ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

The Role of Advanced Anomaly Detection in Transforming Program Management in Government with Scikit-Learn, A Machine Learning Library in Python

Chiedozie M. Okafor[1]*, Owolabi Ogunse[2] , Mercy Nneoma Iheke[3] , Dickson O. Oseghale[4] , Imuetinyan Ogiehor[5] , Ebuka Emmanuel Aniebonam[6]

[1]Position/Role: Financial Analyst | Independent Researcher | Certified Information Systems Auditor (CISA)
Affiliation (Institution/Department): Independent Researcher & Information Systems Audit and Control Association (ISACA) – Abuja Chapter

[2]Position/Role: IT Manager Affiliation: Independent Researcher,

[3]Position/Role: Assistant Manager, Grants Assurance | Independent Researcher Affiliation (Institution/Department): Independent Researcher

[4]Position/Role: Budget Analyst Affiliation (Institution/Department): Division of Global HIV & TB, U.S Centers for Disease Control and Prevention

[5]Position/Role: Senior Stakeholder Engagement Manager

[6]Position/Role: MBA Graduate Student, North Star Mutual School of Business Department of Business, Innovation and Strategy. Southwest Minnesota State University

***Corresponding Author**

DOI: https://doi.org/10.51244/IJRSI.2025.120500014

Received: 21 April 2025; Accepted: 25 April 2025; Published: 26 May 2025

ABSTRACT

Government agencies increasingly face challenges in managing programs efficiently, especially in preventing fraud, abuse, and resource waste. Traditional oversight techniques often struggle to detect early signs of problems such as resource misallocation, project delays, and budget overruns. This paper explores how anomaly detection, powered by Scikit-learn, a machine learning library in Python, can improve government program management. Using Isolation Forest, One-Class Support Vector Machine (SVM), and Local Outlier Factor models, we illustrate how advanced anomaly detection can monitor budgets, timelines, and resource utilization. Our findings show that these approaches can enhance program efficiency, enable agile risk management, and support data-driven decisions through a hypothetical example focused on government project data.

Beyond improving technical oversight, this research explores how the integration of advanced anomaly detection techniques, enabled by Scikit-learn's machine learning capabilities, can fundamentally transform program management practices within government agencies. Recognizing that program managers may not necessarily be programmers, the work highlights practical pathways for them to either collaborate closely with data scientists or develop foundational skills in interpreting machine learning outputs, fostering a stronger, analytics-driven management culture. This shift encourages a new governance model where data-driven insights enhance budgeting, scheduling, and resource allocation decisions. Furthermore, the reviewed literature provides a foundation for positioning machine learning anomaly detection as a strategic instrument for strengthening oversight, enhancing operational performance, and promoting proactive decision-making across government initiatives.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Keywords: Program Management, Anomaly Detection, Machine Learning, Scikit-learn, Data Analysis.

Overview And Challenges In Government Program Management And The Role Of Anomaly Detection

Government program management often involves large, complex projects with significant budgets, extended timelines, and diverse stakeholders (Hopmere et al., 2020). Due to this complexity, issues such as budget overruns, missed deadlines, and resource misallocations are common, compounded by a lack of real-time monitoring that limits early detection of emerging problems (Schrage, 2018).

Anomaly detection offers a powerful solution by identifying deviations in key project indicators such as spending patterns, milestone completion, and resource utilization (Barrados & Blain, 2013). Early detection enables agencies to assess risks swiftly and intervene before issues escalate. Proven successful in domains like fraud detection, healthcare, and cybersecurity, anomaly detection demonstrates versatility in analyzing complex datasets to generate actionable insights (Davis et al., 2020).

When integrated into program management, advanced machine learning algorithms—especially through platforms like Scikit-learn allow agencies to monitor vast datasets in real time, identify irregularities, and strengthen proactive risk management (Coglianese & Lehr, 2019; Emery-Xu et al., 2024). These systems not only flag anomalies but also uncover root causes behind inefficiencies, enabling smarter resource allocation and improved accountability (Ashkanani & Franzoi, 2022; Shah et al., 2023; Wang et al., 2017).

Furthermore, techniques such as clustering and predictive modelling allow agencies to forecast potential delays and detect systemic resource misalignments, reinforcing proactive oversight and enhancing overall program success (Hopmere et al., 2020).

Problem Statement

Despite advancements in data science, many government agencies still rely on traditional project management and oversight tools that are manual, retrospective, and slow (Foorthuis, 2018). These outdated approaches are not scalable for the growing complexity of modern public programs and often fail to detect emerging risks before they escalate (Khalid et al., 2018).

Fixed performance metrics, while useful for post-event analysis, lack the agility needed for real-time anomaly detection. As a result, critical issues such as budget overruns, resource misallocations, and timeline delays frequently remain hidden until they undermine project outcomes (Al-Jibouri, 2002; Othman et al., 2018).

To address these challenges, agencies must transition toward data-driven oversight systems that leverage machine learning models capable of identifying deviations from expected behavior in real time. These models enable proactive risk management, informed decision-making, and more efficient resource use capabilities essential in today's fast-paced, data-intensive government environments where traditional methods fall short (Dunning & Friedman, 2014).

Objective of the Study

This study examines how machine learning–based anomaly detection, powered by the Scikit-learn library, can transform program management practices within government agencies. Scikit-learn, a widely used Python library, provides efficient implementations of classification, clustering, and anomaly detection algorithms suited for analyzing complex public sector datasets (Hao & Ho, 2019).

The primary objective is to evaluate how models such as Isolation Forest, One-Class Support Vector Machine (SVM), and Local Outlier Factor can identify anomalies in government project budgets, timelines, and resource utilization. Using an illustrative example, the study demonstrates how these techniques enable real-time risk detection and support timelier, data-driven decisions.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Ultimately, this research aims to show how machine learning–enabled anomaly detection can enhance program oversight, optimize resource allocation, strengthen accountability, and drive more effective public service delivery.

Research Questions

This study is guided by the following research questions:

How can machine learning–based anomaly detection improve the efficiency and effectiveness of government program management?

This question examines how data-driven anomaly detection can enhance oversight, strengthen accountability, and improve operational agility in public sector projects.

What types of anomalies commonly occur in government program management and how can machine learning detect and address them?

This question identifies typical challenges, such as cost irregularities, schedule delays, and resource misallocations, and explores how machine learning can provide early warning signs and diagnostic insights.

How well do Scikit-learn's anomaly detection tools: Isolation Forest, One-Class SVM, and Local Outlier Factor meet the needs of real-time, data-driven oversight?

This question evaluates the technical strengths and limitations of these algorithms, focusing on their scalability, interpretability, and responsiveness in the public sector context.

Together, these questions guide the investigation into how Scikit-learn's anomaly detection algorithms can transform oversight practices in government program administration.

Thesis Statement

Integrating machine learning–based anomaly detection into government program management represents a transformational shift in addressing inefficiencies, resource misallocations, and performance risks. Through Scikit-learn's powerful algorithms, agencies can detect deviations in budgets, timelines, and resource utilization in real time, advancing oversight from reactive monitoring to proactive, data-driven governance. By equipping program managers with advanced analytical tools, this approach strengthens transparency, optimizes resource allocation, and reinforces accountability across public sector programs.

This study establishes a data-driven framework for enhancing operational efficiency, responsiveness, and governance outcomes in complex, high-volume program environments.

