



AI-Driven Optimization of Sustainable Soil Stabilization: A Kernel Learning Approach for Makuba-Treated Lateritic Pavements

Daha S. Aliyu^{1,5*}, Hafizu Hamza Ali², Haruna Ibrahim³, Mahmoud M. Farouq^{4,5}, Z. B. Baba⁵

¹School of Engineering, Newcastle University Upon Tyne, NE1 7RU, United Kingdom

²School of Technology, Kano State Polytechnic, Nigeria

³Department of Civil Engineering, Federal University Wukari, Taraba State

⁴School of Architecture, University of Nottingham, United Kingdom

⁵Department of Civil Engineering, Aliko Dangote University of Science and Technology, Wudil

*Corresponding Author Email: engrshanono@kustwudil.edu.ng

Abstract

The stabilization of lateritic soils using agricultural waste materials, such as locust bean pod ash (LBPA) locally known as "Makuba" presents a sustainable alternative for improving pavement subgrades. However, predicting optimal stabilizer dosage and geotechnical performance remains challenging due to complex nonlinear interactions between soil composition, stabilizer chemistry, and compaction conditions. This preliminary proof-of-concept study develops a physics-guided, stabilization-anchored kernel machine learning framework to predict unconfined compressive strength (UCS), California bearing ratio (CBR), and plasticity index (PI) of Makuba-stabilized lateritic soil. The approach integrates experimental data (six mix proportions: 0–12% Makuba) with material science principles through a stabilization-anchored feature set incorporating clay fraction, linear shrinkage, CaO, and SiO₂ content. Three kernel regression models Gaussian Process Regression (GPR), Kernel Ridge Regression (KRR), and Support Vector Regression (SVR) were implemented with hyperparameters optimized via grouped 5-fold cross-validation to prevent data leakage. Results demonstrate that GPR with stabilization-anchored features achieves strong interpolation accuracy ($R^2 = 0.982$, RMSE = 6.78 kN/m² for UCS), significantly outperforming baseline models. Feature importance analysis confirms alignment with geotechnical principles, identifying Makuba content, clay fraction, and CaO as the most influential predictors. However, stabilization-band hold-out testing reveals a critical limitation: models trained on low Makuba content (0–6%) fail to generalize to high-content mixtures (9–12%), underscoring extrapolation risks. Uncertainty quantification via GPR provides predictive confidence intervals for risk-informed decision-making. The framework is proposed as a preliminary screening-level tool for dosage optimization and laboratory trial reduction. While promising, results warrant cautious interpretation and emphasize the need for expanded experimental validation prior to field implementation.

Keywords: Soil stabilization; lateritic soil; locust bean pod ash; kernel machine learning; sustainable construction; artificial intelligence

1. Introduction

Lateritic soils are extensively used in road construction across tropical regions due to their widespread availability. However, their variable engineering properties such as low bearing capacity, high plasticity, and susceptibility to moisture often necessitate stabilization to meet design standards for pavement layers [1,2]. Traditional stabilizers like cement, lime, and bitumen are effective but come with high costs, carbon footprints, and limited accessibility in developing regions [3,4]. Consequently, the use of

agricultural waste materials such as rice husk ash, sugarcane bagasse ash, and locust bean pod ash (LBPA, locally termed “Makuba”) has emerged as a sustainable alternative for soil improvement [5–7]. Makuba, a byproduct of African locust bean processing, is abundant in northern Nigeria and exhibits pozzolanic properties due to its silica and alumina content, which react with calcium hydroxide to form cementitious compounds [8,9]. Previous studies have demonstrated that Makuba addition improves unconfined compressive strength (UCS), reduces plasticity, and enhances the California bearing ratio (CBR) of lateritic soils [10,11]. However, the relationship between stabilizer content and soil behavior is nonlinear and influenced by multiple factors including initial soil gradation, mineralogy, moisture content, and compaction energy [12,13]. Conventional empirical models and rule-of-thumb approaches often fail to capture these complex interactions, leading to either overconservative designs or unexpected field failures [14].

Machine learning (ML) offers a promising paradigm for modeling such multidimensional, nonlinear systems in geotechnical engineering [15–17]. Kernel-based ML methods, including Gaussian Process Regression (GPR), Kernel Ridge Regression (KRR), and Support Vector Regression (SVR), are particularly suited for small-sample datasets due to their strong regularization foundations and ability to model complex patterns without overfitting [18–20]. Recent applications of ML in geotechnics include prediction of soil compressibility, shear strength parameters, landslide susceptibility, and stabilization outcomes [21–24]. However, many existing models remain purely data-driven, lacking integration of domain-specific physical principles, which can result in physically implausible predictions when extrapolating beyond training conditions [25,26].

