IJIRSET
International Journal of Innovative Research in Science | Engineering | Technology

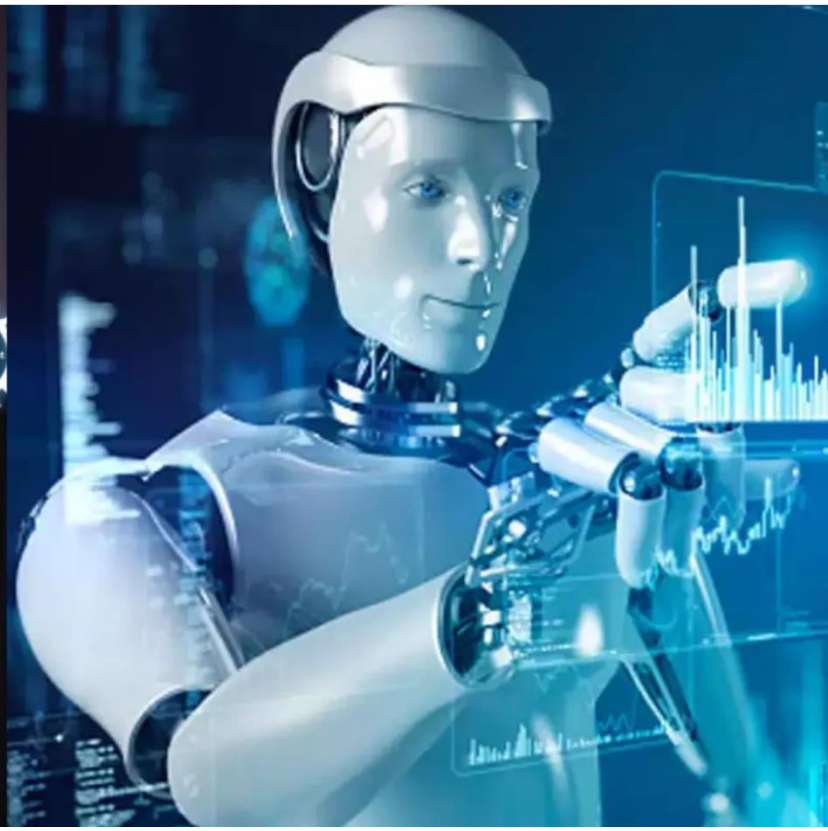
ISSN(O): 2319-8753
ISSN(P): 2347-6710



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)





Impact Factor: 9.935

Volume 15, Issue 10, October 2026

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

Renal Cancer Detection Using Machine Learning and Deep-Learning Image Analysis with ESP32-Based Vital-Sign Monitoring

Sudharshan M K, Nikhilraj G, Nikhilkumar A K, Yashwanth S

Professor, Department of Electronics and Communication Engineering, S.J.M. Institute of Technology, Affiliated to Visvesvaraya Technological University (VTU), Belagavi, Chitradurga, Karnataka, India

Final Year Students, Department of Electronics and Communication Engineering, S.J.M. Institute of Technology, Affiliated to Visvesvaraya Technological University (VTU), Belagavi, Chitradurga, Karnataka, India

ABSTRACT: Renal cell carcinoma (RCC) is the most common form of adult kidney cancer, and a suspicious renal mass is normally investigated through clinical assessment, imaging and, when indicated, pathology. This paper presents a research-oriented decision-support prototype that combines three independent pipelines: a convolutional neural network (CNN) or transfer-learning classifier for labelled kidney CT/MRI images; an optional comparison of conventional machine-learning models (Logistic Regression, Random Forest, Support Vector Machine and gradient boosting) on structured clinical or radiomic data; and an ESP32-based module that displays heart rate, SpO_2 and temperature on an OLED. A Python Flask interface accepts an allowed image, shows a dataset-defined class score with its model version and limitations, and presents live vital readings in a separate panel; vital signs are never merged into the image prediction because they are not specific to renal cancer. The design uses patient-level data splitting, training-only augmentation and a single evaluation on a held-out test set to limit data leakage. The paper also summarizes the emerging, largely associative evidence on the tumour microbiome and the genetic alterations relevant to RCC. The work is a proposed design; no performance figures are reported and quantitative evaluation is left to future experiments. The prototype is not a diagnostic device.