LITERATURE REVIEW

This literature review critically examines three foundational areas essential to understanding the integration of machine learning–based anomaly detection in government program management. First, it explores the current challenges in public sector program oversight, highlighting persistent issues of inefficiency, resource misallocation, and delayed risk identification. Second, it reviews the application of anomaly detection techniques across various sectors, illustrating their versatility and effectiveness in identifying irregularities within complex datasets. Finally, it focuses on the role of machine learning particularly the Scikit-learn library in advancing anomaly detection capabilities for real-time oversight, positioning these tools as strategic enablers of more transparent, accountable, and data-driven governance practices.

Overview of Program Management in Government Agencies

Government program management encompasses large, multiyear initiatives that demand significant financial resources, complex logistical coordination, and engagement with diverse stakeholder groups. These programs

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

often operate within critical sectors such as infrastructure, national defense, healthcare, and social welfare, where effective oversight is essential but frequently undermined by regulatory complexity, political pressures, and leadership instability (DeGroff & Cargo, 2009; Willoughby, 1927).

Persistent issues such as budget overruns, project delays, and scope expansions continue to erode public trust and program performance. Yet many agencies still rely on manual tracking methods and traditional project management tools that are poorly suited for today's data-intensive environments. These outdated systems emphasize retrospective reviews rather than real-time risk detection, leaving agencies unable to intervene early during program execution.

As government programs generate increasingly complex and voluminous datasets, there is a pressing need to transition toward data-driven oversight models. Integrating anomaly detection techniques enables early identification of emerging risks, improves resource management, and supports greater accountability and transparency in public sector operations (West, 2023).

Exploring Anomaly Detection Methods for Operational Risk Identification.

Anomaly detection is a core element of data analysis, widely applied across cybersecurity, fraud detection, and industrial process monitoring. Fundamentally, anomaly detection seeks to find groups of data that do not conform to expected behavior. This is especially useful in program management where any variance from expected budget outcome, time or allocation of resources may signal a potential challenge with project delivery(Xu et al., 2019).

Some of the types of anomalies which are addressed in this literature include:

Point Anomalies: These are individual points which differ from other values in the data. For example, a sudden and unexpected increase in spending within a specific department can be classified as a point anomaly.

Contextual Anomalies: These anomalies are context-dependent and may only be considered anomalous when viewed in a particular setting. For instance, while a delay in a specific project phase may be normal, a delay during a critical milestone could indicate a significant issue.

Collective Anomalies: This involves a collection of related data points that together form an anomalous pattern. In the context of program management, multiple projects within a program showing resource constraints simultaneously could be a collective anomaly signalling broader systemic issues.

Various techniques for anomaly detection exist, including statistical methods, proximity-based methods, and clustering approaches. Statistical methods rely on probability distributions and thresholds, while proximity-based methods evaluate the distance between data points to identify outliers. Clustering approaches, such as K-means clustering (K-Means) and Density-Based Spatial Clustering of Applications with Noise (DBSCAN), group data points into clusters and identify anomalies based on their proximity to the cluster centroids (Braei & Wagner, 2020).

Advanced Machine Learning Algorithms for Detecting Anomalies in Public Program Management.

Machine learning (ML) has emerged as a critical tool for anomaly detection, offering advanced capabilities to identify irregular patterns within large and complex datasets. Unlike traditional systems, ML algorithms adapt dynamically by learning from historical data, enabling the detection of both subtle and significant deviations (Agyemang, 2024; Oprea et al., 2021). This adaptability is particularly valuable for public sector agencies, which must monitor diverse and rapidly changing data streams such as budgeting, procurement, and resource utilization in real time.

Three notable machine learning algorithms often used for anomaly detection are:

Isolation Forest:

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Isolation Forest detects anomalies by breaking down data into smaller partitions and isolating records that behave differently from the majority. Anomalous points tend to be separated quickly, making this method efficient and scalable, especially for large datasets with many variables.

In the United States, the Department of Health and Human Services (HHS) could use Isolation Forest to identify unusual increases in Medicaid payments across states, which might signal fraud or administrative errors.

In Nigeria, the Budget Office of the Federation might apply the algorithm to detect sudden spikes in capital project spending within ministries, prompting audits or deeper reviews.

One-Class Support Vector Machine (SVM):

One-Class SVM focuses on learning the characteristics of "normal" data and flags any new observations that do not fit this pattern. This approach is particularly useful when examples of abnormal behavior are rare or unknown in advance, a common challenge in public sector oversight.

In Nigeria, the National Social Investment Programme (NSIP) could use One-Class SVM to detect duplicate or suspicious registrations in its cash transfer programs.

In the United States, the Internal Revenue Service (IRS) could deploy the algorithm to flag tax refund claims that deviate from normal filing patterns.

Local Outlier Factor (LOF):

Local Outlier Factor evaluates how different a data point is compared to its immediate neighbours, focusing on local patterns rather than global ones. This makes it especially useful for detecting context-specific anomalies.

For example, the Nigerian Federal Ministry of Education could use LOF to spot schools with unusually high reported meal deliveries or attendance, which could indicate misreporting within the school feeding program.

Similarly, in the United States, the Department of Transportation (DOT) could apply LOF to detect infrastructure projects facing unusual delays compared to similar projects nearby, enabling early interventions.

Table 1: Visual Comparison of Anomaly Detection Algorithms for Public Sector Program Management

Algorithm	Strengths	Weaknesses	Ideal Use Cases
Isolation Forest	- Fast and scalable to large datasets - Handles high-dimensional data well	- May perform poorly with small datasets - Less effective for context-specific anomalies	- Monitoring large budget datasets for irregular spending - Detecting procurement anomalies across ministries
One-Class SVM	- Effective with scarce anomaly examples - Learns the profile of "normal" behavior without explicit labels	- Sensitive to kernel and parameter settings - Computationally intensive on very large datasets	- Identifying fraudulent entries in welfare programs - Spotting unusual vendor registrations in procurement systems
Local Outlier Factor (LOF)	- Detects local deviations in density - Excellent for context-dependent or localized anomalies	- Struggles with high-dimensional data - May misclassify when clusters are not uniform	- Finding localized project timeline delays - Detecting underperformance in specific service delivery regions

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Machine learning algorithms offer transformative opportunities to modernize public sector oversight by enabling real-time anomaly detection within large, complex datasets. These tools can uncover irregularities such as resource misallocations, procurement anomalies, service delivery gaps, and project execution delays that traditional monitoring systems often miss (Goh et al., 2021).

In Nigeria, where public institutions often grapple with manual processes, fragmented data systems, and limited technical capacity, machine learning could significantly strengthen operational oversight. Agencies such as the Bureau of Public Procurement (BPP), the Office of the Accountant General of the Federation, and the National Social Investment Programme (NSIP) could integrate algorithms like Isolation Forest, One-Class SVM, and Local Outlier Factor (LOF) to flag unusual spending patterns, detect duplicate welfare entries, and monitor deviations from project milestones. Embedding anomaly detection into operational workflows would reduce audit burdens, enhance financial reporting accuracy, and promote more transparent and effective service delivery.

In contrast, U.S. federal agencies, already supported by advanced digital infrastructure, could leverage machine learning to optimize operational efficiency further. The General Services Administration (GSA) could use Isolation Forest to identify anomalies in procurement costs, while the Environmental Protection Agency (EPA) and Department of Veterans Affairs (VA) could deploy LOF to monitor grant distributions and healthcare services, ensuring resources are allocated effectively and performance bottlenecks are minimized.

By enabling proactive risk detection across diverse operational contexts, machine learning transitions both Nigerian and U.S. institutions from reactive, audit-heavy oversight models to predictive, performance-driven governance (Coglianese & Lehr, 2019; Grover et al., 2019). This shift enhances transparency, strengthens accountability, and improves the agility and responsiveness of public sector programs, ultimately building greater public trust.

