IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

From Fragmented Capability to Unified Investigation: A Critical Review of Artificial Intelligence Across the Police Investigation Lifecycle

Goury S[1], Anjana V[2], Nythal A[3], Roobini R[4], Girija V R[5], Rini V R[6]

Undergraduate Student, Department of Computer Science & Engineering, College of Engineering Kottarakkara, Kollam, Kerala, India[1-4]

Head of the Department, Department of Computer Science & Engineering, College of Engineering Kottarakkara, Kollam, Kerala, India[5]

Assistant Professor, Department of Computer Science & Engineering, College of Engineering Kottarakkara, Kollam, Kerala, India[6]

Abstract: Criminal investigation still relies on manual complaint registration, hours of Closed-Circuit Television (CCTV) review and hand-drafted case files. Artificial intelligence (AI) has been applied to parts of this work, such as crime prediction, criminal network analysis, Automatic Number Plate Recognition (ANPR), face recognition and language-based reporting, but these tools remain disconnected. This paper reviews twenty studies published between 2022 and 2026 and identifies nine recurring gaps, including the absence of end-to-end integration, limited explainability, dataset bias and weak coupling with legal procedure. It then proposes Kaval AI, an officer-in-the-loop investigation assistant that connects seventeen modules through one shared case record, from Malayalam and English voice intake to Bharatiya Nyaya Sanhita (BNS) section prediction and charge sheet drafting, with every suggestion confirmed by an officer. The requirements, architecture, database design and a planned evaluation are presented; as a design-stage proposal, the paper reports no experimental results.

Keywords: Police Investigation, Crime Prediction, Graph Neural Networks, Number Plate Recognition, Face Recognition, Knowledge Graphs, Explainable AI, Bharatiya Nyaya Sanhita, Kaval AI

I. INTRODUCTION

Criminal investigation is slow and labour-intensive: an officer must register each complaint, examine statements, documents and footage, link them to people, vehicles and earlier cases, decide which legal provisions apply and draft the First Information Report (FIR) and charge sheet, largely by hand. Crime analytics has moved from K-Means and K-Nearest Neighbour (KNN) toward Long Short-Term Memory (LSTM) and other deep spatio-temporal forecasting [1][2]; graph frameworks such as DccGraph expose hidden criminal communities [3]; YOLO and Optical Character Recognition (OCR) pipelines read number plates [4][5][6][7][8]; face recognition has become lightweight and privacy-preserving [9][10][11]; and Natural Language Processing (NLP) has been applied to crime classification and to detecting crime reports in citizens' text [12][13].

Despite this progress, a recognised vehicle must still be linked by hand to a person, a person to a case, and a case to comparable incidents and the applicable law, because the reviewed tools stop at their own output. Most are also evaluated on English or foreign data, none produces the FIR or charge sheet, and the two surveys among the reviewed works [14][15] do not examine the full investigation lifecycle, which is the view taken in this paper.

To close this gap, we propose **Kaval AI**, an AI-powered police investigation assistant designed for police stations in Kerala, India, which connects seventeen modules covering the whole investigation. A complainant's spoken statement in Malayalam or English is to be converted into a structured First Information Statement (FIS), a draft record of the complainant's account from which the officer prepares the FIR. The system is also designed to digitise documents, translate statements, analyse CCTV footage, verify number plates, assist sketch drafting, match missing persons, link all outputs in one shared case record, suggest related cases and Bharatiya Nyaya Sanhita (BNS) sections, and draft the FIR,

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

case summary and charge sheet. Every suggestion will enter the case only after the officer accepts or edits it, so Kaval AI supports rather than replaces the officer's judgement.

The contributions of this paper are a comparison of twenty studies across five themes, nine research gaps traced to those studies, and the design of Kaval AI, which combines one shared case record, provenance from each observation to its evidence and model version, confidence carried through the record, and officer confirmation before any output becomes a case fact. None of the reviewed studies combines these elements; the AI techniques are taken from the literature, and no module has yet been implemented or tested.

Section II reviews the literature, Section III discusses it, Section IV lists the research gaps, Section V presents the proposed system, Section VI sets out the planned evaluation and Section VII concludes.

II. LITERATURE REVIEW AND ANALYSIS

A. Review Methodology

This review covers twenty peer-reviewed studies published between 2022 and 2026, drawn mainly from IEEE Xplore and supplemented by open-access journals, selected to span the investigation lifecycle from complaint registration to documentation. Search terms combined crime prediction, criminal network analysis, graph neural network, number plate recognition, face recognition for criminal identification, forensic sketch synthesis, knowledge graphs for crime reports and language models for policing. A study was included if it addressed a defined investigation stage, described its method in enough detail for comparison, reported a quantitative or structured qualitative evaluation, appeared in English in a peer-reviewed venue, and contributed a technique, modality or evaluation not duplicated by another included study; preprints, technical blogs and vendor material were excluded. The review is narrative and purposive rather than systematic or exhaustive, because the studies report measures that cannot be compared across tasks.

The studies were organised into five themes: crime analytics, prediction and predictive policing [1][2][14][15]; graph-based criminal network and evidence analysis [3][16][17]; vehicle identification through number plate recognition [4][5][6][7][8]; suspect identification, privacy and sketch synthesis [9][10][11][18][19][20]; and language models, reporting and investigation support [12][13]. Each study is compared on its contribution, reported results and remaining gap (Table I), with attention to integration with casework, explainability, data origin, robustness, privacy and coupling with legal procedure. Results attributed to a study are those reported by its authors; assessments of investigative relevance are this review's interpretation.