KEYWORDS: convolutional neural network, deep learning, ESP32, kidney cancer, machine learning, MAX30102, medical image analysis, OLED, renal cell carcinoma, SpO_2.

I. INTRODUCTION

The kidneys filter blood, regulate fluid and electrolyte balance and produce urine. Renal cell carcinoma (RCC) accounts for most adult kidney cancers, and some patients have no early symptoms [1]. A suspicious renal mass is typically evaluated through history, examination, laboratory tests, ultrasound, CT or MRI, and tissue pathology when indicated, and staging then describes tumour extent and spread [2], [3]. Interpreting images and records requires trained specialists and can be time-consuming. Machine learning can support research by learning patterns from labelled images or clinical datasets, but reported performance depends on data quality, representative samples, sound validation and clinical review.

This paper presents a research-oriented decision-support prototype with three deliberately separate components: (1) classification of kidney CT/MRI images by deep learning, (2) optional classification of structured data by conventional machine learning, and (3) acquisition and display of heart rate, SpO_2 and temperature by an ESP32-based unit. Keeping these outputs separate avoids implying that general vital signs identify a kidney tumour. A Flask web interface lets a user upload an image, view the model output with its version and limitations, and see live vital readings in a separate panel.

The main contributions are: (i) a leakage-aware design for image-based renal classification using patient-level splitting, training-only augmentation and a single held-out evaluation; (ii) an optional, clearly bounded comparison with conventional machine-learning models; (iii) a low-cost ESP32 vital-monitoring module with explicit handling of invalid or disconnected sensor readings; and (iv) a transparent interface that carries a research-only disclaimer and never merges vital signs into the prediction. A short review of tumour-microbiome and genetic evidence in RCC states

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

what that literature does and does not establish. The system is an academic prototype and not a medical diagnostic device.

The remainder of this paper is organized as follows. Section II reviews related work, Section III defines the problem and requirements, Sections IV–VI describe the architecture, hardware and software design, Section VII gives the validation plan, Section VIII discusses ethics and limitations, and Sections IX–XI cover expected outcomes, future work and conclusions.

II. RELATED WORK

Studies in medical image analysis show that CNNs and transfer-learning models can learn visual patterns from labelled radiology or pathology datasets. Reported performance can, however, be inflated by duplicate images, images from the same patient appearing in more than one data split, class imbalance, or testing only on data from a single institution. Validation should therefore split by patient, preserve an untouched test set and, where possible, include external data. These lessons shape the methodology of this work.

RCC includes several subtypes. Clear-cell RCC is the most common adult subtype and is often associated with disruption of the VHL pathway; papillary and chromophobe RCC are distinct subtypes with their own molecular and histological characteristics, and other kidney malignancies such as collecting-duct carcinoma and renal-pelvis urothelial carcinoma are not interchangeable with RCC. Clinical guidelines describe a diagnostic pathway that combines imaging, laboratory tests and pathology [4]–[6]; a model prediction cannot replace this pathway. CT and MRI show anatomy and lesion characteristics, whereas histopathology shows tissue architecture under a microscope, so these modalities must not be mixed as if they were equivalent.

Researchers are studying microbial signals in the gut, urine and tumour tissue and their possible relationships with inflammation, immune activity and response to immunotherapy [7]–[11]. In RCC, current evidence supports the investigation of associations; it does not establish any single bacterial species as a routine cause, screening marker or diagnostic test. Microbiome measurements are sensitive to contamination, sequencing method, antibiotics, diet and cohort characteristics, and low-biomass tumour samples are especially vulnerable to artefacts. General bacterial morphology (cocci, bacilli, curved or spiral forms) and laboratory growth phases are microbiology concepts and are not stages of kidney cancer.

