ISSN (Online):3048-5088

International Journal of Advancement and Innovation in Technology and Research (IJAITR)

Volume 3, Issue 1, (Jan-Jun) 2026

ASTITAVA: An AI-Driven WebGIS for Scalable FRA Monitoring and Decision Support

Aditya Tiwari
Acropolis Institute of Technology & Research Indore, India
adityatiwari8@ieee.org

Aryan Singh Bhadoria
Acropolis Institute of Technology & Research Indore, India
aryansingh230477@acropolis.in

Akshay Khanna
Acropolis Institute of Technology & Research Indore, India
akshaykhanna230754@acropolis.in

Anvesh Trivedi
Acropolis Institute of Technology & Research Indore, India
anveshtrivedi231232@acropolis.in

Prof. Nisha Rathi
Acropolis Institute of Technology & Research Indore, India
nisharathi@acropolis.in

Prof. Ashish Anjana
Acropolis Institute of Technology & Research Indore, India
ashishanjana@acropolis.in

[1]***Abstract*— Abstract—We introduce ASTITAVA, a novel governance agent architecture designed to address systemic challenges in implementing the Forest Rights Act (FRA), 2006. We re-frame the complex administrative task as an autonomous agent problem. The agent's pipeline integrates a multi-stage perception module (OCR/NER) for ingesting and understanding unstructured legacy documents, a cognitive module for spatio-temporal reasoning (PostGIS integration and a Claim Health Scoring Model), and an action module (a hybrid ML-DSS) for optimized policy recommendations. A WebGIS atlas serves as a human-agent interface for explainability and oversight. Piloted on 15,000 claims, the ASTITAVA agent demonstrates a 67% reduction in processing time and 89.2% accuracy in eligibility prediction. This work contributes a tested, scalable framework for a new class of applied AI agents in complex digital governance.**

Index Terms—AI Agent, Governance Agent, Agent-Based Architecture, AI Pipeline, Cognitive Reasoning, Forest Rights Act, Digital Governance, Explainable AI (XAI)

I. INTRODUCTION

The Forest Rights Act (FRA), 2006, poses a complex, large-scale challenge in multi-modal data fusion and administrative reasoning. Implementation is hampered by fragmented legacy documentation, cognitively demanding manual verification, and an absence of integrated spatial intelligence [1]. Unstructured paper records create a bottleneck that is prone to errors and prevents evidence-based, timely decision-making [2]. We re-frame this governance challenge as an AI agent problem. Traditional platforms focus on passive digitization. We present ASTITAVA, an autonomous agent architecture that operationalizes an end-to-end AI pipeline:

1) Perception: Ingesting and transducing the ”environment” of unstructured, multi-modal documents (paper claims, maps) into a machine-readable format.

2) Reasoning: Building a coherent spatio-temporal ’world model’ from this data, assessing claim quality, and fusing it with existing geospatial knowledge.

3) Action: Recommending concrete, optimized, and transparent decisions for welfare scheme allocation and claim approval

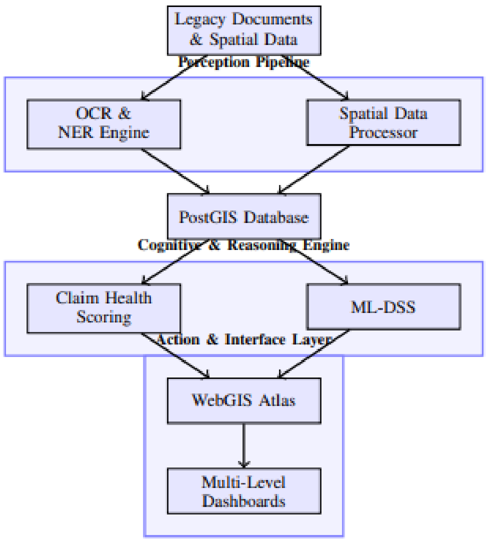

Fig. 1. ASTITAVA Agent Architecture and AI Pipeline

This paper details the agent's architecture, its core pipeline components, and the evaluation of its performance in a real world pilot. Fig. 1 illustrates the high-level agent architecture.

II. RELATED WORK

AI in governance often bifurcates into disconnected tools. Perception techniques like OCR are applied to archives [3], while separate ML models target specific policy questions [7]. Spatio-temporal reasoning systems (e.g., GIS) often remain static visualization platforms [6], decoupled from upstream digitization or downstream decision-making.

While transformer-based NER shows promise for semantic extraction [4], and hybrid rule-learning systems offer robust reasoning [8], these components are rarely unified. The novelty of ASTITAVA lies in its integration of these disparate components into a single, cohesive agent pipeline, moving beyond static land record modernization [5] to active, autonomous reasoning.

