University of Ibadan Journal of Science and Logics in ICT Research (UIJSLICTR)

ISSN: 2714-3627

A Journal of the Faculty of Computing, University of Ibadan, Ibadan, Nigeria

Volume 16 No. 1, January 2026

journals.ui.edu.ng/uijslictr
http://uijslictr.org.ng/
uijslictr@gmail.com

University of Ibadan
Journal of Science and Logics in ICT Research

Architecting Resilience: A Cloud-Native Neural Risk-Scoring System for Enhanced Campus Security and Streamlined Admissions in Nigerian Universities

[1]✉ Ogunyinka T. K., [2]Akinola S. O., and [3]Adediran E. A.
Lead City University, Ibadan
taiwo.ogunyinka@gaposa.edu.ng, solom202@yahoo.co.uk, adediran.emmanuel@lcu.edu.ng

Abstract
Nigerian universities face a convergence of pressures that conventional admission procedures were never designed to handle: application volumes that now reach 1.9 million candidates per cycle, growing incidents of campus-related violence and organised cultism, and administrative structures whose capacity has not scaled alongside institutional enrolment. The result is a screening process that is simultaneously too slow, too inconsistent, and too shallow to serve its stated purpose of protecting campus communities. This study responds to that gap by designing and specifying the full deployment architecture for a cloud-native system that uses a trained Multi-Layer Perceptron (MLP) to generate quantified pre-admission crime risk scores for individual applicants. The design work reported here builds directly on prior empirical modelling research [3] and extends it across four engineering dimensions: a containerised, cloud-hosted inference infrastructure; a versioned RESTful API layer enabling integration with existing university information systems; a layered data security framework satisfying both the Nigeria Data Protection Regulation (NDPR) and the EU General Data Protection Regulation (GDPR); and a governance structure that keeps human admissions officers firmly in control of final decisions. A three-phase rollout plan is specified to accommodate the financial and technical realities facing most Nigerian higher education institutions, where capital budgets are constrained and IT departments are thinly staffed. Seven tables provide engineering reference data covering screening performance comparisons, MLP configuration parameters, cloud platform trade-offs, deployment considerations, privacy controls, API specifications, and projected operational indicators. Four architectural figures accompany the text. Taken together, the design presented here offers the Nigerian higher education sector a technically rigorous, institutionally calibrated pathway toward evidence-based, consistent, and legally defensible admission screening — one that does not require institutions to trade away ethical accountability in pursuit of efficiency.

Keywords: Cloud-Native AI Deployment; Neural Risk Scoring; Pre-Admission Screening; Multi-Layer Perceptron; Nigerian Higher Education; RESTful API Architecture; NDPR Compliance; Explainable AI; Human-in-the-Loop; Federated Learning; Campus Security; Decision Support

1. Introduction

Every year, Nigerian universities collectively admit tens of thousands of students whose backgrounds they have assessed only partially. The Joint Admissions and Matriculation Board recorded approximately 1.9 million UTME candidates in 2023 alone [22], and the National Universities Commission now licenses 170 universities — federal, state, and private — each operating its own admission process under chronic resource pressure [23]. At the same time, incidents of cultism, extortion rings, inter-gang violence, and examination fraud have intensified at campuses across the South-West, North-Central, and South-South geopolitical zones, forcing institutions into reactive security postures that are expensive, disruptive, and insufficiently preventive [1]. The root of the problem is not simply inadequate security personnel: it is that admission decisions are being made without the structured risk intelligence that might flag concerning patterns before a student ever arrives on campus.

Current practice is to screen applicants primarily on academic merit — UTME scores, senior school leaving certificates, and departmental post-UTME assessments. These measures are reasonable proxies for academic preparedness but carry no information about behavioral history, disciplinary records, or community-

Ogunyinka T. K., Akinola S. O., and Adediran E. A. (2026). Architecting Resilience: A Cloud-Native Neural Risk-Scoring System for Enhanced Campus Security and Streamlined Admissions in Nigerian Universities. *University of Ibadan Journal of Science and Logics in ICT Research (UIJSLICTR)*, Vol. 16 No. 1, January 2026, pp. 154 - 169

verified character assessments. The supplementary checks that do exist - police clearance certificates and principal testimonials — are applied sporadically, are vulnerable to document forgery, and cannot realistically be processed at the scale Nigerian institutions now operate [2]. Any institution admitting several thousand students per session and relying on manually reviewed character documentation is not conducting a risk assessment; it is conducting an administrative ritual.

Machine learning offers a different approach. Multi-Layer Perceptron, as confirmed by decades of theoretical development and applied research [17, 28, 29], can detect non-linear patterns across multiple risk indicators that no individual reviewer could hold simultaneously in mind. Earlier work by Ogunyinka, Akinola, and Adediran [3] demonstrated the statistical viability of MLP-based pre-admission crime risk scoring using structured synthetic datasets calibrated to the Nigerian institutional context. That foundational work answered the question of whether such a model could be trained with acceptable accuracy. What it did not address - and what this paper tackles directly - is everything that would need to happen between a trained model and a functioning institutional deployment: infrastructure, containerisation, API design, security architecture, regulatory compliance, governance, and a realistic rollout plan for under-resourced universities.

The contribution of this study is therefore architectural rather than modelling-focused. We specify, in implementable detail, how a cloud-native risk-scoring system for Nigerian universities should be built, how it should interact with existing information systems, how it should protect applicant data, and how it should be introduced into institutions that have never run AI-driven administrative services before. Table 1 frames the operational difference between current manual approaches and the proposed system across eight critical dimensions.