Genomic studies have identified recurrent molecular pathways in RCC. In clear-cell RCC, loss or inactivation of VHL commonly contributes to dysregulated hypoxia-inducible-factor signalling, and changes in chromatin-remodelling genes such as PBRM1, SETD2 and BAP1 can also occur; the pattern differs by subtype and by patient, and inherited variants are associated with hereditary syndromes [12]. Mutation status is established by validated genetic assays with clinical interpretation. It cannot be inferred from heart rate, temperature, SpO_2 or a generic image classifier, and it can be predicted from images only in studies with linked, validated genomic labels.
Low-cost physiological monitors based on pulse oximetry and temperature sensing are widely available, but they measure general physiology and are not cancer-specific. The practical gap addressed here is a transparent, leakage-aware student prototype that separates image classification, optional tabular machine learning and vital-sign monitoring, and states the limitations of each. The prototype makes no claim to detect mutations or bacteria; that would require validated genomic or microbiology data and suitable analysis pipelines.

III. PROBLEM STATEMENT AND REQUIREMENTS

A. Problem Definition

The problem is to build a reproducible and transparent prototype for classifying labelled kidney images that avoids common validation pitfalls, while demonstrating real-time physiological monitoring without implying that such monitoring detects renal cancer. The data modality (CT, MRI, histopathology or structured records) must be selected per experiment, labels must come only from the dataset, and all images from one patient must remain in one split.

B. Functional Requirements

The system shall support (1) loading labelled images and metadata with checks for class counts, duplicates, corrupted files and patient identifiers; (2) patient-level partitioning into training, validation and held-out test sets before

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

augmentation; (3) resizing, normalization and training-only augmentation; (4) training and evaluation of a CNN/transfer-learning model with saved weights, preprocessing configuration, class mapping and version information; (5) optional comparison of Logistic Regression, Random Forest, SVM and gradient boosting on suitable tabular data; (6) sampling of heart rate, SpO_2 and temperature at a defined interval, rejecting invalid or disconnected readings; (7) display of these values with a sensor-status indicator on the OLED; (8) validated upload of an allowed CT/MRI image through Flask; (9) presentation of the predicted class, model score, model version, dataset scope and a research-only disclaimer; (10) a vital panel separate from the prediction; and (11) generation of evaluation metrics with only minimal, de-identified experimental records.

C. Non-Functional Requirements and Constraints

Experiments shall be reproducible through fixed random seeds and saved, versioned configuration. Preprocessing, imputation and feature selection shall be fitted on training data only. Data shall be licensed and de-identified, uploads secured and DICOM patient identifiers never exposed. Outputs shall be called model scores rather than probabilities of cancer unless calibration has been assessed, and false negatives and false positives shall be reported, not only overall accuracy. Invalid or missing sensor values shall be flagged rather than displayed as valid, and image analysis and vital monitoring shall remain independent modules. Constraints include the academic scope (the prototype is not a clinically validated device), the dependence of performance on the dataset, which cannot be stated before experiments are run, and the 3.3 V logic of the ESP32. The open-source toolchain (Python, TensorFlow/Keras or PyTorch, scikit-learn, Flask and the Arduino toolchain), low-cost hardware and optional GPU use make the work technically and economically feasible for a student project.

IV. SYSTEM ARCHITECTURE

The architecture consists of three independent pipelines that can be shown together in one interface but are deliberately not merged into a single cancer score: (A) image analysis, (B) optional structured-data machine learning and (C) vital monitoring through the ESP32. Fig. 1 shows the deployment view. A user uploads a CT/MRI image to the Flask server, which validates file type, size and dimensions, applies the saved modality-specific preprocessing and passes the image to the CNN/transfer-learning classifier. The traditional-ML model contributes only when its clinical or radiomic inputs are available. The result page shows the predicted class, model version and limitations. In parallel, the ESP32 reads the MAX30102 and the temperature sensor, drives the OLED and feeds a separate vital-sign panel; this panel is shown together with, but not combined into, the result.