Integrating Scikit-Learn Algorithms Into Public Sector Oversight Systems.

Scikit-learn, a widely used open-source machine learning library in Python, offers efficient implementations of classification, clustering, and anomaly detection algorithms optimized for large, complex datasets (PedregosaFabian et al., 2011). For public sector oversight, models such as Isolation Forest, One-Class Support Vector Machine (SVM), and Local Outlier Factor (LOF) are particularly valuable for detecting deviations in financial, operational, and service delivery data streams. Scikit-learn's efficiency, scalability, and user-friendly interface make it a practical tool for embedding real-time anomaly detection into existing government systems with minimal integration complexity (Hoag et al., 2023; Coglianese & Lehr, 2019; Pi, 2021).

In Nigeria, where digital infrastructure is still developing, Scikit-learn provides a cost-effective pathway to strengthening public sector oversight. The National Primary Health Care Development Agency (NPHCDA) could apply anomaly detection models to track irregularities in vaccine distribution across rural centers, while the National Bureau of Statistics (NBS) could identify inconsistencies in subnational poverty or employment reporting. Embedding such models would enable earlier interventions, reduce audit burdens, and improve data reliability.

In the United States, where agencies operate within mature digital ecosystems, Scikit-learn supports scaling proactive oversight efforts. The Environmental Protection Agency (EPA) could deploy these models to detect compliance anomalies in environmental reporting, the General Services Administration (GSA) could identify irregular procurement patterns, and the Department of Housing and Urban Development (HUD) could monitor anomalies in housing assistance disbursements and project timelines. These applications would enhance risk management, optimize resource use, and strengthen transparency across complex operational networks.

Research increasingly supports the role of machine learning–based anomaly detection in public sector analytics. Techniques embedded in Scikit-learn have proven effective in areas such as social welfare fraud detection and public health data quality assurance (Davis et al., 2020; Novita & Anissa, 2022). As governments expand data-driven governance strategies, Scikit-learn offers a robust, adaptable framework for modernizing oversight, improving accountability, and delivering more efficient public services.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Data Preparation and Model Evaluation

Before deploying machine learning models for anomaly detection, careful data preparation is essential. Key preprocessing steps include cleaning datasets to handle missing values, normalizing features for consistent scaling, and removing irrelevant variables that could introduce noise. Choosing the right features is also important to ensure models are both accurate and easy to interpret. Additionally, it is critical to watch for biases in the data because if historical records are unbalanced, models might mistakenly classify normal behaviors as anomalies or miss true problems.

In this study, we used simulated datasets representing government programs to test model performance. To assess the effectiveness of each anomaly detection model, standard evaluation metrics like **precision**, **recall**, and **F1 score** were used. These measures are especially important when working with public sector data, where true anomalies are rare. Providing clear documentation on how the data was prepared and how the models were tested is crucial for building credibility and ensuring trust in the findings, both of which are essential when applying machine learning tools in government settings.

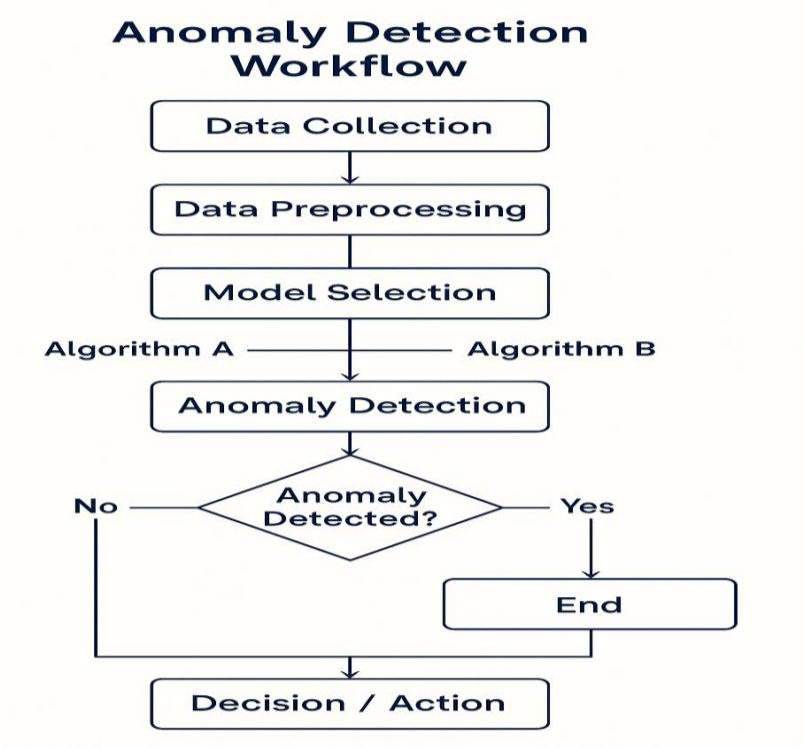

Figure 1: Anomaly Detection Workflow in Public Sector Program Management

Figure 1 presents a flowchart that outlines the full process for deploying machine learning–based anomaly detection systems within public program management. The workflow begins with **Data Collection**, where both structured and unstructured data are gathered from operational, financial, and administrative systems across government agencies. Examples of such data include budget execution records, procurement transactions, beneficiary registries, and project milestone reports.

Next, in the **Data Preprocessing** stage, raw inputs are cleaned, normalized, and transformed. These steps handle missing values, correct inconsistencies, and enhance the relevance of features, all of which are critical for improving the accuracy and performance of the machine learning models.

During the **Model Selection** phase, agencies evaluate which algorithms to use such as Isolation Forest, One-Class SVM, or Local Outlier Factor (LOF) based on the types of anomalies they need to detect (point, contextual, or collective), the size and complexity of the dataset, and the level of interpretability required. This decision-making process is illustrated in the flowchart by the branching paths leading to Algorithm A or Algorithm B, allowing agencies to tailor their model choices to specific operational needs.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Once an algorithm is selected, the system moves to the **Anomaly Detection** stage, where the model analyses the input data and assigns anomaly scores to individual records. At the **Decision Point**, the system asks, "Anomaly Detected?"

If an anomaly is identified, the process moves to the **Decision / Action** stage, prompting auditors, analysts, or program managers to investigate the flagged issue, validate it, and take corrective or preventive measures.

If no anomaly is detected, the system completes the cycle but remains ready for periodic or real-time re-evaluation.

This modular, logic-driven workflow supports scalable, explainable, and responsive oversight mechanisms across various areas of public administration. It helps agencies move from reactive governance to proactive, data-informed decision-making.

To further demonstrate the practical feasibility of machine learning–based anomaly detection for government program oversight, this study applies the discussed models to a simulated dataset designed to reflect typical budgetary, timeline, and resource utilization scenarios in public sector projects. Using Scikit-learn within a Python environment, the models are tested on synthetic data to visualize anomalies, assess detection performance, and illustrate real-world applicability. The following section presents the methodology, results, and insights generated from this illustrative application.

Illustrative Application: Simulated Anomaly Detection In Government Program Data

Data Collection

This study employed an AI-generated dataset designed to simulate critical features of government program management. The dataset comprises the following elements:

Budget Data: Captures overall spending, actual versus planned costs, and monthly expenditure trends.

Project Milestone Data: Tracks planned versus actual dates for key project phases.

Resource Utilization: Includes personnel hours, equipment use, and material costs across departments.

