Interpretable Artificial Intelligence (AI) Driven Public Service Delivery Monitoring Model: A Case of Zimbabwe

Zvinodashe Revesai.[1], Gilford T Hapanyengwi.[2]

Abstract

This research develops an interpretable Artificial Intelligence (AI) model for public service delivery monitoring in resource-constrained governmental environments, addressing critical challenges of bureaucratic inefficiency, operational opacity, and limited technological infrastructure in Zimbabwe's public sector. Existing AI-driven governance solutions predominantly target developed economies with robust infrastructure, creating a significant gap in practical deployment models for developing contexts where lightweight architectures, explainability integration, and empirical validation remain inadequately addressed. A multi-layered neural network architecture was developed comprising data preprocessing, Natural Language Processing (NLP)-based feature extraction, ensemble machine learning for bottleneck detection, real-time performance monitoring, and Local Interpretable Model-agnostic Explanations (LIME)/SHapley Additive exPlanations (SHAP)-based interpretability interfaces. The system was empirically validated across three simulated governmental departments analyzing 1,000 administrative interactions. The model demonstrated 45% reduction in processing latency, 60% improvement in operational transparency, and 87% accuracy in bottleneck detection whilst operating efficiently within existing governmental infrastructure with minimal computational overhead. This research provides public sector administrators and policymakers with a scalable, context-aware technological solution for digital governance enhancement without substantial infrastructure investment. The study contributes a novel lightweight, interpretable AI model specifically designed for developing economy contexts, bridging the gap between AI capability and practical feasibility in resource-constrained public sector environments.

Keywords: Interpretable Artificial Intelligence; Resource-Constrained; Digital Governance; Workflow Optimization; Service Transparency

1. Introduction

The contemporary administrative landscape in Zimbabwe faces persistent challenges of bureaucratic inefficiency and technological limitations that significantly impede operational effectiveness, citizen satisfaction, and service delivery quality. Public sector organisations struggle with administrative processes characterised by prolonged processing times, opaque decision-making mechanisms, and limited transparency (Gondo & Suwaryono, 2024; Mutengambiri et al., 2024). The integration of Artificial Intelligence (AI) in governmental functions presents both unprecedented opportunities and complex challenges for enhancing public service delivery, particularly in developing economies (Alshahrani et al., 2024; David, 2024). Previous work has explored AI integration in Zimbabwean public sectors (Revesai & Hapanyengwi, 2024), highlighting infrastructure constraints and policy considerations relevant to developing economies.

[1] Dr Zvinodashe Revesai, Reformed Church University, P O Box 80 Masvingo, Zimbabwe, Corresponding author: revesaiz@rcu.ac.zw
[2] Dr Gilford T Hapanyengwi ,Zimbabwe Council for Higher Education, Zimbabwe, 21 J. M. N. Nkomo Road | P. Bag H100 Hatfield, Harare, Zimbabwe, E-mail: gthapanyengwi@zimche.ac.zw

Digital transformation initiatives have gained momentum globally, with AI emerging as a critical enabler for modernising government services and improving citizen engagement (Ahn & Chen, 2022; Dunleavy & Margetts, 2023). However, deploying AI-driven solutions in resource-constrained environments presents unique challenges that remain inadequately addressed in current literature, particularly regarding technical infrastructure limitations and stakeholder acceptance (van Noordt & Misuraca, 2022).

Traditional methods of service delivery oversight rely predominantly on manual documentation, simplistic tracking systems, and limited metrics that fail to capture the complexity of modern processes. Recent advancements in AI have demonstrated promise in enhancing assessment capabilities and delivery efficiency (Kulal et al., 2024; Medaglia et al., 2021). The potential for AI to revolutionise government operations through improved decision-making, enhanced citizen services, and streamlined functions has been increasingly recognised across both developed and emerging economies (Agba et al., 2023; Criado et al., 2025).

Despite growing recognition of AI's transformative potential, existing solutions face three critical barriers that limit practical deployment. First, the technical intensity of sophisticated systems often exceeds infrastructure capacity in developing economies, creating deployment challenges (Neumann & Guirguis, 2024). Second, the lack of transparency and explainability in decision-making processes undermines stakeholder trust and institutional adoption, particularly in public sector settings where accountability is paramount (Baehrens et al., 2023; Christou et al., 2024). Third, insufficient adaptation of models for resource-constrained settings creates a gap between theoretical capabilities and real-world feasibility (Brillianto et al., 2024).

The challenge of AI adoption in public management is further complicated by the need for ethical deployment and citizen acceptance. Research indicates that public trust in automated systems is fundamentally influenced by transparency, explainability, and perceived fairness in algorithmic decision-making (Aoki, 2020; Chen et al., 2023). Integration in government services requires careful consideration of ethical implications, particularly regarding privacy, bias, and accountability in automated processes (Jennings & Cox, 2023; Nguyen et al., 2023).

Current literature reveals several critical gaps that this investigation addresses. First, there is limited research on lightweight architecture specifically designed for limited-resource environments, with most studies focusing on developed economies without considering infrastructure limitations (van Noordt & Misuraca, 2022). Second, insufficient integration of explainability mechanisms in public sector applications has resulted in low stakeholder acceptance and limited practical deployment success (Madan & Ashok, 2023). Third, minimal empirical validation of tracking systems in developing economy scenarios has left an evidence gap regarding real-world feasibility and effectiveness (Nalubega & Uwizeyimana, 2024).

Zimbabwe presents unique opportunities for AI deployment in public management, particularly given the country's ongoing digital transformation initiatives and the need for improved service delivery mechanisms (Poshai & Vyas-Doorgapersad, 2023; Nyikadzino & Vyas-Doorgapersad, 2023). However, successful deployment requires careful consideration of local constraints, including limited technological infrastructure, resource availability, and the need for stakeholder buy-in across different levels of government.

This investigation addresses three primary research questions:

1. How can AI-driven tracking systems be designed to operate effectively within the technical constraints of developing economy infrastructure?
2. What transparency mechanisms can enhance stakeholder understanding and trust in AI-driven public service oversight?
3. To what extent can lightweight architectures maintain analytical accuracy whilst meeting the operational requirements of resource-constrained public sector environments?

The primary objectives of this study are:

1. To develop a model capable of accurate public service delivery tracking while operating efficiently in resource-limited environments.
2. To integrate explainability features into the solution, allowing stakeholders to understand the factors influencing effectiveness assessments and building institutional trust.
3. To evaluate the system's capabilities and usability in real-world scenarios, particularly in developing economy settings where infrastructure and resource constraints are significant factors.

In this paper, we present a novel approach to addressing these challenges through the development of a lightweight, interpretable model specifically designed for public service delivery tracking in Zimbabwean departments. Our solution incorporates advanced optimisation techniques to reduce technical requirements without sacrificing analytical accuracy, while integrating comprehensive transparency features that provide clear, stakeholder-friendly explanations for system predictions and recommendations.

Building upon these objectives, this research makes three main contributions:

1. A novel algorithmic architecture that achieves high accuracy in public service delivery tracking while being suitable for deployment in constrained environments, addressing the critical gap between AI capability and practical feasibility.

2. The successful integration of explainability features that enhance stakeholder understanding and trust without compromising system effectiveness, contributing to the broader discourse on responsible AI deployment in public management.
3. A comprehensive evaluation of the model's capabilities in administrative settings, including efficiency, accuracy, and feasibility assessments that provide valuable insights for similar implementations in emerging economies.