Why these twenty studies: each study adds something the others do not, and together they cover every stage from complaint to documentation. Rajesh et al. [1] is one of the few reviewed studies on Indian data, Shahmoradi et al. [2] show what spatio-temporal deep models add to fine-grained hotspot forecasting, and [14] and [15] are surveys of the wider field. DccGraph [3] reports large community-detection gains over eight baselines, and [16] and [17] are the only studies that build evidence or motive knowledge graphs. The plate studies [4][5][6][7] cover different detectors and conditions, including fog, rain and low light, and Gayaki et al. [8] alone checks plate authenticity against Indian registration records. Among suspect studies, [9] and [18] represent accuracy-focused and applied identification, [10] and [11] the two privacy-preserving approaches, and [19] and [20] the only witness-driven retrieval and sketch methods. Sarzaeim et al. [12] is the only study testing large language models on crime records, and Palma-Preciado et al. [13] the only study that detects crime reports in citizens' free text. Studies that repeated one of these techniques without a new modality, condition or evaluation were left out.

B. Crime Analytics, Prediction and Predictive Policing

Rajesh et al. reduced Indian metropolitan crime records to 5,038 occurrences with thirty-five attributes and applied K-Means clustering and K-Nearest Neighbour classification [1]. Shahmoradi et al. combined a spatio-temporal residual network (ST-ResNet) with LSTM and weather, calendar and park covariates to forecast daily theft hotspots in Chicago (242,770 records, 2014–2016), reaching a hit rate of 88.7% for the top 20% of 500 m cells and a Predictive Accuracy Index (PAI) of 4.44 [2]. The survey of Ersöz et al. analysed 142 crime prediction studies, found Random Forest the most used and best performing technique, and reported that 60.0% of datasets come from the United States and only 7.78% from India [14]; Monika and Kumar reviewed an architecture combining records, Geographic Information System (GIS) data, social media and surveillance [15]. These systems indicate where and when crime is likely but do not assist with an individual registered case.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

C. Graph-Based Criminal Network and Evidence Analysis

DccGraph, a Graph Neural Network (GNN) framework, builds an augmented criminal network, balances community distribution and applies a criminal-oriented Siamese graph encoder, outperforming eight baselines by an average of 35.01% in accuracy and 61.59% in F1 and recovering minority groups that modularity-based methods miss [3]. Mazepa et al. built evidence knowledge graphs from homicide reports using Named Entity Recognition (NER), coreference resolution, Support Vector Machine (SVM) classification and ontology construction [16], and Ugai extended a criminal-motive ontology with 1,069 Wikidata classes and used TransE link prediction to score a property motive at 0.9 on a Sherlock Holmes narrative [17]. All three studies assume that a graph has already been built from transactions, parsed reports or a curated story.

D. Vehicle Identification Through Number Plate Recognition

Number plate recognition is the most mature theme. Khan et al. used YOLOv4 for mixed Asian and European plates, reporting 98.42% detection mean average precision (mAP) and character accuracy of 98% (AlexNet), 96% (R-CNNL3) and 90.94% (OCR) [4]. Patil et al. paired YOLOv8 with EasyOCR at about 86% real-time accuracy [5], Pabba et al. reported about 99% validation accuracy with contrast enhancement [7], and Mustafa and Karabatak added make-and-model recognition and Grad-CAM, reporting 97.5% on about 1,000 images in fog, rain and low light [6]. Gayaki et al. checked plate text and vehicle colour against registration records to detect forged Indian plates, with 92% accuracy on fake plates falling to 85% in rain [8]. Robustness is claimed more often than measured, and no pipeline links a vehicle to a case.

E. Suspect Identification, Privacy and Sketch Synthesis

Face-related work covers accuracy, privacy and reconstruction. Munusamy and Senthilkumar distilled a VGG16, ResNet50 and DenseNet121 ensemble into a compact student reaching 96.78% [9], and Teja et al. matched CCTV faces against a criminal database without reporting accuracy [18]. MinusFace stores only the difference between an original and a regenerated face, resisting recovery attacks [10], and SynthDistill trained lightweight models on StyleGAN faces, reaching 99.52% on the Labeled Faces in the Wild (LFW) benchmark without real identities [11]. FaIRCoP lets a witness's selections steer the search toward the remembered suspect [19], and HiFiSketch uses Generative Adversarial Network (GAN) inversion and CLIP text guidance for editable sketches [20]. However, the applied system of Teja et al. [18] adopts none of these privacy mechanisms.

F. Language Models, Reporting and Investigation Support

Sarzaeim et al. compared BART, GPT-3 and GPT-4 for crime classification; fine-tuned GPT-3 reached 97% on San Francisco data but the best Los Angeles result was 68%, showing strong dataset dependence [12]. Palma-Preciado et al. separated crime reports from other posts on X in Mexico City; the Spanish BERT model BETO reached an F1 of 0.86 against 0.78 for the best SVM or Random Forest baseline on term frequency–inverse document frequency (TF-IDF) or embedding features, and the mapped reports agreed only partly with official records (Spearman's $\rho = 0.37$) [13]. It is closest to citizen reporting but uses Spanish social-media posts, not police complaints. Neither study handles visual or relational evidence or maps facts onto statutory provisions.

G. Comparative Analysis

Table I summarises each study, its contribution and the gap it leaves open. Work has moved from perception and text extraction in 2022 toward graph learning, privacy-preserving recognition and language models from 2023 onward, yet nine of the twenty studies concentrate on perception, only one works with citizen crime reports, none addresses complaint registration, and none addresses documentation.