Simulating this data enables the controlled evaluation of anomaly detection models while addressing privacy and confidentiality concerns often associated with real government program data. The dataset was created and analysed within a Jupyter Notebook environment, facilitating real-time adjustments, documentation, and visualization.

Analytical Framework

To demonstrate the practical value of machine learning in public program oversight, this study used Python's Scikit-learn library to implement three widely used anomaly detection models:

Isolation Forest

One-Class Support Vector Machine (SVM)

Local Outlier Factor (LOF)

Each model was applied to detect anomalies across three key dimensions: budgets, timelines, and resource utilization. These models help flag deviations that may signal financial inefficiencies, scheduling delays, or resource misallocations.

The results were visualized using color-coded scatter plots, which enable clear interpretation of anomalous points and support proactive interventions by program managers.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

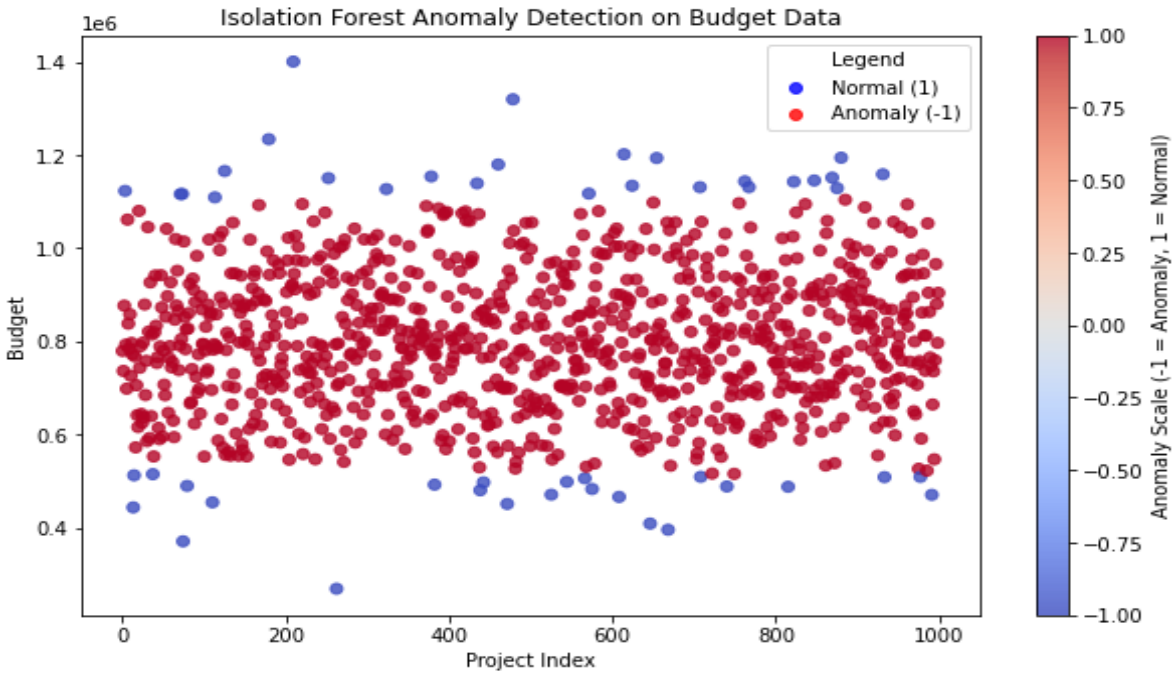

Figure 2: Isolation Forest Anomaly Detection on Budget Data

This figure illustrates the application of the **Isolation Forest** model in identifying budget anomalies. The output was generated using Scikit-learn within a Python-based Jupyter Notebook environment. Red data points (labelled as -1) denote anomalies such as budget overspending or unauthorized expenses while blue points (labelled as 1) indicate typical spending patterns.

By isolating these outliers, the model provides a powerful early warning tool to enhance financial oversight and accountability in government budgeting.

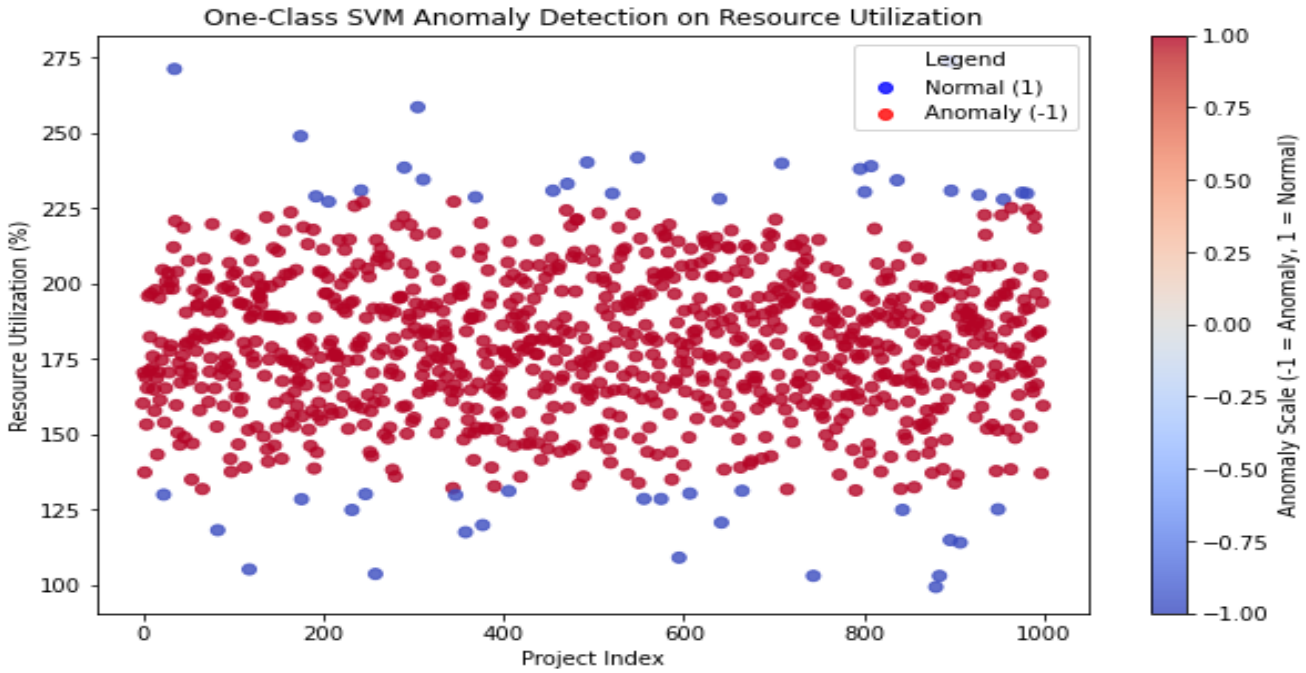

Figure 3: One-Class SVM Anomaly Detection on Resource Utilization

This figure demonstrates how the **One-Class SVM** model detects anomalies in resource utilization data. Implemented via Scikit-learn in Python, the model identifies abnormal patterns in how labor, equipment, or materials are used across projects.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Min-max normalization was applied to standardize feature ranges. Red points (labelled as -1) highlight projects with unusually high or low resource use, while blue points (labelled as 1) represent typical utilization. These insights are critical for minimizing waste and improving operational efficiency.