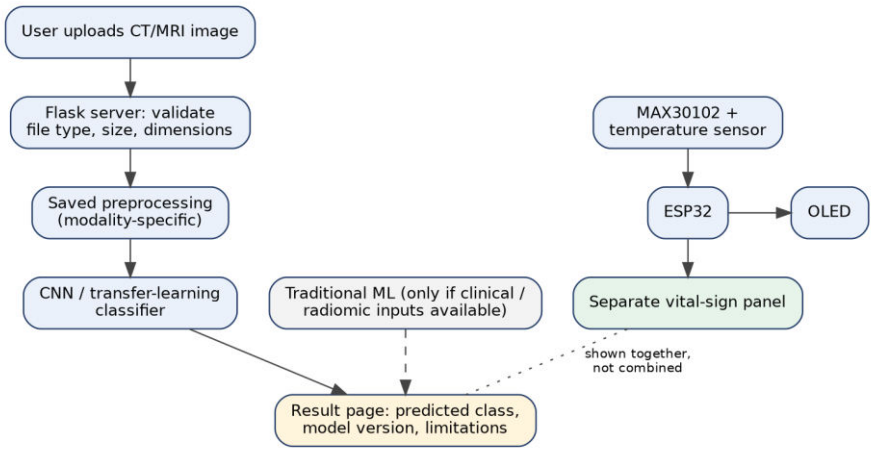


Fig. 1. Project architecture for CT/MRI classification, optional traditional ML, ESP32 vital monitoring and Flask

The system can be viewed in four layers: a data layer (labelled images, metadata and optional clinical or radiomic features), a model layer (CNN/transfer-learning and traditional ML), a hardware layer (ESP32, MAX30102, temperature sensor and OLED) and an application layer (Flask interface with separate result and vital panels). Inputs are a kidney CT or MRI image (a PNG/JPEG export, or DICOM with identifiers removed, with the modality recorded

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

for every experiment), optional structured variables, and sensor readings. Outputs are the dataset-defined class with model score, version and scope; the optional traditional-ML result; heart rate, SpO_2 and temperature with sensor status; and evaluation reports (confusion matrix, sensitivity, specificity, precision, F1-score and ROC-AUC).

V. HARDWARE DESIGN

A. ESP32 Controller and OLED Display

An ESP32 development board [13] is the controller of the vital-monitoring unit; it uses 3.3 V logic, and the exact board variant is to be confirmed. A 5 V signal must never be connected directly to an ESP32 GPIO pin unless the input is explicitly 5 V tolerant. The firmware, written with the Arduino toolchain, initializes the I^2C bus, the OLED, the pulse-oximetry module and the temperature sensor, and then samples at a defined interval. A commonly used I^2C SSD1306-compatible OLED shows heart rate (bpm), SpO_2 (%) and temperature (°C) with a sensor-status indicator; its resolution and address are to be confirmed.

B. Pulse and Oxygen Sensor

The MAX30102 module [14] estimates pulse rate and SpO_2 by photoplethysmography (PPG). Its readings are affected by motion, poor contact, perfusion, ambient light and device implementation, and pulse-oximetry accuracy can also vary between patient groups [15]. The outputs are therefore treated as estimates that require validation against a reference device before any accuracy claim is made.

C. Temperature Sensor

The exact part number of the temperature sensor is to be documented, because contact and non-contact sensors have different limitations. Whether it measures skin, ambient or body temperature must also be recorded, and a skin or ambient measurement must not be represented as core body temperature.