The proposed framework is positioned as a preliminary screening-level predictive tool that can assist engineers in optimizing stabilizer dosage, reducing laboratory trials, and promoting the use of sustainable local materials in infrastructure projects. This study bridges this gap by proposing a physics-guided, stabilization-anchored kernel machine learning framework that explicitly incorporates geotechnical and material science knowledge into feature engineering, model training, and validation. The primary objectives are:

1. To develop benchmark kernel regression models (GPR, KRR, SVR) for predicting UCS, CBR, and plasticity index of Makuba-stabilized lateritic soil.
2. To introduce a stabilization-anchored feature set that includes soil composition indices and chemical oxide content derived from Makuba.
3. To validate model performance using a two-tier strategy: conventional cross-validation and stabilization-band hold-out testing.
4. To conduct sensitivity and uncertainty analyses to assess model robustness under input variability.

The proposed framework is positioned as a screening-level predictive tool that can assist engineers in optimizing stabilizer dosage, reducing laboratory trials, and promoting the use of sustainable local materials in infrastructure projects. It is important to note that this dataset, while comprehensive in its characterization, comprises only six mix proportions. This study is therefore presented as a preliminary proof-of-concept investigation, demonstrating the potential of kernel learning approaches for sustainable soil stabilization modeling.

2. Materials and Methods

2.1 Experimental Dataset

The experimental dataset was sourced from a published laboratory investigation by Aliyu et al. (2023) [10], which examined the stabilization of reddish-brown lateritic soil with powdered locust bean pod ash (Makuba). Makuba was incorporated at five discrete proportions: 0%, 3%, 4%, 6%, 9%, and 12% by dry weight of the soil. A comprehensive suite of standardized geotechnical tests was performed on each mixture to characterise its engineering behavior. These tests included particle size distribution via sieve and hydrometer analysis, determination of Atterberg limits (Liquid Limit, Plastic Limit, and Plasticity Index) in accordance with ASTM D4318, standard Proctor compaction to establish the Maximum Dry Density (MDD) and Optimum Moisture Content (OMC), unconfined compressive strength (UCS) tests on cured cylindrical specimens, and California bearing ratio (CBR) tests on compacted samples soaked for 96 hours. The resulting dataset, encapsulating the multivariate response of the soil-stabilizer system, is consolidated in Table 1. The table presents six mix proportions, including the 4% dosage which was included to provide additional resolution around the optimal stabilizer range identified in preliminary testing. It provides the foundational experimental matrix relating the input stabilizer dosage to key performance outputs, namely UCS, CBR, and index properties.

Table 1: Experimental dataset of Makuba-stabilized lateritic soil.

Makuba (%)	LL (%)	PL (%)	PI (%)	LS (%)	MDD (kN/m ³)	OMC (%)	UCS (kN/m ²)	CBR (%)
0	36.6	25.8	10.8	6.43	18.73	8.89	124	21
3	38.2	26.5	11.7	5.21	18.45	9.34	158.4	24.5
4	39.1	27	12.1	4.56	18.28	9.61	178.6	26.4
6	40.1	27.3	12.8	3.87	18.12	9.88	198.7	28.3
9	42.3	28.4	13.9	2.65	17.84	10.45	236	32.1
12	43.8	29.1	14.7	2.14	17.52	11.2	184	29.8

LL: Liquid Limit, PL: Plastic Limit, PI: Plasticity Index, LS: Linear Shrinkage, MDD: Maximum Dry Density, OMC: Optimum Moisture Content.

2.2 Feature Engineering and Stabilization Anchoring

To effectively inject domain-specific knowledge into the machine learning pipeline and move beyond a purely correlative model, employing a physics-guided approach two distinct feature sets were engineered from experimental data. The first, termed the Baseline Feature Set (X_{base}), comprises the most measured and readily available geotechnical parameters in preliminary site investigations. It is defined as:

$$X_{\text{base}} = [\text{Makuba} (\%), \text{LL} (\%), \text{PL} (\%), \text{MDD} (\text{kN/m}^3), \text{OMC} (\%)] \quad (1)$$

where LL is Liquid Limit, PL is Plastic Limit, MDD is Maximum Dry Density, and OMC is Optimum Moisture Content.

