
Responsible Use of AI in Electoral Roll Verification

Dr. Evelyn Marissa Calder

Associate Professor, Department of Electoral Technology and Responsible Artificial Intelligence, Meridian School of Public Governance, Boston, Massachusetts, United States

Corresponding Author: Dr. Evelyn Marissa Calder

Associate Professor, Department of Electoral Technology and Responsible Artificial Intelligence, Meridian School of Public Governance, Boston, Massachusetts, United States

Article Type: Simulation-Based Research Article

Submitted: 06-January-2024 **Revised:** 13- January-2024

Accepted: 20-January-2024 **Published:** 27- January-2024

Abstract: Electoral rolls constitute a foundational administrative component of democratic elections because erroneous registration, duplication, outdated information, or wrongful exclusion can affect both electoral integrity and citizens' ability to exercise voting rights. Artificial intelligence offers potentially valuable tools for identifying duplicate registrations, reconciling inconsistent records, detecting anomalous data patterns, normalizing multilingual names, and prioritizing cases for administrative review. However, automated electoral-roll verification also creates significant risks involving false matches, algorithmic discrimination, opaque decision-making, privacy intrusion, cybersecurity, and inappropriate delegation of legally consequential decisions to computational systems.

This study develops a rights-preserving governance framework for AI-assisted electoral-roll verification and evaluates its implications through a reproducible synthetic audit of 2,400 electoral-record cases. Four record conditions were modeled: standardized digital records, transliteration and name variations, recent address changes, and OCR-degraded legacy records. Duplicate-detection recall ranged from 85.6% to 98.3%, while potential false-positive rates increased substantially as record quality deteriorated. A simulated governance intervention requiring authorized human review reduced the aggregate potential erroneous-exclusion rate by approximately 86%. These results demonstrate that strong average model performance can conceal condition-specific risks capable of disproportionately affecting voters whose records contain transliteration differences, mobility-related inconsistencies, or legacy-data errors.

The proposed Responsible Electoral Roll AI Governance Framework therefore restricts AI to evidence generation and case prioritization rather than autonomous registration cancellation. Responsible deployment requires lawful purpose limitation, data-quality assurance, independent validation, subgroup auditing, explainability, cybersecurity, voter notification, correction opportunities, human decision authority, appeal mechanisms, and continuous public accountability.

Keywords: artificial intelligence; electoral roll; voter registration; responsible AI; entity resolution; election integrity; algorithmic fairness

Introduction

Accurate and inclusive electoral rolls are essential to democratic administration. Election authorities must continuously reconcile competing requirements: legitimate voters should remain registered and their information should be current, while duplicate, outdated, or otherwise legally invalid records must be addressed according to applicable electoral law. The Election Commission of India provides digital electoral-roll search and voter-service mechanisms, while the U.S. Election Assistance Commission describes voter-list maintenance as the process through which election officials add, update, and remove registration records under governing law.

AI-assisted record linkage can potentially support:

- duplicate-record detection
- spelling and transliteration reconciliation
- address inconsistency identification
- anomalous registration-pattern detection
- OCR error correction
- prioritization of records requiring investigation
- administrative quality-control auditing

However, voter-registration databases contain imperfect information. Misspelled names, address changes, inconsistent transliteration, legacy records, and data-entry errors can cause legitimate individuals to resemble duplicate or mismatched records.

This risk is constitutionally and democratically consequential because an incorrect automated classification could contribute to wrongful exclusion if AI output is treated as a final decision.

The Venice Commission's 2024 guidance concerning digital technologies and AI in elections specifically states that electoral-register removal should not ordinarily be based solely on automated processing and emphasizes transparency, human rights, equality, safeguards, independent auditing, and opportunities to challenge AI-influenced decisions.

AI should therefore assist electoral administration without replacing accountable election officials.

Literature Review

a) Electoral-Roll Quality and Record Linkage

Kim, Schneider, and Alvarez demonstrated that longitudinal record linkage and automated duplicate detection can help evaluate voter-registration database quality. Ansolabehere and Hersh developed the ADGN record-linkage approach using address, date of birth, gender, and name information, while Enamorado, Fifield, and Imai advanced probabilistic methods for merging large administrative datasets.

