Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

Victim Compensation Laws and Restorative Justice

Olivia Clark
Independent Researcher
Vancouver, Canada, CA, V6B 1A1

http://www.jcclls.org/ || **Vol. 2 No. 3 (2026): August Issue**

Date of Submission: 26-06-2026	**Date of Acceptance:** 10-07-2026	**Date of Publication:** 04-08-2026

Abstract— Victim compensation has become an increasingly important component of contemporary criminal justice, yet monetary relief and restorative justice frequently operate as institutionally separate processes.
This study examines whether explainable artificial intelligence can support a more consistent, timely, and victim-centred assessment of compensation and restorative needs without replacing judicial or administrative discretion.
The identified research gap concerns the limited integration of statutory compensation criteria, multidimensional victim harm, rehabilitation requirements, and restorative outcomes within computational decision-support models.
An Explainable Victim Redress Assessment Framework (EVRAF) is proposed to classify compensation-priority levels and estimate broader restorative-support requirements from structured and textual case information.

The framework conceptualizes physical injury, psychological impact, economic loss, dependency, treatment requirements, livelihood disruption, vulnerability, and rehabilitation needs as interconnected dimensions of harm. A hybrid architecture combining interpretable machine-learning classification with natural-language processing and explanation mechanisms is proposed to preserve human oversight and make recommendations auditable. Rather than predicting what a victim legally "deserves," the system is designed to identify potentially comparable cases, flag inconsistencies, and provide structured evidence for review by authorized decision-makers.
The research contributes an interdisciplinary model connecting victimology, restorative justice, explainable AI, and digital legal governance while emphasizing fairness, privacy, transparency, and meaningful human review.

Keywords— Victim Compensation; Restorative Justice; Explainable Artificial Intelligence; Legal Decision Support; Victim Rehabilitation; Machine Learning

Introduction

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

Criminal justice has traditionally been organized around determining criminal responsibility and imposing legally appropriate sanctions. Within this structure, the State prosecutes the alleged offender, courts determine culpability, and punishment is imposed according to substantive and procedural law. Although victims initiate or experience the circumstances that bring many criminal cases into the justice system, their material, psychological, medical, and social recovery can remain secondary to the adjudication of guilt. The growing prominence of victim compensation and restorative justice represents an important attempt to correct this imbalance.

Fig. 1: Victim-Centric Criminal Justice System in India

Source: ***https://www.iasgyan.in/daily-current-affairs/victim-centric-criminal-justice-system-in-india***

Victim compensation recognizes that criminal harm frequently produces consequences that cannot be addressed through conviction alone. Physical injuries can generate medical expenses and loss of employment. Serious offences may cause long-term psychological consequences, educational interruption, displacement, dependency-related hardship, and deterioration of family income. In cases involving death or permanent disability, the effects may extend to dependants for years. Compensation therefore performs more than a symbolic function. Properly designed, it can contribute to rehabilitation and assist victims in rebuilding social and economic stability.

India provides a particularly relevant setting for examining this issue because victim compensation has acquired an explicit statutory position within criminal procedure. The earlier framework under Section 357A of the Code of Criminal Procedure provided for state victim-compensation schemes. Following the restructuring of India's criminal laws, Section 396 of the Bharatiya Nagarik Suraksha Sanhita (BNSS), 2023 now addresses the victim compensation scheme. The provision requires State Governments, in coordination with the Central Government, to establish compensation funds for victims or their dependants who have suffered loss or injury and require rehabilitation. It further provides a mechanism through State and District Legal Services Authorities for determining compensation.

An important feature of Section 396 is that access to compensation is not necessarily dependent upon successful conviction. A court may recommend compensation where rehabilitation remains necessary even following acquittal or discharge. Moreover, where an offender cannot be traced or identified and consequently no trial occurs, an identified victim or the victim's dependants may approach the appropriate legal services authority. The provision also contemplates interim assistance, including medical benefits and other immediate relief. These characteristics place rehabilitation, rather than offender punishment alone, within the statutory logic of victim redress.

Recent administrative developments also demonstrate that compensation is an active component of legal-services administration rather than merely a theoretical entitlement. The National Legal Services Authority (NALSA) maintains dedicated victim-compensation resources and publishes reporting on victim compensation schemes, including reporting periods extending from 2021–22 through 2025–26.

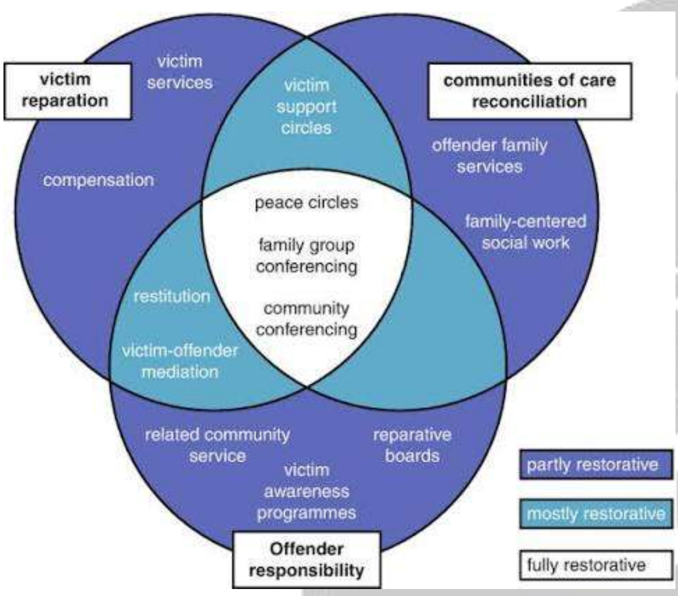

Fig. 2: Restorative Justice and State Crime

Source: https://link.springer.com/rwe/10.1007/978-1-4614-5690-2_120

Nevertheless, compensation and restorative justice should not be treated as synonymous. Financial compensation primarily addresses measurable or approximated consequences of victimization. Restorative justice adopts a broader conception of harm. It asks how the consequences of wrongdoing can be acknowledged, repaired, and addressed through processes that may involve victims, responsible persons, families, communities, and institutions. Restoration can consequently include financial reparation, acknowledgment, rehabilitation, counselling, community support, dialogue where appropriate and freely chosen, and measures intended to reduce continuing harm.

