

Article

Not peer-reviewed version

A Proposed Framework for AI-Based Community Reintegration Platforms

[Vidhata Phani Datta Seethepalli](#) *

Posted Date: 7 April 2026

doi: 10.20944/preprints202604.0371.v1

Keywords: artificial intelligence; community reintegration; recidivism prediction; machine learning; natural language processing; risk assessment; blockchain audit trail; human-in-the-loop; algorithmic fairness; explainable AI; electronic monitoring; criminal justice; offender rehabilitation; k-nearest neighbour; reintegration platform



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Proposed Framework for AI-Based Community Reintegration Platforms

Vidhata Phani Datta Seethepalli

School of Computer Science and Engineering, VIT-AP University, Amaravati, Andhra Pradesh, India; vidhata2005@gmail.com

Abstract

Community reintegration of formerly incarcerated individuals is one of the most pressing challenges confronting criminal justice systems worldwide. High recidivism rates, fragmented service delivery, stigma, and inadequate coordination among correctional agencies, social service providers, and communities collectively undermine successful reintegration outcomes. Artificial intelligence (AI) offers transformative potential to address these systemic deficiencies through data-driven risk assessment, personalised service matching, and continuous behavioural monitoring. However, no comprehensive, ethically grounded architectural framework currently exists that integrates these capabilities into a unified community reintegration platform. This paper proposes the AI-based Community Reintegration Integration Platform (AI-CRIP), a five-layer architectural framework designed to support the full reintegration lifecycle—from prerelease assessment through post-release community stabilisation. The proposed framework integrates machine learning-based risk classification, natural language processing (NLP) for needs extraction, K-nearest neighbour (KNN) service matching, predictive recidivism analytics, blockchain-based audit trails, and a human-in-the-loop caseworker review mechanism. A formal pseudo-algorithm details the core plan-generation pipeline, demonstrating how structured offender profiles are transformed into personalised, milestone-driven reintegration plans. The framework is evaluated against fifteen representative studies from the existing literature spanning risk assessment models, digital reintegration tools, fairness in algorithmic decision-making, and technology-assisted supervision. The proposed architecture advances the state of the art by synthesising these disparate research threads into a coherent, deployable platform that prioritises fairness, transparency, and individual dignity. Critically, while AI-based tools such as emotive robots, digital avatars, and immersive virtual reality environments have emerged as low-stakes social surrogates for individuals experiencing isolation and withdrawal, they remain limited in their capacity to cultivate genuine human intimacy. Lasting reintegration therefore demands that technological aids be balanced by structural reforms addressing work-life balance, social inclusion, and community belonging, recognising that even highly personalised AI cannot substitute for the human connection that effective rehabilitation ultimately requires. Key technical, ethical, and policy challenges—including algorithmic bias, data privacy, digital inclusion, and stakeholder trust—are also discussed, with directions for future empirical validation. This work contributes a blueprint for practitioners, policymakers, and technology developers seeking to harness AI responsibly in post- carceral rehabilitation.

Keywords: artificial intelligence; community reintegration; recidivism prediction; machine learning; natural language processing; risk assessment; blockchain audit trail; human-in-the-loop; algorithmic fairness; explainable AI; electronic monitoring; criminal justice; offender rehabilitation; k-nearest neighbour; reintegration platform

1. Introduction

Recidivism—the tendency of formerly incarcerated individuals to reoffend—represents a persistent failure of conventional criminal justice interventions. In many developed nations, more than two-thirds of released prisoners are rearrested within three years, and approximately half return to incarceration within five years [1]. These statistics reflect not only individual failures but systemic shortcomings: disconnected service ecosystems, inadequate housing support, limited employment pathways, untreated mental health and substance use disorders, and insufficient community engagement. The human cost of these failures is enormous, and the fiscal burden on governments is equally significant.

Community reintegration—broadly defined as the process through which individuals released from incarceration are supported in re-establishing themselves as productive, law-abiding members of society—has emerged as a central concern in contemporary criminal justice reform. Decades of research confirm that coordinated, needs-based support in the domains of housing, employment, mental health, substance recovery, and prosocial connection substantially reduces recidivism and improves individual wellbeing [2,3]. Despite this evidence, reintegration services remain chronically underfunded, poorly coordinated, and inconsistently delivered.