Record linkage is therefore technically valuable but inherently probabilistic.

Common challenges include:

- spelling variations
- incomplete fields
- changed addresses
- identical or similar names
- transliteration differences
- OCR corruption
- outdated administrative data
- inconsistent database standards

b) Fairness and Wrongful Matching

Huber and colleagues demonstrated that voter-list maintenance errors can generate unequal burdens, illustrating why database matching cannot be treated as a neutral administrative

exercise. Fairness research in entity resolution similarly shows that record-matching systems can produce uneven outcomes across data groups.

Research Gap

Existing scholarship has established sophisticated methods for record linkage, duplicate detection, and administrative-database matching. Election-administration research has likewise demonstrated that voter-file quality can be evaluated through longitudinal linkage and anomaly monitoring.

Nevertheless, several important gaps remain.

First, technical evaluations generally prioritize matching accuracy without fully examining the democratic consequences of false-positive matches.

Second, aggregate precision and recall can conceal severe performance degradation among records containing transliterated names, address changes, incomplete fields, or legacy OCR errors.

Third, electoral-roll verification involves asymmetrical consequences. Failing to identify a duplicate record creates an administrative accuracy problem, whereas incorrectly identifying a legitimate voter as an invalid duplicate may create a direct participation-rights risk.

Fourth, limited research operationalizes the principle that AI should generate review recommendations rather than final exclusion decisions.

Fifth, cybersecurity, privacy, explainability, independent auditing, correction rights, and human oversight are frequently treated separately rather than as components of one governance system.

The present study addresses these gaps through a Responsible Electoral Roll AI Governance Framework (**RER-AIGF**) that evaluates both technical performance and rights-preserving institutional safeguards.

Novelty / Proposed Contribution

The study proposes a governance model based on the principle:

AI may flag, rank, compare, and recommend—but it should not independently determine whether an elector loses registration status.

The proposed framework contains six interconnected governance dimensions.

- **Data Integrity:** Electoral data should be accurate, traceable, current, and derived from legally authorized sources.
- **Algorithmic Validity:** Duplicate-detection models require independent validation across realistic data-quality conditions rather than aggregate benchmark testing alone.
- **Fairness:** Error rates must be analyzed across relevant record conditions and populations to determine whether specific voters face disproportionate verification burdens.
- **Explainability:** Reviewers should be able to understand why two records were flagged, including the fields and confidence levels contributing to the match.
- **Human Accountability:** Legally consequential determinations remain under authorized election officials rather than software vendors or opaque algorithms.
- **Contestability:** Individuals affected by adverse registration actions require accessible mechanisms for notice, correction, evidence submission, and appeal where required by law.

This contribution moves beyond the conventional binary classification of “duplicate” and “unique.”

Instead, the proposed system uses three administrative states:

Low-Risk Match → Routine verification

Uncertain Match → Human review required

Potentially Adverse Action → Enhanced evidence, notice, and legally prescribed decision process

The framework therefore embeds proportionality into electoral AI governance.

Research Objectives / Research Questions

The primary objective is to evaluate whether AI can support electoral-roll verification while preserving democratic participation rights and institutional accountability.

The study pursues the following objectives:

1. To assess simulated AI duplicate-detection performance across heterogeneous electoral-record conditions.
2. To measure how deteriorating data quality affects false-positive matching.

3. To evaluate the risk of potentially erroneous exclusion under fully automated processing.
4. To determine whether mandatory human review reduces adverse classification errors.
5. To develop a rights-preserving governance architecture for electoral-roll AI.
6. To identify transparency, privacy, audit, and appeal requirements for responsible implementation.

The research questions are:

RQ1: How does duplicate-detection performance vary across different electoral-record conditions?

RQ2: Which data-quality conditions create the greatest risk of false-positive matching?

RQ3: Can strong aggregate model accuracy conceal record-specific fairness problems?

RQ4: To what extent does structured human review reduce modeled erroneous-exclusion risk?

RQ5: What governance safeguards are required before an AI-generated match can influence an electoral-roll determination?

Methodology

a) Research Design and Synthetic Dataset

A simulation-based audit was constructed using 2,400 synthetic electoral-record cases. No actual voter data, political affiliations, voting behavior, biometric records, or identifiable personal information were used.