TABLE I HARDWARE COMPONENTS AND INTERFACES

Component	Role	Interface / note
ESP32 development board	Controller; sampling and display logic	3.3 V logic; board variant to be confirmed
OLED display	Shows bpm, SpO_2, °C and status	Commonly I^2C SSD1306-compatible
MAX30102 module	Heart-rate and SpO_2 estimates (PPG)	I^2C; verify module voltage and pin labels
Temperature sensor	Temperature measurement	Part number to be documented
Accessories	Power and wiring	USB cable, jumper wires, breadboard; optional enclosure

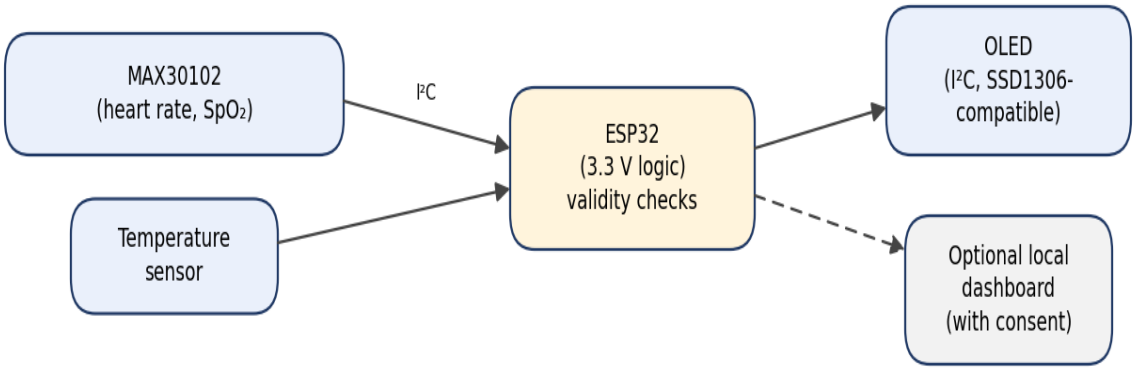


Fig.2. ESP32 vital-monitoring module (block diagram; pin assignment to be verified per module).

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

D. Sampling and Failure Handling

In each sampling cycle the firmware reads heart rate and SpO_2 from the MAX30102 and the temperature value, and checks that they are within a plausible range and that the sensor has contact. If a check fails, the OLED shows a sensor-check status and the value is not displayed as valid; if a sensor is disconnected, the status changes to invalid or disconnected and no stale value is shown as valid. Valid values are shown with an OK status and may optionally be sent to a local dashboard, only with user consent and appropriate privacy safeguards. Exact GPIO pins cannot be specified safely without the exact modules: many I²C OLEDs and MAX30102 boards use SDA/SCL, but assignments vary by ESP32 board and library, so module labels, operating voltage, pull-ups and datasheets must be verified. Fig. 2 is a block diagram and not a wiring diagram.

VI. SOFTWARE AND MODEL DESIGN

A. Dataset Selection and Preprocessing

One primary imaging modality is selected together with a dataset that has documented labels, class definitions, patient identifiers where available and permitted research use. Labels such as normal/benign/malignant or RCC subtype are used only if the dataset supplies them, and unrelated datasets are not combined without harmonization. For DICOM, orientation and spacing metadata are preserved while patient identifiers are removed. As Fig. 3 shows, the data are audited for missing labels, duplicates, corrupted files and class imbalance; split at patient level into training, validation and held-out test sets before augmentation; resized and normalized with modality-appropriate preprocessing, since CT and MRI differ in intensity characteristics; and augmented conservatively (for example by small rotations or flips where medically appropriate) on the training set only.