TABLE I DETAILS ABOUT PREVIOUS STUDIES ON AI FOR POLICE INVESTIGATION AND THEIR GAPS

Title of Study	Author and Year	Research Gain	Gaps in Study
Machine Learning Based Crime Identification System Using Data Analytics	Rajesh et al. (2023) [1]	K-Means clustering, KNN identification and map visualisation over 5,038 Indian crime records with 35 attributes	Structured historical records only; cannot ingest complaint narratives
Hybrid ST-ResNet and LSTM Approach for Precise Crime Hotspot Prediction	Shahmoradi et al. (2025) [2]	ST-ResNet with LSTM and weather, calendar and park covariates; 88.7% hit rate (top 20% of 500 m cells), PAI 4.44	Hotspot forecasting for one crime type in one city; no support for an individual case

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

Title of Study	Author and Year	Research Gain	Gaps in Study
DccGraph: Detecting Criminal Communities with Augmented Criminal Network Construction and Graph Neural Network	Yang et al. (2023) [3]	Augmented criminal network and Siamese graph encoder; beat eight baselines by 35.01% ACC and 61.59% F1 on average; finds minority groups	Presupposes an existing network; applied case reports no outcome measures
A Novel Deep Learning Based ANPR Pipeline for Vehicle Access Control	Khan et al. (2022) [4]	YOLOv4 detection for mixed plate layouts (98.42% mAP); character accuracy 98% (AlexNet), 96% (R-CNNL3) and 90.94% (OCR)	Access control only; sensitive to angle and speed
Vehicle Number Plate Detection Using YOLOv8 and EasyOCR	Patil et al. (2023) [5]	Filtering and edge preprocessing before YOLOv8 and EasyOCR; about 86% real-time accuracy on live video	Lighting and blur errors; a misread plate can misdirect an inquiry
Real Time Car Model and Plate Detection System by Using Deep Learning Architectures	Mustafa and Karabatak (2024) [6]	Make-and-model plus plate recognition with Grad-CAM; 97.5% on about 1,000 real images in fog, rain and low light	No link from vehicle to case records
Automatic Number Plate Detection Using Deep Learning and OCR	Pabba et al. (2024) [7]	Contrast enhancement with a convolutional neural network (CNN) and OCR; about 99% validation accuracy	Near-saturated accuracy suggests easy data; no degraded-condition results
Fake Number Plate Detection Using Vehicle Image Analysis: A Robust Verification System	Gayaki et al. (2025) [8]	Plate OCR and vehicle colour checked against registration; 95% on authentic and 92% on fake plates	Falls to 85% in rain; plate verification only
Leveraging Lightweight Hybrid Ensemble Distillation (HED) for Suspect Identification with Face Recognition	Munusamy and Senthilkumar (2025) [9]	Three-network ensemble distilled into a compact student; 96.78% at lower cost	Low-resolution CCTV performance unknown
Privacy-Preserving Face Recognition Using Trainable Feature Subtraction	Mi et al. (2024) [10]	MinusFace residual representation with channel shuffling resists recovery attacks	Recognition component, not an investigative workflow
Knowledge Distillation for Face Recognition Using Synthetic Data with Dynamic Latent Sampling	Shahreza et al. (2024) [11]	Lightweight models trained on StyleGAN faces; 99.52% on LFW without real identities	No case integration; transfer to CCTV untested
A Framework for LLM-Assisted Smart Policing System	Sarzaeim et al. (2024) [12]	BART, GPT-3 and GPT-4 crime classification; fine-tuned GPT-3 reached 97% on San Francisco data	Only 68% on Los Angeles data; English only; no statutory mapping
Exploring Crime Reports from Social Media Through Transformer-Based Text Classification and Geovisualization	Palma-Preciado et al. (2026) [13]	Crime-report detection in X posts; BETO F1 0.86 against 0.78 for the best SVM baseline; reports mapped against official records	Spanish social-media text, not police complaints; only partial agreement with official records (ρ = 0.37)
Artificial Intelligence in Crime Prediction: A Survey with a Focus on Explainability	Ersöz et al. (2025) [14]	Analysed 142 studies; Random Forest most used and best; SHAP the most common post-hoc method, followed by LIME	Explainable AI underdeveloped; 60.0% of datasets from the US, 7.78% from India
Advancements in AI-Based Crime Detection and Prediction	Monika and Kumar (2024) [15]	Reviews a multi-source architecture of records, GIS, social media and surveillance	No implementation; bias, privacy and transparency unresolved

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

Title of Study	Author and Year	Research Gain	Gaps in Study
Relationships Knowledge Graphs Construction Between Evidence Based on Crime Reports	Mazepa et al. (2022) [16]	NER, coreference and SVM build evidence knowledge graphs from homicide reports; >92% text-analysis target	Text evidence only; accuracy stated as a target, not an investigative outcome
Criminal Investigation with Augmented Ontology and Link Prediction	Ugai (2023) [17]	Ontology extended with 1,069 Wikidata classes; TransE scored a motive at 0.9	Single fictional (Sherlock Holmes) graph
Face Detection and Recognition for Criminal Identification	Teja et al. (2023) [18]	Haar Cascade and CNN embeddings match CCTV faces to a criminal database	No quantitative accuracy; no privacy safeguard
A Suspect Identification Framework Using Contrastive Relevance Feedback	Gupta et al. (2023) [19]	Witness feedback with Separating Cluster Loss converges on the suspect and lowers witness burden	Needs an existing gallery; not linked to cases
HiFiSketch: High Fidelity Face Photo-Sketch Synthesis and Manipulation	Peng et al. (2023) [20]	GAN inversion and CLIP text guidance give editable, high-fidelity sketches	Sketch only; not fed back into recognition

III. DISCUSSION

A. Accuracy versus Investigative Value

Reported accuracy is highest for narrow tasks such as plate and face recognition [4][6][7][11], while studies of investigative questions, such as which individuals form a group or which motive is plausible, report lower or less comparable figures [3][17]. A plate reader with 99% accuracy that cannot say which case a plate belongs to contributes less than its accuracy suggests, so components should be chosen by how well their outputs can be linked and verified.