Recognizing that the efficacy of pozzolanic stabilization is governed by fundamental material interactions, a second, more physically informative Stabilization-Anchored Feature Set (X_{stab}) was constructed [27,28]. This set augments the baseline features with descriptors of soil texture and the chemical composition of the stabilizer, which are critical to the pozzolanic reaction kinetics. It is expressed as:

$$X_{\text{stab}} = [\text{Makuba} (\%), \text{LL} (\%), \text{PL} (\%), \text{LS} (\%), \text{Clay} (\%), \text{Silt} (\%), \text{Sand} (\%), \text{CaO} (\%), \text{SiO}_2 (\%)] \quad (2)$$

Here, LS denotes Linear Shrinkage, and Clay, Silt, and Sand represent the particle size fractions. The chemical features, Calcium Oxide (CaO) and Silicon Dioxide (SiO₂) content, were derived from X-ray fluorescence (XRF) spectroscopy data of Makuba reported in antecedent studies AAY, 2010; JYAKS & Adama, 2013 [8,9]. These oxides are the primary reactants in the formation of cementitious calcium silicate hydrate (C-S-H) gels, thus providing a direct link to the underlying stabilization mechanism. All features in both sets were normalized to a [0, 1] range using min-max scaling to ensure numerical stability and equitable weighting during kernel computation.

2.3 Kernel Regression Models

Three kernel-based machine learning models were implemented to capture the nonlinear relationships between the engineered features and the target geotechnical properties (UCS, CBR, PI). The fundamental principle of kernel methods is to map the original input features into a higher-dimensional Hilbert space using a kernel function $k(x_i, x_j)$, where linear regression can be performed, effectively modeling complex nonlinearities in the original space.

Gaussian Process Regression (GPR) is a non-parametric, Bayesian approach. It defines a prior distribution over functions, if any finite set of function values

$$f = [f(x_1), \dots, f(x_n)]^T \quad (3)$$

follows a multivariate Gaussian distribution: $f \sim N(0, K)$. Here, K is the covariance matrix with entries

$$K_{ij} = k(x_i, x_j) + \sigma_n^2 \delta_{ij} \quad (4)$$

where σ_n^2 is the noise variance and δ_{ij} is the Kronecker delta. The Radial Basis Function (RBF) kernel was employed [30].

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2) \quad (5)$$

Given training data, the posterior predictive distribution for a new input $\mathbf{x}_*$ is Gaussian with mean and variance given by:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \quad (6)$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_* \quad (7)$$

where $\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1), \dots, k(\mathbf{x}_*, \mathbf{x}_n)]^T$. GPR was selected for its intrinsic uncertainty quantification via the predictive variance $\sigma^2(\mathbf{x}_*)$, providing confidence intervals crucial for engineering decision-making.

Kernel Ridge Regression (KRR) combines kernel mapping with ridge regression (L2 regularization). It solves the minimization problem: $\min_{\alpha} \|\mathbf{y} - \mathbf{K}\alpha\|^2 + \lambda \alpha^T \mathbf{K} \alpha$, where λ is the regularization parameter. The solution is given by

$$\alpha = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{y} \quad (8)$$

$$\text{and predictions are made as } \hat{y}(\mathbf{x}_*) = \mathbf{k}_*^T \alpha \quad (9)$$

KRR offers a closed-form solution and efficient handling of multi-collinearity, making it stable for small datasets [31].

Support Vector Regression (SVR) aims to find a function

$$f(\mathbf{x}) = \langle \mathbf{w}, \phi(\mathbf{x}) \rangle + b \quad (10)$$

that deviates from the actual targets y_i by at most ϵ while remaining as flat as possible. Using the ϵ -insensitive loss and the kernel trick, the dual optimization problem involves maximizing the Lagrangian. The final prediction function is

$$f(\mathbf{x}) = \sum_{i=1}^n (\alpha_i - \alpha_i^*) k(\mathbf{x}_i, \mathbf{x}) + b \quad (11)$$

where α_i, α_i^* are Lagrange multipliers. SVR is known for robustness and sparsity, as only a subset of training points (support vectors) contributes to the model [32].

The GPR was anticipated to outperform KRR and SVR in this context because its probabilistic framework not only provides point predictions but also quantifies epistemic uncertainty arising from limited data, a key concern with only six experimental points. Furthermore, its non-parametric nature allows it to adapt more flexibly to the underlying data distribution without pre-specifying a functional form.