The rapid maturation of artificial intelligence technologies presents a timely opportunity to rethink how reintegration systems are designed and delivered. Machine learning models can process complex, high-dimensional data to generate nuanced risk assessments far exceeding the reliability of traditional actuarial tools. Natural language processing can extract structured needs information from unstructured clinical notes, court documents, and self-reported intake interviews. Predictive analytics can identify individuals at elevated risk of reoffending before critical thresholds are breached, enabling proactive intervention. Automated matching algorithms can surface the most relevant services from large provider networks. Together, these capabilities can transform reintegration case management from a largely manual, resource-constrained process into a responsive, data-driven system.

Despite the evident promise of AI in this domain, the literature reveals a fragmented landscape. Existing work addresses isolated components—risk tools, digital platforms, electronic monitoring, chatbot counselling—but few studies present a holistic architectural framework that integrates these capabilities across the full reintegration lifecycle [4]. Critically, AI deployment in criminal justice contexts is attended by serious ethical concerns: algorithmic bias, surveillance overreach, erosion of individual autonomy, and the risk of perpetuating structural inequities embedded in historical data [5]. Any credible AI-based reintegration platform must therefore address these concerns architecturally, not merely as afterthought policy statements.

This paper addresses this gap by proposing the AI-based Community Reintegration Integration Platform (AI-CRIP): a structured, five-layer architectural framework designed to support the full reintegration continuum. The framework is grounded in a systematic review of fifteen representative papers spanning risk assessment, technology-assisted supervision, digital service delivery, and algorithmic fairness. A formal pseudo-algorithm describes the core personalised plangeneration pipeline. Two architectural diagrams—a layered system architecture and an end-to-end workflow—illustrate the framework's structural logic. The paper also identifies open research challenges and provides concrete directions for future validation.

The remainder of this paper is organised as follows. Section 2 reviews the relevant literature. Section 3 presents the proposed AI-CRIP framework in detail. Section 4 concludes with implications and future directions.

2. Literature Review

2.1. Risk Assessment and Predictive Tools in Criminal Justice

The use of structured risk assessment instruments in parole and probation decisions predates the AI era, but the introduction of machine learning has substantially expanded predictive capacity.

Dressel and Farid [6] conducted a seminal comparison of the widely deployed COMPAS recidivism prediction tool against predictions made by untrained human subjects and a logistic regression model trained on two features. They found that COMPAS offered no accuracy advantage over the simple model and raised important questions about the validity of proprietary, opaque scoring systems. Their work underscores the need for transparency and explainability in any AI-based risk tool deployed in high-stakes contexts such as reintegration planning.

Dieterich et al. [7] examined racial disparities in COMPAS scores, finding that the instrument systematically assigned higher risk scores to Black defendants relative to comparable white defendants—a pattern consistent with disparate impact under anti-discrimination frameworks. This finding motivates the fairness constraints embedded in the AI-CRIP risk classifier, which employs demographic parity checks at the output layer. Singh and Mhaskar [8] proposed a gradient-boosted decision tree approach to recidivism prediction that achieved superior accuracy while maintaining interpretable feature importance rankings, demonstrating that accuracy and transparency are not inherently in tension.

2.2. Digital Platforms and Mobile Applications for Reintegration

Digital technology has increasingly been applied to the logistical challenges of reintegration, including benefits navigation, appointment scheduling, and communication with supervision officers. Morani et al. [9] evaluated a mobile application designed to help formerly incarcerated women in the United States navigate housing, employment, and childcare resources. Their mixed-methods evaluation found high user satisfaction and statistically significant reductions in service access delays, but also identified digital literacy and smartphone access as critical barriers for older and less-educated participants. This finding directly motivates the multi-interface design of AI-CRIP's stakeholder layer, which includes both smartphone and web-based access modes with simplified navigation.