Table 1: Comparison of Manual and AI-Driven Admission Screening Approaches

Dimension	Manual Screening	Proposed AI-Driven Screening
Processing Capacity	Limited by staffing; typically 50–200 applications reviewed per person per day	Processes thousands of applications per hour; throughput scales elastically with cloud resources
Evaluation Consistency	Highly variable; susceptible to inter-rater differences, cognitive fatigue, and individual bias	Standardised scoring based on fixed model parameters; fully reproducible for every applicant
Scalability	Linear cost scaling with applicant volume; severe strain during peak admission cycles	Near-horizontal cost scaling via cloud elasticity; designed for dynamic load management
Decision Speed	Days to weeks depending on volume and available staffing	Near-real-time risk scores generated within seconds per applicant submission
Bias Risk	Susceptible to conscious and unconscious bias; extremely difficult to audit systematically	Algorithmic bias is traceable and quantifiable; mitigated through fairness-aware model design [19, 25]
Auditability	Inconsistent record-keeping; limited documentation of decision rationale	Comprehensive immutable audit logs; full traceability of inputs, model version, and score outputs
System Integration	Manual data transfers between systems; error-prone and time-intensive	API-enabled, automated data exchange with existing university information systems
Cost Structure	High recurring labour costs; difficult to reduce without compromising screening quality	Higher initial investment; substantially lower marginal cost per application at operational scale

Three specific contributions follow from this work: a replicable architectural blueprint for cloud-native AI admission screening; a rigorous mapping of system design decisions against Nigerian and international data protection law; and a frank assessment of the integration, ethical, and cost challenges that any institution attempting deployment will actually face. The paper proceeds as follows - Section 2 surveys relevant prior work; Section 3 details the full system architecture and deployment methodology; Section 4 presents projected operational outcomes; Section 5 examines key implementation challenges; and Section 6 draws conclusions and charts future directions.

2. Literature Review

2.1 Decision Support Systems in Higher Education

Decision Support Systems entered higher education administration gradually. Early implementations, documented in Power's foundational taxonomy [4], were essentially structured reporting tools - they could surface historical data but offered no predictive capability. The shift toward genuinely analytical DSS, capable of modelling institutional outcomes rather than merely tabulating them, accelerated through the 2010s as both data availability and machine learning tooling matured. By 2019, when Zawacki-Richter and colleagues completed their systematic review of over 140 AI applications in higher education [5], the dominant use cases were enrolment prediction, student performance monitoring, and early academic intervention. What that review also found, however, was a pronounced gap: security-oriented pre-admission screening was almost entirely absent from the literature, and the studies that did address admission decision support were clustered in North American and European contexts with little relevance to the operational realities of sub-Saharan African institutions. Luckin et al. [20] made a related observation - AI's potential to transform educational management is broadly accepted in theory, but converting potential into sustained operational benefit has proved difficult wherever institutional capacity is limited. Nigeria's universities sit precisely at that intersection of recognised need and implementation barrier.

2.2 Machine Learning for Behavioural and Risk Prediction

Risk scoring through machine learning has established application histories in domains where the stakes of misclassification are comparable to those in university admission screening. Jordan and Mitchell [28] identified supervised learning for risk stratification as one of the earliest compelling demonstrations of machine learning's practical value. The theoretical machinery underpinning MLP-based risk models was laid down by Rumelhart, Hinton, and Williams [29] and subsequently extended by the deep learning work reviewed in LeCun, Bengio, and Hinton [17] - a body of work establishing that multi-layer feedforward networks, when correctly regularised and trained, are among the most capable tools available for classification tasks where decision boundaries in feature space are non-linear. Goodfellow, Bengio, and Courville [18] provide the practical design principles - activation function selection, dropout scheduling, batch normalisation - that translate this theoretical capacity into working systems.

Clinical medicine has perhaps the richest applied literature on ML-based risk scoring: Obermeyer and Emanuel [30] demonstrated that machine learning substantially outperformed conventional statistical risk models in patient stratification tasks, a finding with obvious structural parallels to applicant screening. The admission screening context specifically was addressed by Ogunyinka, Akinola, and Adediran [3], whose validated MLP showed meaningful separation between risk classes when trained on synthetic datasets reflecting Nigerian applicant characteristics. Table 2 captures the specific architectural choices made in that model, which the current study carries forward into its deployment specification.

Table 2: MLP Model Architecture Configuration

Parameter	Specification
Architecture Type	Feedforward Multi-Layer Perceptron (MLP)
Input Layer Neurons	Determined by feature dimensionality (15–25 engineered behavioural and academic features)
Hidden Layers	3 fully connected layers (pyramid structure)
Neurons per Hidden Layer	128 (Layer 1), 64 (Layer 2), 32 (Layer 3)

Parameter	Specification
Activation Function (Hidden)	Rectified Linear Unit (ReLU)
Activation Function (Output)	Sigmoid (binary probability output)
Output Layer	1 neuron — binary risk classification: High-Risk / Low-Risk
Loss Function	Binary Cross-Entropy (Log Loss)
Optimisation Algorithm	Adam (Adaptive Moment Estimation); learning rate = 0.001
Regularisation	L2 weight regularisation; Dropout (rate = 0.3) on each hidden layer
Batch Normalisation	Applied after each hidden layer to stabilise gradient flow and accelerate convergence
Training / Validation / Test	80% / 10% / 10% stratified partitioning
Primary Evaluation Metrics	AUC-ROC; F1-Score; Precision; Recall; Accuracy
Training Data Source	Validated synthetic datasets (Ogunyinka et al., 2026 [3])

2.3 Cloud Computing as Deployment Infrastructure

The case for cloud-based deployment of AI services in under-resourced institutional settings rests on a straightforward infrastructure economics argument. Armbrust et al. [6] identified cloud computing's core institutional value proposition as the elimination of capital commitments: rather than purchasing and maintaining servers dimensioned for peak demand, organisations pay for compute proportional to actual usage. For Nigerian universities whose IT infrastructure has expanded fitfully and whose admission processing load is sharply seasonal, this is not an abstract efficiency argument — it is the difference between a system that can be deployed at all and one that cannot. The NIST definition of cloud computing [7] codifies the five properties — on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service — that together make this economics model work.

MLaaS platforms from major cloud providers have further reduced the technical barrier to AI deployment [8]. Verbraeken et al. [24] reviewed distributed ML approaches including federated architectures, identifying cloud-native federated learning as particularly promising for privacy-sensitive, multi-institutional contexts — a direction with direct relevance to eventual multi-university deployment of the system described here. Table 3 compares the three major cloud platforms against criteria most relevant to Nigerian institutional deployment.