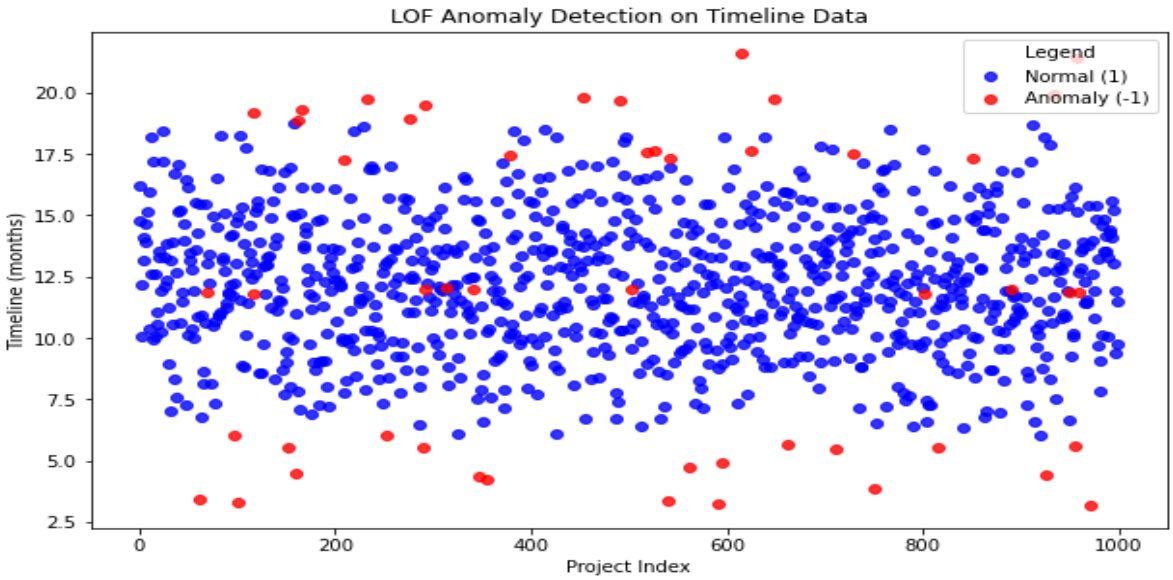

Figure 4: LOF Anomaly Detection on Timeline Data

This figure presents the output of the **Local Outlier Factor (LOF)** algorithm applied to project timeline data. The model, implemented using Scikit-learn in Jupyter Notebook, assesses local deviations to detect scheduling anomalies.

Red points (labelled as -1) signify projects with timelines that diverge significantly from their neighbours, potentially indicating delays or inefficiencies. Blue points (labelled as 1) fall within expected ranges. Because LOF focuses on localized patterns, it is well-suited for identifying risks in diverse project portfolios.

Performance Comparison of Anomaly Detection Models

To evaluate how each anomaly detection model performs, we assessed their accuracy using standard metrics sumarized below:

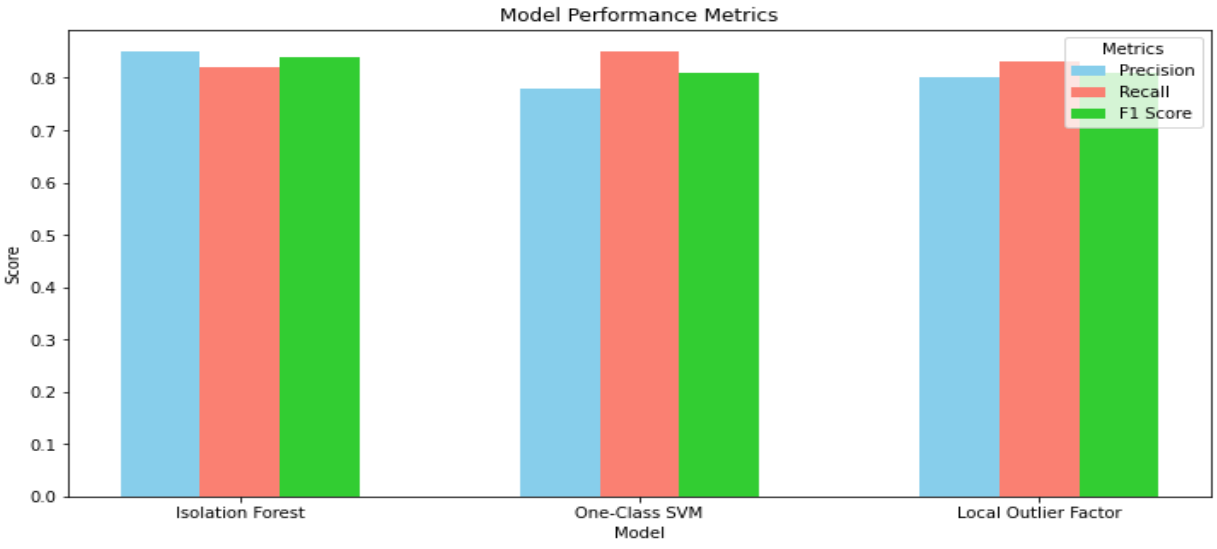

Figure 5.

The bar chart in **Figure 5** compares the precision, recall, and F1 score of each anomaly detection model.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Isolation Forest showed the highest precision in F1 score, making it ideal for **budget oversight**, where minimizing false positives is crucial. In financial monitoring, a **type-1 error** flagging a normal transaction as anomalous can waste valuable time and resources. High precision ensures that flagged anomalies are truly problematic, supporting focused corrective action.

The One-Class SVM model demonstrated the highest **recall**, meaning it successfully detected more real anomalies in **resource utilization**. Even if it generates more false alarms, its strength lies in ensuring that few real issues go unnoticed, making it particularly useful for programs where missed anomalies could have costly consequences.

Finally, Local Outlier Factor (LOF) delivered **balanced scores** across precision, recall, and F1 metrics. This balance makes it well-suited for monitoring **project timelines**, where local deviations from expected milestones are important to detect early before delays cascade into broader project risks.

From Theory To Practice: Machine Learning Applications In Government Program Management.

Integrating machine learning-based anomaly detection into government program management provides major advantages in budget oversight, resource optimization, and schedule management. Each model discussed: Isolation Forest, One-Class SVM, and Local Outlier Factor targets specific types of risks and inefficiencies, giving agencies the ability to intervene early and strategically (Kanksha et al., 2021).

The use of machine learning technologies is expanding across industries, including the public sector, making complex analytical methods more accessible to a broader range of users (Sehatbakhsh et al., 2020). Machine learning empowers governments to adopt data-driven strategies for smarter and more responsive decision-making (Alexopoulos et al., 2019).

Strengthening Financial Oversight

Isolation Forest has proven effective at identifying irregular financial patterns, such as overspending or unexpected deviations from budget allocations. By enabling the early detection of material budget variances, the model allows program managers to investigate and address issues before they escalate. Implementing Isolation Forest can promote greater financial discipline, reduce the likelihood of budget overruns, and strengthen accountability for public spending (Prasad et al., 2023).

Improving Resource Allocation and Efficiency

One-Class SVM is particularly strong in identifying anomalies related to resource utilization. It can flag departments or projects that are either overusing or underutilizing resources. These insights allow managers to redistribute labor, equipment, or funding to where it is most needed, preventing inefficiencies and ensuring resources are allocated more effectively across programs.

Enhancing Timeline Management

Local Outlier Factor is especially useful for detecting anomalies in project timelines. It helps uncover early warning signs of delays or prolonged phases across similar projects. Early identification allows managers to adjust schedules, reassign responsibilities, or intervene before delays become serious. For time-sensitive programs such as healthcare, infrastructure, or emergency response, early detection is critical to maintaining service delivery standards.

Limitations And Implementation Of Machine Learning Models.