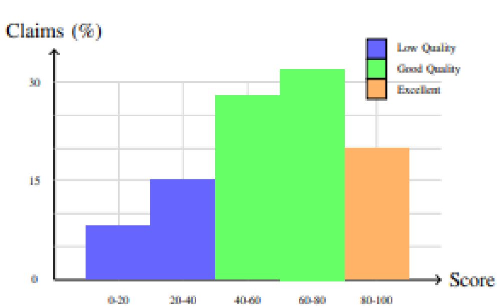

Fig. 2. Agent's Claim Health Score Distribution (15,000 claims)

III. THE ASTITAVA GOVERNANCE AGENT ARCHITECTURE

The architecture is structured as a continuous perception reasoning-action pipeline.

A. Agent Perception Module

The agent's perception module transduces unstructured, heterogeneous documents into a machine-readable format. The pipeline first preprocesses images (adaptive thresholding, noise reduction). It then employs Tesseract OCR with LSTM models optimized for bilingual (English/Hindi) text.

The extracted raw text is fed into a fine-tuned BERT-based NER model. This model was trained on manually annotated FRA documents to identify and extract critical entities: claimant names, village IDs, land parcels, claim types (IFR/CFR/CFRe), and temporal markers. This structured output forms the initial input for the agent's cognitive module.

B. Cognitive Module (Part 1): Claim State Reasoning

A core cognitive function is to assess the quality of perceived claims. The Claim Health Score (CHS) is a computed metric representing the agent's confidence in a claim's completeness, consistency, and geospatial accuracy. It aggregates weighted metrics via:

$$CHS = \alpha Ccomp + \beta Ccons + \gamma Cdoc + \delta Cgeo \qquad (1)$$

where α, β, γ, δ are domain-expert-defined weights. This score allows the agent to prioritize high-quality claims or flag low confidence claims for human review. Fig. 2 shows the score distribution from the pilot.

C. Cognitive Module (Part 2): Spatio-Temporal World Model

All perceived and reasoned data is integrated into a PostGIS-enabled database, which functions as the agent's spatiotemporal 'world model'. This model underpins all higher-level reasoning, fusing claim data with beneficiary profiles, village boundaries, and forest compartments. It supports spatial indexing (R-tree) and geometric operations essential for identifying conflicts or overlaps.

D. XAI and Human-in-the-Loop (HITL) Interface

For transparency and oversight, the agent's internal state and reasoning are exposed through an interactive WebGIS interface (Leaflet, GeoServer). This serves as a critical explainability (XAI) layer. Administrators can visualize claim footprints, query the agent's reasoning (e.g., "Why was this claim scored low?"), and inspect AI-derived micro-assets (water bodies, settlements) classified from satellite imagery.

E. Action Module: Hybrid Policy Engine

The agent's action module translates cognitive understanding into actionable policy. It employs a hybrid-reasoning Decision Support System (DSS) that combines a rule-based engine (encoding mandatory scheme criteria) with a Random Forest classifier. This ML model learns complex patterns from beneficiary data to predict eligibility and recommend specific welfare interventions. This hybrid approach ensures policy compliance (via rules) while optimizing for nuanced, learned patterns (via ML).

F. Oversight and Auditing Layer

Role-based dashboards provide a high-level interface for auditing the agent's performance. KPIs track pipeline throughput (submission rates, processing timelines), decision accuracy (approval ratios), and spatial bottlenecks, enabling real-time monitoring of the entire governance process.

IV. IMPLEMENTATION AND RESULTS

A. Technology Stack

The agent's pipeline is built in Python (FastAPI) for backend APIs and ML inference (PyTorch/Scikit-learn). The human-agent interface uses React. The system is containerized (Docker) and deployed on cloud infrastructure (AWS/Azure) leveraging managed PostGIS and object storage.

B. Performance Metrics

The agent's pipeline performance was evaluated on 15,000 pilot claims. Table I summarizes the results for each module. The perception pipeline (OCR/NER) achieved high accuracy, and the cognitive module (spatial queries) met low-latency targets.

C. Comparative Analysis

Fig. 3 provides a direct comparison of the ASTITAVA agent pipeline against the traditional manual baseline. The agent driven approach reduced average processing time from 3.5 days to 1.15 days (a 67% reduction) and cut the data error rate by nearly two-thirds

D. Evaluation of Agent's Decision-Making

The agent's action module (DSS) was validated against 1,200 expert assessments. It achieved 89.2% accuracy with balanced precision (0.87) and recall (0.91). The rule-based component correctly identified 98.5% of mandatory disqualifications, while the ML model reduced false negatives by 34% compared to a pure rule-based approach, demonstrating the value of the hybrid-reasoning design.