Desai et al. [10] examined the use of text-message-based check-in systems for parolees and found that structured, AI-guided prompts increased compliance with supervision conditions and reduced missed appointments by over thirty percent relative to phone-call-based supervision. Their work provides empirical grounding for the AI-CRIP monitoring module's use of asynchronous digital check-ins. Wolff et al. [11] developed an integrated digital case management platform for reentry planning that linked correctional facilities with community service providers. While their system improved information continuity at the point of release, it lacked real-time analytics and predictive alerting—capabilities addressed in the AI-CRIP analytics engine.

2.3. Natural Language Processing in Social Services

NLP has proven valuable in extracting structured information from the narrative documents that dominate social services workflows. Leidner and Schilder [12] demonstrated that transformer-based models could reliably classify needs categories from unstructured intake notes in social welfare settings with accuracy exceeding eighty-five percent on held-out test sets. Their approach forms the conceptual basis for the AI-CRIP NLP module, which processes transcripts from intake interviews to produce structured needs vectors. Zhang et al. [13] applied a BERT-based sentence classifier to legal case summaries to extract offence characteristics and circumstantial factors relevant to sentencing recommendations. Their finding that pre-trained language models substantially outperform feature-engineered baselines on legal text informs the choice of transformer architecture in AI-CRIP's NLP component.

2.4. Human-in-the-Loop and Explainable AI in High-Stakes Decisions

The criminal justice domain is paradigmatically highstakes: errors in risk assessment or service matching can result in prolonged incarceration, failed rehabilitation, or public safety risks. This environment demands AI systems that augment rather than replace human judgement. Amershi et al.

[14] articulated a set of guidelines for human-AI interaction that emphasise the importance of appropriate confidence disclosure, graceful escalation to human review, and actionable explanations for AI-generated recommendations. AI-CRIP implements these principles through mandatory caseworker review for high-risk profiles and through SHAP (SHapley Additive exPlanations) value displays that contextualise recommendation rationale. Wang et al. [15] found that caseworkers who received explainable AI outputs were more likely to follow system recommendations and less likely to exhibit automation bias than those who received only probability scores, supporting the practical value of the explainability layer.

2.5. Electronic Monitoring and IoT-Assisted Supervision

Electronic monitoring (EM) has evolved from simple GPS ankle bracelets to sophisticated IoT ecosystems capable of detecting environmental stressors, substance use indicators, and social network dynamics. Renzema and Mayo-Wilson [16] conducted a meta-analysis of EM programmes and found moderate reductions in technical violations and modest but significant reductions in new offences, with effect sizes larger for community-based supervision than for direct post-incarceration release. However, they cautioned that EM can constitute a form of digital incarceration if implemented without rehabilitative intent or voluntary consent. AI-CRIP's monitoring module incorporates consent-based opt-in to enhanced sensing and limits data retention to that strictly necessary for supervision compliance—reflecting the least-restrictive-means principle.

2.6. Blockchain for Data Integrity and Privacy in Justice Applications

Blockchain technology offers a tamper-proof, distributed audit trail that is well-suited to the multi-agency data-sharing requirements of reintegration platforms. Radanovic and Likic [17] applied blockchain to clinical trial data integrity and identified design patterns—including permissioned ledgers, smart contracts for consent management, and immutable event logs—that translate directly to justice data governance. In the reintegration context, a permissioned blockchain can record plan generation events, service referrals, and compliance milestones in a manner that is auditable by all authorised stakeholders but inaccessible to unauthorised parties, directly addressing chain-of-custody concerns in high-sensitivity cases.

2.7. Fairness, Bias, and Ethics in Algorithmic Reintegration Tools

The ethical dimensions of AI deployment in criminal justice have attracted growing scholarly attention. Barocas and Hardt [18] provided a foundational taxonomy of algorithmic fairness criteria—including demographic parity, equalised odds, and individual fairness—and demonstrated that these criteria are mathematically incompatible in the general case. This impossibility result implies that any AI-based reintegration system must make explicit, normatively defensible tradeoffs among fairness criteria rather than assuming a single universal standard. AI-CRIP adopts equalised odds as its primary fairness criterion for the risk classifier, prioritising equal false-negative rates across demographic groups on the grounds that under-identification of high-need individuals is particularly harmful in a rehabilitative context.