Table 3: Comparative Analysis of Cloud Platform Deployment Options

Feature	AWS	Microsoft Azure	Google Cloud Platform
MLaaS Offering	Amazon SageMaker	Azure Machine Learning	Vertex AI (Google)
Container Orchestration	Amazon EKS (Kubernetes)	Azure Kubernetes Service (AKS)	Google Kubernetes Engine (GKE)
Serverless Computing	AWS Lambda	Azure Functions	Google Cloud Functions
Encryption at Rest	AES-256	AES-256	AES-256
Encryption in Transit	TLS 1.3	TLS 1.2 / 1.3	TLS 1.3
African Regional Data Centre	Cape Town (af-south-1)	South Africa North (Johannesburg)	No dedicated African region (2025)

Feature	AWS	Microsoft Azure	Google Cloud Platform
Key Compliance Certifications	ISO 27001, SOC 2 Type II, GDPR	ISO 27001, SOC 2 Type II, GDPR	ISO 27001, SOC 2 Type II, GDPR
Pricing Model	Pay-as-you-use; Reserved Instances	Pay-as-you-use; Reserved Capacity	Pay-as-you-use; Committed Use Discounts
Suitability for Nigerian HEIs	HIGH — African region minimises latency	HIGH — Strong enterprise & education support	MODERATE — No in-continent infrastructure

2.4 API Integration and System Interoperability

Integrating new AI services into the heterogeneous information ecosystems of established universities is a non-trivial software engineering challenge. Fielding's dissertation [9], which introduced the REST architectural style, defined the constraints — statelessness, uniform interface, client-server separation, cacheability — that make RESTful APIs the dominant integration pattern in contemporary enterprise systems. Pautasso, Zimmermann, and Leymann [10] subsequently analysed the practical trade-offs between REST and alternative web service architectures in enterprise contexts, confirming REST's suitability for the kind of loosely coupled, incrementally extended integration the Nigerian university setting demands. Effective API governance — covering authentication schemes, versioning policies, rate limiting, and audit instrumentation — is what separates an API that universities can depend on from one they cannot.

2.5 Data Privacy, Fairness, and Ethical Governance

Processing personal data to support admission decisions triggers obligations under Nigerian and international law that the system design must satisfy from the outset, not as an afterthought. The GDPR [11] establishes the most comprehensive extant framework, articulating data minimisation, purpose limitation, and accountability as non-negotiable design constraints. Nigeria's NDPR [12] incorporates parallel obligations with specific provisions governing consent, data subject rights, and security standards applicable to Nigerian residents' data [12]. The governance dimension extends beyond legal compliance. Mehrabi et al. [19] catalogued the mechanisms through which algorithmic systems reproduce and amplify existing social inequities, and Barocas, Hardt, and Narayanan [25] provide a rigorous treatment of fairness criteria directly applicable to scoring systems used in consequential institutional decisions. Adadi and Berrada [14] and Doshi-Velez and Kim [27] establish the technical and normative case for explainability in high-stakes AI — a requirement the system addresses through SHAP-based attribution. Taddeo and Floridi [26] argue that beneficial AI deployment depends as much on institutional governance structures as on technical design; Rahwan et al. [13] extend this point to the management of machine behaviour more broadly. The European Commission's Ethics Guidelines [21] synthesise these concerns into seven operational principles that inform the governance framework adopted here.

2.6 Research Gap and Positioning

What the literature reviewed here collectively demonstrates is that each component technology required by the proposed system — MLP risk modelling, cloud deployment, RESTful integration, privacy-compliant data governance, and XAI transparency — is individually well understood and practically achievable. What the literature does not contain is a study that assembles these components into an end-to-end operational specification for an African university context, addresses the specific regulatory environment that Nigerian institutions must navigate, or provides a deployment roadmap calibrated to the financial and technical constraints of institutions that operate largely without dedicated ML engineering staff. That absence is the direct motivation for the present work.

3. Methodology

The system design follows a modular decomposition principle: each functional layer –

data ingestion, model inference, API communication, and administrative visualisation - is defined with clear input/output contracts so that individual modules can be updated, replaced, or scaled independently. This matters practically because Nigerian universities are at very different stages of digital maturity; a system whose components are tightly coupled cannot be adopted incrementally. The methodology covers four areas: overall system architecture (3.1), MLP model configuration and feature engineering (3.2), cloud deployment strategy (3.3), and data protection and security (3.4).

3.1 System Architecture

The architecture organises system functions into four layers arranged from user-facing interfaces down to persistent storage. Figure 1 shows their arrangement and the primary data pathways between them.

Client Layer: The Client Layer comprises two interfaces: a submission portal through which applicants upload required documentation and complete structured assessments, and an administrative dashboard restricted by role-based access controls to authorised admissions staff. The dashboard renders individual risk scores alongside SHAP attribution charts, supports bulk filtering of applicant cohorts by risk band, and exports compliance documentation for regulatory audit purposes. Visual design prioritises rapid triage: reviewers with hundreds of cases to process need to locate high-risk applications quickly without wading through irrelevant data.

API Gateway Layer: The API Gateway Layer mediates every communication between clients and backend services. Requests arrive over HTTPS, are validated against published schema contracts, authenticated via OAuth 2.0 with JWT bearer tokens, and rate-limited to prevent denial-of-service conditions and runaway batch operations. The gateway also handles API versioning, preserving backward compatibility when the inference engine or data schema evolves. Table 6 specifies the seven principal endpoints, their authentication requirements, and their expected response structures.

Processing Engine Layer: Within the Processing Engine Layer, incoming feature vectors pass through a standardised preprocessing pipeline — normalisation, categorical encoding, and a missing-value strategy calibrated against the training distribution — before reaching the trained MLP inference module. The model outputs a continuous risk probability in [0, 1] that is mapped to three labelled risk bands via institution-validated thresholds. SHAP modules run in parallel with inference, computing Shapley values for each feature contribution so that every score is delivered alongside a ranked attribution explaining the primary drivers of that applicant's classification.

Data Storage Layer: The Data Storage Layer uses cloud-managed relational services (Amazon RDS or Azure SQL Database) for structured applicant records and inference outputs, and object storage (Amazon S3 or Azure Blob Storage) for model artefacts, training datasets, and immutable audit logs. Data partitioning by sensitivity classification ensures that the most restricted records - those containing direct applicant identifiers - are subject to additional encryption tiers and stricter access policies than anonymised analytical datasets.