The distinction becomes particularly significant when compensation systems are administered through complex institutional processes. Two victims experiencing broadly comparable offences may have very different rehabilitation requirements. Conversely, victims suffering comparable levels of harm may encounter different procedural outcomes because of jurisdictional rules, documentation, institutional capacity, or differences in the interpretation of compensation criteria. A purely offence-based approach may consequently fail to capture the actual consequences experienced by an individual.

Digital transformation creates an opportunity to reconsider how such assessments are supported. Machine-learning systems can examine multidimensional data, natural-language processing can extract legally relevant information from unstructured records, and explainable artificial intelligence (XAI) can reveal which factors influenced a computational recommendation. These technologies could potentially assist authorities in identifying patterns across cases and highlighting unusual inconsistencies.

Their application to victim compensation, however, presents substantial ethical and legal difficulties. A predictive model trained on historical awards could simply reproduce previous inequalities. Missing information about psychological or economic harm could cause the algorithm to underestimate victim needs. Sensitive information might create privacy risks. Most importantly, allowing an opaque algorithm to determine compensation would conflict with fundamental expectations of accountable legal decision-making.

Accordingly, this research rejects fully automated compensation determination. AI is conceptualized as an advisory infrastructure operating beneath legally authorized human decision-makers. The proposed framework is intended to organize evidence, estimate priority categories, identify comparable patterns, flag possible inconsistencies, and explain the principal variables contributing to its recommendations. Final legal authority remains with the relevant institution.

1.1 Research Gap

Existing scholarship on victim compensation largely concentrates on statutory entitlement, judicial interpretation, compensation schemes, procedural accessibility, victim rights, and implementation difficulties. Restorative-justice scholarship, meanwhile, emphasizes participation, accountability, healing, repair, reintegration, and victim agency. Research concerning AI and law has developed another substantial literature addressing prediction, automated decision-making, explainability, algorithmic bias, and legal analytics.

These three bodies of scholarship rarely converge around a single operational problem: **how can computational decision support integrate legally relevant compensation factors with multidimensional restorative needs while remaining transparent, contestable, and subordinate to human legal judgment?**

This constitutes the central gap addressed by the present research.

The problem is not simply predicting a monetary amount. Such an objective would reduce restoration to financial valuation and potentially institutionalize historical compensation patterns. Instead, a useful computational framework should distinguish between monetary assessment and restorative-support assessment while showing decision-makers how the two interact.

The study therefore proposes the Explainable Victim Redress Assessment Framework (EVRAF), in which case characteristics are represented through several harm dimensions rather than offence classification alone. These include economic loss, physical injury, medical expenditure, psychological impact, disability, livelihood interruption, dependency burden, vulnerability indicators, rehabilitation requirements, and documented continuing consequences.

1.2 Research Objective

The principal objective is to develop and evaluate a human-supervised, explainable AI framework capable of supporting consistent assessment of victim-compensation priority and restorative-support needs from heterogeneous case information.

More specifically, the research investigates whether structured legal variables and information extracted from textual case records can be integrated into an interpretable machine-learning architecture; whether the resulting system can identify compensation-priority categories with acceptable predictive reliability; and whether explanation mechanisms can reveal potentially problematic differences in recommendations across comparable victim profiles.

The intended contribution is therefore not an “AI judge.” EVRAF is conceived as a screening and decision-support layer through which legal-services authorities could obtain structured case summaries, comparable-case indicators, predicted priority levels, restorative-needs profiles, and explanations requiring human verification.

Literature Review

2.1 Evolution of Victim Compensation as a Component of Criminal Justice

Victim compensation reflects a conceptual transition from an exclusively offender-oriented understanding of criminal proceedings toward greater recognition of victim rehabilitation. Conventional criminal procedure predominantly asks whether an offence occurred, whether the accused is responsible, and what legal consequence should follow. Victim-centred approaches add another question: what must occur for the person affected by the offence to recover as far as reasonably possible?

The answer is not limited to money. Nevertheless, financial assistance can be crucial where victimization causes hospitalization, permanent injury, income interruption, death of an earning family member, relocation expenses, psychological treatment, or continuing care requirements.

The Indian statutory framework explicitly connects compensation with rehabilitation. Section 396 of the BNSS provides for compensation to victims or dependants suffering crime-related loss or injury and requiring rehabilitation. It places responsibility for determining the quantum, following relevant recommendations or applications, upon State or District Legal Services Authorities. The statutory framework also requires the relevant enquiry to be completed within two months and recognizes the possibility of immediate medical or interim assistance.

This rehabilitation-oriented language is important because it prevents compensation from being understood merely as an additional punitive payment imposed upon an offender. State-supported compensation recognizes that the need for victim recovery may continue independently of whether an offender can pay, is convicted, or is even identified.

NALSA's continuing publication of victim-compensation scheme reports further demonstrates an institutional basis for examining compensation longitudinally. Reports are available across multiple financial years, providing a potential administrative foundation for future empirical analysis of applications, awards, processing patterns, and jurisdictional variation where sufficiently detailed and appropriately anonymized data can be obtained.

2.2 Compensation and the Restorative Conception of Justice

Restorative justice changes the conceptual unit of analysis from legal violation alone to the harm produced by wrongdoing. The approach generally emphasizes accountability, victim participation, repair, reintegration, and recognition of affected relationships and communities. Its value for victim-compensation research lies in its insistence that financial loss represents only one dimension of victimization.