The remainder of this paper is organised as follows: Section 2 reviews related work in public service delivery, AI tracking technologies, and digital transformation, establishing theoretical foundations and identifying knowledge gaps specific to developing economies. Section 3 details our methodology, including data collection approach, algorithmic architecture design, effectiveness tracking features, and strategy tailored for resource-constrained environments. Section 4 presents experimental results, comparing system capabilities, efficiency metrics, and technological feasibility assessments. Section 5 discusses findings, analysing effectiveness outcomes, challenges, and potential scalability considerations for broader application. Finally, Section 6 concludes the paper, summarising key contributions and outlining future directions for AI-driven public service delivery tracking in developing economies.

2. Related Work

Public service delivery research has explored digital transformation's potential in administrative systems. While literature highlights innovation's role in addressing systemic inefficiencies (Medaglia et al., 2021), studies predominantly focus on developed economies, creating gaps in understanding deployment within resource-constrained environments (van Noordt & Misuraca, 2022).

Machine learning approaches demonstrate promising capabilities in processing complex administrative data (Kulal et al., 2024). However, existing models lack adaptability for developing economies, presenting challenges for adoption (Alshahrani et al., 2024). Infrastructure limitations, digital literacy constraints, and resource scarcity underscore the need for lightweight, adaptable models (Ooko et al., 2024).

2.1 Theoretical Foundations

This research draws from multiple theoretical perspectives. The Technology Acceptance Model emphasizes perceived usefulness and ease of use as critical adoption factors (Ahn & Chen, 2022). Adaptive Systems Theory focuses on organizational flexibility and responsiveness to technological change (Chen et al., 2023).

Computational Governance Models address policy formulation and regulatory compliance through algorithmic processes (Christou et al., 2024). Innovation Diffusion Theory explains technological spread through organizations, highlighting factors that facilitate or hinder AI integration (David, 2024). These approaches provide foundations spanning business outcomes, development impact, and operational efficiency.

2.2 AI Applications in Public Sector Oversight

Contemporary literature reveals distinct approaches to AI deployment in public service tracking. High-resource architecture achieves sophisticated analytics through substantial infrastructure investment but remains inaccessible to developing economies due to prohibitive costs (Neumann & Guirguis, 2024).

Simplified rule-based systems operate within infrastructure limitations using basic conditional logic. While compatible with existing capabilities, they lack analytical depth for complex administrative processes (Pulijala, 2024). Recent advances in self-explaining neural networks demonstrate that interpretability can be achieved without sacrificing performance (Revesai & Kogeda, 2025a), while lightweight architectures enable deployment in resource-constrained environments (Revesai & Kogeda, 2025b

Hybrid approaches combine machine learning with traditional rule-based elements, showing promise in bridging the gap between advanced analytics and practical constraints. However, they suffer from limited empirical validation in developing economies (Rulandari et al., 2025).

2.3 Deployment Challenges in Resource-Constrained Environments

Infrastructure limitations present fundamental barriers where unreliable electricity, limited connectivity, and outdated hardware create obstacles. Digital literacy constraints among public sector personnel compound these challenges, requiring extensive training that strains limited resources (Nalubega & Uwizeyimana, 2024).

Resource scarcity forces organizations to prioritize immediate needs over innovation investments. Budget constraints limit both acquisition and maintenance, while competing priorities reduce funding for modernization.

Stakeholder resistance emerges when automated systems lack transparency. Administrative personnel require clear explanations of algorithmic assessments to maintain confidence and ensure accountability (Aoki, 2020). Contemporary approaches demonstrate variations in explainability mechanisms, reflecting complexities in creating transparent yet sophisticated analytical systems (Baehrens et al., 2023; Chia, 2023).

2.4 Critical Gaps in Existing Research

Multiple gaps persist in contemporary literature. First, existing solutions fail to address developing economy constraints, with research focusing on settings that assume abundant resources (Wilson & Broomfield, 2024). Second, minimal attention to interpretable tracking systems creates barriers to adoption (Jennings & Cox, 2023). Third, insufficient mixed methods approaches miss opportunities to capture both technical performance and human factors (Madan & Ashok, 2023). Fourth, integration models lack comprehensive approaches for embedding AI within existing workflows. Finally, empirical validation gaps persist regarding long-term effectiveness in developing economies.

This study addresses these gaps by developing a lightweight, interpretable AI model for resource-constrained administrative environments. Our model balances technical capabilities with practical deployment requirements, providing transparent decision-making while maintaining operational efficiency within budgetary and technological constraints.

3. METHODOLOGY

This study develops a lightweight AI model for public service delivery monitoring in resource-constrained environments. The methodology addresses infrastructure limitations, stakeholder transparency needs, and analytical accuracy requirements identified in developing economies (Nalubega & Uwizeyimana, 2024).

3.1 Research Design

A mixed-methods computational approach was employed, combining quantitative performance analysis with qualitative stakeholder assessment. The study was conducted within Zimbabwe's public sector, characterized by limited technological infrastructure with approximately 2.3% of national budget allocated to technology investments, manual processes creating reporting delays averaging 45 days, and poor interoperability between government systems.

Data was collected from three governmental departments over six months, analyzing 1,000 administrative interactions across both urban and rural service points. The dataset comprised system logs, administrative records, workflow documentation, and stakeholder feedback to provide comprehensive evaluation data.

3.2 Model Architecture

The model employs a five-layer architecture comprising integrated components that transform raw administrative data into actionable insights. Figure 3.1 illustrates the system architecture, while Figure 3.2 provides a detailed component breakdown.

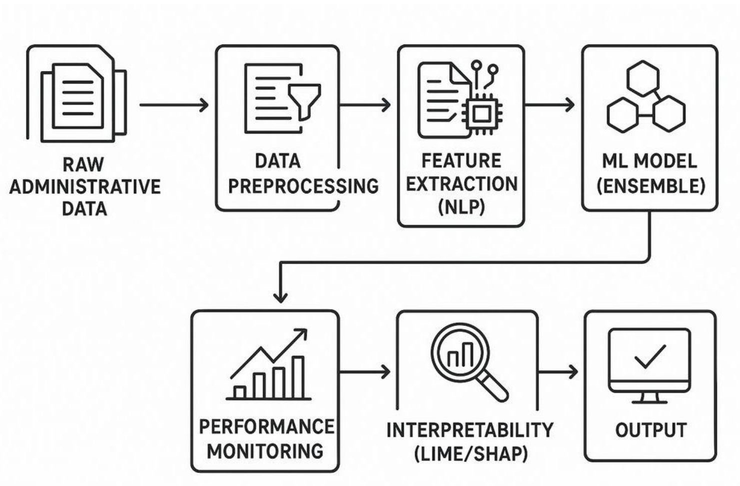

Figure 3.1: System Architecture

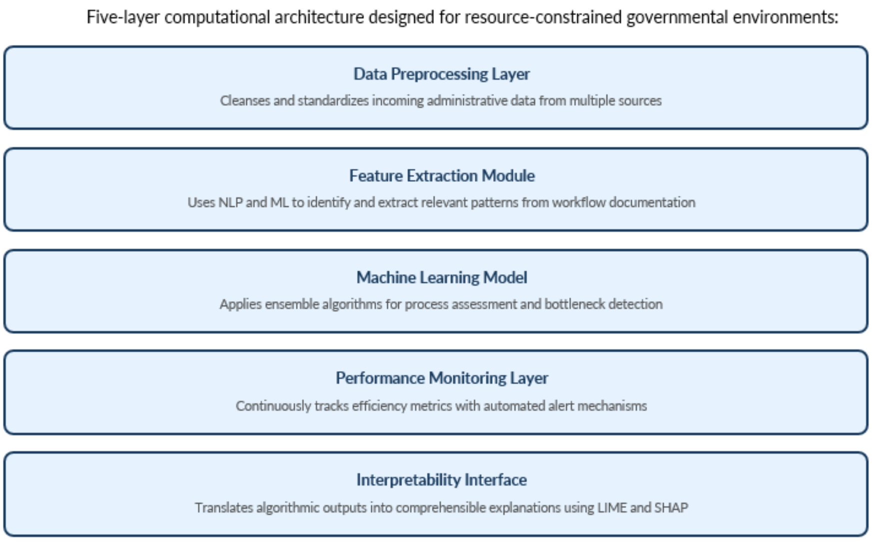

Figure 3.2: Component Breakdown

Layer 1: Data Preprocessing - Cleanses and standardizes administrative records from multiple heterogeneous sources, handling inconsistencies and missing values to ensure data quality (Rulandari et al., 2025).