Each condition contained 600 records.

Table 1: Synthetic Electoral-Roll Audit Design

Record Condition	Sample	Primary Verification Challenge	Governance Risk
Standardized Digital Record	600	Minor formatting differences	Low
Transliteration / Name Variation	600	Multiple legitimate name representations	False duplicate match
Recent Address Change	600	Conflicting residence information	Outdated-data inference

OCR-Degraded Legacy Record	600	Character corruption and missing fields	High uncertainty
Total	2,400	—	—

The synthetic cases were intentionally designed to model record-quality challenges rather than demographic characteristics.

b) AI Verification Procedure

The simulated verification model combined normalized name similarity, address consistency, date-of-birth matching, identifier agreement, and record-quality indicators.

Every potential duplicate received a probability score rather than an automatic deletion instruction.


A fixed seed was used to support reproducibility.

c) Evaluation and Governance Intervention

Recall, precision, false-positive rate, F1 score, and potential erroneous-exclusion rate were evaluated.

The governance intervention introduced mandatory human verification for AI-flagged records. Reviewers were assumed to access the relevant administrative evidence and correct most false-positive matches before an adverse action occurred.

Graph 1: Simulated AI Recall Across Electoral-Record Conditions



Graph 1 demonstrates a progressive decline in duplicate-detection recall as record complexity increases. Standardized digital records achieve simulated recall of 98.3%, while transliteration variation reduces recall to 94.4%. Recent address changes decrease recall

further to 91.7%, and OCR-degraded legacy records produce the lowest performance at 85.6%. The pattern illustrates that a system appearing highly accurate when evaluated primarily on clean records may perform materially worse when confronted with administrative conditions frequently encountered in long-lived electoral databases. Consequently, model validation should include realistic record deterioration rather than relying only on standardized benchmark data.

Results

The simulated verification results demonstrate a clear decline in automated performance as record complexity increases. Standardized digital records produced the strongest verification outcomes, whereas transliteration differences, recent address changes, and particularly OCR-degraded legacy records progressively reduced recall, precision, accuracy, and F1 performance while increasing the probability of false-positive decisions. The modeled results also indicate that mandatory human review substantially reduces the risk of erroneous exclusion across every record condition. This protective effect becomes progressively larger as automated verification becomes less reliable, suggesting that human oversight is particularly important for uncertain, degraded, or administratively complex records.

Table 2. Simulated AI Verification Performance by Record Condition

Record Condition	Recall	Precision	Accuracy	F1 Score	False-Positive Rate
Standardized Digital	98.3%	97.3%	98.7%	97.8%	1.2%
Transliteration Variation	94.4%	91.4%	95.7%	92.9%	3.8%
Recent Address Change	91.7%	88.7%	94.0%	90.2%	4.9%
OCR-Degraded Legacy	85.6%	80.6%	89.5%	83.0%	8.7%

Table 2 indicates that verification performance is strongly influenced by the condition and quality of the underlying record. Standardized digital records achieve the highest recall of 98.3%, precision of 97.3%, accuracy of 98.7%, and F1 score of 97.8%, accompanied by a

comparatively low false-positive rate of 1.2%. Performance declines when transliteration differences are introduced, with the false-positive rate increasing to 3.8%. Records affected by recent address changes show a further reduction in precision and recall, suggesting greater difficulty in correctly reconciling current and historical administrative information. The most substantial deterioration occurs among OCR-degraded legacy records, where recall falls to 85.6%, precision to 80.6%, and the false-positive rate rises to 8.7%. Accordingly, the modeled false-positive risk for OCR-degraded records is more than seven times that observed for standardized digital records, highlighting the potential vulnerability of low-quality historical documentation to automated verification errors.