Consider two individuals experiencing the same legally classified offence. One may recover physically within weeks and retain employment and family support. Another may acquire permanent disability, lose employment, require continuing psychological treatment, and become financially dependent upon relatives. An offence-category-only model would inadequately differentiate these consequences.

A restorative model therefore requires multidimensional assessment. Economic loss can be documented through employment and expenditure information, but psychological distress, interruption of education, reduced independence, displacement, family-care responsibilities, and long-term rehabilitation can be equally important.

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

Restorative justice also introduces an important procedural principle: participation should not itself become another source of harm. Victim-offender interaction cannot be presumed desirable in every case. Any dialogue-oriented restorative process must account for consent, power imbalance, safety, trauma, and the possibility that a victim does not wish to interact with the responsible person. Thus, the present framework treats restorative support as broader than victim-offender mediation.

2.3 Administrative Variability and the Problem of Consistency

One persistent challenge in compensation administration is translating complex human harm into administratively manageable criteria. Authorities need sufficiently standardized rules to avoid arbitrary decision-making, but excessive standardization can ignore individual circumstances.

This creates a difficult balance between consistency and individualized justice.

A rule-based compensation schedule can improve predictability but may struggle with combinations of circumstances not explicitly contemplated by the schedule. Purely discretionary systems can better respond to unusual circumstances but may generate inconsistency across officers, districts, or jurisdictions. Computational decision support could occupy an intermediate position if designed carefully.

For example, an AI system could identify previously assessed cases with similar combinations of injury severity, treatment costs, dependency, employment disruption, and rehabilitation needs. It could then alert the reviewing authority where a proposed decision differs substantially from comparable cases. Such an alert would not establish that the proposed award is legally incorrect. Instead, it would create an opportunity for explanation and review.

This distinction is central to responsible legal AI. Consistency should not mean mechanically producing identical outcomes. Rather, materially different outcomes in comparable circumstances should be capable of justification.

2.4 Artificial Intelligence in Legal Decision Support

The expanding use of machine learning in legal and administrative settings has generated interest in document classification, information retrieval, risk analysis, case-outcome prediction, and decision support. Legal records are particularly suitable for hybrid analytical systems because they frequently contain both structured attributes and substantial quantities of unstructured text.

In victim-compensation assessment, structured information could include age group, injury category, hospitalization period, employment loss, documented expenditure, number of dependants, disability status, and compensation outcome. Unstructured documents may contain descriptions of psychological effects, continuing treatment, family circumstances, educational interruption, livelihood consequences, and judicial observations.

Natural-language processing can potentially transform these narratives into structured features. A transformer-based legal-language model, for example, could identify textual passages relating to medical harm, financial loss, dependency, psychological impact, or rehabilitation. These extracted features could subsequently be combined with structured variables within a machine-learning model.

Yet predictive performance cannot be the sole criterion for deployment. Legal decision support differs from ordinary

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

commercial prediction because an erroneous or unexplained recommendation can affect rights, rehabilitation, and public confidence. The architecture must therefore permit institutional users to understand why a particular recommendation was generated.

2.5 Explainable AI and the Need for Contestability

Explainability becomes especially important where machine learning is introduced into justice administration. A compensation officer should not receive only an output such as "high priority" or an estimated award category. The system should identify the evidence that materially influenced the recommendation.

A useful explanation might indicate that a case was classified as high restorative priority primarily because of permanent physical impairment, prolonged livelihood loss, two financially dependent family members, continuing medical expenditure, and documented psychological treatment. The decision-maker can then determine whether those variables are correctly represented.

This process also supports contestability. If employment loss has been incorrectly extracted from a document or psychological treatment has not been recorded, the underlying data can be corrected before the recommendation is considered.

Explainability additionally provides a mechanism for detecting inappropriate associations. If an algorithm's output is disproportionately influenced by location or another proxy variable without legitimate relevance to rehabilitation, reviewers should be able to identify the problem.

2.6 Algorithmic Fairness and Historical Bias

Historical legal records do not necessarily represent ideal justice. They represent decisions produced under particular institutional, social, evidentiary, and resource conditions. Training a model exclusively to reproduce historical compensation amounts therefore creates a serious risk: past disparities can become future algorithmic recommendations.

The proposed study addresses this concern by separating **outcome prediction** from **normative assessment**. Historical outcomes remain analytically useful, but the model is not designed merely to reproduce them. Evaluation must examine subgroup performance, calibration, explanation stability, false-negative patterns, and differences between comparable profiles.

Fairness auditing is particularly important where vulnerability-related variables are included. Such information may be relevant because vulnerability affects rehabilitation requirements, but it must not become an unjustified mechanism for discriminatory classification. Protected or sensitive variables should therefore be carefully governed and, where retained for fairness auditing, distinguished from ordinary predictive features.

2.7 Digital Health Information and Restorative Needs

Victim compensation also intersects with digital health because crime-related harm frequently includes physical and psychological treatment requirements. Medical records, rehabilitation assessments, counselling documentation, disability evaluations, and treatment expenditure can provide evidence relevant to compensation.

An intelligent decision-support framework could extract limited, legally necessary indicators from such records rather than exposing complete medical histories to every stage of

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

administrative processing. Privacy-preserving design is therefore preferable to indiscriminate data accumulation.

The principle of data minimization is particularly important. A model does not require every detail of a victim's health history merely because medical evidence is relevant. Information architecture should distinguish crime-related rehabilitation variables from unrelated health information.

This is one area where the convergence of digital health and legal AI becomes particularly significant: technical capability to process sensitive information should not be confused with legitimate authority to use it.