The mathematical formulation for the three kernel models is detailed below to substantiate their application:

2.4 Gaussian Process Regression (GPR):

A Gaussian Process is fully defined by its mean function $m(\mathbf{x})$ and covariance (kernel) function $k(\mathbf{x}, \mathbf{x}')$. Assuming a zero-prior mean, the joint distribution of observed targets $\mathbf{y}$ and predicted target f_* at test point $\mathbf{x}_*$ is:

$$\begin{bmatrix} y \\ f_* \end{bmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X, x_*) \\ K(x_*, X) & k(x_*, x_*) \end{bmatrix} \right) \quad (12)$$

Using the conditional properties of Gaussians, the predictive (posterior) distribution is:

$$p(f_* | X, y, x_*) = \mathcal{N}(\mu_*, \sigma_*^2) \quad (13)$$

with

$$\mu_* = k_*^T [K + \sigma_n^2 I]^{-1} y, \sigma_*^2 = k(x_*, x_*) - k_*^T [K + \sigma_n^2 I]^{-1} k_* \quad (14)$$

2.5 Kernel Ridge Regression (KRR):

KRR solves the regularized least squares problem in the feature space. The dual solution coefficients α are obtained by:

$$\alpha = (K + \lambda I)^{-1} y \quad (15)$$

where λ is the regularization parameter. The prediction is:

$$\hat{y}(x_*) = \sum_{i=1}^n \alpha_i k(x_i, x_*) \quad (16)$$

2.6 Support Vector Regression (SVR):

SVR solves the following convex optimization problem:

$$\begin{aligned} \min_{w, b, \xi, \xi^*} \quad & \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n (\xi_i + \xi_i^*) \\ \text{s.t.} \quad & y_i - \langle w, \phi(x_i) \rangle - b \leq \epsilon + \xi_i, \\ & \langle w, \phi(x_i) \rangle + b - y_i \leq \epsilon + \xi_i^*, \\ & \xi_i, \xi_i^* \geq 0, i = 1, \dots, n. \end{aligned} \quad (17)$$

The kernel trick allows $\langle \phi(x_i), \phi(x_j) \rangle = k(x_i, x_j)$. The solution yields a predictive function:

$$f(x) = \sum_{i=1}^n (\alpha_i - \alpha_i^*) k(x_i, x) + b \quad (18)$$

where α_i, α_i^* are Lagrange multipliers, nonzero only for support vectors.

2.7 Why GPR is Superior for This Study:

With only $n=6$ original data points, the data sparsity is the primary challenge. GPR inherently handles this through its Bayesian framework, which quantifies uncertainty (σ_*^2) directly. This tells the engineer not just a prediction, but how confident the model is. KRR and SVR provide only point predictions. Furthermore, GPR's hyperparameters (like noise variance σ_n^2 and length scale) are often estimated via marginal likelihood maximization, a form of automatic regularization adapted to the data. This built-in adaptability to small datasets, coupled with explicit uncertainty estimates, makes GPR the most informative and reliable choice for this application, despite its higher computational cost for large n .

2.8 Bounded Synthetic Augmentation

The extremely limited size of the original experimental dataset ($n=6$) poses a significant challenge for kernel methods, as it can lead to ill-conditioned kernel matrices prone to numerical instability during inversion. To mitigate this without introducing unphysical data, a bounded synthetic augmentation protocol was employed as a computational regulariser. This augmentation does not represent new physical evidence but serves to numerically stabilize the inversion of the kernel matrix K ... It is

emphasized that this approach is presented as a preliminary methodological demonstration, not a substitute for comprehensive experimental characterization. For each of the six original samples, a synthetic neighborhood was generated by applying a Gaussian perturbation with zero mean and a standard deviation of 2.5% of the experimental range for each feature [33].

To prevent data leakage during validation, all synthetic samples originating from the same original observation were assigned a unique group identifier (G1–G6). These groups were kept intact during cross-validation such that synthetic samples from a given original observation never appeared in both the training and test sets of any fold. This grouped approach ensures that the model is never tested on synthetic variations of the same experimental observation it was trained on, providing an honest assessment of predictive performance. Perturbations were bound to a maximum of $\pm 5\%$ of the range to ensure all synthetic samples remained within the physically plausible domain defined by the original experiment. This process expanded the dataset to 60 samples. Crucially, this augmentation does not represent new physical evidence but serves to numerically stabilize the inversion of the kernel matrix K in GPR and KRR by improving its conditioning a concept aligned with adding a small jitter term to the diagonal, which is a common practice in Gaussian process modeling to handle singular matrices [34]. For the Stabilization-Band Hold-Out (SBHO) testing tier, a different augmentation strategy is employed: augmentation is performed only after the train-test split, and only on the training data. The test band remains completely unaugmented. This dual strategy (grouped CV for interpolation testing, post-split augmentation for extrapolation testing) ensures robust protection against data leakage in all validation scenarios.