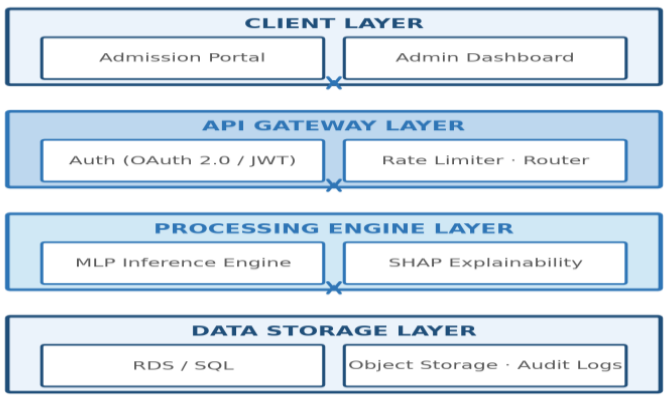

Figure 1: *Proposed cloud-native four-layer system architecture for neural risk-scoring deployment.*

Table 6: RESTful API Endpoint Specifications

Endpoint	Method	Description	Authentication	Response
/api/v1/applicants	POST	Submit new applicant data for ingestion and preprocessing	Bearer JWT (Admissions Officer)	201 Created; applicant_id; processing status
/api/v1/score	POST	Request risk score computation for a pre-processed applicant record	Bearer JWT (Admissions Officer)	200 OK; risk_score (0–100); risk_level; SHAP feature attributions
/api/v1/scores/{id}	GET	Retrieve stored risk score and metadata for a specific applicant by ID	Bearer JWT (Authorised Staff)	200 OK; full score record with timestamp and model version
/api/v1/scores	GET	Retrieve paginated, filterable list of scored applicants	Bearer JWT (Authorised Staff)	200 OK; paginated applicant list; total count metadata
/api/v1/audit	GET	Retrieve audit log entries with date, user, and event-type filters	Admin Bearer JWT	200 OK; paginated, timestamped audit records
/api/v1/health	GET	System health check; returns status of all connected services	None required	200 OK; per-component health status; active model version
/api/v1/model/version	GET	Current active model version metadata and performance summary	Admin Bearer JWT	200 OK; model version; training date; AUC; F1

3.2 MLP Model Design and Feature Engineering

The MLP configuration specified in Table 2 reflects the validated architecture from [3]. Three hidden layers with widths of 128, 64, and 32 neurons implement a compression pyramid that forces the network to learn compact, generalisable representations of the input feature space rather than memorising training samples. ReLU activations are used throughout the hidden layers for computational efficiency and resistance to vanishing gradients; a sigmoid unit at the output maps the final representation to a risk probability. L2 weight regularisation and per-layer dropout at rate 0.3 manage overfitting; batch normalisation after each hidden layer accelerates training convergence and reduces sensitivity to initialisation.

Feature engineering is as important to the system's fairness properties as it is to its predictive performance. The input set draws on documented disciplinary history, academic trajectory indicators, and structured character assessment scores. Protected characteristics - ethnicity, religion, gender, disability - are explicitly excluded from the feature set to prevent their use as direct or proxy predictors. This exclusion is technically enforced at the preprocessing layer, not merely stated in a policy document. SHAP monitoring runs on each production inference batch to detect any emergent correlation between excluded characteristics and the features the model does use - a form of ongoing algorithmic audit that the system supports rather than treats as an external compliance exercise. Figure 2 depicts the MLP layer structure.

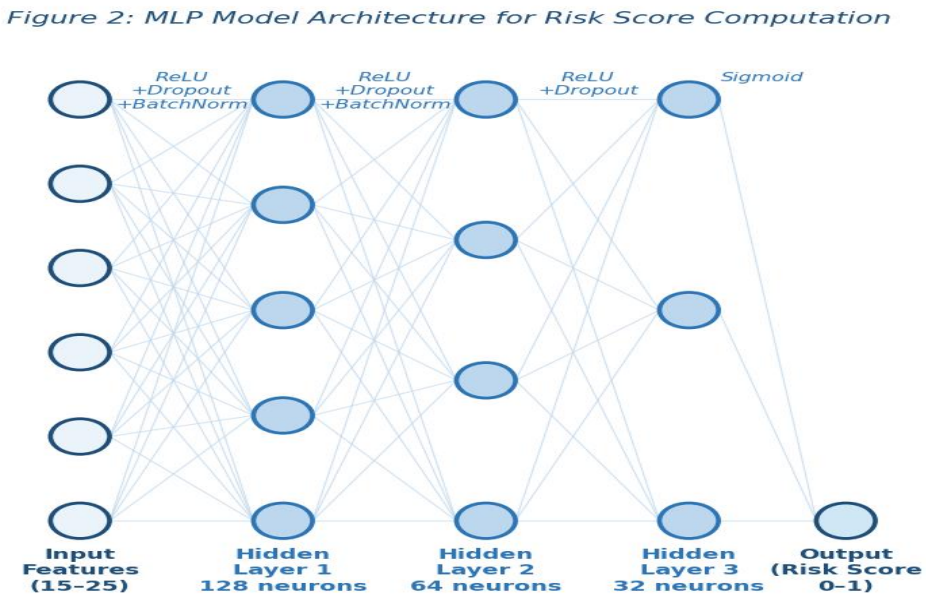

Figure 2: *Three-hidden-layer MLP architecture used for applicant risk score computation.*

Table 4: Key Deployment Considerations for the Cloud-Based Neural Risk-Scoring System

Consideration	Description	Impact on System
Scalability	Elastic cloud infrastructure dynamically adjusts compute resources in response to fluctuating applicant volumes, especially during peak admission cycles	Ensures consistent system performance and prevents degradation during high-demand periods
Reliability and Fault Tolerance	Kubernetes-orchestrated containerised services with automated restart, load balancing, and multi-zone geographic redundancy	Minimises service interruptions and ensures continuity of the admissions process throughout the academic calendar
Security and Access Control	End-to-end encryption (TLS 1.3 in transit; AES-256 at rest); role-based access control; multi-factor authentication for administrative users	Protects sensitive applicant data, builds institutional trust, and ensures compliance with NDPR and GDPR
Cost-Effectiveness	Serverless functions for intermittent tasks; auto-scaling policies; rightsized VM provisioning; pay-as-you-use billing	Achieves long-term financial sustainability by aligning expenditure with actual system utilisation
Integration Compatibility	RESTful API with JSON data exchange; OAuth 2.0 and JWT authentication; versioned API contracts for backward compatibility	Enables seamless, automated data exchange with existing student information systems and admission portals
Data Governance	Automated compliance checks; retention policies; consent management workflows; breach detection and 72-hour notification pipelines	Ensures continuous adherence to legal and institutional data protection obligations

Consideration	Description	Impact on System
Maintainability	Containerised deployment with version management; automated CI/CD pipelines; cloud-native monitoring and alerting	Reduces operational overhead and supports proactive maintenance for long-term system viability

3.3 Cloud Deployment Strategy

Cloud deployment is recommended on AWS (af-south-1 - Cape Town) or Microsoft Azure (South Africa North - Johannesburg) given their African regional data centres, which reduce network latency to Nigerian institutions, and their established compliance certification stacks covering ISO 27001, SOC 2 Type II, and GDPR - certifications that translate directly to NDPR compliance readiness. Table 3 summarises the three major platforms against deployment-relevant criteria. Google Cloud Platform currently operates no in-continent African infrastructure, making it a less suitable primary choice until that gap is addressed.