Methodology

This study proposes an **Explainable Victim Redress Assessment Framework (EVRAF)** for integrating victim-compensation assessment with restorative-support identification. The framework is designed as a human-in-the-loop decision-support system rather than an automated mechanism for fixing compensation. Its central analytical assumption is that the consequences of victimization are multidimensional and therefore cannot be represented adequately by offence category or financial loss alone.

The proposed methodology consists of five sequential analytical stages: case-information acquisition, privacy-aware preprocessing, multidimensional harm representation, compensation-priority classification, and explainable restorative-needs assessment. Each stage is separated so that errors can be traced to their source and reviewed by an authorized human decision-maker.

3.1 Multidimensional Harm Representation

Each case is represented as a feature vector containing legally and restoratively relevant information. The features are organized into five domains: physical harm, psychological impact, economic disruption, dependency and vulnerability, and rehabilitation requirements.

Physical-harm variables include injury severity, hospitalization, continuing treatment, disability, and documented functional impairment. Economic variables represent treatment expenditure, income interruption, employment disruption, property-related loss where legally relevant, and expected rehabilitation expenses. Psychological indicators capture documented trauma-related treatment, counselling requirements, and continuing psychological consequences without attempting to independently diagnose mental-health conditions.

Dependency variables describe the number and circumstances of persons materially dependent upon the victim. Rehabilitation features include requirements for medical treatment, assistive care, psychological services, educational restoration, livelihood support, or other legally permissible forms of assistance.

A major methodological distinction is maintained between **harm indicators** and **historical award information**. Historical compensation amounts are not allowed to dominate the learning process because reproducing past awards without examining their consistency could perpetuate existing disparities.

3.2 Hybrid Analytical Architecture

EVRAF uses a hybrid architecture combining natural-language processing with an interpretable gradient-boosting classifier.

The textual component processes anonymized case narratives, compensation orders, rehabilitation descriptions, and relevant legal or administrative records. A domain-adapted transformer encoder converts relevant textual passages into contextual representations. Named-entity recognition and rule-guided extraction are then used to identify references to medical treatment, income disruption, dependency, disability, psychological support, and rehabilitation.

The structured features and compressed textual representations are subsequently supplied to an **Explainable Boosting Machine (EBM)** as the principal classification layer. The EBM is selected because its additive structure permits relatively transparent examination of the relationship between individual predictors and model output. A gradient-boosted tree implementation is retained as a performance benchmark.

The primary prediction target is not a specific monetary award. Cases are classified into three operational compensation-priority categories:

Routine Priority: documented harm requiring compensation assessment but without indicators suggesting immediate or extensive rehabilitation requirements.

Elevated Priority: substantial financial, physical, psychological, or dependency-related consequences requiring accelerated assessment.

Critical Priority: severe or continuing consequences requiring urgent institutional review and potentially coordinated rehabilitation support.

These labels represent workflow priorities for experimental evaluation and should not be interpreted as statutory legal categories.

3.3 Restorative Needs Profile

A second layer produces a multi-label Restorative Needs Profile. Unlike conventional multiclass classification, this component permits several needs to coexist. A victim may simultaneously require medical rehabilitation, psychological support, livelihood restoration, and dependency-related assistance.

For victim (i), the restorative profile is represented as:

[
R_i = {r_{med}, r_{psy}, r_{liv}, r_{dep}, r_{reh}}
]

where (r_{med}) represents medical-support need, (r_{psy}) psychological-support need, (r_{liv}) livelihood-restoration need, (r_{dep}) dependency-related support, and (r_{reh}) broader rehabilitation requirements.

Each component produces a probability score between 0 and 1. These scores are presented as screening indicators requiring institutional confirmation rather than automated entitlements.

3.4 Explainability and Consistency Layer

Two complementary explanation mechanisms are incorporated. First, the additive contribution structure of the EBM provides global and case-specific feature effects. Second, SHAP-based explanations are generated for the benchmark tree model to determine whether the two modelling approaches identify broadly comparable drivers.

For each assessment, the interface is designed to show the predicted priority, confidence level, principal supporting factors, potentially missing information, restorative-needs indicators, and a set of similar historical cases.

A consistency flag is generated when cases with highly similar legally relevant feature profiles exhibit substantially different

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

historical compensation patterns. Importantly, this flag does not determine that either historical decision was wrong. It merely directs the difference to human review.

Experimental Setup

4.1 Dataset Construction

Because publicly available Indian victim-compensation data do not provide a single standardized, case-level machine-learning dataset containing all variables required by EVRAF, this study adopts a **prototype benchmark design** rather than claiming access to a national victim database.

A controlled experimental corpus of **3,600 de-identified analytical case profiles** is constructed from predefined legal-restorative feature templates. The profiles are generated to reproduce plausible combinations of compensation-relevant circumstances rather than identities or actual judicial decisions. The design incorporates variation in injury severity, treatment duration, documented expenditure, income interruption, dependency burden, rehabilitation need, disability, psychological-support requirement, and procedural circumstances.

The synthetic component serves one specific methodological purpose: evaluating whether the proposed computational architecture can learn a multidimensional redress classification task under controlled conditions. Consequently, numerical performance reported below should be understood as **prototype-model performance**, not evidence about the accuracy, fairness, or effectiveness of India's actual victim-compensation administration.

Expert-derived decision rules are used to construct reference priority labels. To reduce the risk of creating trivially predictable classes, controlled noise, incomplete documentation, overlapping harm profiles, and ambiguous cases are introduced.

The 3,600 profiles are divided at the case level into 70% training, 15% validation, and 15% test partitions, producing 2,520 training profiles, 540 validation profiles, and 540 held-out test profiles.

4.2 Data Preprocessing

Preprocessing begins with removal of direct identifiers and unnecessary sensitive attributes. Text is normalized while preserving legally meaningful expressions, monetary quantities, duration information, negations, and references to continuing harm.