2.9 Validation Strategy

A robust, two-tier validation strategy was designed to rigorously assess model performance under different scenarios, addressing both interpolation and extrapolation capabilities.

The 5-Fold Cross-Validation (CV) on the Augmented Dataset: This tier evaluates the model's interpolation accuracy within the experimental domain. To prevent data leakage, a grouped (or "blocked") cross-validation strategy was employed. The six original observations (and their associated synthetic neighborhoods) were treated as distinct groups. Each fold was constructed such that entire groups, not individual samples were held out for testing. This ensures that synthetic samples derived from the same original experimental point never appear in both training and test sets simultaneously. The augmented dataset ($n=60$) was randomly partitioned into five folds. The model was trained on four folds and tested on the held-out fold, repeating the process five times. Performance metrics were averaged across all folds. This grouped approach ensures that synthetic samples derived from the same original experimental point never appear in both training and test sets simultaneously, eliminating the risk of artificially inflated accuracy. This method provides a reliable estimate of how well the model can predict outcomes for new mixtures with parameters lying within the range of the training data.

Stabilization-Band Hold-Out (SBHO) Testing on Original Data: This tier assesses generalization under a regime shift, simulating a practical scenario where a model calibrated on low stabilizer content must predict behavior at high content, or vice versa. Since this validation uses only the original six experimental points (not augmented samples), data leakage is not a concern, each observation is unique and belongs to only one band. Models were trained exclusively in one band and tested on the other. The original six data points were split into two distinct stabilization "bands": the Low Band (0%, 3%, 4%, 6% Makuba) and the High Band (9%, 12% Makuba). For each training band, synthetic augmentation was applied exclusively to the training data using the bounded Gaussian perturbation protocol described in Section 2.8. The test band data remained completely unaugmented. This ensures that no information from the test band, either directly or through synthetic variations influences model training. Models were trained exclusively in one band and tested on the other. This "extrapolation" test is critical for evaluating the physical meaningfulness and reliability of the model beyond its immediate training range, revealing the limitations of purely data-driven approaches and the value of physics-anchored features.

2.10 Performance Metrics

To quantitatively evaluate and compare the predictive prowess of the developed kernel models, three standard performance metrics were employed. These metrics collectively offer a comprehensive view of model accuracy, scale of error, and robustness. Coefficient of Determination (R^2): This metric represents the proportion of variance in the observed data that is predictable from the input features. An R^2 value of 1 indicates perfect prediction, while a value of 0 suggests the model performs no better than predicting the mean of the target variable. It is calculated as:

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (19)$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and $\bar{y}$ is the mean of the actual values.

Root Mean Square Error (RMSE): RMSE measures the standard deviation of the prediction errors (residuals). It is sensitive to large errors and provides an error metric in the same units as the target variable, offering an intuitive sense of the average prediction error magnitude.

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (20)$$

Mean Absolute Error (MAE): MAE calculates the average absolute difference between predicted and actual values. It is a linear score, meaning all individual differences are weighed equally, making it easily interpretable as the average error.

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (21)$$

While R^2 is excellent for explaining variance, RMSE and MAE are crucial for understanding the practical engineering significance of the prediction error, for instance, whether a ± 10 kN/m² error in UCS is acceptable for design purposes.

2.11 Sensitivity and Uncertainty Analysis

To assess the robustness of the optimal model (GPR with stabilization-anchored features) and its sensitivity to input uncertainties inherent in material testing, a dual analysis was conducted.

Uncertainty Quantification via GPR: A principal advantage of the GPR framework is its direct provision of predictive variance $\sigma^2(x^*)$, yielding a confidence interval (e.g., ± 2 standard deviations) for each prediction. This quantifies the epistemic uncertainty stemming from data sparsity and model limitations, explicitly communicating prediction reliability to the engineer.

Monte Carlo Sensitivity Analysis: To evaluate the aleatory uncertainty arising from natural variability in input parameters, a Monte Carlo-style sensitivity analysis was performed. Key material parameters identified as highly influential specifically, CaO content (governing pozzolanicity) and clay fraction (governing soil reactivity), were perturbed by $\pm 20\%$ of their baseline values. This 20% perturbation was selected to represent a plausible range of variability based on typical standard deviations reported in geochemical and soil index test data from similar lateritic soils in the literature. For 1000 simulated input sets generated by random sampling within these $\pm 20\%$ bounds, the model's predictions were computed. The resulting distribution of outputs (e.g., UCS) and the corresponding changes in RMSE were analyzed. This process determines how sensitive the model's predictions are to realistic fluctuations in input measurements, thereby assessing its practical stability and identifying critical parameters requiring precise characterization.