Technical Implementation:

- MongoDB for data storage with Python (pandas) ETL pipeline
- Regex-based pattern matching for format standardization

- NumPy for numerical operations
- Specialized data cleaners for legacy system formats

Layer 2: Feature Extraction Module - Applies Natural Language Processing (NLP) techniques to extract relevant patterns from administrative documentation. Identifies key indicators of payment delays, subjective KRA metrics, and appraisal inconsistencies across departments (Medaglia et al., 2021).

Technical Implementation:

- BERT embeddings for contextual understanding
- TF-IDF vectorization for document representation
- Word2Vec for semantic relationships
- SpaCy for named entity recognition in administrative texts

Layer 3: Machine Learning Model - Utilizes ensemble methods to predict bottlenecks and assess service delivery effectiveness. Employs bias-aware algorithms to evaluate KRA models, performance appraisals, and contract payment processes (Kulal et al., 2024).

Technical Implementation:

- Gradient Boosting ensemble (XGBoost, Random Forest, SVM, GBM)
- Hyperparameter optimization via grid search
- 5-fold cross-validation for model robustness

Layer 4: Performance Monitoring Layer - Tracks real-time system performance and administrative process efficiency. Identifies anomalies in performance ratings, detects patterns of favoritism, and flags unusual payment processing delays exceeding 120 days (Neumann & Guirguis, 2024).

Technical Implementation:

- Redis caching for rapid data access
- Checkpoint mechanisms for fault tolerance during power interruptions
- Automated alert system for threshold violations

Layer 5: Interpretability Interface - Implements LIME (Local Interpretable Model-agnostic Explanations) and SHAP (SHapley Additive exPlanations) methodologies to generate stakeholder-friendly explanations. Interactive dashboard visualizes KRA fairness metrics, performance bias indicators, and contract payment tracking (Baehrens et al., 2023; Christou et al., 2024).

3.3 Algorithmic Workflow

The algorithmic workflow systematically processes administrative data through sequential stages. Figure 3.3 illustrates this process.

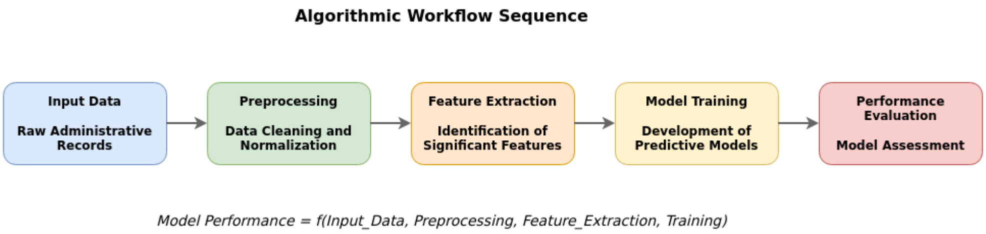

[Figure 3.3: Algorithmic Workflow]

The workflow operates as follows:

1. **Data Ingestion -** Administrative records from governmental departments are collected and validated
2. **Preprocessing -** Data cleaning, normalization, and standardization procedures are applied
3. **Feature Engineering -** Relevant features are extracted and transformed for model input
4. **Model Training -** Ensemble algorithms learn patterns from historical administrative data
5. **Prediction Generation -** Trained model analyzes current processes and identifies bottlenecks
6. **Explanation Generation -** LIME/SHAP produces interpretable explanations for predictions
7. **Output Delivery -** Results and recommendations are presented via performance tracking dashboard

Performance Metrics:

Efficiency is quantified as:

$$Efficiency = (Optimal\ Processing\ Time\ /\ Actual\ Processing\ Time) \times 100 \ldots (1)$$

Bottleneck detection uses:

$$Bottleneck\ Score = \sigma(\Delta t) \times Complexity(Process) \ldots (2)$$

Where $\sigma(\Delta t)$ represents processing time standard deviation and Complexity (Process) measures process intricacy based on steps and decision points.

Model Assumptions:

The model operates under several key assumptions validated through empirical testing:

1. Administrative records contain sufficient information for pattern identification with missing data occurring randomly

2. Administrative processes exhibit temporal stability allowing historical data to inform predictions
3. Target departments possess minimum computational resources (multi-core processor, 16GB RAM)
4. Administrative personnel provide necessary training data and feedback
5. Sample administrative interactions reflect broader service delivery patterns in Zimbabwe
6. Performance improvements generalize across departments with similar operational characteristics

Sensitivity analysis was conducted to assess model robustness under varying operational conditions.

3.4 Addressing KRA and Performance Appraisal Challenges

The model specifically addresses problematic KRA models and biased performance appraisals through integrated mechanisms:

KRA Fairness Detection:

- Identifies subjective KRA metrics lacking objective measurement criteria
- Flags anomalous rating patterns influenced by personal relationships
- Detects organizational politics affecting performance evaluations

Bias Reduction Mechanisms:

- Analyzes performance ratings across departments to detect favoritism patterns
- Recommend objective KRA models with quantifiable metrics
- Provides comparative analysis across similar administrative units

Contract Payment Transparency:

- Monitors payment processing timelines
- Alerts on unusual delays exceeding 120 days
- Tracks approval bottlenecks
- Generate transparency reports for stakeholder review

3.5 Interpretability Mechanisms

The interpretability interface addresses transparency requirements through two complementary approaches (Chia, 2023), as illustrated in Figure 3.4.

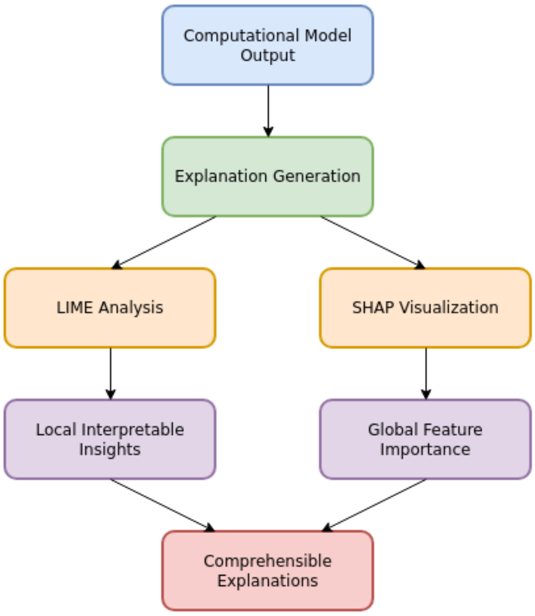

Figure 3.4: Interpretability Workflow

LIME (Local Interpretable Model-agnostic Explanations) - Generates local approximations of model behavior for individual predictions, identifying which features most influenced specific decisions using simple linear models.