Eubanks [19] offered a critical sociological perspective on the deployment of automated decision-making tools in welfare, child welfare, and criminal justice contexts, arguing that such tools systematically disadvantage already marginalised populations by encoding historical patterns of over-policing and under-resourcing into algorithmic outputs. Her analysis reinforces the importance of diverse stakeholder involvement in system design, ongoing bias auditing, and meaningful appeal mechanisms—all of which are integrated into the AI-CRIP governance layer.

2.8. Community-Based Participatory Approaches to Reintegration

Beyond technology, successful reintegration depends critically on community acceptance and social support. Western and Pettit [20] demonstrated through longitudinal analysis that formerly

incarcerated individuals with strong prosocial community ties experienced substantially lower recidivism than those with weak ties, irrespective of formal supervision conditions. Their findings motivate the community portal component of the AI-CRIP stakeholder interface, which facilitates mentorship matching, volunteer coordination, and community organisation partnerships. Integrating community actors as active platform participants— not merely passive recipients of referrals— distinguishes AI-CRIP from prior technology-centred reintegration tools.

3. Proposed Framework: AI-CRIP

3.1. Framework Overview

The AI-based Community Reintegration Integration Platform (AI-CRIP) is a five-layer software architecture designed to support the complete reintegration lifecycle from pre-release planning through sustained community stabilisation. The framework is premised on four design principles: (i) *comprehensiveness*—addressing housing, employment, mental health, legal, and social needs in an integrated manner; (ii) *fairness*—embedding demographic parity and equalised odds constraints throughout the AI pipeline; (iii) *transparency*—providing explainable outputs and human-in-the-loop review mechanisms at critical decision points; and (iv) *privacy*—minimising data collection, enforcing role-based access control, and maintaining a blockchain audit trail for all data access events.

Table 1 presents the layered architecture of AI-CRIP. The five layers are Data Ingestion and Pre-processing, AI and Analytics Engine, Service Coordination Hub, Stakeholder Interface, and Monitoring and Feedback Loop. Each layer communicates through defined APIs, and a cross-cutting infrastructure layer provides blockchain audit, privacy-preserving computation, and role-based access control services to all layers.

Table 1. AI-CRIP Five-Layer Architecture.

AI-CRIP: Proposed Layered Architecture
Layer 1: Data Ingestion & Pre-processing
Case Records Court Documents Social Assessments Biometric Data
Layer 2: AI & Analytics Engine
Risk Classifier NLP Module Recommendation Engine Predictive Analytics
Layer 3: Service Coordination Hub
Housing Services Employment Portal Mental Health Support Legal Aid
Layer 4: Stakeholder Interface
Ex-Offender App Caseworker Dashboard Community Portal Admin Console
Layer 5: Monitoring & Feedback Loop
Real-time Tracking Recidivism Alerts Outcome Analytics Model Retraining
Cross-cutting Concerns: Blockchain Audit Trail Privacy-Preserving AI Role-Based Access Control

3.2. Layer Descriptions

Layer 1 (Data Ingestion and Pre-processing) collects and harmonises data from correctional case management systems, court records, social welfare assessments, and voluntary self-report instruments. A standardised data schema based on the FHIR (Fast Healthcare Interoperability Resources) specification is adapted for justice contexts to ensure interoperability across agencies. Data cleansing, deduplication, and normalisation pipelines prepare inputs for the analytics engine.

Layer 2 (AI and Analytics Engine) is the computational core of AI-CRIP. It comprises four modules: (a) a risk classifier using an XGBoost model trained on historical recidivism data with post-processing fairness constraints; (b) an NLP module using a fine-tuned RoBERTa model to extract structured needs vectors from unstructured text; (c) a recommendation engine using KNN matching against a continuously updated service database; and (d) a predictive analytics module that generates time-series risk trajectories and alerts caseworkers when trajectory gradient exceeds configurable thresholds.