B. Realism versus Evaluation Rigour

The best-evaluated studies use public benchmarks [11][20], while those closest to policing rely on private or local data [1][8][18], because investigative data cannot be shared freely. Evaluation-set sizes are reported unevenly, Teja et al. [18] report no quantitative accuracy, and none of the studies reports whether its output helped an officer reach a correct conclusion, so governed, anonymised Indian benchmarks would be a valuable contribution.

C. Privacy, Explainability and Accountability

Explainability and privacy remain weak in applied systems (gaps 2 and 5) [6][10][11][14][15][18], and predictive systems trained on recorded crime may learn where policing occurred as much as where crime occurred, so human confirmation and auditability must be designed in from the start.

D. What Works and What Remains Fragmented

On benchmark and locally collected data, the reviewed systems handle single tasks well, but tasks that cross modalities remain fragmented: linking a vehicle to a person and a case, reasoning over accumulated evidence and mapping findings onto legal categories. What is missing is a shared case representation through which these outputs can be exchanged.

IV. RESEARCH GAPS

Nine gaps follow from the twenty studies.

1) Absence of End-to-End Integration: No study carries a case from registration to documentation; prediction [1][2][14][15], number plate [4][5][6][7][8], face [9][10][11][18] and graph [3][16][17] systems each stop at their own output.

2) Limited Explainability and Accountability: Explainable AI remains unresolved in crime prediction [14][15], and apart from Grad-CAM in [6], applied systems such as that of Teja et al. [18] do not justify their output in a reviewable form.

3) Dataset and Geographic Bias: Only 7.78% of the datasets surveyed in [14] come from India, against 60.0% from the United States, and few of the reviewed studies, such as Rajesh et al. [1] and Gayaki et al. [8], report results on Indian data.

4) Robustness Under Deployment Conditions: Low light, blur, occlusion and weather degrade results [4][5][9][18], yet such results are rarely reported separately; Gayaki et al. report a fall from 92% to 85% in rain [8].

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

5) Privacy Safeguards in Isolation: Privacy-preserving recognition [10][11] has not reached applied identification systems [18].
6) Minimal Multilingual and Multimodal Capability: Each text framework works in a single language, English in [12] and Spanish in [13], and none handles Malayalam, spoken or handwritten input or combines text, visual and relational evidence.
7) Absence of Field Validation: No study reports validation in live police casework; most are evaluated once on retrospective or benchmark data, and the real-case application of DccGraph [3] reports no outcome measures.
8) Limited Attention to Evidence Integrity: Only forged-plate detection [8] and resistance to recovery attacks [10] consider adversarial conditions; tampered media and duplicate complaints are not examined.
9) Weak Coupling with Legal Procedure: None of the twenty reviewed studies maps findings onto statutory provisions or drafts an FIR or charge sheet.

V. PROPOSED SYSTEM: KAVAL AI

A. Problem Statement and Scope

Current tools leave the officer to connect complaints, evidence, earlier cases and the law by hand. Kaval AI is an officer-in-the-loop assistant designed to register complaints and evidence, turn them into structured observations linked in one case record, suggest related cases and applicable Bharatiya Nyaya Sanhita (BNS) sections, and draft case documents, with every suggestion traceable to its evidence and confirmed by an officer before it enters the case. This section describes the proposed design; no module has yet been implemented.

The system is delivered through a web interface for police stations and covers the seventeen modules in Table II. Digital and cyber forensics, speaker and gait analysis, aerial surveillance, autonomous decision-making and live deployment in police networks are outside the scope of this work.

B. System Architecture

Fig. 1 shows the architecture of Kaval AI. The input is a complaint, given as a spoken or written statement, together with any documents and CCTV footage. It passes through six stages:

1) Acquisition: Each statement, document and video is registered with its source, time and custodian.
2) Perception: Footage, number plates, sketches and faces are turned into structured observations.
3) Shared Case Record: These observations are linked into one case graph of persons, vehicles, places and evidence.
4) Reasoning: The system suggests related cases, applicable BNS sections and likely hotspots, each ranked by confidence.
5) Officer Review: Every suggestion is shown with its supporting evidence, so that the officer can accept, edit or reject it. The decision is logged, and rejected or edited suggestions are fed back to the module that produced them.
6) Documentation: The FIR, case summary and charge sheet are drafted from confirmed facts only.

Three shared components support all six stages. *Governance and Audit* applies role-based access, audit logging and retention limits at every stage; *External Sources* supply the BNS text, crime and missing-person records and vehicle registration details; and the *Case Database* stores cases, entities, evidence, drafts and the audit log.

C. Module Description

Table II lists the seventeen modules by stage; Officer Review (stage 5) is an interface step rather than an AI module, and "partial" in the Basis column means the cited study covers only part of the task. Each module connects to the shared case record through a fixed interface, and the three documentation modules are new.