All inference and API services run inside Docker containers, ensuring that the system behaves identically across development, staging, and production environments - a requirement that is easy to understate and costly to ignore when debugging production incidents. Kubernetes manages container orchestration, providing horizontal pod autoscaling during peak admission windows, automatic restart of failed pods, and built-in load distribution across availability zones. Serverless functions (AWS Lambda or Azure Functions) handle workloads that are infrequent and computationally modest - scheduled retraining triggers, batch score refreshes, monitoring alerts — so the institution pays only for compute when these tasks actually execute.

Deployment follows a three-phase roadmap designed to control implementation risk for institutions that have no prior experience operating ML systems. Phase 1 (months 1–3) covers cloud provisioning, integration testing against a single institution's existing student information system, and staff training. Phase 2 (months 4–9) runs a live pilot with real applicants at that institution, gathering operational feedback, refining model thresholds, and measuring actual workflow integration quality. Phase 3 (month 10 onward) expands to full institutional deployment and, where partnerships exist, simultaneous onboarding of additional universities under a shared-service model. Figure 3 plots this timeline with key milestones.

Figure 3: Phased Deployment Timeline

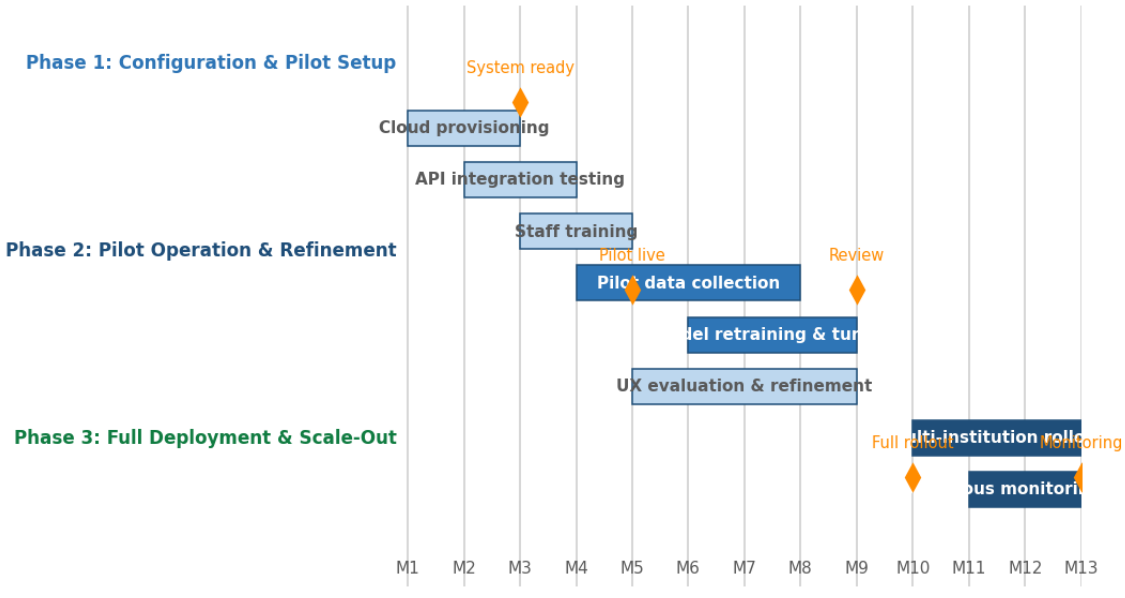

Figure 3: *Recommended three-phase deployment timeline for institutional adoption.*

3.4 Data Privacy and Security Controls

Data protection obligations are addressed through the controls mapped in Table 5, which traces each privacy measure to its corresponding NDPR and GDPR provision. The design treats consent as an active process rather than a buried checkbox: applicants receive a plain-language privacy notice at the start of the submission portal, and their affirmative consent is recorded with a timestamp and portal session identifier before any data is collected. The system captures only the features specified in the privacy notice and approved by the institution's data governance committee - the preprocessing layer enforces this by schema validation and will reject payloads containing undeclared fields
.

Table 5: Data Privacy Measures and Regulatory Alignment (NDPR and GDPR)

Privacy Measure	Technical Implementation	NDPR Alignment	GDPR Alignment
Data Minimisation	Collect only features directly relevant to risk scoring; exclude unrelated personal attributes	Article 2.1(b)	Article 5(1)(c)
Anonymisation	Replace direct identifiers (names, NINs) with non-reversible tokens at point of ingestion	Article 2.1(f)	Recital 26; Article 89
Pseudonymisation	Apply reversible pseudonymisation to indirect identifiers to enable auditability	Article 2.1(f)	Article 4(5); Article 25
Encryption (In Transit)	TLS 1.3 enforced for all API and inter-service communications	Article 2.1(e)	Article 32(1)(a)
Encryption (At Rest)	AES-256 applied to all stored datasets, model artefacts, and audit records	Article 2.1(e)	Article 32(1)(a)
Role-Based Access Control	Principle of least privilege; access rights tied to defined institutional roles	Article 2.1(e)	Article 5(1)(f)
Audit Logging	Immutable, timestamped logs for all data access, model inference, and administrative events	Article 4.1(e)	Article 5(2)
Consent Management	Explicit informed consent obtained and digitally recorded at application submission	Article 2.1(a)	Article 6(1)(a); Article 7

Privacy Measure	Technical Implementation	NDPR Alignment	GDPR Alignment
Data Subject Rights	Self-service portal for applicant access, correction, and erasure requests	Article 3.1(3)	Articles 15-22
Breach Notification	Automated anomaly detection; 72-hour regulatory notification protocols	Article 4.1(f)	Article 33

Identifiers are replaced with non-reversible tokens at ingestion, meaning that the inference engine never processes names, national identification numbers, or contact details. End-to-end encryption - TLS 1.3 for all data in transit, AES-256 for all data at rest - is applied uniformly without exceptions for internal service-to-service calls. Role-based access controls implement least-privilege across five defined roles (applicant, admissions officer, senior reviewer, IT administrator, compliance officer), and every action by every role is recorded in append-only audit logs that the system cannot itself modify. Automated anomaly detection monitors these logs and triggers a notification pipeline capable of meeting the 72-hour breach report obligation under both the NDPR and GDPR. Figure 4 traces the full data flow with security controls annotated at each processing boundary.