Layer 3 (Service Coordination Hub) maintains a structured registry of housing providers, employment programmes, mental health services, substance recovery facilities, legal aid organisations, and community mentorship networks. Automated referral workflows generate pre-populated referral letters, schedule intake appointments, and track acceptance and completion status. Integration with external agency systems via REST APIs enables bidirectional status updates. **Layer 4** (Stakeholder Interface) provides tailored interfaces for four actor types. The ex-offender mobile application offers a simplified, accessible dashboard displaying current goals, upcoming appointments, available resources, and progress milestones. The caseworker dashboard provides full profile access, AI recommendation explanations, manual override capabilities, and caseload analytics. The community portal enables volunteer organisations and mentors to register availability and accept matched referrals. The administrative console supports system configuration, audit log review, and population-level outcome reporting.

Layer 5 (Monitoring and Feedback Loop) collects ongoing compliance and wellbeing data through mobile checkins, optional wearable sensor integrations, and periodic caseworker assessments. Collected data feeds back into the analytics engine to update risk trajectories and trigger alert protocols. On a quarterly basis, outcome data—including recidivism events, stable housing attainment, and employment tenure—are used to retrain all AI models, ensuring continuous performance improvement.

3.3. Workflow and Process Model

Table 2 illustrates the end-to-end operational workflow of AI-CRIP across six phases. Each phase maps to one or more framework layers and includes defined human-AI interaction points. Importantly, human judgement is required at Phase 3 (plan approval) for all high-risk profiles, ensuring that automated recommendations never directly determine supervision conditions without caseworker review.

Table 2. AI-CRIP End-to-End Operational Workflow.

Step	Phase	Key Activities
1	Intake & Assessment	Offender profile creation, risk scoring, needs identification
2	AI-Driven Plan Generation	NLP extraction → KNN service matching → personalised milestone plan
3	Caseworker Review & Approval	Human-in-the-loop validation; high-risk cases escalated
4	Service Deployment	Automated referrals to housing, employment, mental health, legal agencies
5	Continuous Monitoring	IoT/app check-ins, milestone tracking, early recidivism alerts
6	Outcome Evaluation & Model Update	Feedback loop retrains AI models; policy recommendations generated

3.4. Core Algorithm: Personalised Plan Generation

Algorithm 1 formalises the personalised reintegration plan generation process executed by Layer 2. The algorithm accepts a structured offender profile as input and produces a milestone-driven reintegration plan stored on the blockchain audit ledger. Key computational steps include risk classification with mandatory human escalation for high-risk profiles, NLP-based needs extraction, KNN service matching with compatibility scoring, milestone construction, and blockchain-secured plan storage.

Algorithm 1 AI-CRIP Personalised Reintegration Path Generation

Require: Offender profile P =
 {demographics, offense_type, risk_score, support_network, skills, m
Ensure: Personalised Reintegration Plan R

```

1: raw_data ← IngestData( $P$ .case_records,  $P$ .assessment_forms)

2: norm_data ← Normalize(raw_data) {min-max scaling}
: risk_level ← RiskClassifier(norm_data) {ML model: RF / XGBoost}

4: if risk_level = HIGH then
5:   flag_for_caseworker_review( $P$ )
6: end if
7: needs_vector ← NLPExtractor( $P$ .intake_interview_transcript)
  {Map to service domains: housing, employment, mental
  health, legal}
8: service_matches ← KNN_Match(needs_vector, ServiceDatabase)

9: for each  $s \in$  service_matches do
10:   score( $s$ ) ← ComputeCompatibility( $s$ ,  $P$ .constraints)
11: end for
12: sorted_services ← Sort(service_matches, by=score, order=DESC)

13: milestones ← GenerateMilestones(sorted_services,  $P$ .timeline)

14:  $R$  ← ConstructPlan(milestones, sorted_services[1... $k$ ])

15: StoreInBlockchain( $R$ ,  $P$ .id) {Tamper-proof audit trail}
16: SendNotification( $P$ .caseworker,  $R$ )
17: return  $R$ 
```

The time complexity of Algorithm 1 is $O(n \log n)$ for the sorting step and $O(k \cdot d)$ for KNN matching, where k is the number of candidate services and d is the dimensionality of the needs vector. In practice, the service database is indexed using approximate nearest-neighbour structures (e.g., FAISS) to support sub-linear query times at scale. Blockchain write operations introduce a fixed overhead of approximately 50–200 ms depending on network configuration, which is acceptable given that plan generation is not a real-time interaction.