TABLE II MODULES OF KAVAL AI AND THE LITERATURE THAT INFORMS THEM

Module	Function	Basis
Acquisition		
1. Voice-to-FIS Generation	Speech in Malayalam or English is converted to text, and the complainant, accused, time, place and incident are extracted into FIS fields.	New (speech recognition and FIS extraction are proposed components)
2. OCR Document Digitisation	Printed and handwritten documents are read by OCR with layout parsing, and the text is stored as evidence with its source.	New ([4][6] use OCR for plates only)

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

Module	Function	Basis
3. Multilingual Translation of Witness Statements	Witness statements are translated between Malayalam and English by neural machine translation, and the original is kept beside the translation.	New
Perception		
4. CCTV Video Analysis	Persons and vehicles in footage are detected with YOLO, and time-stamped observations are added to the case.	Partial [5][6] (vehicles only)
5. Vehicle and Number Plate Recognition	Plates are detected with YOLO and read by OCR, and the registered colour and model are checked to flag fake plates.	[4][5][6][7][8]
6. AI-based Facial Sketch Generation	Proposed extension: a sketch is drafted from a witness description by text-guided sketch editing as in [20] and refined by witness feedback as in [19]; direct generation from a description is not shown in the reviewed studies and needs development.	Partial [19][20] (feedback retrieval; photo-to-sketch with text editing)
7. Missing Person Identification	Face embeddings are matched against reported missing persons, and facial data are held in protective representations as in [10][11].	[9][10][11]
Shared Case Record		
8. Criminal Community and Hidden Network Analysis	Communities are detected with a GNN over the case graph of persons, contacts and vehicles.	[3]
9. Evidence Classification	Evidence items are classified from their text and images, and each item is linked to the entities it concerns.	Partial [16] (text only)
10. Duplicate FIR Detection	Registered complaints are compared by text similarity to flag repeated registrations.	New ([13] classifies report text only)
Reasoning		
11. BNS Section Prediction and Legal Guidance	Case facts are classified onto BNS sections by NLP, and each section is shown with the facts that support it.	Partial [12] (crime-type classification)
12. Crime Hotspot Prediction	Crime hotspots are predicted from spatio-temporal crime data with Random Forest and LSTM-based models.	[2][14]
13. Crime Analytics	Crime patterns and trends are found by clustering and shown on dashboards.	[1][15]
14. Intelligent Search for Related Cases and Evidence	Related cases and evidence are found by semantic search and graph queries across cases.	Partial [3][17] (graph links only)
Documentation		
15. Automatic FIR Generation	An FIR draft is generated from confirmed FIS fields using templates, for officer review and the informant's signature.	New
16. Case Summary Generation	A summary of the case record and its evidence is generated.	New
17. Charge Sheet Generation	A draft police report (charge sheet) is generated from confirmed facts and BNS sections, for officer review.	New

© IJARCCE This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

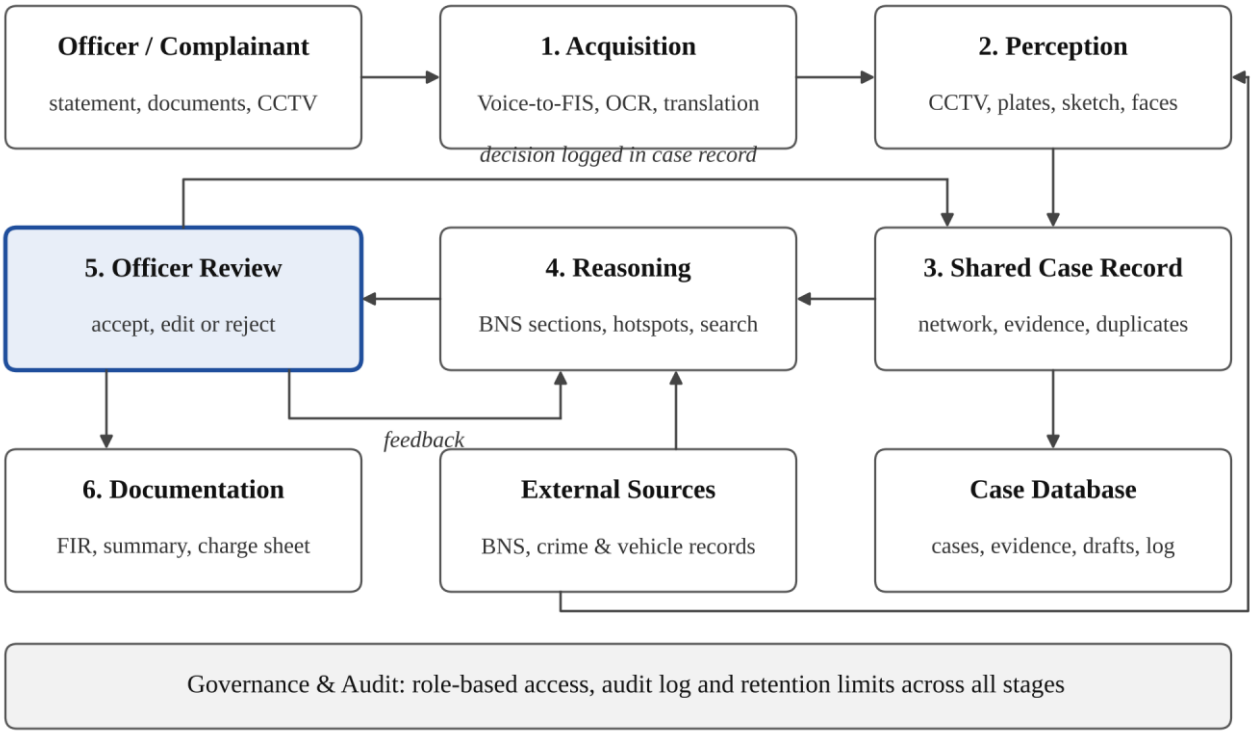


Fig. 1 Architecture of Kaval AI

D. Officer Confirmation, Uncertainty and Provenance

Every suggestion carries a status (pending, accepted, edited or rejected), a confidence value and its model version. Suggestions below a configurable threshold appear only in a separate low-confidence list and never pre-fill a document, and only accepted or edited suggestions become case facts, which responds to the accountability concern raised in [14] and [15].