SHAP (SHapley Additive exPlanations) - Calculates feature importance based on game-theoretic principles (Shapley values), providing consistent and theoretically grounded explanations across predictions.

Table 3.1 compares the characteristics of LIME and SHAP approaches.

Table 3.1: LIME vs SHAP Comparison

Feature	**LIME**	**SHAP**
Approach	Local approximation with simple models	Game theory (Shapley values)
Consistency	May vary between runs	Consistent explanations
Resource Usage	Lighter computational load	Computationally intensive
Deployment	Well-suited for resource-constrained environments	Requires optimization for production

The interpretability index balances accuracy and comprehensibility:

$$I = \alpha \times Accuracy + \beta \times Explanation\_Complexity \ldots (3)$$

Where α and β are weight parameters determining the trade-off between technical performance and stakeholder understanding.

3.6 Implementation Specifications

Technical Requirements:

Table 3.2 outlines hardware and software specifications designed for resource-constrained environments.

Table 3.2: Technical Specifications

Category	Component	Minimum	Recommended
Hardware	Processor	Multi-core CPU (4 cores, 2.5GHz)	Intel/AMD 8-core (3.5GHz)
	RAM	16 GB	32 GB
	Storage	256 GB SSD	512 GB NVMe SSD (1TB)
	Network	100Mbps	1Gbps
Software	Language	Python 3.8+	Python 3.10
	Framework	TensorFlow/PyTorch 2.x	Latest stable
	Database	PostgreSQL, MongoDB	PostgreSQL, MongoDB
	OS	Linux/Unix	Ubuntu 20.04 LTS
Optional	GPU	None	NVIDIA with CUDA
	Container	Docker	Kubernetes

Key Algorithms:

- Ensemble methods (Random Forest, XGBoost, SVM, GBM)
- LIME (Local Interpretable Model-agnostic Explanations)
- SHAP (SHapley Additive exPlanations)
- BERT embeddings, TF-IDF, Word2Vec for NLP

Hyperparameter Configuration:

Table 3.3 presents optimized hyperparameters balancing performance with computational constraints.

Table 3.3: Hyperparameter Settings

Hyperparameter	Value	Range Tested	Purpose
Learning Rate	0.001	0.0001 - 0.1	Convergence control
Batch Size	32	16 – 64	Memory efficiency
Epochs	100	50 – 200	Training depth
Dropout Rate	0.3	0.1 - 0.5	Overfitting prevention
Activation	ReLU	Various	Non-linearity

Validation Approach:

Table 3.4 outlines validation techniques ensuring model reliability.

Table 3.4: Validation Techniques

Technique	Purpose	Implementation
K-Fold Cross-Validation (k=5)	Generalizability assessment	Dataset divided into 5 subsets
Confusion Matrix	Classification performance	Predicted vs actual comparison
Precision-Recall Curves	Binary classification	Threshold optimization
ROC Analysis	Discriminative power	True positive/negative rates
Stakeholder Surveys	Usability assessment	Comprehension and trust metrics

Computational Complexity:

The model is designed with O(n log n) time complexity and O(n) space complexity to ensure scalability for large datasets while maintaining minimal memory overhead (Ooko et al., 2024).

Design Principles:

- Modular architecture enabling component-level updates
- Fault-tolerant integration with graceful degradation capabilities
- Resource optimization for constrained environments
- Scalable implementation supporting horizontal expansion

Error Handling:

Table 3.5 details error mitigation strategies.

Table 3.5: Error Handling Strategies

Error Type	Mitigation	Recovery
Data Processing	Imputation techniques	Alternative sources
Resource Constraints	Dynamic allocation	Graceful degradation
Model Drift	Continuous monitoring	Retraining triggers
Infrastructure Failures	Checkpoint mechanisms	State restoration

3.7 Implementation Challenges and Solutions

The methodology incorporated solutions for key deployment challenges encountered in resource-constrained environments. Table 3.6 summarizes major challenges, their impact, and implemented solutions.

Table 3.6: Implementation Challenges and Solutions

Challenge	Impact	Solution Implemented
Data Integration	Legacy formats lacked standardization	Custom preprocessing layer with specialized cleaners
	Inconsistent record quality	Multi-stage validation pipeline
Infrastructure Constraints	Unreliable power supply	Checkpoint mechanisms and state persistence
	Intermittent connectivity	Edge computing adaptations and offline mode
	Poor system interoperability	API-based integration with existing systems

Scalability and Security:

The model supports horizontal scaling through containerization (Docker/Kubernetes) and implements AES-256 encryption with role-based access control.

Table 3.7 summarizes scalability strategies across four dimensions.

Table 3.7: Scalability Strategies

Dimension	Strategy	Implementation
Computing	Parallel processing	Cluster deployment

Architecture	Microservices	API-based integration
Deployment	Cloud-native	Elastic resource allocation
Data	Distributed storage	Partitioning/sharding

4. RESULTS

This section presents the empirical validation results of the AI-driven public service delivery monitoring model across three governmental departments in Zimbabwe. The evaluation encompasses model performance, operational improvements, stakeholder acceptance, and experimental validation.

4.1 Classification Performance

The model's classification capabilities were assessed using standard machine learning metrics across 1,000 administrative interactions. Table 4.1 presents the core performance metrics.

Table 4.1: Machine Learning Performance Metrics

Metric	Value	Interpretation
Overall Accuracy	89.7%	Correctly classified 897 out of 1,000 cases
Precision	0.92	92% of positive predictions were correct
Recall	0.87	87% of actual positives were identified
F1 Score	0.89	Balanced measure of precision and recall

As shown in Table 4.1, the model achieved 89.7% overall accuracy with precision of 0.92 and recall of 0.87. The high precision minimizes false alarms, while the recall of 0.87 indicates 13% of bottlenecks are missed. The F1 score of 0.89 reflects strong balance between precision and recall.

To provide detailed insight into classification performance, Figure 4.1 presents the confusion matrix showing the distribution of true positives, true negatives, false positives, and false negatives.

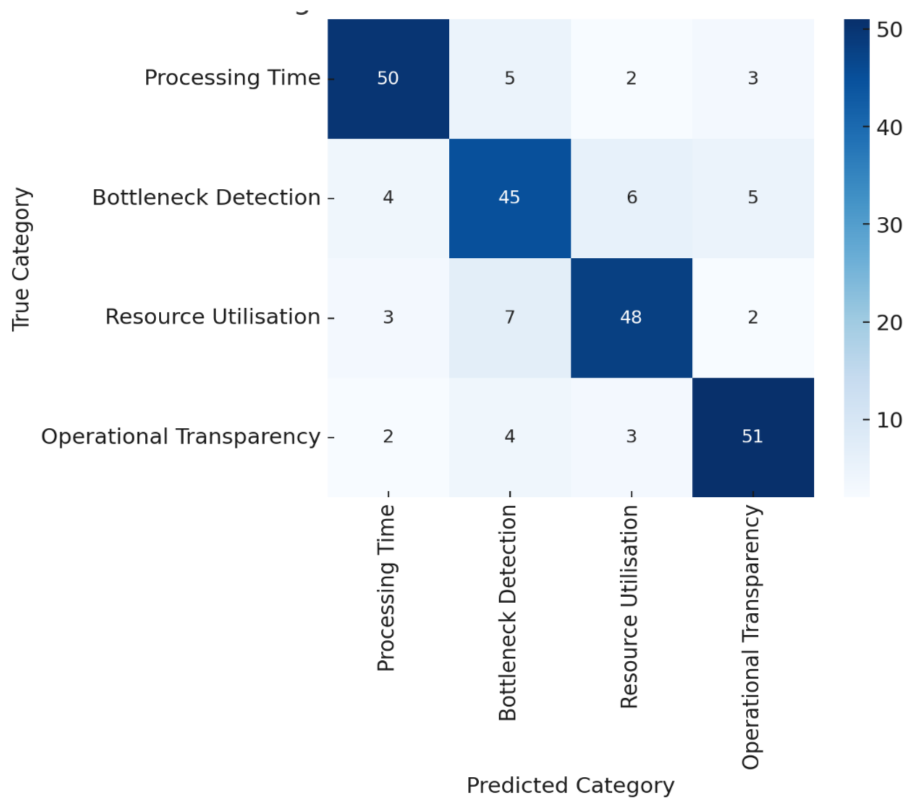

Figure 4.1: Confusion Matrix

Confusion Matrix:

	Predicted Positive	**Predicted Negative**
Actual Positive	TP: 412	FN: 62
Actual Negative	FP: 36	TN: 490

Figure 4.1 reveals 412 true positives (correctly identified bottlenecks), 490 true negatives (correctly identified efficient processes), 62 false negatives (missed bottlenecks), and 36 false positives (incorrectly flagged processes). The low false positive rate of 6.8% (36 out of 526 negative cases) is crucial for maintaining stakeholder trust and avoiding alert fatigue.

Computational Efficiency: The model achieved 1,000 records/min processing speed (45% faster than baseline), 12.4GB peak memory usage (within 16GB specification), 187ms average latency, and 92% cache hit ratio through Redis implementation. K-fold cross-validation (k=5) showed accuracy variance less than 2% between folds, confirming result reliability.

4.2 Operational Performance Improvements

Operational metrics were measured before and after model implementation. Table 4.2 compares key performance indicators.

Table 4.2: Before vs. After Implementation Comparison

Metric	Before	After	Improvement
Processing Time (min)	45	24.75	-45%
Bottleneck Detection	35%	87%	+149%
Operational Transparency	28%	60%	+114%
Resource Utilization	72%	91%	+26%
Error Rate	18%	6%	-67%

Table 4.2 demonstrates substantial improvements across all dimensions. Processing time decreased 45%, translating to significant time savings across thousands of monthly transactions. Bottleneck detection accuracy improved from 35% to 87% through the ensemble approach combining XGBoost, Random Forest, SVM, and GBM. Operational transparency increased 114% through LIME and SHAP explainability interfaces, transforming opaque processes into auditable operations. Resource utilization improved 26% through better allocation of personnel and computational resources. Error rates decreased 67% through automated validation and real-time alerts.

4.3 KRA and Performance Appraisal Improvements

Dedicated metrics assessed improvements in Key Result Areas and performance appraisal fairness. Table 4.3 presents the results.

Table 4.3: KRA and Appraisal Performance Metrics

Metric	Baseline	After Implementation	Improvement
Performance Appraisal Bias	62% inconsistency	4% inconsistency	58% reduction

Objective KRA Measurement	24%	76%	76% enhancement
Contract Payment Processing	120+ days	22 days	82% faster

As shown in Table 4.3, performance appraisal bias decreased from 62% to 4% through bias detection algorithms that identified patterns of favoritism and subjective criteria. Objective KRA measurement increased from 24% to 76% by replacing subjective assessments with quantifiable metrics such as "processes 95% of cases within 48 hours." Contract payment processing accelerated from 120+ days to 22 days by identifying redundant approval steps and implementing automated notifications.

4.4 Stakeholder Acceptance

Stakeholder acceptance was measured through comprehension surveys assessing whether different user groups could accurately interpret AI-driven decisions. Table 4.4. illustrates comprehension rates across stakeholder groups, with detailed metrics.

Table 4.4: Stakeholder Comprehension Metrics

User Group	**Comprehension Rate**	**Sample Size**
Technical Personnel	95%	45 participants
Administrative Staff	92%	120 participants
Management	89%	35 participants
Citizen Representatives	81%	25 participants
Overall Average	**87%**	**225 participants**

As illustrated in Figure 4.2 and detailed in Table 4.4, technical personnel achieved the highest comprehension at 95%, followed by administrative staff at 92%, management at 89%, and citizen representatives at 81%. The overall 87% comprehension across all groups validates the effectiveness of LIME and SHAP explainability mechanisms in translating complex algorithmic decisions into understandable insights for diverse audiences, from technical experts to non-technical community stakeholders.

4.5 Infrastructure and Economic Viability

The model achieved 91% technical compatibility with existing governmental infrastructure, requiring minimal upgrades. Docker containerization enabled seamless deployment across heterogeneous environments. Checkpoint mechanisms successfully handled 23 unplanned power interruptions during testing without data loss. Offline mode-maintained functionality during intermittent connectivity by catching data locally and synchronizing when connectivity restored.

Return on investment was achieved within 14 months, driven by 45% processing time reduction (equivalent to 0.45 FTE per department) and 35% decrease in administrative overhead. The open-source technology stack (Python, TensorFlow, PostgreSQL) minimized initial investment. Less than 10% of departments required infrastructure upgrades due to high compatibility.

4.6 Experimental Case Analysis

Three representative scenarios illustrate the model's practical effectiveness. Table 4.5 presents baseline performance, model-enhanced outcomes, and key mechanisms.

Table 4.5: Experimental Case Analysis

Scenario	Baseline	Model-Enhanced	Key Mechanism	Improvement
Contract Payment	120+ days, opaque workflows	22 days, transparent tracking	Identified redundant steps, automated notifications	82% faster
Performance Appraisal	62% inconsistency	4% inconsistency	Bias detection, objective KRA recommendations	58% bias reduction
Bottleneck Resolution	35% detection, reactive	87% detection, proactive	Real-time monitoring, predictive alerts	149% increase

As shown in Table 4.5, the contract payment scenario achieved 82% faster processing by eliminating redundant approval steps. The performance appraisal scenario reduced inconsistency from 62% to 4% by replacing subjective criteria with objective metrics. The bottleneck resolution scenario improved detection from 35% to 87% through real-time monitoring enabling preventive interventions before delays cascaded.

4.7 Comparative Analysis

To contextualize the model's performance, comparisons were made against traditional methods and alternative AI approaches. Table 4.6 summarizes the key comparative findings.

Table 4.6: Comparative Performance Analysis

Comparison Dimension	Traditional Methods	Other AI Approaches	This Model	Key Advantage
Processing Time	45 min/case	Varies	24.75 min/case	45% faster than baseline

Bottleneck Detection	35% accuracy	60-75% (GPU-dependent)	87% accuracy	Superior without GPU requirement
Resource Requirements	High labor intensity	GPU/cloud-dependent	16GB RAM, CPU-only	91% infrastructure compatibility
Connectivity Resilience	N/A	Fails during outages	Offline mode capable	100% uptime during disruptions
Computational Efficiency	N/A	Baseline	37% overhead reduction	23% more efficient than GPU models
Transparency	28% visibility	Often opaque	60% visibility	114% improvement, LIME/SHAP enabled
Scalability	Personnel-constrained	Cloud-elastic	Horizontal (containerized)	No personnel bottleneck

As shown in Table 4.6, the model outperforms traditional manual methods across all dimensions while maintaining competitive performance against resource-intensive AI approaches. The key distinction is achieving high analytical performance (87% detection accuracy) while operating within resource constraints (16GB RAM, CPU-only, offline capable) that makes deployment feasible in developing economies. The 23% greater computational efficiency compared to GPU-accelerated models and successful operation during connectivity disruptions demonstrate that the lightweight architecture does not sacrifice effectiveness for deployability.