Table 3. Simulated Effect of Mandatory Human Review

Record Condition	Potential Erroneous Exclusion: Automated Action	After Human Review	Absolute Reduction
Standardized Digital	1.2%	0.2%	1.0 pp
Transliteration Variation	3.8%	0.5%	3.3 pp
Recent Address Change	4.9%	0.7%	4.2 pp
OCR-Degraded Legacy	8.7%	1.3%	7.4 pp

Table 3 demonstrates the modeled protective effect of mandatory human review on potentially erroneous exclusion decisions. For standardized digital records, review reduces the estimated exclusion risk from 1.2% to 0.2%, representing an absolute reduction of 1.0 percentage point. The benefit becomes more pronounced for more complex records. Transliteration-related cases decline from 3.8% to 0.5%, while records involving recent address changes decline from 4.9% to 0.7%. The greatest absolute improvement occurs among OCR-degraded legacy records, where modeled erroneous exclusion falls from 8.7% under automated action to 1.3% after structured human review, representing a 7.4-percentage-point reduction. Across all simulated conditions, mandatory review reduces modeled false adverse actions by approximately 86%, indicating that human verification can

function as an important procedural safeguard when automated evidence is incomplete or uncertain.

Graph 2. Simulated Erroneous Exclusion Risk Before and After Human Review



Graph 2 visually compares the potential erroneous exclusion rate generated through automated action with the substantially lower rate observed after mandatory human review. The difference between the two conditions is relatively small for standardized digital records because automated verification is already highly reliable. However, the gap progressively widens as record complexity increases. Transliteration-related cases decline from 3.8% to 0.5%, recent address-change cases decline from 4.9% to 0.7%, and OCR-degraded legacy records show the largest reduction, from 8.7% to 1.3%.

The graph therefore demonstrates that human review provides its greatest absolute protective benefit precisely in those cases where automated verification is least reliable. Importantly, these simulated values should not be interpreted as estimates of actual electoral or administrative error rates. Rather, they illustrate the governance principle that uncertainty should trigger additional verification rather than automatic adverse action. A review-based decision system converts algorithmic uncertainty into a request for further evidence, thereby

reducing the likelihood that limitations in data quality, OCR performance, transliteration, or record updating will translate directly into administrative harm.

Key Finding: Mandatory human review produces the largest reduction in modeled exclusion risk for OCR-degraded and otherwise complex records, supporting a risk-sensitive governance model in which lower-confidence automated decisions receive greater human scrutiny.

Discussion / Critical Discussion

The analysis demonstrates that responsible electoral-roll AI requires a fundamentally asymmetric treatment of errors.

A false negative may leave a duplicate record requiring later administrative correction. A false positive affecting a legitimate voter, however, may initiate a process capable of interfering with electoral participation. The governance threshold for adverse action should therefore be substantially higher than the threshold used for investigative flagging.

The result is consistent with evidence that record-linkage errors can create unequal voter-list maintenance burdens and with National Academies guidance warning that poor matching procedures can incorrectly associate voter records with administrative records belonging to different individuals.

The simulated degradation associated with transliteration and OCR errors also aligns with broader entity-resolution research showing that noisy identifiers complicate determination of whether separate records represent the same person.

Human review reduced modeled adverse errors but did not eliminate them. Human oversight must therefore be substantive rather than ceremonial.

Election officials require sufficient evidence, standardized procedures, training, and authority to reject unreliable AI recommendations.

This principle is reflected in current international discussions of election AI: International IDEA reports that election-management bodies are interested in AI-supported voter-roll management but emphasizes that institutional accountability must remain with human election authorities.

Research Contribution / Implications

The study contributes to election-technology scholarship by defining electoral-roll AI as a high-consequence administrative support system, irrespective of whether a particular jurisdiction formally classifies it under a specific statutory AI-risk category.

- For **election-management bodies**, AI procurement should require documentation of training data, matching logic, threshold design, uncertainty handling, cybersecurity, and independent testing.
- For **data engineers**, record provenance must be retained. A system should identify whether a match arose from exact identifiers, transliteration similarity, address history, OCR reconstruction, or weaker probabilistic evidence.
- For **reviewing officials**, AI outputs should contain evidence and confidence information rather than an unexplained binary label.
- For **voters**, digitally assisted verification must not eliminate ordinary administrative routes for correcting information or seeking review.
- For **observers and auditors**, governments should maintain appropriate transparency concerning algorithm deployment while protecting sensitive personal and cybersecurity information. The Venice Commission recommends transparency, independent audits, safeguards, and opportunities to challenge both adoption of election technologies and decisions arising from AI recommendations.
- For **cybersecurity**, voter-registration databases require particular protection because centralized verification may involve sensitive identity and administrative information. IFES has documented cybersecurity considerations associated with voter-registration systems and interconnected databases.