Provenance is recorded at every step. Each registered item stores its SHA-256 hash, source, time and custodian; each derived observation, such as a plate string or a face match, stores the ID of its evidence item and the model version; and each link in the case graph cites the observations that support it, so any fact in a draft can be traced back to the raw evidence. Confidence is carried along these chains: a derived link is never given a higher confidence than the weakest link supporting it, so a single misread plate character cannot silently create a false suspect. The case graph is stored in the relational database itself, with entities as rows of the ENTITY table and links as rows of the RELATION table (Table IV).

Entity resolution merges observations of the same person, vehicle or place automatically only on exact keys, such as a registration number, while fuzzy matches on names or faces are proposed to the officer. A canonical ID links the same entity across cases without merging the cases themselves. Conflicting observations, such as two different reads of one plate, are both kept and flagged for the officer. A confirmed suggestion keeps its original model confidence for the audit trail, and a rejected one leaves the case facts but stays in the log.

Officer edits and rejections are kept as feedback for retraining, but because such labels may be noisy or biased, feedback will be reviewed before use, retrained models will be validated on held-out data and versioned for rollback, and only anonymised, confirmed records will be used. If no suggestion is acceptable, the officer can continue manually.

E. Requirements

1) Functional Requirements:

The functional requirements follow from the seventeen modules in Table II, each of which defines one required function. Two further functions apply to all modules: every suggestion must be shown with its evidence and confidence so that the officer can accept, edit or reject it, and every query, suggestion and decision must be logged.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

2) Non-Functional Requirements:

The non-functional requirements are listed in Table III.

TABLE III NON-FUNCTIONAL REQUIREMENTS OF KAVAL AI (PROPOSED TARGETS)

Category	Requirement
Correctness	No suggestion enters a case record or document without officer confirmation. The target plate-reading accuracy is at least 90% on daytime footage, against a reviewed range of about 86% to 99% [4][5][6][7][8].
Performance	The target is an FIS draft within one minute of the end of a statement. CCTV footage should be processed at least at real-time speed on a single-GPU station.
Security and privacy	Access is role-based, storage is encrypted and case data are processed on premises. Facial data are held in protective representations [10][11] with retention limits.
Legal alignment	Outputs are aligned with the BNS, the Bharatiya Nagarik Suraksha Sanhita and the Bharatiya Sakshya Adhiniyam. Personal data are handled in line with the Digital Personal Data Protection Act, 2023 (Section V-G).
Explainability	Each suggestion shows its supporting facts, graph path or image region, confidence and model version.
Usability	The interface is available in Malayalam and English, officers without technical training can use it, and all drafts are editable.
Reliability	Each module fails independently, and the officer can continue manually if a module is unavailable.
Maintainability	Each module can be replaced behind its fixed interface to the shared case record.
Auditability	An append-only, tamper-evident log records every query, suggestion and officer decision.
Portability	The web client runs on station workstations, and the server uses Python 3.10 or later with PyTorch.

F. Data Flow and Database Design

Statements, documents and footage flow from the complainant and officer into the case record, external sources supply registration, crime and BNS data, suggestions flow to officer review, confirmed facts flow to documentation, and rejected suggestions return to the modules that produced them.

For example, in a vehicle theft reported in Malayalam, Voice-to-FIS Generation drafts the statement, Duplicate FIR Detection checks for an earlier registration, plate recognition flags a suspect vehicle whose plate does not match its registered colour and model, Intelligent Search finds earlier thefts involving it, network analysis links suspects who share contacts, and BNS Section Prediction suggests sections; once the officer confirms them, the FIR, case summary and charge sheet are drafted for editing.

A relational database with nine tables, summarised in Table IV, holds the case data and the audit trail. Entities describing the same real-world person or vehicle share a canonical_id across cases, each draft lists in fact_refs the confirmed facts it uses, and each audit-log entry stores the hash of the previous entry, so any later change to the log can be detected.

TABLE IV DATABASE TABLES OF KAVAL AI

Table	Purpose	Key fields
CASE	One investigation	case_id, station, officer_id, status, created_at
COMPLAINT	Registered statement (FIS)	complaint_id, case_id, language, audio_ref, fis_text
EVIDENCE	Documents, footage and images	evidence_id, case_id, type, source, custodian, sha256
ENTITY	Persons, vehicles and places	entity_id, case_id, canonical_id, entity_type, attributes, confidence
RELATION	Links in the case graph	relation_id, from_entity, to_entity, relation_type, evidence_id
BNS_SECTION	Statutory provisions	section_no, title, description

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510115

TABLE IV (CONTINUED)

Table	Purpose	Key fields
SUGGESTION	Module outputs for review	suggestion_id, case_id, module, confidence, model_version, status
DOCUMENT	FIR, summary and charge sheet drafts	document_id, case_id, doc_type, version, fact_refs, approved_by
AUDIT_LOG	Queries and decisions	log_id, user_id, action, target_id, timestamp, prev_hash

G. Legal and Ethical Safeguards

BNS section prediction is advisory: each section is shown with its supporting facts and confidence, the officer may record their own assessment first to avoid anchoring, suggestions are checked against the BNS text in force, and special and local laws are outside the current scope. Drafting must also follow the Bharatiya Nagarik Suraksha Sanhita (BNSS). Under Section 173, information given orally must be reduced to writing, read over to the informant and signed, and information given by electronic communication is taken on record once the informant signs it within three days, so the FIS is only a draft that the officer reads back to the informant for correction and signature. Separately, the charge sheet is a draft of the police report under Section 193, which must state the sequence of custody in the case of an electronic device. All AI-generated drafts, suggestions and matches require officer review, and the officer and prosecution remain responsible for the legal characterisation of an offence.