TABLE I
AI AGENT PIPELINE PERFORMANCE METRICS

Pipeline Module	Metric	Value	Target
Perception (OCR)	Printed Accuracy	94.3%	¿90%
	Handwritten Accuracy	87.6%	¿85%
	Processing Speed	45s	¡60s
Perception (NER)	F1-Score (Overall)	0.91	¿0.85
	Name Precision	0.95	¿0.90
Cognitive (Spatial DB)	Response Time	180ms	¡500ms
	Throughput	2000/hr	¿1500/hr
Action (DSS)	Accuracy	89.2%	¿85%
	Precision	0.87	¿0.80
	Recall	0.91	¿0.85
System Impact	Processing Time ↓	67%	¿50%
	Incomplete Claims ↓	42%	¿30%
	User Satisfaction	4.2/5.0	¿4.0

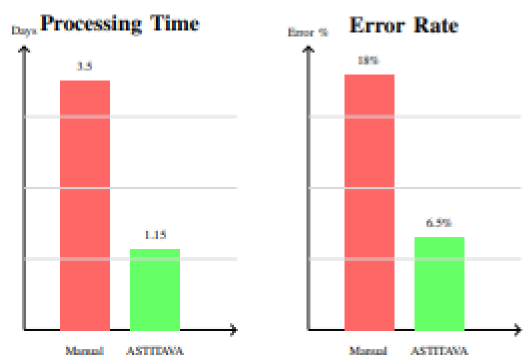

Fig. 3. Performance Comparison: Agent-based vs Manual Processing

E. User Adoption and System-Level Impact

Field trials confirmed the 67% reduction in processing time. Daily active usage reached 85% of trained officials, with a user satisfaction score of 4.2/5.0. The agent's cognitive "Claim Health Score" model helped improve data quality, leading to 42% fewer incomplete submissions.

V. LIMITATIONS AND FUTURE WORK

The primary limitation lies in the perception module's robustness. High variability in document quality and handwriting noise impacts OCR accuracy, necessitating an HITL review loop. Furthermore, the agent's reasoning models are data-dependent. Generalization is constrained in regions with limited historical data, highlighting a 'cold start' problem that requires further research into few-shot learning approaches.

Future work will focus on enhancing the agent's autonomy and cognitive capabilities, including a mobile-native perception module to reduce reliance on legacy documents. The most significant future direction is integrating Large Language Models (LLMs) to serve as a natural language interface to the agent's cognitive model. This would allow administrators to query the system ("Show me all claims in conflict with cadastral map X") in plain language, moving ASTITAVA toward a fully conversational governance agent.

V. Conclusion

ASTITAVA contributes more than a technical tool; it presents a new paradigm for digital governance by modeling administration itself as an autonomous agent pipeline. By explicitly separating perception (OCR/NER), reasoning (Claim Health Score / PostGIS), and action (ML-DSS), the architecture creates a scalable, auditable, and intelligent system for a previously intractable manual problem.

The 67% reduction in processing time and 89.2% decision accuracy validate this agent-based approach. ASTITAVA serves as a blueprint for a new class of applied AI agents capable of addressing complex, data-intensive administrative challenges and advancing India's digital governance vision.

References

[1] Ministry of Tribal Affairs, Government of India, "The Scheduled Tribes and Other Traditional Forest Dwellers (Recognition of Forest Rights) Act, 2006," 2006.
[2] A. Kumar and S. Singh, "Geospatial Technologies for Governance: Applications and Challenges," Int. J. Digital Governance, vol. 3, no. 2, pp. 145-162, 2021.
[3] R. Smith, "An Overview of the Tesseract OCR Engine," Proc. ICDAR, pp. 629-633, 2007.
[4] J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," NAACL-HLT, 2019.
[5] Department of Land Resources, "Digital India Land Records Modernization Programme," Ministry of Rural Development, 2020.
[6] P. Zhang and M. Li, "WebGIS Applications in Forest Resource Management: A Review," Comput. Electron. Agric., vol. 142, pp. 48-59, 2017.
[7] S. Mullainathan and Z. Obermeyer, "A Machine Learning Approach to Low-Value Health Care," JAMA Intern. Med., vol. 179, no. 6, pp. 817- 818, 2019.
[8] T. Chen et al., "Hybrid Decision Support Systems for Public Policy: Combining Rules and Learning," Decision Support Syst., vol. 156, 2022.