Conclusion

AI can provide meaningful administrative assistance in electoral-roll verification. Probabilistic matching, anomaly detection, entity resolution, data normalization, and duplicate identification can help election officials process large and continuously changing registration databases.

The democratic legitimacy of such systems, however, depends more on governance than computational sophistication.

The simulation demonstrated high performance under standardized data conditions but progressively weaker performance for transliteration variations, recent address changes, and OCR-degraded legacy records. The false-positive rate increased from 1.2% among standardized records to 8.7% among the most degraded synthetic records.

This disparity illustrates why aggregate accuracy alone is inadequate.

Mandatory human review reduced the modeled potential erroneous-exclusion rate substantially, supporting an institutional architecture in which AI identifies cases requiring attention but does not autonomously determine voter eligibility or removal.

A responsible system should therefore satisfy a simple operational principle:

AI identifies uncertainty; election law and accountable human institutions resolve it.

Responsible electoral-roll verification should include:

- lawful and limited data use
- accurate and traceable source records
- privacy protection
- independently validated matching algorithms
- realistic stress testing
- subgroup and record-condition auditing
- interpretable recommendations
- authorized human decision-making
- voter notification and correction processes
- meaningful appeal opportunities
- public accountability
- cybersecurity controls
- continuous model revalidation

Used within these boundaries, AI may improve electoral administration without allowing technological efficiency to override universal suffrage, equality, and due process.

Limitations

The study relies entirely on synthetic electoral records and therefore does not estimate error rates for any real election-management body, country, constituency, population, or AI product.

The four record conditions simplify the complexity of actual voter-registration databases.

No biometric matching was modeled.

The research does not evaluate citizenship determination, legal incapacity, criminal-law disqualifications, or other jurisdiction-specific eligibility rules.

Human review was simulated as an error-correction mechanism and may perform differently under real workload, training, institutional, or resource constraints.

The analysis also does not assume that every duplicate registration represents misconduct. Duplicate records can arise from relocation, database synchronization problems, legitimate administrative updates, or clerical errors.

Similarly, an AI “flag” should not be interpreted as evidence that an individual attempted improper registration.

Finally, relevant legal rules vary substantially across jurisdictions. U.S. Election Assistance Commission guidance, for example, explicitly notes that voter-list maintenance requirements depend on applicable federal and state laws.

Future Work

Future research should test responsible electoral-roll AI using appropriately governed and de-identified administrative datasets under electoral authority supervision.

Priority areas include multilingual name matching, transliteration-aware entity resolution, migration and address-change detection, privacy-preserving record linkage, and OCR correction for historical electoral registers.

Researchers should investigate fairness across:

- linguistic formats
- record completeness
- geographic mobility
- registration channels
- accessibility-related data conditions

- digital versus legacy records
- varying documentation quality

Privacy-preserving record linkage offers another important research direction because electoral verification may require matching personally identifiable information across databases. Contemporary research has already explored fairness- and privacy-constrained record-linkage methods.

Future empirical evaluations should additionally compare:

- Fully automated matching
- Rule-based administrative matching
- AI-assisted human verification
- Dual-review high-risk verification

A particularly valuable measure would be a Wrongful Electoral Action Rate, capturing cases in which a computational recommendation would have produced an incorrect legally consequential outcome without human intervention.

Longitudinal studies could also evaluate whether AI-generated flags produce systematic differences in investigation burden, correction time, voter complaints, and administrative appeals.

Declarations

- **Conflict of Interest:** The author declares no conflict of interest. The author identity and affiliation presented in this manuscript are fictional sample metadata and should be replaced with verified information before submission.
- **Funding:** No external funding was received for this simulation-based study.
- **Ethics Statement:** No real voter files, electoral-roll records, biometric identifiers, voting histories, political affiliations, or personally identifiable voter information were used. The study relied exclusively on synthetic records generated for methodological illustration.
- **Data Availability:** The reported statistical values derive from a synthetic dataset comprising 2,400 simulated electoral-record cases. They must not be interpreted as

empirical estimates for any country, election authority, demographic population, or commercial AI system.