False identification is the most serious risk: a face or plate match is a lead to be checked, never the sole basis for an arrest or a charge, and it is shown with its confidence and the image regions it relied on. Before any use, error rates will be compared across demographic groups, regions and languages, and a module with widely differing error rates will not be used until the difference is understood. Hotspot predictions will be reviewed for patrol-feedback bias, in which more heavily patrolled areas record more crime.

Under Section 17(1)(c) of the Digital Personal Data Protection Act, 2023, processing in the interest of the prevention, detection, investigation or prosecution of an offence is exempt from most of Chapter II, from Chapter III and from Section 16, but Section 8(1) and the duty of reasonable security safeguards under Section 8(5) still apply. Kaval AI nevertheless adopts data minimisation, purpose limitation, role-based access, encryption, on-premises processing and retention limits on facial data as design choices, and how the exemption applies to each module requires qualified legal review. Every query, suggestion and decision enters an append-only, tamper-evident audit log, and drafts stay marked as AI-assisted until signed. Original evidence is stored unchanged under its hash, and processed versions such as enhanced frames are stored separately. These measures support traceability and integrity, but they do not by themselves make evidence admissible: admissibility of electronic records under Section 63 of the Bharatiya Sakshya Adhiniyam (BSA) depends on the conditions and certificate that section prescribes, which also requires qualified legal review.

H. How Kaval AI Addresses the Gaps

The design responds to each gap without yet closing it: the shared case record (gap 1), evidence-linked suggestions (gap 2), BNS categories, Indian plates and Malayalam intake (gap 3), degraded-condition testing (gap 4), the safeguards in Section V-G (gap 5), multilingual and multimodal intake (gap 6), a mock-case officer study that covers field validation only in part (gap 7), duplicate detection, fake-plate checks and hashed evidence (gap 8), and BNS prediction with the drafting modules (gap 9). Gaps 1, 3, 6 and 9 are targeted directly, gaps 2, 5 and 8 in part, and gaps 4 and 7 depend on evaluating the built system.

VI. PLANNED EVALUATION AND FUTURE DIRECTIONS

Table V sets out the planned evaluation; none of its measures has yet been taken.

TABLE V PLANNED EVALUATION OF KAVAL AI

Module Group	Technical / Investigative Measures	Baseline	Data	Proposed Pass Criterion
Speech, OCR and translation	Word error rate, FIS field F1, translation quality (character n-gram F-score, chrF); FIR registration time	Manual FIR registration	Mock complaints in Malayalam and English, recorded with consent	Draft within one minute (Table III); faster than manual



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.1510115

Module Group	Technical / Investigative Measures	Baseline	Data	Proposed Pass Criterion
CCTV and number plates	Detection mAP, character-level and whole-plate accuracy, normal and degraded; correct and false leads per case	YOLOv8 and EasyOCR as in [5]	Annotated Indian plate footage with low-light and rain subsets	At least 90% daytime plate accuracy (Table III); degraded subsets reported separately
Sketch and missing person	Retrieval rank; true-match rate at a fixed false-match rate, by group; witness effort	Feedback retrieval as in [19]; embeddings as in [9]	LFW (web photos, not CCTV) and consented local images	False-match rate and group gap fixed before testing
Network and evidence analysis	Accuracy, normalised mutual information (NMI), F1; officer-confirmed associates	Modularity detection vs. GNN as in [3]	Anonymised cases with officer-labelled communities	Higher F1 than the modularity baseline
Hotspot prediction	Hit rate and Predictive Accuracy Index (PAI); officer-judged usefulness	ST-ResNet–LSTM [2]; Random Forest [14]	Historical crime records, time-based split	At least the best baseline on held-out months
BNS prediction and drafts	Micro- and macro-F1, top-3 recall; edits before acceptance	Proposed TF-IDF text classifier (planned baseline)	Past FIRs labelled by legal experts	Higher top-3 recall than TF-IDF; officer reviews every draft

Test data will be kept separate from training data through time-based, case-level and, for faces, identity-level splits, with degraded-condition results reported separately. Mock cases will compare officer time and correct leads with and without the shared case record. Splits, seeds, model versions and scripts will be recorded; local data will be anonymised and used only with institutional approval; and the officer study will have a pre-set sample size, ethics approval and 95% confidence intervals.

The review is purposive, excluding digital forensics, cybercrime, speaker and gait analysis and aerial surveillance, and its figures are taken as reported. The targets in Tables III and V may change once data are collected, and ethical risks are addressed in Section V-G. The shared case record will be built first, then the Malayalam language and OCR modules, then the reasoning and drafting modules.

VII. CONCLUSION

This paper reviewed twenty studies on AI for police investigation published between 2022 and 2026 and identified nine gaps, chief among them the absence of end-to-end integration and weak coupling with legal procedure. Kaval AI responds with a proposed officer-in-the-loop design of seventeen modules linked through one shared case record, together with an evaluation plan of baselines and proposed pass criteria. Its contribution is this architecture and plan, not a validated or deployed system; building and testing it with practising officers is the next step.

REFERENCES

[1]. S. M. Rajesh, I. Chiranmai and N. JayaPandian, "Machine learning based crime identification system using data analytics," in *Proc. Int. Conf. Sustainable Communication Networks and Application (ICSCNA)*, 2023, pp. 951–956, doi: 10.1109/ICSCNA58489.2023.10370717.