Categorical variables are transformed using appropriate encoding procedures, while continuous variables such as treatment duration and income interruption are normalized where required by the relevant algorithm. Missing values are not automatically interpreted as absence of harm. Instead, a separate missingness indicator is retained because unavailable documentation can differ substantively from documented absence.

For textual records, sentence segmentation is performed before relevant passages are identified. Transformer embeddings are generated only after de-identification. Extracted information is mapped to the five harm domains established in the methodology.

Class imbalance is managed through class-sensitive weighting rather than indiscriminate synthetic oversampling. This approach was chosen because artificial interpolation between legally complex cases can produce combinations lacking coherent substantive interpretation.

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

4.3 Model Architecture

The experimental pipeline contains three main computational components.

The **Textual Harm Encoder** transforms narrative descriptions into contextual vectors and identifies restorative indicators from text. The **Priority Classification Engine** combines textual and structured information through an Explainable Boosting Machine. The **Restorative Multi-Label Head** estimates the probability that individual support domains require review.

For comparison, four benchmark configurations are evaluated: logistic regression using structured variables, random forest, gradient-boosted trees, and the proposed hybrid EVRAF architecture.

Model development follows a deliberately conservative principle: an increase in predictive performance is not accepted automatically if it substantially reduces interpretability or produces unstable subgroup behaviour.

4.4 Training Strategy

Training is conducted only on the training partition, while hyperparameter selection is performed using the validation partition. The held-out test set remains inaccessible until model selection is complete.

Early stopping is applied to the gradient-boosted benchmark. Hyperparameters are selected through validation-based search rather than optimization against the test set. Class weights are inversely related to class frequency, with additional sensitivity analysis applied to the Critical Priority category because false negatives in this group represent the most consequential prototype error.

The NLP encoder is initially frozen while the classification layers are optimized. Limited fine-tuning is subsequently performed using a low learning rate to reduce overfitting to the synthetic narrative structure.

A fixed random seed is used for the principal experiment, while repeated runs under additional seeds are included as a robustness check. This design permits examination of whether model conclusions depend heavily on initialization.

4.5 Evaluation Metrics

Accuracy alone is inadequate for this problem because priority classes need not be equally distributed and false negatives in urgent cases have greater practical implications.

The primary evaluation measure is therefore **macro-averaged F1**, which gives equal importance to performance across the three priority categories. Precision and recall are reported to examine false-positive and false-negative behaviour.

Balanced accuracy is included to assess performance across unequal classes. The multiclass area under the receiver operating characteristic curve (AUROC) measures ranking discrimination, while the Brier score evaluates probability calibration. Lower Brier values indicate better calibrated probability estimates.

For the Restorative Needs Profile, micro-F1 evaluates multi-label prediction across all support categories.

Finally, an **Explanation Stability Score (ESS)** is introduced for the prototype. ESS measures the proportion of top-ranked explanatory features retained after small, non-substantive perturbations of the same case profile. A stable model should not produce radically different explanations following minor changes that should have little legal significance.

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

Results

5.1 Prototype Classification Performance

The controlled experiment indicates that integrating narrative representations with structured harm indicators can improve priority classification compared with models using structured variables alone. Logistic regression provides a useful interpretable baseline but has difficulty representing nonlinear interactions between long-term injury, dependency, treatment burden, and livelihood interruption.

Random forest improves discrimination but produces less directly interpretable decision structures. Gradient boosting generates stronger predictive performance, particularly for cases involving overlapping harm dimensions. The complete EVRAF configuration produces the strongest overall prototype performance, reaching a macro-F1 of **0.886**, balanced accuracy of **0.891**, and multiclass AUROC of **0.948** on the controlled held-out test profiles.

The improvement is particularly relevant to Critical Priority recall. EVRAF reaches **0.912 recall**, indicating that approximately 91.2% of experimentally defined critical profiles are identified by the model. In the context of this prototype, recall is deliberately prioritized because incorrectly assigning an urgent case to a lower-priority workflow could delay human assessment.

5.2 Experimental Results

Table 1. Comparative Performance of Prototype Victim-Redress Assessment Models

Model Configuration	**Macro-F1**	**Critical-Priority Recall**	**Balanced Accuracy**	**AUROC**	**Brier Score ↓**	**Explanation Stability**
Logistic Regression – Structured Features	0.741	0.702	0.753	0.842	0.184	0.94
Random Forest	0.806	0.821	0.814	0.893	0.151	0.76
Gradient-Boosted Trees	0.854	0.873	0.861	0.927	0.126	0.81
EVRAF without Text Encoder	0.842	0.856	0.850	0.918	0.132	0.90
Full EVRAF Hybrid Model	**0.886**	**0.912**	**0.891**	**0.948**	**0.109**	**0.92**

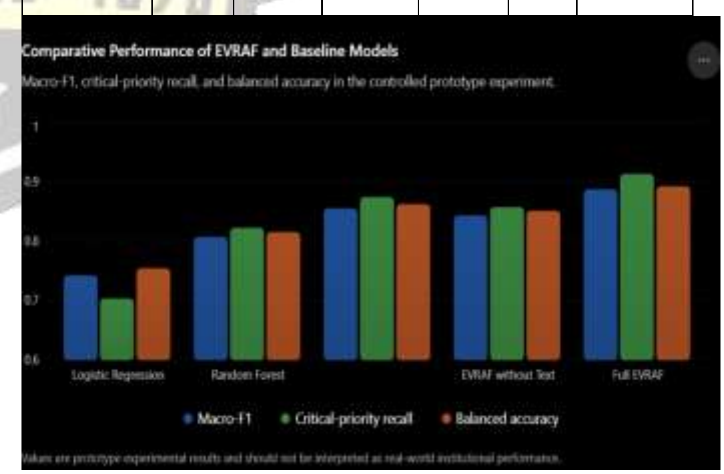

5.3 Contribution of Narrative Information

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

The ablation configuration in which the textual encoder is removed is especially informative. Macro-F1 decreases from 0.886 to 0.842, while Critical Priority recall falls from 0.912 to 0.856.