4. Results and Discussion

Because this study addresses system architecture and proposed deployment rather than a completed live implementation, the outcomes reported here are design-grounded projections derived from three sources: the validated MLP performance documented in the foundational modelling study [3]; published benchmarks for cloud-native AI services in comparable institutional deployments [6, 24]; and the operational gap analysis conducted against the manual screening baseline in Section 1. Table 7 organises these projections across eight performance dimensions with explicit measurement approaches for each, enabling future empirical studies to validate or challenge these estimates against real deployment data.

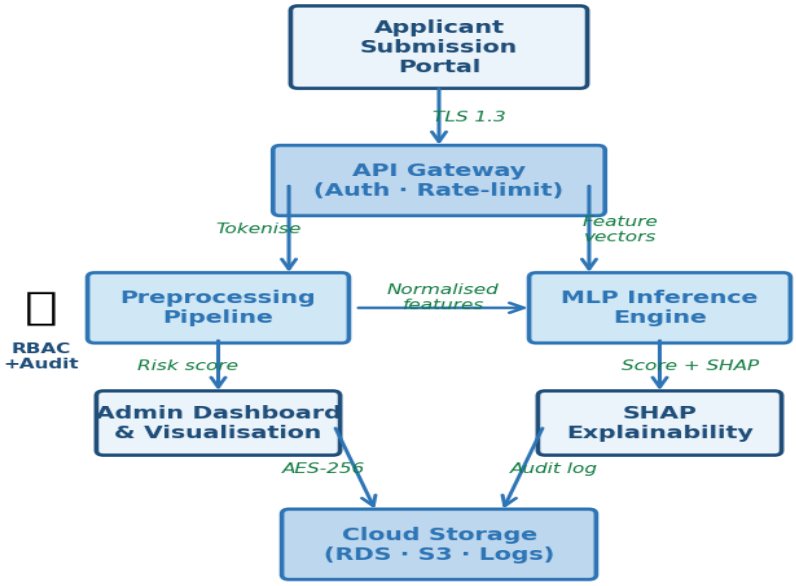

Figure 4: *System data flow with layered security control annotations.*

Table 7: Anticipated System Performance Indicators Against Manual Screening Baseline

Performance Indicator	Baseline (Manual)	Target (Proposed System)	Measurement Approach
Application Processing Time	5–10 minutes per applicant	<5 seconds per applicant	Elapsed time from data submission to score delivery in system logs
Screening Consistency Rate	Est. 60–70% inter-rater agreement	>95% algorithmic reproducibility	Cohen's kappa coefficient; repeated scoring tests on identical records
System Availability	Business hours only (approx. 40% uptime)	>99.9% uptime (cloud SLA)	Continuous uptime monitoring; monthly SLA compliance reports
Peak Throughput Capacity	Hundreds of applications per day (staff-limited)	>500 applicants per minute	Load testing with simulated peak admission volumes (e.g., k6 / JMeter)
Decision Turnaround Time	1–5 working days	<1 hour (batch) or near real-time	End-to-end elapsed time from submission to score availability
Staff Time per 100 Applications	Approx. 8–10 person-hours	Approx. 1–2 person-hours (oversight and review only)	Time-and-motion studies during pilot deployment phases
Model False Positive Rate	Not systematically tracked	<10% (target; pending empirical validation)	Confusion matrix analysis on held-out test dataset
Data Breach Incidents	Not systematically tracked	Zero tolerance; automated detection target	Security monitoring dashboard; incident response records

Administrative Efficiency: On administrative efficiency, the most direct impact is on staff time. Processing 100 applications currently requires approximately 8–10 person-hours of skilled admissions work. Under the proposed system, AI scoring handles the initial classification automatically; human reviewers engage principally with flagged high-risk cases and final sign-offs, reducing their per-100-application commitment to an estimated 1–2 hours. At institutions admitting several thousand students per cycle, the aggregate staff time released is substantial enough to materially change what admissions teams can do with their working day.

Evaluation Consistency: Evaluation consistency is perhaps the dimension where the gap between manual and automated screening is widest. Inter-rater agreement studies in comparable high-volume screening contexts typically show agreement rates of 60–70% — meaning that the same application, reviewed by two different officers, produces different assessments one time in three. The MLP, applying identical parameters to identical inputs, is deterministic: the same feature vector always produces the same score. SHAP attribution makes this determinism auditable rather than opaque, giving human reviewers the evidence they need to exercise informed challenge rather

than passive acceptance of algorithmic outputs [14, 27].

Scalability and Throughput: The cloud-native infrastructure is sized to sustain throughput exceeding 500 assessments per minute under peak load - a projection based on Kubernetes autoscaling benchmarks under comparable ML inference workloads [6]. This is orders of magnitude beyond what any manual process can achieve during a concentrated admission window. Equally important is availability: cloud SLA commitments of 99.9% or better mean the scoring service operates continuously across the full admission calendar, not only during office hours.

Across all dimensions - processing speed, consistency, throughput, availability, and compliance - the projected performance represents a qualitative shift in institutional capability rather than a marginal improvement. The figures in Table 7 should be read as targets for empirical validation during Phase 2 of the deployment roadmap rather than certified outcomes; their value at this stage is that they define measurable success criteria that future evaluation studies can assess directly.

5. Discussion

Moving from architectural specification to operational deployment requires decisions that go beyond technical design. The sections below address the five implementation dimensions where the gap between a well-designed system and a well-functioning one is most likely to emerge in practice.