While machine learning models offer powerful tools for anomaly detection, they also present challenges that must be addressed to ensure successful implementation.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Challenges in Detecting Contextual Anomalies

One major limitation is the difficulty in identifying contextual anomalies. Data points that appear unusual may be justified under specific conditions, such as emergency procurements or seasonal fluctuations, leading to false positives (Yuan & Wu, 2022). Overcoming this challenge requires integrating statistical modelling with domain expertise to refine anomaly definitions within the agency's operational context.

Issues with Interpretability

Interpretability remains another significant concern. Models like Isolation Forest and Local Outlier Factor (LOF) generate anomaly scores but often lack transparent, easily understandable explanations. Without clear interpretability, decision-makers may hesitate to act on model outputs. Developing explainable AI (XAI) tools such as visual dashboards and attribution methods can improve transparency and build stakeholder trust (Kostopoulos et al., 2024; Busuioc, 2020).

Institutional Constraints and Technical Readiness

Implementing machine learning systems also depends on the institutional and technical environment. Real-time monitoring requires strong digital infrastructure, secure data pipelines, and trained personnel which may be limited, especially in developing settings.

Moreover, agencies must comply with data privacy regulations, such as GDPR and national privacy laws, requiring secure storage, anonymization, and restricted data access protocols. Continuous model maintenance and recalibration are also necessary to adapt to changing project dynamics and prevent model performance from degrading over time (Gruschka et al., 2018).

Strategic Pathways for Successful Integration

To successfully integrate anomaly detection systems, agencies should start by investing in secure, scalable data infrastructure capable of supporting real-time analytics. Building internal capacity among program managers and data analysts is critical to ensuring that machine learning outputs are correctly interpreted and acted upon.

Partnerships with research institutions and private technology firms can further help tailor solutions to agency-specific needs and foster innovation in public sector data science (Dunleavy & Margetts, 2023).

Rather than attempting large-scale rollouts immediately, agencies are encouraged to begin with pilot programs that allow controlled testing of machine learning applications. A phased approach helps minimize risks, enables fine-tuning, and builds stakeholder confidence.

Throughout implementation, transparency and explainability must remain core principles, ensuring that machine learning insights are not only technically sound but also aligned with public accountability standards (Coglianese & Lehr, 2019).

Ethical Considerations And Data Privacy

While machine learning–based anomaly detection offers significant potential for enhancing government program oversight, it also raises important ethical and data privacy concerns that must be addressed to ensure responsible implementation.

One major concern is the risk of bias in the underlying data. If historical datasets reflect systemic inequalities or incomplete reporting, anomaly detection models may unintentionally reinforce or amplify these biases. As a result, certain regions, departments, or demographic groups could be disproportionately flagged for investigation, even when deviations are justifiable in context. To minimize this risk, technical modelling must be complemented by human oversight and domain expertise during the interpretation of results.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Data privacy is another critical issue, particularly when models analyse sensitive financial, personnel, or citizen service information. Compliance with relevant data protection regulations, such as the General Data Protection Regulation (GDPR) and country specific or national privacy laws, are essential. Privacy safeguards including data anonymization, encryption, and restricted access protocols must be embedded into the system design from the outset.

In addition, public trust hinges on transparency and explainability. When anomaly detection models are deployed by public institutions, it is vital that stakeholders can understand how decisions are made. Agencies must prioritize the use of explainable AI models, document the rationale behind anomaly flags, and create clear mechanisms for appeals or independent reviews of automated decisions.

Finally, ethical deployment of these systems should go beyond punitive measures. Anomaly detection should be used not only to identify irregularities but also to support early interventions, capacity-building, and constructive support for struggling programs or resource-challenged areas.

By embedding principles of fairness, transparency, and accountability into both the technical design and operational deployment of anomaly detection systems, public sector institutions can ensure these tools align with democratic values and protect citizens' rights.

Recommendation For Government Agencies.

To meet the growing demand for improved oversight and performance management in government programs, the adoption of machine learning–based anomaly detection must be approached with strategic planning and institutional readiness. The following recommendations aim to guide agencies toward effective implementation:

Invest in Modern Data Infrastructure

Agencies should prioritize building scalable, secure, and interoperable data systems that support the continuous collection, integration, and analysis of program-related information. Establishing robust data pipelines and adopting cloud-based platforms can facilitate real-time data processing and streamline the deployment of machine learning models.

Furthermore, modern infrastructure must include strong data governance frameworks to ensure data quality, consistency, and accessibility across departments forming the foundation for effective model training and deployment (Adesina et al., 2024).

Develop Internal Technical Capacity

Sustainable use of machine learning tools requires investment in workforce development. Agencies should focus on upskilling program managers, data analysts, IT specialists, and decision-makers in data literacy, algorithmic thinking, and the practical use of anomaly detection systems.

Building internal technical capacity reduces dependency on external consultants, strengthens institutional knowledge, and empowers staff to interpret model outputs, provide feedback, and refine analytical frameworks as operational needs evolve (Chalapathy & Chawla, 2019).

Foster Interdisciplinary Collaboration

Successful implementation requires collaboration across technical, operational, and policy domains. Agencies should create cross-functional teams comprising data scientists, project managers, procurement officers, and regulatory stakeholders to co-design use cases and validate model outputs (Chen et al., 2021; Kreuter et al., 2019).

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

In addition, partnerships with research institutions, civic technology groups, and private **sector** vendors can accelerate innovation, support pilot testing, and facilitate the customization of models for diverse operational environments.

Prioritize Model Explainability and Transparency

Trust and accountability are essential in public sector innovation. Agencies must prioritize explainable AI (XAI) tools that clarify how anomalies are flagged, making it easier for decision-makers to act confidently on model outputs (Coglianese & Lehr, 2019; Warner & Sloan, 2021). Clear documentation, visual dashboards, and feedback mechanisms should accompany all model deployments to ensure transparency both internally and with external oversight bodies.

Adopt an Incremental Implementation Strategy

To minimize operational disruptions and encourage stakeholder buy-in, agencies should initially implement machine learning models in pilot projects targeting high-impact areas such as budgeting, procurement, or project milestone tracking. These pilots can serve as proofs of concept for wider adoption (Lee et al., 2019).

A phased rollout allows agencies to calibrate models, test technical and managerial feasibility, and build internal momentum, with lessons from early phases informing broader policy and procedural adjustments necessary for sustainable scaling.

Future Research Directions

Anomaly detection techniques offer transformative potential for strengthening government program oversight. However, achieving their full impact at scale requires further research to advance the technical, operational, and institutional aspects of machine learning applications in the public sector. Future work should focus on improving methodologies, validating models in real-world settings, enhancing interpretability, and addressing ethical considerations to ensure reliable and responsible deployment (Alexopoulos et al., 2019).

To fully realize the potential of anomaly detection in public sector oversight, several key research areas below must be explored:

Advancing Anomaly Detection Models

Future research should explore more advanced machine learning architectures, including deep learning, hybrid models, and ensemble methods. Neural network–based techniques, such as autoencoders and recurrent neural networks (RNNs), are particularly effective at capturing complex, high-dimensional, and time-sequenced patterns making them well suited for government program data.

Hybrid approaches that combine supervised and unsupervised learning, especially in semi-supervised environments with limited labelled data, can significantly improve anomaly detection performance. Integrating multiple models such as combining Isolation Forest with autoencoders or clustering methods—can further enhance the detection of both point and contextual anomalies, supporting greater adaptability across different program management settings (Palacio, 2018; Chalapathy & Chawla, 2019; Yuan & Wu, 2022).