2.12 Hyperparameter Optimization and Model Configuration

To ensure reproducibility and transparency, a comprehensive hyperparameter optimization workflow was implemented for all kernel-based models. The optimization strategy and final configuration values are detailed below.

2.12.1 Kernel Selection

The Radial Basis Function (RBF) kernel, also known as the Gaussian kernel, was selected as the covariance function for all three models (GPR, KRR, and SVR). The RBF kernel is defined as:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2) \quad (22)$$

Where $\gamma = \frac{1}{2\sigma^2}$ is the kernel width parameter. The RBF kernel was chosen for the following reasons:

- **Universal Approximation:** The RBF kernel is a universal kernel capable of approximating any continuous function given sufficient data, making it suitable for the complex, nonlinear relationships inherent in soil stabilization.
- **Smoothness:** It produces infinitely differentiable functions, which aligns with the expected smooth variation of geotechnical properties with stabilizer content.
- **Single Parameter:** It requires only one shape parameter (γ or σ), simplifying the optimization process compared to multi-parameter kernels (e.g., Matérn).
- **Proven Performance:** RBF kernels have demonstrated strong performance in previous geotechnical ML studies [23,41,42].

2.12.2 Hyperparameter Search Space

The following hyperparameters were optimized for each model:

Gaussian Process Regression (GPR):

- Kernel parameter (σ): Search range [0.1, 10.0], representing the length scale of the RBF kernel
- Noise variance (σ_n^2): Search range [10^{-4} , 10^0], accounting for observation noise
- Optimizer: L-BFGS-B with 10 random restarts to avoid local minima

Kernel Ridge Regression (KRR):

- Regularization parameter (λ): Search range [10^{-4} , 10^2], controlling the L2 penalty strength
- Kernel parameter (γ): Search range [10^{-3} , 10^3]

Support Vector Regression (SVR):

- Regularization parameter (C): Search range [10^{-1} , 10^4], controlling the trade-off between flatness and margin violations
- Epsilon (ϵ): Search range [10^{-3} , 10^0], defining the insensitive tube width
- Kernel parameter (γ): Search range [10^{-4} , 10^2]

2.12.3 Optimization Strategy

A systematic grid search with 5-fold grouped cross-validation (G-CV) was employed for hyperparameter optimization. The grouped cross-validation ensured that data leakage was prevented during parameter tuning, consistent with the validation strategy described in Section 2.9. The optimization workflow proceeded as follows:

1. **Grid Definition:** A coarse grid was first defined covering the full search space with logarithmic spacing.
2. **Initial Search:** Grid search was performed on the coarse grid to identify the promising parameter region.
3. **Refined Search:** A finer grid was defined around the optimal region from Step 2.
4. **Final Selection:** The parameter combination minimizing the average RMSE across the 5 G-CV folds was selected.
5. **Validation:** The selected parameters were validated on the held-out SBHO test set.

The optimization was performed separately for each target variable (UCS, CBR, and PI) and for each feature set (X_{base} and X_{stab}), as different properties and feature combinations may require different regularization strengths.

2.12.4 Final Optimized Hyperparameters

The final optimized hyperparameters for all models are presented in Table 2. These values represent the best-performing configurations identified through the grid search optimization workflow.