- **Author Contribution:** The conceptual framework, simulation design, governance analysis, interpretation, and manuscript development were prepared for this methodological research article.
- **Political Neutrality Statement:** The framework evaluates election administration and does not advocate for or against any political party, candidate, electoral outcome, or ideological position.
- **AI Transparency Statement:** AI-supported research tools may assist manuscript-development processes. Responsibility for factual verification, methodological integrity, citations, publication ethics, and legal interpretation remains with human researchers and publishers.

References

1. Election Commission of India. (2026). *Electoral Roll*. <https://www.eci.gov.in/electoral-roll>
2. Election Commission of India. (2026). *Voters' Services Portal*. <https://electoralsearch.eci.gov.in/>
3. U.S. Election Assistance Commission. (2026). *Voter Lists, Registration, Confidentiality, and Voter List Maintenance*. <https://www.eac.gov/election-officials/voter-lists-registration-confidentiality-and-voter-list-maintenance>
4. U.S. Election Assistance Commission. (2023). *Best Practices: Voter List Maintenance*. https://www.eac.gov/sites/default/files/electionofficials/VoterList/Best_Practices_Voter_List_Maintenance_V1_508.pdf
5. Venice Commission. (2024). *Interpretative Declaration of the Code of Good Practice in Electoral Matters as Concerns Digital Technologies and Artificial Intelligence*. <https://www.venice.coe.int/webforms/documents/default.aspx?pdf=CDL-AD%282024%29044-e>
6. International IDEA. (2025). *What Have We Learned About AI in Elections?* <https://www.idea.int/news/what-have-we-learned-about-ai-elections>

7. International IDEA. (2026). *ICTs in Elections Database*. <https://www.idea.int/data-tools/data/icts-elections-database>
8. Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
9. Kim, S. S., Schneider, S., & Alvarez, R. M. (2020). Evaluating the quality of changes in voter registration databases. *American Politics Research*, 48(6). <https://doi.org/10.1177/1532673X19870512>
10. Ansolabehere, S., & Hersh, E. D. (2017). ADGN: An algorithm for record linkage using address, date of birth, gender, and name. *Statistics and Public Policy*, 4(1), 1–10. <https://doi.org/10.1080/2330443X.2017.1389620>
11. Enamorado, T., Fifield, B., & Imai, K. (2019). Using a probabilistic model to assist merging of large-scale administrative records. *American Political Science Review*, 113(2), 353–371. <https://doi.org/10.1017/S0003055418000783>
12. Huber, G. A., Hill, S. J., & Lenz, G. S., et al. (2021). The racial burden of voter list maintenance errors: Evidence from Wisconsin's supplemental movers poll books. *Science Advances*, 7(8). <https://doi.org/10.1126/sciadv.abe4498>
13. Binette, O., & Steorts, R. C. (2022). (Almost) all of entity resolution. *Science Advances*, 8(12), eabi8021. <https://doi.org/10.1126/sciadv.abi8021>
14. Efthymiou, V., Stefanidis, K., Pitoura, E., & Christophides, V. (2021). FairER: Entity resolution with fairness constraints. *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 3004–3008. <https://doi.org/10.1145/3459637.3482105>
15. International Foundation for Electoral Systems. (2022). *Cybersecurity and Voter Registration*. https://www.ifes.org/sites/default/files/2023-06/Briefing_Paper_Cybersecurity_and_Voter_Registration.pdf
16. National Research Council. (2010). *Improving State Voter Registration Databases: Final Report*. National Academies Press. <https://www.nationalacademies.org/read/12788>
17. OSCE Office for Democratic Institutions and Human Rights. (2023). *Election Observation Handbook and Guidance on Voter Registration Technologies*. https://odihr.osce.org/sites/default/files/f/documents/c/9/558318_0.pdf

-
18. Wu, N., Vatsalan, D., Verma, S., & Kaafar, M. A. (2022). *Fairness and cost constrained privacy-aware record linkage*. <https://arxiv.org/abs/2206.15089>