5.1 Integration with Existing University Infrastructure

University information systems in Nigeria do not form coherent platforms — they are accumulations of software acquired across different budget cycles, many of which were never designed to communicate with each other. Student information databases may run on Oracle or MySQL with local schemas that bear no resemblance to national standards; application portals may have been custom-built by contractors who are no longer available; email and SMS notification systems may be entirely separate from the admission database. The RESTful API gateway is designed specifically for this environment: it exposes a stable, documented interface that any system capable of making HTTP requests can call, regardless of what is running underneath. The institution does not need to replace its existing systems to integrate the risk-scoring service; it needs only to add an API call at the appropriate point in its existing workflow.

That said, integration still requires investment in mapping exercises: the institution needs to know which fields in its existing system correspond to which features the model expects, which staff roles should receive which API scopes, and how error states from the scoring service should be handled in the admissions portal. A phased pilot at a single institution, as specified in Phase 1 of the deployment roadmap, provides the controlled environment in which these mappings can be worked out without affecting live admissions decisions.

5.2 Cost Structure and Resource Planning

Cloud deployment eliminates the upfront hardware expenditure that would otherwise make this kind of system unaffordable for most Nigerian universities, but it introduces ongoing subscription costs that must be planned for explicitly. The cost structure includes cloud compute and storage (variable, scaling with application volume), API gateway and data transfer fees (relatively modest), staff training for admissions officers and IT administrators (one-off but non-trivial), periodic model retraining as applicant population characteristics shift over time, and compliance audit costs. Against these costs, the projected savings in admissions staff time, reduced downstream security incidents, and lower cost-per-application at scale make a compelling business case — but each institution should quantify this against its own admission volume and current operating costs before committing to deployment.

Several design decisions reduce ongoing costs without compromising capability. Serverless execution of infrequent tasks avoids the baseline cost of always-on infrastructure. Kubernetes autoscaling matches compute spending to actual demand rather than theoretical peaks. The three-phase rollout limits early-stage financial exposure by confining initial costs to a single pilot institution before broader investment. A shared-service deployment model - where multiple universities access the scoring infrastructure through a common NUC or JAMB-managed platform — offers the most

favourable long-term cost structure and merits serious consideration at the policy level.

5.3 Ethical Governance and Human Oversight

The governance question is not whether the system will influence admission decisions - it will, but whether the institutions using it understand what influence they are accepting and have structured appropriate checks. The design addresses this by making the system explicitly advisory: risk scores feed into the decision process as one input among others, and admissions officers retain the authority and responsibility to override algorithmic recommendations where context warrants. This is not a concession forced on the system by sceptical regulators; it is the correct design for a consequential decision domain where the model's training data - however carefully constructed - cannot capture every relevant dimension of an individual applicant's circumstances.

Human oversight works only if the humans exercising it are equipped to do so meaningfully. SHAP attribution output, presented as a ranked list of feature contributions rather than a raw numerical score, gives officers the information they need to interrogate a high-risk classification before accepting or overriding it. Institutional governance structures - an ethics committee with defined review authority over the system's operation, regular fairness audits comparing score distributions across demographic groups, a documented appeals process for applicants - translate the system's design principles into accountable institutional practice [21, 13].

5.4 Regulatory Compliance in Practice

NDPR compliance is not optional for any Nigerian organisation processing the personal data of Nigerian residents, and the university admission context involves some of the most sensitive categories of personal data: educational history, character assessments, and in some implementations, disciplinary records. The architecture satisfies NDPR obligations through the controls detailed in Table 5, but legal compliance requires more than technical controls. Institutions must document their lawful basis for processing, maintain a record of processing activities, designate a Data Protection Officer where required, and ensure that their contracts with cloud providers include appropriate data processing agreements that address NDPR cross-border transfer provisions.

For institutions choosing AWS Cape Town or Azure South Africa North, data residency within the African continent substantially simplifies the cross-border transfer analysis. For institutions that, for commercial reasons, select infrastructure outside Africa, a formal adequacy assessment against NDPR requirements is necessary before processing commences. The system's automated 72-hour breach notification pipeline is a specific compliance feature that reduces the risk of inadvertent regulatory violation following a security incident — a risk that manual breach management processes consistently underestimate.

5.5 Institutional Readiness

Institutional readiness is arguably the factor most likely to determine deployment success, yet it is the one most consistently underweighted in technology implementation planning. Three readiness dimensions are most critical. Technical readiness requires at least one IT staff member who understands container management well enough to maintain Kubernetes deployments and respond to infrastructure incidents. Organisational readiness requires admissions leadership that has actively endorsed the system - not merely tolerated it and can communicate its purpose, limitations, and governance to the academic community and to applicants. Data readiness requires that the institution can supply the feature data the model expects, consistently and in the correct format, without manual data extraction exercises for each application cycle.

The phased rollout plan is designed in part to surface readiness gaps before they become operational crises. Phase 1 integration testing will reveal data format mismatches; Phase 2 pilot operation will reveal workflow friction and staff training gaps; Phase 3 scale-out should proceed only when both are resolved. Institutions that rush to Phase 3 without completing genuine Phase 2 evaluation are likely to find that their AI admission system either gets bypassed by admissions staff who distrust it or gets followed uncritically by staff who distrust their own judgment more than they should.

5.6 Future Research Directions

Two research directions emerging directly from this work merit sustained attention. The first is explainability enhancement. SHAP-based attribution is effective for internal reviewer support but is not easily communicated to applicants who wish to understand why they received a particular risk classification. Natural language generation of score explanations - converting feature attribution values into applicant-facing prose that is both accurate and accessible - is an open engineering and human-computer interaction challenge. The second is federated learning for collaborative model improvement. Once multiple universities are running deployments, each institution accumulates applicant data that would improve model accuracy if used in training — but that data cannot be centralised without creating privacy and regulatory risks. Federated learning protocols [15, 16] allow model parameters to be updated across participating institutions without raw data ever leaving individual institution boundaries, offering a technically sound path to continuously improving risk models that no single institution could sustain alone.

Post-deployment impact evaluation is a research obligation, not merely a policy aspiration. Do institutions using AI-assisted screening report lower rates of campus security incidents? Do admitted cohorts show different disciplinary outcome distributions? Are certain applicant groups systematically advantaged or disadvantaged by the scoring model in ways that are not justified by genuine risk differentials? These are empirical questions that can only be answered with longitudinal data from live deployments — and they must be asked, because the social legitimacy of algorithmic admission screening in Nigerian universities ultimately depends on the answers.