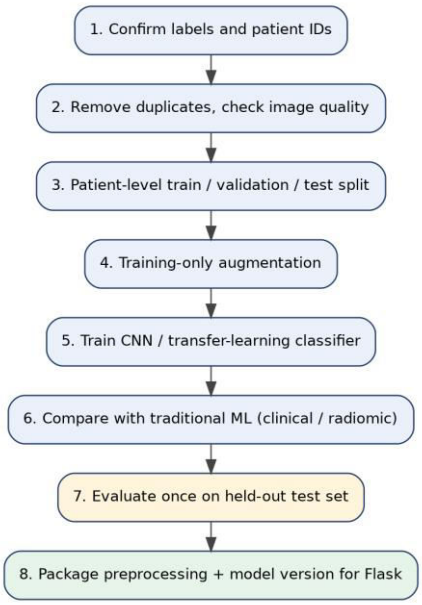


Fig.3.CT/MRI model development and deployment workflow

B. Deep-Learning Classification

A CNN/transfer-learning backbone such as ResNet [16] or EfficientNet [17] is used for the CT/MRI task. If scans are 3D volumes, the approach is defined explicitly as 2D slices, 2.5D neighbouring slices or a 3D network, and slices from one patient are never treated as independent samples. The final classification layer is replaced for the dataset's classes; the model is trained on the training set, tuned on validation data and evaluated once on the held-out test set. Class weights or suitable sampling address class imbalance. The model, preprocessing configuration, class mapping and version information are saved for Flask inference.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

A. Traditional Machine Learning

Logistic Regression, Random Forest, Support Vector Machine and gradient boosting are compared using scikit-learn [18] only when inputs and target labels are defined, using structured clinical variables or radiomic features from a documented segmentation and feature-extraction method. Imputation, scaling, feature selection and hyper-parameter tuning are fitted only within training folds. Flattened pixels are not fed to a classical model without a justified baseline, and a tabular model is never trained on arbitrary vital-sign readings and described as a renal-cancer detector.

B. Flask Interface and Inference Flow

The interface uses Flask [19] with HTML, CSS and JavaScript, OpenCV or Pillow for image handling, and TensorFlow/Keras or PyTorch for inference. At inference, (1) the user uploads an allowed CT/MRI image; (2) the server validates file type, size and dimensions, rejects malformed input and avoids exposing DICOM identifiers; (3) the same saved preprocessing used during validation is applied; (4) the CNN produces a dataset-defined class score, and the traditional-ML result is shown only if its inputs are available; (5) the page shows the predicted class, model version, dataset scope and limitations, calling the output a model score and not a probability of cancer unless it has been calibrated; and (6) a separate vital panel displays the ESP32 readings, which are not combined into the prediction. Every result page carries a prominent research-only disclaimer. An optional SQLite database stores minimal, de-identified experimental records with access control, and each run generates metrics and configuration files for reproducibility.

VII. VALIDATION PLAN

This paper describes a proposed design and does not claim measured performance. Results will be reported only after the experiments have been run; no accuracy values are assumed or taken from examples. Splitting is done at patient level to prevent leakage between slices of the same scan or patient, and the held-out test set is used once. Table II lists the planned unit and integration tests, and Table III lists the evaluation items to be reported.

TABLE II PLANNED TEST CASES

ID	Test	Expected outcome
TC-01	Dataset audit: duplicate, corruption, missing-label and class-count checks	Issues detected and reported
TC-02	Patient-level split: each patient ID checked across splits	Zero overlap between splits
TC-03	Preprocessing: output shape and range; augmentation scope	Correct shape and range; no augmentation of validation/test data
TC-04	CNN training on a small subset	Loss decreases; weights and class mapping saved
TC-05	Traditional-ML preprocessing fitted within training folds	No leakage from validation/test data
TC-06	Sensor read with and without finger contact; disconnected sensor	Invalid readings rejected and flagged
TC-07	Upload of wrong-type, oversized and malformed files	Inputs rejected with a clear message
TC-08	Valid CT/MRI upload through Flask; traditional-ML inputs unavailable	Saved preprocessing applied; class, model version and limitations shown; ML result omitted without error
TC-09	ESP32 module running with the web page open	OLED values readable with correct status; vital panel separate from model result
TC-10	Sensor disconnected during operation	Status becomes invalid/disconnected; no stale value shown as valid
TC-11	DICOM input with identifiers; review of every result page	Identifiers not exposed; research-only disclaimer present