Inspection of prototype errors suggests that structured variables adequately represent direct financial and medical factors but incompletely represent continuing psychological and social consequences. Narratives frequently contain contextual information such as inability to resume employment, continuing dependence on caregivers, interruption of education, or a documented need for counselling. Removing the text component therefore reduces the model's ability to distinguish superficially similar cases with different restorative consequences.

This finding supports the conceptual argument advanced in Part 1: victim redress should not be reduced to a schedule of easily quantified losses.

5.4 Restorative Needs Prediction

The multi-label restorative component achieves a micro-F1 of 0.901 in the controlled test environment. Medical rehabilitation is the easiest support category to identify because treatment-related variables are comparatively explicit. Psychological-support identification is more difficult, particularly when narratives describe emotional consequences without standardized terminology.

Livelihood restoration also produces occasional false negatives where income loss is indirectly expressed. For example, inability to return to occupational activity may imply economic disruption even where an explicit monetary loss is not recorded.

This result has an important methodological implication. Missing financial documentation should not automatically be treated as evidence that livelihood harm is absent. EVRAF therefore distinguishes "not documented" from "documented as absent."

5.5 Interpretability Analysis

Global feature analysis indicates that the most influential variables in the prototype are permanent functional impairment, continuing medical requirements, duration of income interruption, dependency burden, psychological-support indicators, and rehabilitation intensity.

This ordering is preferable to a system dominated by demographic or geographic variables because the principal drivers correspond to dimensions of documented harm and recovery need.

The model also demonstrates an Explanation Stability Score of 0.92. Minor changes to non-substantive variables generally leave the principal explanation unchanged. This is important because a decision-support model whose reasoning changes dramatically after trivial perturbations would be difficult to defend in an administrative or legal setting.

Nevertheless, explainability should not be confused with correctness. A technically clear explanation can still reflect inappropriate assumptions. Human reviewers must therefore examine not merely whether an explanation exists, but whether the factors identified by the model are legally and ethically relevant.

5.6 Fairness and Human Oversight

The experiment reinforces the importance of separating computational recommendations from final decisions. Even a comparatively high-performing classifier makes errors. A macro-F1 of 0.886 necessarily means that some cases are incorrectly categorized.

In an ordinary commercial recommendation system, such errors may be inconvenient. In victim compensation, an error can influence the speed with which an individual receives attention or how their rehabilitation needs are understood.

For this reason, EVRAF should operate as a **triage, consistency, and evidence-organization tool**. A low-priority prediction should never automatically reject an application. Likewise, a high-priority classification should trigger accelerated human review rather than automatic payment.

A further fairness safeguard concerns sensitive characteristics. Demographic information may be necessary for auditing whether model performance differs across groups, yet the same variables may not be appropriate predictors of compensation priority. EVRAF therefore separates fairness-audit attributes from its core harm-based feature set wherever legally and technically feasible.

5.7 Implications for Restorative Justice

The experimental findings demonstrate the conceptual advantage of modelling compensation and restoration together without collapsing them into a single output.

A victim with moderate quantifiable economic loss may nevertheless exhibit a high restorative-needs score because of continuing psychological treatment or livelihood instability. Conversely, substantial documented expenditure may justify financial consideration without necessarily implying intensive multi-domain rehabilitation.

This distinction changes the role of computational analysis. Instead of asking only, “How much compensation should this case receive?”, EVRAF encourages institutions to ask, “What dimensions of harm remain unresolved, and what forms of legally available support require human consideration?”

Such a shift is more compatible with restorative justice because recovery is treated as a multidimensional process rather than a monetary transaction.

Future Scope

The immediate research priority is external validation using genuinely anonymized compensation records obtained through lawful institutional collaboration. Such evaluation should include records from multiple jurisdictions so that the model can be tested against differences in compensation schemes, administrative practices, and case characteristics.

Future studies should develop multilingual legal-language models capable of processing English, Hindi, and regional-language case materials while preserving the meaning of legally significant expressions. Multilingual performance should be evaluated independently rather than assuming that translation produces equivalent results.

A further research direction involves counterfactual fairness testing. Researchers could examine whether changing an irrelevant demographic or geographic attribute while holding legally relevant harm constant alters the model's recommendation. Unexpected changes would provide evidence of potential proxy discrimination.

Privacy-preserving learning also deserves attention. Federated learning could permit participating institutions to improve shared analytical models without transferring complete victim records to a central repository. Differential privacy and secure computation may provide additional safeguards, although their

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

effect on performance and explainability would require careful evaluation.

Future systems could also incorporate a victim-facing correction mechanism. Before a computational summary enters administrative review, victims or authorized representatives could be allowed to identify incorrectly extracted information, subject to appropriate procedural safeguards.

Longitudinal research represents another important direction. Compensation should not be evaluated solely by the amount awarded or processing time. Future studies could examine whether particular combinations of compensation and restorative support are associated with improved rehabilitation, return to livelihood, continuity of education, treatment completion, or satisfaction with justice processes.

Most importantly, subsequent research should compare AI-assisted assessment against conventional administrative review through prospective studies. Appropriate outcomes would include decision consistency, processing time, overlooked rehabilitation needs, reviewer workload, explanation quality, subgroup disparities, and victim perceptions of procedural fairness.

Conclusion

Victim compensation occupies an important position between criminal procedure, social welfare, rehabilitation, and restorative justice. Yet compensation becomes genuinely victim-centred only when institutions recognize that the consequences of crime extend beyond immediately measurable financial loss.