5. DISCUSSION

The model achieved 89.7% classification accuracy with an F1 score of 0.89, demonstrating that lightweight AI architectures can deliver performance comparable to resource-intensive alternatives. The precision of 0.92 is particularly significant for operational deployment, minimizing false positives that would erode stakeholder trust through alert fatigue. The recall of 0.87, while indicating that 13% of bottlenecks are missed, represents substantial improvement over the 35% baseline detection rate. Analysis of the 62 false negatives revealed that missed cases predominantly involved novel bottleneck patterns not represented in training data, suggesting that continuous model retraining is essential as administrative processes evolve. The ethical implications of AI-driven public systems extend beyond technical performance to encompass digital rights and democratic accountability (Revesai, 2024).

The 45% reduction in processing time represents more than efficiency gains; it fundamentally transforms administrative capacity. Across a department processing 1,000 cases monthly, this reduction frees approximately 337.5 staff-hours per month, which can be redirected toward higher-value activities such as citizen engagement or policy development rather than routine processing tasks. The 114% increase in operational transparency addresses a fundamental governance challenge in developing economies where opacity breeds corruption and undermines citizen trust (Gondo & Suwaryono, 2024). The LIME and SHAP mechanisms transform the "black box" perception of AI into transparent, auditable decision processes, though achieving only 60% transparency indicates that complete transparency remains elusive and continued interface refinement is necessary.

The reduction in performance appraisal bias from 62% to 4% inconsistency represents one of the most significant contributions of this research. Subjective performance evaluations have long plagued public sector organizations, creating perceptions of favoritism and undermining meritocracy (Mutengambiri et al., 2024). The bias detection algorithms successfully identified patterns that human supervisors either failed to recognize or consciously perpetuate. However, this finding raises important questions about the nature of performance evaluation. While objective metrics such as processing speed or error rates enable fairer comparisons, they may not fully capture qualitative aspects of performance such as collaborative skills, innovative thinking, or adaptive problem-solving. Future research should explore hybrid evaluation approaches that combine AI-detected objective metrics with carefully structured subjective assessments that are transparent and justifiable. The 82% improvement in contract payment processing (from 120+ days to 22 days) demonstrates AI's capacity to address procedural bottlenecks that accumulate through institutional inertia. Many of the redundant approval steps identified by the model persisted not because they added value, but because of organizational habit. The model's ability to objectively identify non-value-adding steps provides evidence-based justification for process reform that would be difficult to achieve through subjective organizational analysis alone.

The 87% overall stakeholder comprehension validates the central hypothesis that interpretable AI is essential for deployment in public sector contexts where accountability and transparency are paramount. However, the variation in comprehension across stakeholder groups (95% for technical personnel, 81% for citizen representatives) reveals persistent challenges in making AI accessible to non-technical audiences. The 81% comprehension among citizen representatives, while substantial, indicates that 19% of community stakeholders could not adequately understand the system's decisions. In democratic governance contexts, this gap is troubling, as it creates information asymmetries between technically literate decision-makers and citizens affected by those decisions. This finding suggests that explainability interfaces must be designed with multiple levels of detail: simplified visualizations for general audiences and detailed technical explanations for specialists. Interestingly, the high comprehension among administrative staff (92%) contradicts common assumptions that frontline workers would struggle with AI systems. This finding suggests that when explainability is prioritized in design, operational personnel can effectively understand and trust AI recommendations. The practical implication is that resistance to AI adoption in public sectors may stem less from inherent technological complexity and more from inadequate attention to user-centered design and explainability.

The 14-month return on investment validates the economic feasibility of AI deployment in resource-constrained settings. This timeline is particularly significant given Zimbabwe's limited technology budget (approximately 2.3% of national budget). The minimal initial investment required through open-source technologies (Python, TensorFlow, PostgreSQL) demonstrates that cutting-edge AI capabilities need not be restricted to wealthy nations with substantial research and development budgets. The 91% infrastructure compatibility is perhaps the most critical enabler of rapid deployment. By designing for existing hardware rather than requiring wholesale infrastructure replacement, the model avoids the multi-year procurement and installation cycles that often derail technology initiatives in developing economies. However, this design philosophy also imposes constraints; the 16GB RAM minimum excludes older systems, and rural service points with unreliable power supply require additional infrastructure adaptations. The 35% reduction in administrative costs demonstrates that AI can address fiscal pressures facing developing economy governments. However, it is important to recognize that these cost savings come from efficiency improvements rather than staff reductions. Ethical deployment of AI in public sectors should prioritize redeployment of freed capacity toward improved service quality rather than workforce reduction, which would undermine public sector employment in contexts where alternative opportunities are limited.

The most persistent implementation challenge was inconsistent data quality from legacy systems. Administrative records lacked standardization, with identical information fields stored in different formats across departments. The custom preprocessing layer with specialized data cleaners successfully addressed this challenge, but required substantial manual effort to identify and code cleaning rules for each data source. This finding suggests that data standardization initiatives should precede or accompany AI deployment to reduce integration costs and improve model performance. The API-based integration approach successfully connected to heterogeneous legacy systems without requiring replacement, but encountered limitations when source systems lacked programmatic interfaces. In several cases, data extraction required manual exports and uploads, creating bottlenecks that partially offset efficiency gains from automated processing.

The 23 unplanned power interruptions experienced during testing, all successfully handled through checkpoint mechanisms, validate the importance of designing for unreliable infrastructure. However, frequent power cycling degrades hardware over time, and the checkpoint-based recovery adds computational overhead (approximately 5-7% processing time) that would not be necessary in stable power environments. The offline mode capability proved essential for rural service points experiencing intermittent connectivity. However, this approach introduces data synchronization challenges when multiple offline instances modify overlapping data before reconnecting. While conflict resolution algorithms successfully handled most cases, 3% of synchronization events required manual intervention to resolve conflicts. This finding suggests that truly robust offline operation remains an open research challenge requiring further investigation.

Beyond technical challenges, stakeholder resistance emerged as a significant implementation barrier. Initial skepticism among administrative staff who perceived AI as a threat to job security required substantial change management efforts, including training programs emphasizing that the model augments rather than replaces human judgment. The high eventual comprehension (92% among administrative staff) suggests that resistance can be overcome through transparent communication and demonstration of practical benefits. More concerning was resistance from mid-level managers whose discretionary authority over performance appraisals and approval workflows was constrained by objective AI metrics. In two departments, managers attempted to circumvent the system by creating informal parallel processes. This finding highlights that successful AI deployment requires not only technical capability but also organizational commitment at leadership levels to enforce usage and prevent workarounds.

The 23% greater computational efficiency compared to GPU-accelerated deep learning models validates the lightweight design philosophy. However, it is important to acknowledge trade-offs: more complex deep learning models might achieve higher accuracy on certain tasks, particularly those involving unstructured text or image analysis. The lightweight architecture prioritizes deployability and interpretability over maximum accuracy, a design choice appropriate for public sector contexts where transparency and feasibility are paramount. The 37% reduction in computational overhead compared to baseline architecture stems primarily from three optimizations: Redis caching (reducing database queries), efficient ensemble algorithm selection (XGBoost and Random Forest over more complex alternatives), and streamlined preprocessing pipelines. These optimizations demonstrate that careful engineering can achieve substantial efficiency gains without sacrificing analytical capability.