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

TABLE III EVALUATION ITEMS TO BE REPORTED

Item	What will be reported
Dataset summary	Patients and scans per split, class distribution, imaging modality, missing data
Classification metrics	Class-wise sensitivity, specificity, precision, F1-score, confusion matrix and ROC-AUC; PR-AUC and macro/weighted averages for imbalanced or multi-class tasks
Model comparison	CNN versus a simple baseline and the optional traditional-ML models on the same held-out patients
Reliability	Calibration if probabilities are shown, confidence intervals where feasible, error analysis and external-validation status
Sensor validation	Heart-rate, SpO_2 and temperature error against a reference device, missing-reading rate, update interval and stability

The sensor module is tested against a suitable reference device under controlled conditions. Consumer-grade readings are not clinical-grade measurements unless the specific device has been validated, and MAX30102 outputs remain estimates until they are validated against a reference device.

VIII. ETHICAL CONSIDERATIONS AND LIMITATIONS

The prototype is an academic proof of concept and not a clinically validated medical device. Only clinicians can interpret suspected renal cancer using the full clinical context and accepted diagnostic tests, and the system must not be used to rule out cancer, delay medical assessment or recommend treatment. Data must be appropriately licensed and de-identified, uploads secured, and consent and governance requirements documented. Contamination controls are essential in any low-biomass tumour-microbiome study.

Several limitations apply. CT and MRI are not interchangeable, and preprocessing and performance may vary by scanner, protocol, contrast phase and institution, so a model is valid only for the modality and dataset on which it was trained. Heart rate, SpO_2 and temperature are general measurements and not validated biomarkers for renal cancer. Genetic alterations such as VHL, PBRM1, SETD2 and BAP1 cannot be read from CT/MRI images without a validated linked-genomics study, and microbiome findings are emerging and largely associative, so bacterial morphology or growth phase does not prove causation and is not a cancer stage. Finally, a design description does not prove clinical accuracy or reliability.

IX. EXPECTED OUTCOMES AND APPLICATIONS

The expected outcomes are a documented image-preprocessing and model-training pipeline; a comparative evaluation of selected machine-learning and deep-learning methods on a clearly defined dataset; an ESP32 OLED prototype displaying heart rate, temperature and SpO_2 with sensor-validity checks; a research dashboard that keeps model predictions separate from vital-sign monitoring; and a literature summary of RCC genetics and microbiome evidence with explicit uncertainty. The prototype is suitable as a teaching and research platform for leakage-aware medical-image classification and for demonstrating low-cost embedded sensing, and its modular structure allows image analysis and monitoring to be developed and tested separately. Any clinical use would require specialist oversight, rigorous validation and applicable regulatory review.

X. FUTURE WORK

Planned extensions include external validation on data from other institutions, explainability methods such as Grad-CAM [20] for the image models, uncertainty estimation and calibration of model outputs, and clinically meaningful subgroup analysis. Properly governed genomic or microbiome datasets could allow carefully designed association studies, and any clinical deployment would require specialist oversight, rigorous validation and regulatory review.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

XI. CONCLUSION

This paper presented a research-oriented prototype that combines deep-learning analysis of kidney CT/MRI images, optional structured-data machine learning and ESP32-based physiological monitoring while keeping the three functions separate. Its scientific value depends on a well-documented dataset, patient-level validation, transparent metrics and careful separation of functions; heart rate, SpO_2 and temperature are never used as renal-cancer evidence. The reviewed literature indicates that microbiome evidence in RCC is emerging and largely associative, whereas the genetic pathways, such as VHL, are comparatively well characterized. Future work will focus on implementation, recorded evaluation and external validation.

XII. ACKNOWLEDGMENT

The authors thank the project guide, department faculty and the institution for their guidance and encouragement, and acknowledge the Python, Flask, TensorFlow/PyTorch, scikit-learn and Arduino/ESP32 open-source communities whose documentation supported this work.