[2]. N. Shahmoradi, A. A. Alesheikh, A. Jafari and A. Lotfata, "Hybrid ST-ResNet and LSTM approach for precise crime hotspot prediction," *Sci. Rep.*, vol. 15, Art. no. 40754, 2025, doi: 10.1038/s41598-025-24559-7.

[3]. Y. Yang, L. Yang, L. Li, X. Ma, L. Yu and C. Zuo, "DccGraph: Detecting criminal communities with augmented criminal network construction and graph neural network," in *Proc. Int. Joint Conf. Neural Networks (IJCNN)*, 2023, pp. 1–8, doi: 10.1109/IJCNN54540.2023.10191121.

[4]. M. G. Khan, S. Salma, M. Saeed, A. Zulfiqar, Y. Y. Ghadi and M. Adnan, "A novel deep learning based ANPR pipeline for vehicle access control," *IEEE Access*, vol. 10, pp. 64172–64184, 2022, doi: 10.1109/ACCESS.2022.3183101.

DOI: 10.17148/IJARCCE.2026.1510115

[5]. S. S. Patil et al., "Vehicle number plate detection using YOLOv8 and EasyOCR," in *Proc. 14th Int. Conf. Computing Communication and Networking Technologies (ICCCNT)*, 2023, pp. 1–4, doi: 10.1109/ICCCNT56998.2023.10307420.

[6]. T. Mustafa and M. Karabatak, "Real time car model and plate detection system by using deep learning architectures," *IEEE Access*, vol. 12, pp. 107616–107630, 2024, doi: 10.1109/ACCESS.2024.3430857.

[7]. P. Pabba, A. Ram, L. Saiprasad and M. Sathwik, "Automatic number plate detection using deep learning and OCR," in *Proc. IEEE 4th Int. Conf. ICT in Business Industry & Government (ICTBIG)*, 2024, pp. 1–5, doi: 10.1109/ICTBIG64922.2024.10911385.

[8]. U. Gayaki, O. Nishane, N. Warkad and S. Bansod, "Fake number plate detection using vehicle image analysis: A robust verification system," in *Proc. 2nd Int. Conf. Electronic Circuits and Signaling Technologies (ICECST)*, 2025, pp. 1051–1056, doi: 10.1109/ICECST66106.2025.11307270.

[9]. V. Munusamy and S. Senthilkumar, "Leveraging lightweight hybrid ensemble distillation (HED) for suspect identification with face recognition," *IEEE Access*, vol. 13, pp. 2112–2126, 2025, doi: 10.1109/ACCESS.2024.3523101.

[10]. Y. Mi et al., "Privacy-preserving face recognition using trainable feature subtraction," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 297–307, doi: 10.1109/CVPR52733.2024.00036.

[11]. H. O. Shahreza, A. George and S. Marcel, "Knowledge distillation for face recognition using synthetic data with dynamic latent sampling," *IEEE Access*, vol. 12, pp. 187800–187812, 2024, doi: 10.1109/ACCESS.2024.3505621.

[12]. P. Sarzaeim, Q. H. Mahmoud and A. Azim, "A framework for LLM-assisted smart policing system," *IEEE Access*, vol. 12, pp. 74915–74929, 2024, doi: 10.1109/ACCESS.2024.3404862.

[13]. C. Palma-Preciado, M. Saldana-Perez and G. Sidorov, "Exploring crime reports from social media through transformer-based text classification and geovisualization," *ISPRS Int. J. Geo-Inf.*, vol. 15, no. 9, Art. no. 397, 2026, doi: 10.3390/ijgi15090397.

[14]. F. Ersöz, T. Ersöz, F. Marcelloni and F. Ruffini, "Artificial intelligence in crime prediction: A survey with a focus on explainability," *IEEE Access*, vol. 13, pp. 59646–59674, 2025, doi: 10.1109/ACCESS.2025.3553934.

[15]. E. Monika and T. R. Kumar, "Advancements in AI-based crime detection and prediction," in *Proc. 8th Int. Conf. Electronics, Communication and Aerospace Technology (ICECA)*, 2024, pp. 1557–1562, doi: 10.1109/ICECA63461.2024.10800911.

[16]. S. Mazepa, V. Vysotska, D. Ivanchyshyn, L. Chyrun, V. Schuchmann and Y. Ryshkovets, "Relationships knowledge graphs construction between evidence based on crime reports," in *Proc. IEEE 17th Int. Conf. Computer Sciences and Information Technologies (CSIT)*, 2022, pp. 165–171, doi: 10.1109/CSIT56902.2022.10000587.

[17]. T. Ugai, "Criminal investigation with augmented ontology and link prediction," in *Proc. IEEE 17th Int. Conf. Semantic Computing (ICSC)*, 2023, pp. 288–289, doi: 10.1109/ICSC56153.2023.00056.

[18]. K. P. Teja, G. D. Kumar and T. P. Jacob, "Face detection and recognition for criminal identification," in *Proc. 8th Int. Conf. Communication and Electronics Systems (ICCES)*, 2023, pp. 1431–1435, doi: 10.1109/ICCES57224.2023.10192845.

[19]. D. Gupta et al., "A suspect identification framework using contrastive relevance feedback," in *Proc. IEEE/CVF Winter Conf. Applications of Computer Vision (WACV)*, 2023, pp. 4361–4369, doi: 10.1109/WACV56688.2023.00434.

[20]. C. Peng, C. Zhang, D. Liu, N. Wang and X. Gao, "HiFiSketch: High fidelity face photo-sketch synthesis and manipulation," *IEEE Trans. Image Process.*, vol. 32, pp. 5865–5876, 2023, doi: 10.1109/TIP.2023.3326680.

 This work is licensed under a Creative Commons Attribution 4.0 International License