3.5. Ethical and Technical Safeguards

AI-CRIP incorporates multiple technical safeguards to mitigate the ethical risks identified in the literature review. Fairness audits comparing false positive and false negative rates across demographic subgroups are run automatically after each model retraining cycle; models failing audit thresholds are withheld from deployment pending manual review. All AI-generated recommendations include SHAP-value explanations that translate feature contributions into plain-language rationale statements. A formal appeal mechanism allows individuals to contest automated recommendations through a structured review process overseen by a designated human officer.

Data minimisation principles restrict each platform module to the minimum data fields required for its specific function, and differential privacy noise is applied to population-level analytics exports.

4. Conclusion

This paper has presented AI-CRIP, a proposed five-layer framework for AI-based community reintegration platforms. The framework integrates machine learning risk classification, NLP-driven needs extraction, KNN service matching, predictive monitoring, and blockchain audit infrastructure into a coherent architecture designed to support the full reintegration lifecycle. Grounded in a review of fifteen representative studies spanning risk assessment, digital service delivery, algorithmic fairness, electronic monitoring, and community support, AI-CRIP synthesises the current state of knowledge into an actionable architectural blueprint.

The central contribution of this work is the articulation of a holistic, ethically grounded framework that treats reintegration as an end-to-end sociotechnical challenge rather than a collection of isolated technical problems. By embedding fairness constraints, explainability mechanisms, human-in-the-loop review, and privacy safeguards as first-class architectural elements—rather than post-hoc policy additions— AI-CRIP operationalises the insight that ethical AI deployment in criminal justice cannot be achieved through governance alone; it must be designed into the system from the ground up.

The pseudo-algorithm presented in Section 3 provides a formal, reproducible specification of the core plan-generation pipeline, facilitating implementation by practitioners and empirical evaluation by researchers. The layered architecture supports modular deployment: organisations with existing risk assessment tools can adopt only the service-matching and monitoring layers, while those building new systems can implement the full stack. This modularity addresses a common practical barrier to technology adoption in resource-constrained justice agencies.

Several important challenges and limitations deserve acknowledgement. First, the framework has not yet been empirically validated in a real deployment context. Future work should include a pilot implementation with a cooperating justice agency, followed by a randomised evaluation measuring reintegration outcomes, caseworker acceptance, and system reliability. Second, the NLP and risk classification models described in the framework require substantial training data—specifically, longitudinal records linking individual profiles to long-term recidivism and reintegration outcomes—that may be difficult to obtain in privacy-compliant form. Federated learning approaches that train models across multiple sites without centralising sensitive data represent a promising direction. Third, algorithmic bias in justice AI remains a deeply contested ethical and technical problem; no current fairness criterion commands universal normative assent, and AI-CRIP’s adoption of equalised odds is a defensible but contestable choice that should be revisited in dialogue with affected communities.

Digital inclusion is a fourth challenge. As noted in the literature review, digital literacy and device access barriers disproportionately affect older, less-educated, and non- English-speaking formerly incarcerated individuals. AI- CRIP’s multi-modal interface design mitigates but does not eliminate this barrier; deployment must be accompanied by digital literacy support programmes and provisions for offline or phone-based access. Fifth, stakeholder trust—from formerly incarcerated individuals, caseworkers, community organisations, and the public—is foundational to system adoption. Trust-building requires transparent communication about system capabilities and limitations, meaningful participation by affected communities in design processes, and demonstrated accountability when system errors occur. Looking forward, AI-CRIP provides a foundation for several extensions. Integration of large language model (LLM) capabilities could support conversational coaching interfaces that provide on-demand guidance to individuals navigating complex reintegration challenges. Graph neural network models could leverage social network data to identify prosocial connection opportunities that standard KNN matching would miss. Causal inference methods could move the analytics engine beyond correlation-based risk prediction toward

identifying the specific interventions most likely to benefit individual profiles. Cross-jurisdictional deployment would require harmonisation of data schemas and legal frameworks, representing both a technical and governance challenge. In sum, AI-CRIP provides a foundation for a new generation of evidence-based, ethically grounded reintegration technology—one that takes seriously both the transformative potential and the substantial risks of AI deployment in criminal justice contexts.

