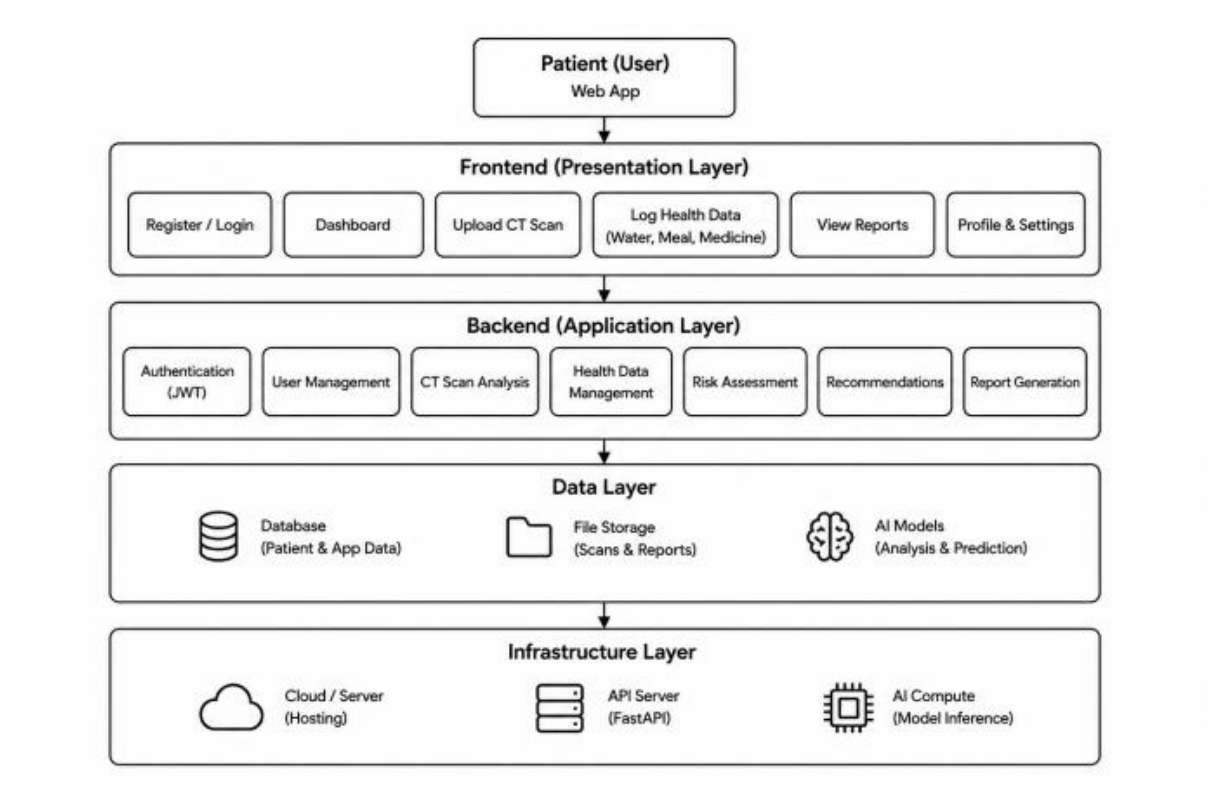


Figure 2. Layered architecture of the RenalCare AI system.

4.2 Security and user management

Authentication uses bcrypt password hashing and JSON Web Token bearer sessions. Ownership checks are applied to protected patient resources, and profile changes can be confirmed through one-time-password flows. HTTPS deployment, scoped access controls, and auditable service boundaries are recommended for any production environment. These measures protect sensitive health-related information while supporting normal patient and clinician workflows.

5 Evaluation and testing

5.1 Preliminary model results

The submitted project report records the held-out metrics shown in Table 1 for its vision-analysis module. Because the report does not specify the test-set size, data partitioning, or external-validation protocol, the values are treated as preliminary prototype results. Future evaluation should report class balance, confidence intervals, calibration, error cases, and performance on independent clinical data. The reported stone-size estimation distribution is summarised in Table 2.

Table 1. Reported held-out test metrics.

Metric	Reported value
Accuracy	1.00
Macro precision	1.00
Macro recall	1.00
Macro F1 score	1.00

Table 2. Reported stone size estimation distribution.

Statistic	Reported value
Median size	5.8 mm
90th percentile	8.9 mm
Maximum ceiling	15.0 mm

5.2 Application testing

Unit, integration, functional, and system tests were carried out with Pytest. The checks covered registration, access control, health logging, CT scan analysis, and rule-based fallback behavior. Representative outcomes from the project report are listed in Table 3. Together, these tests provide evidence of feature-level behavior; they do not validate diagnostic safety or clinical effectiveness.

Table 3. Representative application test cases.

ID	Test case	Expected outcome
TC-01	Valid registration	HTTP 200 and token issued
TC-04	Cross-patient resource access	HTTP 403 forbidden
TC-05	Log water intake	Daily total updated by 500 ml
TC-12	CT scan analysis	Valid prediction and size output
TC-16	Health goals without API	Rule-based fallback returned

6 Discussion

RenalCare AI demonstrates how an image-analysis prototype can be combined with explainability and longitudinal self-management features in a single system. The architecture supports a practical separation between visual inference, health-data management, and guidance generation. Its reported test-split scores are encouraging, but the absence of documented external validation means that no diagnostic-performance conclusion should be generalized beyond the submitted project setting. Independent datasets, clinician review, and prospective usability studies are needed before any clinical deployment.

The design is deliberately extensible. Future work can introduce 3D volumetric CT analysis, PACS integration, clinician review portals, formal audit logs, patient-consent workflows, and regulatory assessments. The multimodal path should also be evaluated against strong single-modality baselines to establish whether cross-attention fusion produces a meaningful benefit.

7 Conclusion

This paper presented RenalCare AI, a protected web system for kidney-stone image analysis and clinical guidance. It combines a transfer-learning vision component, Grad-CAM explanations, health tracking, risk insights, appointment support, and report generation within a layered FastAPI and React architecture. The prototype's reported held-out metrics and functional tests show implementation feasibility. Its future value depends on rigorous external validation, clinical oversight, and careful handling of health information before it can be considered for real-world clinical use.

References

[1] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–444. https://doi.org/10.1038/nature14539

[2] Tan M, Le QV. EfficientNet: Rethinking model scaling for convolutional neural networks. In: *Proceedings of the 36th International Conference on Machine Learning*; 2019. p. 6105–6114. (PMLR vol. 97)

[3] Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision*; 2017. p. 618–626.

[4] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: *Advances in Neural Information Processing Systems 30*; 2017.

[5] Elton DC, Turkbey EB, Pickhardt PJ, Summers RM. A deep learning system for automated kidney stone detection and volumetric segmentation on noncontrast CT scans. *Med Phys*. 2022;49(4):2545–2554. https://doi.org/10.1002/mp.15518

[6] Abdalla PA, et al. Kidney stone detection via axial CT imaging: A dataset for AI and deep learning applications. *Data Brief*. 2025;59:111446.

[7] Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. *Nature*. 2023;620(7972):172–180. https://doi.org/10.1038/s41586-023-06291-2