Real-World Application and Empirical Validation

Pilot studies are essential to move anomaly detection from theoretical promise to practical utility. Future research should prioritize real-world deployments across diverse government sectors such as budget tracking for infrastructure projects, procurement monitoring, or performance evaluations in public health programs. Such empirical validations will refine models, expose operational constraints, and provide insights into data readiness, stakeholder engagement, and institutional capacity (Eddy et al., 2012; Casey et al., 2018).

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

Enhancing Real-Time Monitoring Capabilities

The greatest value of anomaly detection lies in real-time risk identification. Future research should focus on optimizing data pipelines and deploying continuous analytics platforms, such as Apache Kafka and Apache Flink, to enable near-instant detection and response.

Edge computing and real-time dashboards can further enhance responsiveness, especially in decentralized program management. Advancements must not only improve computational efficiency but also establish governance structures that ensure timely action based on system alerts (Pashami et al., 2023).

Improving Explainability and Transparency

A major barrier to public sector adoption of machine learning is the lack of model explainability. Stakeholders including program managers, auditors, and policymakers must be able to understand and trust how anomalies are detected.

Future research should advance explainable AI (XAI) techniques, such as SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-Agnostic Explanations), and visual attribution tools, tailored specifically to anomaly detection systems. Enhanced interpretability not only builds trust but also helps uncover potential biases, ensures fairness, and strengthens compliance with legal and ethical standards (Poursabzi-Sangdeh et al., 2018; Gilpin et al., 2018).

Developing Sector-Specific and Context-Aware Models

Anomalies do not manifest uniformly across public sector domains. For example, what signals a delay in transportation infrastructure may appear routine in social services Future research should focus on creating sector-specific anomaly detection frameworks that incorporate domain expertise, operational constraints, and contextual factors. Collaboration with domain specialists can refine anomaly definitions, optimize detection thresholds, and reduce false positives, making outputs more actionable (Chakraborty et al., 2020).

Addressing Ethical, Legal, and Governance Challenges

As AI adoption widens in public governance, ethical and legal considerations must be embedded at every stage. Future research should examine how anomaly detection intersects with data privacy regulations (e.g., GDPR, FOIA), fairness requirements, and public accountability obligations.

Transparency, data stewardship, and stakeholder consultation must be integral parts of model development and deployment. This includes defining clear roles for oversight, ensuring respect for citizen rights, and establishing mechanisms for appeals or corrections when automated errors occur (Sanderson et al., 2023; Radanliev et al., 2024; Larsson & Heintz, 2020).

REFERENCES

1. Adesina, A. A., Iyelolu, T. V., & Paul, P. O. (2024). Leveraging predictive analytics for strategic decision-making: Enhancing business performance through data-driven insights. World Journal of Advanced Research and Reviews, 22(3), 1927. https://doi.org/10.30574/wjarr.2024.22.3.1961
2. Agyemang, E. F. (2024). Anomaly detection using unsupervised machine learning algorithms: A simulation study. Scientific African. https://doi.org/10.1016/j.sciaf.2024.e02386
3. Alexopoulos, C., Lachana, Z., Androutsopoulou, A., Diamantopoulou, V., Charalabidis, Y., & Loutsaris, M. A. (2019). How Machine Learning is Changing e-Government. https://doi.org/10.1145/3326365.3326412
4. Al-Jibouri, S. H. S. (2002). Monitoring systems and their effectiveness for project cost control in construction. International Journal of Project Management, 21(2), 145. https://doi.org/10.1016/s0263-7863(02)00010-8

5. Ashkanani, S., & Franzoi, R. E. (2022). An overview on megaproject management systems. Management Matters, 19(2), 129. https://doi.org/10.1108/manm-01-2022-0006
6. Barrados, M., & Blain, J. (2013). Improving Program Results Through the Use of Predictive Operational Performance Indicators. American Journal of Evaluation, 34(1), 45. https://doi.org/10.1177/1098214012464426
7. Braei, M., & Wagner, S. (2020). Anomaly Detection in Univariate Time-series: A Survey on the State-of-the-Art. arXiv (Cornell University). https://doi.org/10.48550/arXiv.2004.00433
8. Busuioc, M. (2020). Accountable Artificial Intelligence: Holding Algorithms to Account. Public Administration Review, 81(5), 825. https://doi.org/10.1111/puar.13293
9. Casey, P. C., Wilson, K. H., & Yokum, D. (2018). A Cautionary Tail: A Framework and Case Study for Testing Predictive Model Validity. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1807.03860
10. Chakraborty, S., Shah, S., Soltani, K., Swigart, A., Yang, L., & Buckingham, K. (2020). Building an Automated and Self-Aware Anomaly Detection System. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2011.05047
11. Chalapathy, R., & Chawla, S. (2019). Deep Learning for Anomaly Detection: A Survey. In arXiv (Cornell University). Cornell University. https://arxiv.org/abs/1901.03407
12. Chen, J. C., Rubin, E. A., & Cornwall, G. J. (2021). Building Data Teams. In Springer series in the data sciences (p. 317). Springer International Publishing. https://doi.org/10.1007/978-3-030-71352-2_16
13. Coglianese, C., & Lehr, D. (2019). Transparency and Algorithmic Governance. SSRN Electronic Journal. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3373299_code615352.pdf?abstractid=3293008&mirid=1
14. Davis, N., Raina, G., & Jagannathan, K. (2020). A framework for end-to-end deep learning-based anomaly detection in transportation networks. Transportation Research Interdisciplinary Perspectives, 5, 100112. https://doi.org/10.1016/j.trip.2020.100112
15. DeGroff, A., & Cargo, M. (2009). Policy implementation: Implications for evaluation. New Directions for Evaluation, 2009(124), 47. https://doi.org/10.1002/ev.313
16. Dunleavy, P., & Margetts, H. (2023). Data science, artificial intelligence and the third wave of digital era governance. Public Policy and Administration. https://doi.org/10.1177/09520767231198737
17. Dunning, T., & Friedman, E. (2014). Practical Machine Learning: A New Look at Anomaly Detection. http://ci.nii.ac.jp/ncid/BB18163833
18. Eddy, D. M., Hollingworth, W., Jaime, J., Tsevat, J., McDonald, K. M., & Wong, J. B. (2012). Model Transparency and Validation. Medical Decision Making, 32(5), 733. https://doi.org/10.1177/0272989x12454579
19. Emery-Xu, N., Jordan, R., & Trager, R. (2024). International governance of advancing artificial intelligence. AI & Society. https://doi.org/10.1007/s00146-024-02050-7
20. Foorthuis, R. (2018). A Typology of Data Anomalies. In Communications in computer and information science (p. 26). Springer Science+Business Media. https://doi.org/10.1007/978-3-319-91476-3_3
21. Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., & Kagal, L. (2018). Explaining Explanations: An Overview of Interpretability of Machine Learning. arXiv (Cornell University). https://doi.org/10.48550/arXiv.1806.00069
22. Goh, C., Lee, B., Pan, G., & Seow, P. S. (2021). Forensic analytics using cluster analysis: Detecting anomalies in data. Journal of Corporate Accounting & Finance, 32(2), 154. https://doi.org/10.1002/jcaf.22486
23. Grover, D., Bauhoff, S., & Friedman, J. (2019). Using supervised learning to select audit targets in performance-based financing in health: An example from Zambia. PLoS ONE, 14(1). https://doi.org/10.1371/journal.pone.0211262
24. Gruschka, N., Mavroeidis, V., Vishi, K., & Jensen, M. (2018). Privacy Issues and Data Protection in Big Data: A Case Study Analysis under GDPR. 2021 IEEE International Conference on Big Data (Big Data), 5027. https://doi.org/10.1109/bigdata.2018.8622621
25. Hao, J., & Ho, T. K. (2019). Machine Learning Made Easy: A Review of Scikit-learn Package in Python Programming Language [Review of Machine Learning Made Easy: A Review of Scikit-learn