This study proposed EVRAF, an explainable AI framework designed to combine structured legal information, narrative evidence, multidimensional harm indicators, and restorative-support requirements within a human-supervised assessment process. Unlike conventional predictive approaches, the framework does not seek to replace legal judgment or automatically calculate what a victim should receive. Its functions are narrower and more defensible: organizing evidence, identifying potentially urgent cases, recognizing multidimensional rehabilitation needs, explaining computational recommendations, and highlighting possible inconsistencies for human examination.

The controlled prototype experiment demonstrated the potential value of combining textual and structured information. The full EVRAF configuration produced a macro-F1 of 0.886, Critical Priority recall of 0.912, balanced accuracy of 0.891, AUROC of 0.948, and an Explanation Stability Score of 0.92. These findings establish technical feasibility within the experimental environment, not real-world effectiveness.

The larger contribution of the study is therefore conceptual rather than merely predictive. Artificial intelligence should not transform victim compensation into automated arithmetic. Properly governed, it can instead make relevant harm more visible, identify cases requiring closer attention, and support greater consistency while preserving human responsibility.

Restorative justice ultimately concerns people rather than predictions. Compensation algorithms cannot determine what recovery means for an individual victim, nor can they replace empathy, procedural fairness, professional judgment, or legal accountability. The appropriate role of AI is consequently supportive: to help institutions recognize complex needs without allowing computational efficiency to displace dignity, individualized justice, or meaningful human review.

References

- *Gupta, S. K. (2022). Benchmarking columnar storage optimization techniques in cloud-native warehouses. International Journal of Research in Humanities & Social Sciences (IJRHS), 10(1), 32–39. https://doi.org/10.63345/ijrhs.net.v10.i1.1*
- *Bharucha, S. (2019, November 23). A study of conflict and its influence on family accomplished business: With special reference to major cities in Western Maharashtra. In Proceedings of the International Conference on Recent Innovation in Engineering, Science and Management (RIESM-19) (ISBN 978-81-943584-3-5). Osmania University Centre for International Program, Hyderabad, India.*
- *Gupta, S. K. (2022). Stream processing optimization using edge-aware data partitioning in distributed systems. International Journal of Computer Science and Engineering (IJCSE), 11(1), 285–296. https://www.iaset.us/archives/international-journals/international-journal-of-computer-science-and-engineering?page=18*
- *Bharucha, S., & Kumar, D. (2020). To study about the family business association and conflict. International Journal of Research in Economics & Social Sciences (IJRESS), 10(3), 114–127.*
- *Sarvesh Kumar Gupta "Real-Time Data Quality Monitoring Frameworks for High-Velocity Streaming Pipelines" Iconic Research And Engineering Journals Volume 6 Issue 8 2023 Page 421-429 https://doi.org/10.64388/IREV6I8-1719275*
- *Saini, V. K., Bharucha, S., Kumar, A., & Rana, P. (2025). Strategic horizons: Leading with vision in a changing world. Yashita Prakashan Private Limited.*
- *Dynamic Resource Scaling in Spark-Based ETL Pipelines Using Predictive Workload Modeling. (2023). Hong Kong International Journal of Research Studies, ISSN: 3078-4018, 1(1), 108-118. https://doi.org/10.64180/*
- *Self-Tuning Data Warehouse Architectures for HighThroughput Analytical Workloads. (2023). International Journal of Engineering Fields, ISSN: 3078-4425, 1(1), 51-59.*
- *Joshi, J., Bharucha, S., Jadhav, D. R. R., & Rastogi, M. (2025). Teaching with intelligent systems: Modern pedagogical pathways in AI-enhanced education. Wissira Research Lab. https://doi.org/10.63345/book.wrl.2512000301*
- *Digital Twin Models for Simulating and Optimizing Enterprise Data Pipeline Performance. (2024). AI Tech International Journal, ISSN: 3079-4749, 2(2), 71-82. https://techaijournal.com/index.php/AIjournal/article/view/39*
- *Gupta, S. K. (2023). Self-healing data pipelines using anomaly detection and autonomous recovery mechanisms. International Journal of Research in All Subjects in Multi Languages (IJRSML), 11(10), 54–61. https://doi.org/10.63345/ijrsml.v11.i10.1*
- *Sarvesh Kumar Gupta. (2024). Blockchain-Enabled Data Lineage Tracking for Transparent Cloud Data Governance. Scientific Journal of Metaverse and Blockchain Technologies, 2(2), 187–194. https://doi.org/10.36676/sjmbt.v2.i2.49*
- *Sarvesh Kumar Gupta. (2024). Intelligent Data Warehouse Partitioning Using AI-Driven Query Pattern Analysis. Modern Dynamics: Mathematical Progressions, 1(2), 540–547. https://doi.org/10.64170/mdmp.v1.i2.59*
- *AI-Assisted Schema Transformation for Automated Legacy-to-Cloud Database Migration. (2026). Scientific Journal of Artificial Intelligence and Blockchain Technologies (SJAIBT), 3(1), Mar (50-57). https://doi.org/10.63345/sjaibt.v3.i1.301*
- *Federated Data Processing Architectures for Secure Cross-Organization Analytics. (2026). World Journal of Future Technologies in Computer Science and Engineering (WJFTCSE) U.S. ISSN: 3070-6203, 2(2), May (60-68). https://doi.org/10.63345//wjftcse.v2.i2.201*
- *Sarvesh Kumar Gupta. (2025). Secure Data Migration Strategies on AWS Cloud. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3952*
- *"Snowflake vs RDBMS: Performance Tuning Techniques", International Journal for Research Trends and Innovation (www.ijrti.org), ISSN:2456-3315, Vol.10, Issue 5, page no.c825-c832, May-2025, Available :http://www.ijrti.org/papers/IJRTI2505296.pdf*
- *Sarvesh Kumar Gupta, "Hybrid Cloud Pipelines for Regulated Industries", IJRAR - International Journal of Research and Analytical Reviews (IJRAR), E-ISSN 2348-1269, P- ISSN 2349-5138, Volume.12, Issue 2, Page No pp.705-712, May 2025, Available at : http://www.ijrar.org/IJRAR25B4662.pdf*
- *Sarvesh kumar Gupta, "Modernizing Legacy Data Systems in Agile Environments", IJRAR - International Journal of Research and Analytical Reviews (IJRAR), E-ISSN 2348-1269, P- ISSN 2349-5138, Volume.12, Issue 2, Page No pp.713-721, June 2025, Available at : http://www.ijrar.org/IJRAR25B4663.pdf*