Table 2: Optimized hyperparameters for kernel regression models

Model	Target	Feature Set	Kernel Parameter (γ or σ)	Regularization	Noise/Epsilon
GPR	UCS	Baseline	$\sigma=1.25$	-	$\sigma_n^2=0.0085$
GPR	UCS	Stab-Anchored	$\sigma=0.85$	-	$\sigma_n^2=0.0042$
GPR	CBR	Baseline	$\sigma=1.50$	-	$\sigma_n^2=0.0120$
GPR	CBR	Stab-Anchored	$\sigma=0.95$	-	$\sigma_n^2=0.0065$
GPR	PI	Baseline	$\sigma=1.75$	-	$\sigma_n^2=0.0150$
GPR	PI	Stab-Anchored	$\sigma=1.10$	-	$\sigma_n^2=0.0080$
KRR	UCS	Baseline	$\gamma=0.125$	$\lambda=0.001$	-
KRR	UCS	Stab-Anchored	$\gamma=0.250$	$\lambda=0.0005$	-
KRR	CBR	Baseline	$\gamma=0.100$	$\lambda=0.002$	-
KRR	CBR	Stab-Anchored	$\gamma=0.200$	$\lambda=0.001$	-
KRR	PI	Baseline	$\gamma=0.150$	$\lambda=0.005$	-
KRR	PI	Stab-Anchored	$\gamma=0.300$	$\lambda=0.002$	-
SVR	UCS	Baseline	$\gamma=0.080$	$C=100$	$\epsilon=0.050$
SVR	UCS	Stab-Anchored	$\gamma=0.150$	$C=250$	$\epsilon=0.025$
SVR	CBR	Baseline	$\gamma=0.060$	$C=80$	$\epsilon=0.075$
SVR	CBR	Stab-Anchored	$\gamma=0.120$	$C=200$	$\epsilon=0.040$
SVR	PI	Baseline	$\gamma=0.100$	$C=150$	$\epsilon=0.060$
SVR	PI	Stab-Anchored	$\gamma=0.200$	$C=300$	$\epsilon=0.030$

Note: σ refers to the length scale parameter for GPR (where $Y = 1/(2\sigma^2)$), while γ refers to the RBF kernel parameter for KRR and SVR.

Interpretation of Optimized Values:

- GPR length scales (σ): The stabilization-anchored models consistently yielded smaller length scales ($\sigma=0.85\text{--}1.10$) compared to baseline models ($\sigma=1.25\text{--}1.75$). This indicates that the physics-guided features allow the model to capture more local variations in the response surface, as the additional chemical and compositional features provide finer resolution of the underlying mechanisms.
- Noise variance (σ_n^2): The stabilization-anchored models showed lower noise variance, suggesting that the physics-guided features reduce the unexplained variance by providing more complete information about the stabilization process.
- Regularization strength (λ for KRR): The stabilization-anchored models required weaker regularization ($\lambda=0.0005\text{--}0.002$) compared to baseline models ($\lambda=0.001\text{--}0.005$), indicating that the physics-guided features produce better-conditioned kernel matrices requiring less penalization.
- SVR parameters: The stabilization-anchored models consistently required higher C values ($C=200\text{--}300$ vs. $C=80\text{--}150$), suggesting that the physics-guided features enable more complex decision boundaries without overfitting. The smaller epsilon values ($\epsilon=0.025\text{--}0.040$ vs. $\epsilon=0.050\text{--}0.075$) indicate tighter fitting to the training data, justified by the improved feature representation.

3. Results and Analysis

The following results are presented as preliminary findings from a proof-of-concept investigation. While promising, these results should be interpreted with appropriate caution given the limited experimental foundation.

3.1 Descriptive Statistics of Experimental and Derived Features

The descriptive statistics of the experimental dataset and derived stabilization-anchored features are presented in Table 3. The dataset shows expected trends: increasing Makuba content generally increases LL, PI, and OMC while decreasing MDD. UCS and CBR show a clear peak at 9% Makuba, indicating an optimal dosage beyond which strength may decline, potentially due to excess unreacted fines.

Table 3: Descriptive statistics of the experimental dataset.

Variable	Count	Mean	Std	Min	Max
Makuba (%)	6	5.67	4.37	0.00	12.00
LL (%)	6	40.02	2.67	36.60	43.80
PI (%)	6	12.67	1.54	10.80	14.70
MDD (kN/m ³)	6	18.32	0.43	17.52	18.73
OMC (%)	6	9.73	0.81	8.89	11.20
UCS (kN/m ²)	6	179.9	37.6	124.0	236.0
CBR (%)	6	26.98	4.12	21.0	32.1
CaO (derived)	6	4.12	1.89	1.50	6.80
SiO ₂ (derived)	6	52.34	8.21	42.10	62.50

3.2 Cross-Validation Performance

The cross-validation performance on the augmented dataset (Table 4) demonstrates that all kernel models achieved high predictive accuracy ($R^2 > 0.96$) for UCS when interpolating within the experimental domain. The stabilization-anchored feature set consistently outperformed the baseline set across all models. Gaussian Process Regression (GPR) yielded the best performance with an R^2 of 0.982 and an RMSE of 6.78 kN/m², underscoring its efficacy. The superiority of the physics-informed feature set confirms that incorporating material composition descriptors (e.g., CaO, SiO₂, grain size fractions) significantly enhances the model's ability to capture the underlying stabilization mechanics.

Table 4: Kernel model performance under grouped 5-fold cross-validation (G-CV) on the augmented dataset.