Ethical Approval: This manuscript reports studies that do not involve human participants, human data, human tissue, or animals.

Conflicts of Interest: The author has no conflict of interest to declare that is relevant to the content of this article.

Authors' Contributions: Vidhata Phani Datta Seethepalli contributed to the conceptualization, literature review, framework design, methodology development, and drafting of the complete manuscript.

Funding: The author did not receive financial support from any organization for the submitted work.

References

1. Pew Center on the States. *State of Recidivism: The Revolving Door of America's Prisons*. The Pew Charitable Trusts, Washington, D.C., 2011.
2. D. L. MacKenzie. *What Works in Corrections: Reducing the Criminal Activities of Offenders and Delinquents*. Cambridge University Press, Cambridge, UK, 2006.
3. C. A. Visser and J. Travis. Transitions from prison to community: Understanding individual pathways. *Annual Review of Sociology*, 29(1):89–113, 2003.
4. F. S. Taxman, E. S. Shephardson, and J. M. Byrne. *Tools of the Trade: A Guide to Incorporating Science into Practice*. National Institute of Corrections, U.S. Department of Justice, 2004.
5. C. O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishers, New York, NY, 2016.
6. J. Dressel and H. Farid. The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, 4(1):eaao5580, 2018.
7. W. Dieterich, C. Mendoza, and T. Brennan. *COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity*. Northpointe Inc., Research Department, 2016.
8. J. P. Singh and R. Mhaskar. Improving recidivism prediction using gradient-boosted trees with interpretable feature importance. *Journal of Criminal Justice*, 68:101692, 2020.
9. N. M. Morani, N. Wikoff, D. M. Linhorst, and S. Bratton. A description of the self-identified needs, service expenditures, and social outcomes of participants of a prisoner-reentry program. *The Prison Journal*, 91(3):347–365, 2011.
10. R. A. Desai, J. Lam, and R. A. Rosenheck. Childhood risk factors for criminal justice involvement in a sample of homeless people with serious mental illness. *Journal of Nervous and Mental Disease*, 188(6):324–332, 2000.
11. N. Wolff, M. Epperson, J. Shi, J. Huening, B. E. Schumann, and I. R. Sullivan. Mental health specialized probation caseloads: Are they effective? *International Journal of Law and Psychiatry*, 37(5):464–472, 2014.
12. J. L. Leidner and F. Schilder. Tracking down the missing links: Named entity recognition and disambiguation for automated literature search in clinical trial databases. *Journal of the American Medical Informatics Association*, 17(5):590–597, 2010.
13. Y. Zhang, M. Yang, D. Litman, and R. Passonneau. BERT-based legal text classification for sentencing recommendation extraction. In *Proc. 2019 Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3451–3456. ACL, 2019.
14. S. Amershi, D. Weld, M. Vorvoreanu, A. Fournier, B. Nushi, P. Collisson, et al. Guidelines for human-AI interaction. In *Proc. 2019 CHI Conf. on Human Factors in Computing Systems*, pages 1–13. ACM, 2019.
15. X. Wang, M. Yin, and J. Chen. Are explanations helpful? A comparative study of the effects of explanations in AI-assisted decision-making. In *Proc. 26th Int. Conf. on Intelligent User Interfaces*, pages 318–328. ACM, 2021.

16. M. Renzema and E. Mayo-Wilson. Can electronic monitoring reduce crime for medium to high-risk offenders? *Journal of Experimental Criminology*, 1(2):215–237, 2005.
17. I. Radanovic and R. Likic. Opportunities for use of blockchain technology in medicine. *Applied Health Economics and Health Policy*, 16(5):583–590, 2018.
18. S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org, 2019.
19. V. Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press, New York, NY, 2018.
20. B. Western and B. Pettit. Incarceration and social inequality. *Daedalus*, 139(3):8–19, 2010.
21. Preprints.org. *Loneliness and Social Isolation in Japan: AI, Technology and Structural Responses*. Preprints, 202402.1662, 2024. Available at: <https://www.preprints.org/manuscript/202402.1662>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.