REFERENCES

[1] National Cancer Institute, “Kidney (renal cell) cancer—Patient version.” Accessed: Nov. 1, 2026. [Online]. Available: https://www.cancer.gov/types/kidney

[2] National Cancer Institute, “Renal cell cancer treatment (PDQ®)—Patient version.” Accessed: Nov. 1, 2026. [Online]. Available: https://www.cancer.gov/types/kidney/patient/kidney-treatment-pdq

[3] National Cancer Institute, “Advances in kidney cancer research.” Accessed: Nov. 1, 2026. [Online]. Available: https://www.cancer.gov/types/kidney/research

[4] European Association of Urology, “Renal cell carcinoma: Diagnostic evaluation,” EAU Guidelines. Accessed: Nov. 1, 2026. [Online]. Available: https://uroweb.org/guidelines/renal-cell-carcinoma/chapter/diagnostic-evaluation

[5] T. Powles et al., “Renal cell carcinoma: ESMO Clinical Practice Guideline for diagnosis, treatment and follow-up,” Ann. Oncol., vol. 35, no. 8, pp. 692–706, 2024, doi: 10.1016/j.annonc.2024.05.537.

[6] National Institute for Health and Care Excellence, “Kidney cancer: Diagnosis and management—Diagnosis,” NICE Guideline NG256, 2026. [Online]. Available: https://www.nice.org.uk/guidance/ng256/chapter/Diagnosis

[7] O. Kovaleva, P. Podlesnaya, and A. Gratchev, “Resident microbiome of kidney tumors,” Oncol. Rev., vol. 18, 2024, doi: 10.3389/or.2024.1393664.

[8] “The microbiota and renal cell carcinoma,” 2023. PubMed PMID: 37878209. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/37878209/

[9] H. Xu et al., “Tumor microbiota of renal cell carcinoma affects clinical prognosis by influencing the tumor immune microenvironment,” Heliyon, vol. 10, no. 19, Art. no. e38310, 2024, doi: 10.1016/j.heliyon.2024.e38310.

[10] S. Leonard et al., “Microbiome in RCC—A systematic review of the literature,” 2024, doi: 10.1177/24684570241292246.

[11] S. Elyas, P. Barata, and U. Vaishampayan, “Clinical applications of microbiome in renal cell carcinoma,” Eur. Urol. Focus, 2024, doi: 10.1016/j.euf.2024.12.006.

[12] National Cancer Institute, “Hereditary kidney cancer syndromes.” Accessed: Nov. 1, 2026. [Online]. Available: https://www.cancer.gov/types/kidney/patient/kidney-genetics-pdq

[13] Espressif Systems, “ESP32 series datasheet.” Accessed: Nov. 1, 2026. [Online]. Available: https://www.espressif.com/

[14] Maxim Integrated (Analog Devices), MAX30102 High-Sensitivity Pulse Oximeter and Heart-Rate Sensor for Wearable Health, Datasheet.

[15] M. W. Sjoding, R. P. Dickson, T. J. Iwashyna, S. E. Gay, and T. S. Valley, “Racial bias in pulse oximetry measurement,” N. Engl. J. Med., vol. 383, no. 25, pp. 2477–2478, Dec. 2020, doi: 10.1056/NEJMc2029240.

[16] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, Jun. 2016, pp. 770–778.

[17] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in Proc. 36th Int. Conf. Mach. Learn. (ICML), Long Beach, CA, USA, Jun. 2019, pp. 6105–6114.

[18] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 10, October 2026

|DOI: 10.15680/IJIRSET.2026.1510108|

[19] Pallets Projects, "Flask documentation." Accessed: Oct. 30, 2026. [Online]. Available: https://flask.palletsprojects.com/

[20] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Venice, Italy, Oct. 2017, pp. 618–626.



INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi cross ref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details