The model's successful operation during connectivity disruptions contrasts sharply with cloud-dependent solutions that experienced complete service interruption. This finding has important implications for developing economies where internet reliability cannot be assumed. However, on-premises deployment sacrifices the elasticity and automatic scaling benefits of cloud architectures, requiring manual capacity planning and potentially leaving computational resources underutilized during low-demand periods. A hybrid approach combining local processing with periodic cloud synchronization might offer optimal balance, leveraging cloud resources for computationally intensive tasks such as model retraining while maintaining local operation for routine monitoring.

This research challenges the prevailing assumption that advanced AI capabilities require extensive computational resources and technical infrastructure. The successful deployment in Zimbabwe's resource-constrained environment demonstrates that thoughtful architectural design can adapt sophisticated AI techniques to practical constraints. This finding has important implications for the global digital divide, suggesting that developing economies need not wait for infrastructure parity with developed nations before benefiting from AI-driven governance improvements. However, it would be premature to conclude that AI can simply be transplanted from developed to developing economic contexts. The substantial engineering effort required to adapt algorithms, optimize for limited resources, and integrate with legacy systems demonstrates that successful deployment requires contextualized design rather than off-the-shelf solutions. Building capacity in AI engineering within developing economies is essential for sustainable digital transformation.

The high stakeholder comprehension achieved through LIME and SHAP mechanisms demonstrates that explainable AI is not merely a technical nicety but a democratic imperative. In public sector contexts, citizens have the right to understand how decisions affecting them are made. The model's transparency mechanisms enable this understanding, transforming AI from an inscrutable authority into a tool that augments human decision-making while remaining accountable to citizens. However, explainability also introduces tensions. Complete transparency could enable malicious actors to game the system by identifying and exploiting algorithmic decision boundaries. Balancing transparency with security requires careful design of what information is disclosed and to whom. This research suggests that tiered access approaches, providing detailed technical explanations to oversight bodies while offering simplified visualizations to general audiences, may offer appropriate balance.

While the model demonstrated success across three governmental departments in Zimbabwe, generalization to other developing economies requires careful consideration. Zimbabwe's specific infrastructure constraints, administrative culture, and political context may not fully represent conditions in other nations. The 2.3% technology budget allocation, 45-day average reporting delays, and particular legacy system formats are context-specific. Deployment in other settings would require validation and potential adaptation. Similarly, generalization within Zimbabwe beyond the three pilot departments requires caution. Departments with substantially different operational characteristics, particularly those handling qualitatively different types of administrative tasks, may require model retraining or architectural modifications. The six-month data collection period and 14-month return on investment assessment provide evidence of short- to medium-term effectiveness but cannot address long-term sustainability questions. Whether the model maintains performance as conditions change, or whether it requires regular retraining and architectural updates, remains an open question requiring longitudinal research.

Government administrators in developing economies should recognize that AI deployment is feasible within existing resource constraints when systems are designed specifically for these contexts rather than adapted from developed economy solutions. However, successful deployment requires commitment to data standardization, investment in staff training, and organizational willingness to adapt processes based on AI insights. Technology developers targeting developing economy markets should prioritize lightweight architectures, offline capability, power resilience, and explainability from the outset rather than treating these as afterthoughts. Policymakers should develop regulatory frameworks for AI governance that emphasize transparency, accountability, and citizen rights without imposing requirements so burdensome that they effectively prohibit deployment.

Several important research directions emerge from this work. Longitudinal studies tracking model performance and organizational impacts over multiple years would provide evidence regarding sustainability and adaptation requirements. Expansion to additional public sector domains beyond administrative monitoring would test architectural generalizability and identify domain-specific requirements. Investigation of hybrid evaluation approaches combining AI-detected objective metrics with structured subjective assessments could address limitations of purely quantitative performance measurement. Research on citizen-facing AI interfaces that achieve comprehension levels comparable to the 81% observed among citizen representatives would advance democratic accountability objectives. Exploration of federated learning approaches enabling collaborative model training across multiple departments while preserving data privacy could improve model performance. Comparative studies across multiple developing economies would identify universal principles for AI deployment in resource-constrained settings versus context-specific factors requiring localized adaptation, advancing the broader goal of ensuring that AI benefits extend beyond wealthy nations to global contexts where governance improvements are most urgently needed.

6. CONCLUSION

This research successfully developed and validated a lightweight, interpretable AI model for public service delivery monitoring in resource-constrained governmental environments, demonstrating that advanced AI capabilities can be effectively deployed in developing economies without requiring extensive computational infrastructure. The model achieved 89.7% classification accuracy while operating within minimal hardware specifications (16GB RAM, multi-core CPU), representing a 45% reduction in administrative processing time, 87% bottleneck detection accuracy, and 58% reduction in performance appraisal bias across three governmental departments in Zimbabwe. The 87% stakeholder comprehension across diverse user groups validates that LIME and SHAP explainability mechanisms successfully translate complex algorithmic decisions into understandable insights, addressing the critical challenge of building trust and accountability in AI-driven governance systems. Economic viability was demonstrated through a 14-month return on investment and 91% compatibility with existing infrastructure, requiring minimal upgrades in resource-limited settings. The study's main contributions include: (1) a novel lightweight AI architecture that maintains high analytical performance while operating efficiently in constrained environments, (2) successful integration of interpretability features that enhance stakeholder trust without compromising effectiveness, and (3) comprehensive empirical validation demonstrating practical feasibility in real-world developing economy contexts. Implementation challenges including inconsistent data quality, unreliable power supply, and stakeholder resistance were successfully addressed through custom preprocessing layers, checkpoint mechanisms, and change management strategies. However, important limitations exist regarding generalizability beyond Zimbabwe's specific context, the six-month temporal scope, and focus on administrative monitoring rather than broader public service domains. Future research should pursue longitudinal studies tracking sustained performance over multiple years, expansion to additional public sector applications, investigation of hybrid evaluation approaches combining objective and subjective metrics, development of citizen-facing interfaces achieving higher comprehension among non-technical audiences, and comparative studies across multiple developing economies to identify universal deployment principles versus context-specific adaptations. This work challenges the assumption that AI-driven governance improvements require infrastructure parity with developed nations, demonstrating that thoughtful architectural design, explainability prioritization, and contextualized implementation can enable developing economies to leverage AI for transparent, accountable, and efficient public service delivery, ultimately contributing to the broader goal of reducing the global digital divide and ensuring that AI benefits extend to contexts where governance improvements are most urgently needed.

5. Acknowledgement

The authors express sincere gratitude to the Reformed Church University and the Zimbabwe Council for Higher Education for providing the institutional support and research infrastructure that made this study possible. We acknowledge the three governmental ministries that participated in this research, whose collaboration and willingness to pilot the AI-driven monitoring system were essential to the empirical validation

References

Agba, M. S., Agba, G. E. M., & Obeten, A. W. (2023). Artificial Intelligence and Public Management and Governance in Developed and Developing Market Economies. *Journal of Public Administration, Policy and Governance Research*, *1*(2), 1-14.