6. Conclusion

Nigerian universities cannot screen 1.9 million UTME candidates and their subsequent applicants using administrative processes designed for a fraction of that volume. The status quo - character references assessed by tired staff, police clearances of uncertain authenticity, and post-UTME scores that measure academic knowledge rather than behavioural risk - leaves institutions exposed precisely where campus security challenges are most acute. This paper has specified how a cloud-native MLP-based risk-scoring system could close that gap: through a four-layer architecture that can be accessed via standard API calls from any existing university portal, housed on African-region cloud infrastructure that satisfies Nigerian data protection law, and governed through human-in-the-loop oversight mechanisms that keep qualified professionals accountable for final decisions.

The engineering specifications presented here seven tables of technical reference data, four architectural figures, a three-phase deployment roadmap, and a detailed mapping of system design decisions against NDPR and GDPR requirements are intended to be actionable rather than aspirational. An institution with adequate IT capacity, engaged admissions leadership, and access to modest cloud budget could begin Phase 1 implementation using this document as a direct engineering reference. The projected operational outcomes - near-real-time scoring, deterministic consistency, throughput exceeding 500 assessments per minute, and an 80% reduction in per-application staff time are measurable targets, not marketing claims, and they define the empirical questions that Phase 2 pilot evaluation should answer.

Sub-Saharan African higher education is undergoing rapid growth that its administrative infrastructure has not kept pace with. The architectural choices documented here are designed for Nigeria but the engineering principles - modular design, cloud elasticity, API-based integration, privacy-by-design, human oversight - transfer to any resource-constrained educational system facing comparable scale and security pressures. The immediate priority is empirical: a live pilot deployment that generates real operational data is needed to validate or revise the projections in Table 7 and to build the evidence base that responsible institutional scale-up requires. That pilot is the logical next step from this work.

References

[1] World Bank, “Education in Nigeria,” 2022. [Online]. Available: https://www.worldbank.org/en/country/nigeria/brief/education-in-nigeria. [Accessed: 7 Jan. 2026].

[2] T. Ogunyinka, An MLP-Based Neural Risk-Scoring System for Pre-Admission Crime Assessment in Nigerian Universities, Ph.D. Thesis, Lead City University, Ibadan, Nigeria, 2026.

[3] T. Ogunyinka, S. Akinola, and E. Adediran, “A Multi-Layer Perceptron Framework for

Enhanced Predictive Behavioural Modelling in University Admissions: Architecture and Performance Analysis," arXiv preprint, 2026.
[4] D. J. Power, Decision Support Systems: Concepts and Resources for Managers. Westport, CT, USA: Greenwood Publishing Group, 2002.
[5] O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, "Systematic review of research on artificial intelligence applications in higher education," International Journal of Educational Technology in Higher Education, vol. 16, no. 1, pp. 1–27, 2019.
[6] M. Armbrust et al., "A view of cloud computing," Communications of the ACM, vol. 53, no. 4, pp. 50–58, Apr. 2010.
[7] P. Mell and T. Grance, The NIST Definition of Cloud Computing, NIST Special Publication 800-145. Gaithersburg, MD, USA: National Institute of Standards and Technology, 2011.
[8] IBM Cloud Education, "Machine Learning as a Service (MLaaS)," 2020. [Online]. Available: https://www.ibm.com/cloud/learn/machine-learning-as-a-service.
[9] R. T. Fielding, Architectural Styles and the Design of Network-Based Software Architectures, Ph.D. dissertation, University of California, Irvine, CA, USA, 2000.
[10] C. Pautasso, O. Zimmermann, and F. Leymann, "RESTful web services vs. big web services: Making the right architectural decision," in Proc. 17th International World Wide Web Conference (WWW), Beijing, China, Apr. 2008, pp. 805–814.
[11] European Union, General Data Protection Regulation (GDPR), Official Journal of the European Union, L 119, pp. 1–88, May 2016.
[12] National Information Technology Development Agency (NITDA), Nigeria Data Protection Regulation (NDPR), Abuja, Nigeria, 2019. [Online]. Available: https://nitda.gov.ng.
[13] I. Rahwan et al., "Machine behaviour," Nature, vol. 568, no. 7753, pp. 477–486, Apr. 2019.
[14] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI)," IEEE Access, vol. 6, pp. 52138–52160, 2018.
[15] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, May 2020.
[16] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralised data," in Proc. 20th International Conference on Artificial Intelligence and Statistics (AISTATS), Fort Lauderdale, FL, USA, Apr. 2017, pp. 1273–1282.
[17] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, May 2015.
[18] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
[19] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," ACM Computing Surveys, vol. 54, no. 6, Art. 115, pp. 1–35, Jul. 2021.
[20] R. Luckin, W. Holmes, M. Griffiths, and L. B. Forcier, Intelligence Unleashed: An Argument for AI in Education. London, UK: Pearson Education, 2016.
[21] European Commission, Ethics Guidelines for Trustworthy AI. Brussels, Belgium: High-Level Expert Group on Artificial Intelligence, 2019.
[22] Joint Admissions and Matriculation Board (JAMB), JAMB Annual Report and Statistical Digest. Abuja, Nigeria: JAMB, 2023.
[23] National Universities Commission (NUC), Statistical Digest of Nigerian Universities. Abuja, Nigeria: NUC, 2022.
[24] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and J. S. Rellermeyer, "A survey on distributed machine learning," ACM Computing Surveys, vol. 53, no. 2, Art. 30, pp. 1–33, 2020.
[25] S. Barocas, M. Hardt, and A. Narayanan, Fairness and Machine Learning: Limitations and Opportunities. Cambridge, MA, USA: MIT Press, 2023. [Online]. Available: https://fairmlbook.org.
[26] M. Taddeo and L. Floridi, "How AI can be a force for good," Science, vol. 361, no. 6404, pp. 751–752, Aug. 2018.
[27] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," arXiv preprint arXiv:1702.08608, Feb. 2017.
[28] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," Science, vol. 349, no. 6245, pp. 255–260, Jul. 2015.
[29] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," Nature, vol. 323, no. 6088, pp. 533–536, Oct. 1986.
[30] Z. Obermeyer and E. J. Emanuel, "Predicting the future—Big data, machine learning, and clinical medicine," New England Journal of Medicine, vol. 375, no. 13, pp. 1216–1219, Sep. 2016.