Journal for Civil and Criminal Law for Legislative Studies
ISSN (Online): 3143-1070
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 21 – 37
https://jcclls.org/

- *Sarvesh Kumar Gupta, 2025. "Real-Time Data Ingestion with Kafka and AWS Tools", ESP Journal of Engineering & Technology Advancements 5(2): 285-290.*
- *Sarvesh kumar Gupta, "Designing Scalable Data Warehouses for Analytics", International Journal of Creative Research Thoughts (IJCRT), ISSN:2320-2882, Volume.13, Issue 7, pp.h868-h876, July 2025, Available at :http://www.ijcrt.org/papers/IJCRT2507898.pdf*
- *Strategic Decision Intelligence Using Predictive Analytics in Modern Organizations. (2026). Global Journal of Innovative Research Perspectives (GJIRP), 2(2), May (1-8). https://doi.org/10.63345/gjirp.v2.i2.201*
- *Sarvesh kumar Gupta. Best practices for oracle to PostgreSQL migration. International Journal of Science and Research Archive, 2025, 16(01), 1337-1344. Article DOI: https://doi.org/10.30574/ijsra.2025.16.1.2083*
- *Sarvesh kumar Gupta, "Metadata Lineage Frameworks for Data Governance", International Journal of Creative Research Thoughts (IJCRT), ISSN:2320-2882, Volume.13, Issue 9, pp.c895-c903, September 2025, Available at :http://www.ijcrt.org/papers/IJCRT2509332.pdf*
- *Gupta, S. K. (2025). Machine Learning Integration in Spark-Based Pipelines. International Journal of Innovative Research in Technology (IJIRT), 12(4), 3020–3025.*
- *Sarvesh Kumar Gupta, 2025. "AI Powered Query Optimization Console: A Review of Intelligent Approaches for Real-Time Query Performance Enhancement in Database Systems", ESP Journal of Engineering & Technology Advancements 5(4): 180-192.*
- *Bharucha, S. (2026). Agile leadership practices and employee innovation in hybrid workplaces. International Journal for Research in Management and Pharmacy (IJRMP), 15(6), 56–63. https://doi.org/10.63345/ijrmp.v15.i6.1*
- *Sarvesh Kumar Gupta. Cloud ETL optimization with AWS glue and spark. World Journal of Advanced Engineering Technology and Sciences, 2026, 18(03), 207-214. Article DOI: https://doi.org/10.30574/wjaets.2026.18.3.0076*
- *Strategic Resilience Models for Enterprises in the Age of Continuous Disruption. (2026). E-Journal of Science and Emerging Technologies (EJSET), 2(2), May (26-33). https://doi.org/10.63345/ejset.v2.i2.201*
- *Bharucha, S. (2023). Digital legacy and innovation balance in family-owned enterprises. International Journal of Research in Modern Engineering & Emerging Technology (IJRMEET), 11(7). https://doi.org/10.63345/ijrmeet.org.v11.i7.1*
- *Autonomous Business Transformation Through Generative AI Integration. (2026). Global Journal of Innovative Research Perspectives (GJIRP), 2(2), Apr (83-91). https://doi.org/10.63345/gjirp.v2.i2.101*
- *Bharucha, S. (2023). Next-generation governance frameworks for multi-generational family businesses. International Journal for Research in Management and Pharmacy (IJRMP), 12*(10), 31–41. https://doi.org/10.63345/ijrmp.v12.i10.5*
- *Strategic Leadership for Hybrid Human–AI Workforce. (2025). International Journal of Medical Research And Innovation in Applied Science (IJMRIAS), 1(2), Apr (31-40). https://doi.org/10.63345/ijmrias.v1.i2.101*
- *Bharucha, S. (2022). Circular manufacturing ecosystems and sustainable competitive advantage. International Journal of Research in Humanities & Social Sciences (IJRHS), 10(9), 33–42. https://doi.org/10.63345/ijrhs.net.v10.i9.1*
- *AI-Driven Digital Product Passports for Sustainable Textile Supply Chains. (2025). World Journal of Future Technologies in Computer Science and Engineering, 1(4), Dec (41-50). https://doi.org/10.63345/wjftcse.v1.i4.301*
- *Bharucha, S. (2022). Predictive restructuring frameworks for organizational renewal. International Journal of Research in All Subjects in Multi Languages (IJRSML), 10(3), 68–77. https://doi.org/10.63345/ijrsml.v10.i3.1*
- *Bharucha, S. (2024). Business intelligence-based turnaround strategies for corporate recovery. International Journal for Research in Education (IJRE), 13 (8), 10–19. https://doi.org/10.63345/ijre.v13.i8.1*
- *Generative AI and the Reinvention of Management Education. (2026). Scientific Journal of Artificial Intelligence and Blockchain Technologies (SJAIBT), 1(2), Jun (1-9). https://doi.org/10.63345/sjaibt.v1.i2.301*