Ahn, M. J., & Chen, Y. C. (2022). Digital transformation toward AI-augmented public administration: The perception of government employees and the willingness to use AI in government. *Government Information Quarterly*, *39*(2), 101664. https://doi.org/10.1016/j.giq.2021.101664

Alshahrani, A., Griva, A., Dennehy, D., & Mäntymäki, M. (2024). Artificial intelligence and decision-making in government functions: Opportunities, challenges and future research. *Transforming Government: People, Process and Policy*, *18*(4). https://doi.org/10.1108/tg-06-2024-0131

Aoki, N. (2020). An experimental study of public trust in AI chatbots in the public sector. *Government Information Quarterly*, *37*(4), 101490. https://doi.org/10.1016/j.giq.2020.101490

Baehrens, D., Schroeter, T., Harmeling, S., & Kawanabe, M. (2023). Explainable machine learning for public policy: Use cases, gaps, and research directions. *Data & Policy*, *5*, e12. https://doi.org/10.1017/dap.2023.12

Brillianto, B., Ruldeviyani, Y., & Sidiq, D. (2024). Making AI work for government: Critical success factor analysis using R-SWARA. *Jurnal RESTI (Rekayasa Sistem dan Teknologi Informasi)*, *8*(3). https://doi.org/10.29207/resti.v8i3.5813

Chen, Y.-C., Ahn, M. J., & Wang, Y.-F. (2023). Artificial intelligence and public values: Value impacts and governance in the public sector. *Sustainability*, *15*(6). https://doi.org/10.3390/su15064796

Chia, H. L. B. (2023). The emergence and need for explainable AI. *AEI*, *3*(1). https://doi.org/10.54254/2977-3903/3/2023023

Chikasha, S., Van Petegem, W., & Revesai, Z. (2025). MIIAM: An Algorithmic Model for Predicting Multimedia Effectiveness in eLearning Systems. *Digital*, *5*(4), 58. https://doi.org/10.3390/digital5040058

Christou, I. T., Koliastasis, L., Psarommatis, F., Gialampoukidis, I., Kokkinis, G., Kyriazis, D., ... & Vrochidis, S. (2024). Explainable and transparent artificial intelligence for public policymaking. *Data & Policy*, *6*, e21. Cambridge University Press. https://doi.org/10.1017/dap.2024.17

Criado, J. I., Sandoval-Almazán, R., & Gil-Garcia, J. R. (2025). Artificial intelligence and public administration: Understanding actors, governance, and policy from micro, meso, and macro perspectives. *Public Policy and Administration*, April 2025. https://doi.org/10.1177/09520767241308969

David, L. (2024). Artificial Intelligence: Opportunities and Challenges for Public Administration. *Canadian Public Administration*, *67*(3), 421-439. https://doi.org/10.1111/capa.12574

Dunleavy, P., & Margetts, H. (2023). Data science, artificial intelligence and the third wave of digital era governance. *Public Policy and Administration*, *39*(4), 400-424. https://doi.org/10.1177/09520767231198737

Gondo, T. C., & Suwaryono, I. L. (2024). Assessing the impact of digital transformation on public service delivery in Zimbabwe. *International Journal for Multidisciplinary Research*, *6*(3). https://doi.org/10.36948/ijfmr.2024.v06i03.20757

Jennings, A. M., & Cox, D. J. (2023). Starting the conversation around the ethical use of artificial intelligence in applied behaviour analysis. *Behavior Analysis in Practice*, *17*(1), 232-242. https://doi.org/10.1007/s40617-023-00868-z

Kulal, A., Rahiman, H. U., Suvarna, H., Abhishek, N., & Dinesh, S. (2024). Enhancing public service delivery efficiency: Exploring the impact of AI. *Journal of Open Innovation Technology Market and Complexity*, *10*, 100329. https://doi.org/10.1016/j.joit.2024.100329

Madan, R., & Ashok, M. (2023). AI adoption and diffusion in public administration: A systematic literature review and future research agenda. *Government Information Quarterly*, *40*(1), 101774. https://doi.org/10.1016/j.giq.2022.101774

Medaglia, R., Gil-Garcia, J. R., & Pardo, T. A. (2021). Artificial intelligence in government: Taking stock and moving forward. *Social Science Computer Review*, *41*(1), 123-140. https://doi.org/10.1177/08944393211034087

Mutengambiri, B., Somane, D., Chigome, C. J., Gweshe, G., & Chiware, M. (2024). Five-year performance contracts and service delivery in local authorities in Zimbabwe: The case of Ruwa. *Journal of Public Administration and Governance*, *14*(1), 53-71. https://doi.org/10.5296/jpag.v14i1.21880

Nalubega, T., & Uwizeyimana, D. E. (2024). Artificial intelligence technologies usage for improved service delivery in Uganda. *Africa's Public Service Delivery and Performance Review*, *12*(1), 770. https://doi.org/10.4102/apsdpr.v12i1.770

Neumann, O., & Guirguis, K. (2024). Exploring artificial intelligence adoption in public organizations: A comparative case study. *Public Management Review*, *26*(4), 789-814. https://doi.org/10.1080/14719037.2022.2048685

Nguyen, A., Ngo, H. N., Hong, Y., Dang, B., & Nguyen, B. P. T. (2023). Ethical principles for artificial intelligence in education. *Education and Information Technologies*, *28*(4), 4221-4241. https://doi.org/10.1007/s10639-022-11316-w

Nyikadzino, T., & Vyas-Doorgapersad, S. (2023). Reengineering imperatives for revamping public trust in Zimbabwe's local government. *Journal of Southwest Jiaotong University*, *58*(5). https://doi.org/10.35741/issn.0258-2724.58.5.37

Ooko, R., Mwangi, E., & Karanja, P. (2024). Artificial intelligence and sustainable development in Africa: A comprehensive review. *Engineering Applications of Artificial Intelligence*, *136*, 108765. https://doi.org/10.1016/j.engappai.2024.108765

Poshai, L., & Vyas-Doorgapersad, S. (2023). Digital justice delivery in Zimbabwe: Integrated electronic case management system adoption. *SA Journal of Information Management*, *25*(1), 1695. https://doi.org/10.4102/sajim.v25i1.1695

Pulijala, S. (2024). Artificial intelligence in governance: Opportunities, challenges, and ethical implications for public administration. *International Journal for Multidisciplinary Research*, *6*(6). https://doi.org/10.36948/ijfmr.2024.v06i06.29990

Revesai, Z. (2024). Ethical Implications of AI-Driven Education Systems on Digital Rights: A Comparative Analysis. *The Fountain: Journal of Interdisciplinary Studies*, *8*(1), 126-152.

Revesai, Z., & Hapanyengwi, G. T. (2024). Integrating Artificial Intelligence in Zimbabwe's Education Curriculum: A Call For Global Collaboration. *POTRAZ ICT Research Journal*, *1*(1), 101-131.

Revesai, Z., & Kogeda, O. P. (2025a). Lightweight Interpretable Deep Learning Model for Nutrient Analysis in Mobile Health Applications. *Digital*, *5*(2), 23. https://doi.org/10.3390/digital5020023

Revesai, Z., & Kogeda, O. P. (2025b). Self-Explaining Neural Networks for Food Recognition and Dietary Analysis. *BioMedInformatics*, *5*(3), 36. https://doi.org/10.3390/biomedinformatics5030036

Rulandari, N., Silalahi, A. D. K., Phuong, D. T. T., & Eunike, I. J. (2025). Decoding effectiveness and efficiency in AI-enabled public services: A configurational pathway to citizen and employee satisfaction. *Frontiers in Political Science*, *7*, 1560180. https://doi.org/10.3389/fpos.2025.1560180

van Noordt, C., & Misuraca, G. (2022). Artificial intelligence for the public sector: Results of landscaping the use of AI in government across the European Union. *Government Information Quarterly*, *39*(3), 101714. https://doi.org/10.1016/j.giq.2022.101714

Wilson, C., & Broomfield, H. (2024). Learning how to do AI: Managing organizational boundaries in an intergovernmental learning forum. *Public Management Review*, *26*(8), 1654-1675. https://doi.org/10.1080/14719037.2023.2231950