Package in Python Programming Language]. Journal of Educational and Behavioral Statistics, 44(3), 348. SAGE Publishing. https://doi.org/10.3102/1076998619832248
26. Hoag, C. L., Heyman, J., Asdal, K., Reinertsen, H., & Hull, M. S. (2023). The Government Analytics Handbook: Leveraging Data to Strengthen Public Administration. https://doi.org/10.1596/978-1-4648-1957-5
27. Hopmere, M., Crawford, L., & Harré, M. (2020). Proactively Monitoring Large Project Portfolios. Project Management Journal, 51(6), 656. https://doi.org/10.1177/8756972820933446
28. Kanksha, Bhaskar, A., Pande, S. D., Malik, R., & Khamparia, A. (2021). An intelligent unsupervised technique for fraud detection in health care systems. Intelligent Decision Technologies, 15(1), 127. https://doi.org/10.3233/idt-200052
29. Khalid, H., Noor, A., Iqbal, J., Farid, S., & Chang, V. (2018). Development of public sector information management systems: challenges and promising practices. Information Discovery and Delivery, 46(3), 184. https://doi.org/10.1108/idd-03-2018-0008
30. Kostopoulos, G., Davrazos, G., & Kotsiantis, S. (2024). Explainable Artificial Intelligence-Based Decision Support Systems: A Recent Review [Review of Explainable Artificial Intelligence-Based Decision Support Systems: A Recent Review]. Electronics, 13(14), 2842. Multidisciplinary Digital Publishing Institute. https://doi.org/10.3390/electronics13142842
31. Kreuter, F., Ghani, R., & Lane, J. (2019). Change Through Data: A Data Analytics Training Program for Government Employees. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ed353ae3
32. Larsson, S., & Heintz, F. (2020). Transparency in artificial intelligence. Internet Policy Review, 9(2). https://doi.org/10.14763/2020.2.1469
33. Lee, J., Suh, T., Roy, D., & Baucus, M. S. (2019). Emerging Technology and Business Model Innovation: The Case of Artificial Intelligence. Journal of Open Innovation Technology Market and Complexity, 5(3), 44. https://doi.org/10.3390/joitmc5030044
34. Novita, N., & Anissa, A. I. N. A. (2022). The role of data analytics for detecting indications of fraud in the public sector. International Journal of Research in Business and Social Science (2147-4478), 11(7), 218. https://doi.org/10.20525/ijrbs.v11i7.2113
35. Oprea, S., Bâră, A., Puican, F., & Radu, I. C. (2021). Anomaly Detection with Machine Learning Algorithms and Big Data in Electricity Consumption. Sustainability, 13(19), 10963. https://doi.org/10.3390/su131910963
36. Othman, I., Ghani, S. N. M., Mohamad, H., Alalou, W. l S., & Shafiq, N. (2018). Early Warning Signs of Project Failure. MATEC Web of Conferences, 203, 2008. https://doi.org/10.1051/matecconf/201820302008
37. Palacio, S. (2018). Detecting Outliers with Semi-Supervised Machine Learning: A Fraud Prediction Application. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3165318
38. Pashami, S., Nowaczyk, S., Fan, Y., Jakubowski, J., Paiva, N., Davari, N., Bobek, S., Jamshidi, S., Sarmadi, H., Alabdallah, A., Ribeiro, R. P., Veloso, B., Sayed-Mouchaweh, M., Rajaoarisoa, L., Nalepa, G. J., & Gama, J. (2023). Explainable Predictive Maintenance. arXiv (Cornell University). https://doi.org/10.48550/arXiv.2306.05120
39. PedregosaFabian, VaroquauxGaël, GramfortAlexandre, MichelVincent, ThirionBertrand, GriselOlivier, BlondelMathieu, PrettenhoferPeter, WeissRon, DubourgVincent, VanderplasJake, PassosAlexandre, CournapeauDavid, BrucherMatthieu, PerrotMatthieu, & DuchesnayÉdouard. (2011). Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research. https://doi.org/10.5555/1953048.2078195
40. Pi, Y. (2021). Machine learning in Governments: Benefits, Challenges and Future Directions. JeDEM - eJournal of eDemocracy and Open Government, 13(1), 203. https://doi.org/10.29379/jedem.v13i1.625
41. Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Vaughan, J. W., & Wallach, H. (2018). Manipulating and Measuring Model Interpretability. arXiv (Cornell University). https://doi.org/10.48550/arXiv.1802.07810
42. Prasad, N., Bajpai, M., & Tripathi, A. (2023). The impact of budgetary control on organizational performance. International Journal of Research in Finance and Management, 6(2), 266. https://doi.org/10.33545/26175754.2023.v6.i2c.333
43. Radanliev, P., Santos, O., Brandon-Jones, A., & Joinson, A. (2024). Ethics and responsible AI deployment. Frontiers in Artificial Intelligence, 7. https://doi.org/10.3389/frai.2024.1377011

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XII Issue V May 2025

44. Sanderson, C., Douglas, D. C., & Lu, Q. (2023). Implementing Responsible AI: Tensions and Trade-Offs Between Ethics Aspects. 2022 International Joint Conference on Neural Networks (IJCNN), 1. https://doi.org/10.1109/ijcnn54540.2023.10191274
45. Schrage, B. (2018). Project Portfolio Management Can Ensure Best Use of Time and Resources. Natural Gas & Electricity, 35(1), 21. https://doi.org/10.1002/gas.22069
46. Sehatbakhsh, N., Daw, E., Savas, O., Hassanzadeh, A., & McCulloh, I. (2020). Security and Privacy Considerations for Machine Learning Models Deployed in the Government and Public Sector (white paper). arXiv (Cornell University). https://doi.org/10.48550/arxiv.2010.05809
47. Shah, F. H., Bhatti, O. S., & Ahmed, S. (2023). A Review of the Effects of Project Management Practices on Cost Overrun in Construction Projects [Review of A Review of the Effects of Project Management Practices on Cost Overrun in Construction Projects]. 15, 1. https://doi.org/10.3390/engproc2023044001
48. Wang, Q., Guangping, Z., & Tu, X. (2017). Information Technology Project Portfolio Implementation Process Optimization Based on Complex Network Theory and Entropy. Entropy, 19(6), 287. https://doi.org/10.3390/e19060287
49. Warner, R., & Sloan, R. H. (2021). Making Artificial Intelligence Transparent: Fairness and the Problem of Proxy Variables. Criminal Justice Ethics, 40(1), 23. https://doi.org/10.1080/0731129x.2021.1893932
50. West, D. M. (2023). Using AI and machine learning to reduce government fraud. https://www.brookings.edu/research/using-ai-and-machine-learning-to-reduce-government-fraud/
51. Willoughby, W. F. (1927). Principles of public administration. https://agris.fao.org/agris-search/search.do?recordID=US201300608555
52. Xu, X., Liu, H., & Yao, M. (2019). Recent Progress of Anomaly Detection. Complexity, 2019(1). https://doi.org/10.1155/2019/2686378
53. Yuan, S., & Wu, X. (2022a). Trustworthy Anomaly Detection: A Survey. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2202.07787
54. Yuan, S., & Wu, X. (2022b). Trustworthy Anomaly Detection: A Survey. arXiv (Cornell University). https://doi.org/10.48550/arXiv.2202.07787