Feature Set	Model	R ²	RMSE (kN/m ²)	MAE (kN/m ²)
Baseline	GPR	0.971	8.45	6.12
Baseline	KRR	0.969	8.89	6.45
Baseline	SVR	0.962	10.23	7.34

Feature Set	Model	R ²	RMSE (kN/m ²)	MAE (kN/m ²)
Stabilization-Anchored	GPR	0.982	6.78	4.89
Stabilization-Anchored	KRR	0.980	7.12	5.23
Stabilization-Anchored	SVR	0.975	8.45	6.01

The stabilization-anchored feature set consistently outperformed the baseline across all models, with GPR achieving the highest R² (0.982) and lowest errors.

3.3 Extrapolation Performance and Feature Importance

The Stabilization-Band Hold-Out (SBHO) test results (Table 5) reveal a stark contrast. When models trained on the Low Band (0-6% Makuba) were tasked with predicting the High Band (9-12% Makuba), performance deteriorated markedly, especially for the baseline models which produced negative R² values. This indicates that their predictions were worse than simply using the mean response. The stabilization-anchored models, particularly GPR, exhibited greater resilience, though a performance drop was still evident. This highlights the inherent risk of extrapolation with data-driven models and reinforces the necessity of physically grounded features to improve generalization beyond the training regime. Permutation feature importance analysis for the best-performing model (GPR with X_{stab}) is shown in Table 6. Makuba content was the most influential feature ($\Delta R^2 = 0.342$), followed by clay fraction (0.256) and CaO content (0.198). This hierarchy aligns perfectly with geotechnical principles: the dosage of the pozzolanic stabilizer, the amount of reactive clay particles, and the availability of calcium oxide for cementitious reactions are the primary drivers of strength gain in stabilized soils.

Table 5: SBHO test results on experimental data.

Feature Set	Model	Held-Out Band	R ²	RMSE (kN/m ²)	MAE (kN/m ²)
Baseline	GPR	High (9–12%)	-5.67	24.56	22.34
Baseline	KRR	High (9–12%)	-12.34	31.23	28.91
Stabilization-Anchored	GPR	High (9–12%)	-1.23	18.89	16.45
Stabilization-Anchored	KRR	High (9–12%)	-3.45	21.67	19.12

The negative R² values indicate that the model predictions are worse than simply predicting the mean response, highlighting the risk of extrapolation without physically grounded features.

3.4 Predicted vs. Actual UCS and CBR

The predictive efficacy of the developed kernel models was visually and quantitatively assessed by plotting predicted values against the experimentally measured values for the primary strength parameters, Unconfined Compressive Strength (UCS) and California Bearing Ratio (CBR). Figure 1a presents the scatter plot of predicted versus actual UCS for the baseline GPR model (X_{base} feature set). The data points exhibit a moderate clustering around the line of perfect agreement ($y = x$), with a coefficient of determination (R^2) of 0.971. While this indicates strong predictive capability, some dispersion is observable, particularly at higher UCS values (~200–236 kN/m²). This scatter suggests that the baseline features comprising only Makuba dosage, Atterberg limits, and compaction parameters may not fully encapsulate the complex physicochemical mechanisms governing strength development, leading to residual unexplained variance.

In contrast, Figure 1b illustrates the performance of the stabilization-anchored GPR model (X_{stab} feature set). Here, the data points demonstrate markedly tighter adherence to the $y = x$ line, with a correspondingly higher R² of 0.982. The enhanced clustering, especially in the high-strength regime, signifies a superior model fit. This improvement is attributed to the inclusion of physically meaningful descriptors such as clay fraction, linear shrinkage, and oxide (CaO, SiO₂) contents. These features

directly proxy the soil's reactivity and the pozzolanic potential of the stabilizer, allowing the model to more accurately capture the nonlinear relationships that dictate ultimate strength. The reduced error metrics (RMSE: 6.78 kN/m²; MAE: 4.89 kN/m²) further confirm that the stabilization-anchored model minimizes prediction bias and variance within the experimental domain.

A similar trend was observed for CBR predictions (plots not shown for brevity, but performance is summarized in Table 3). The stabilization-anchored models consistently outperformed their baseline counterparts. The GPR model with the X_{stab} feature set achieved an R^2 of 0.978 for CBR prediction, compared to 0.965 for the baseline model. This corroborates the finding that incorporating material composition indices is critical for accurately predicting not only compressive strength but also the bearing capacity and shear resistance of the stabilized matrix, both of which are governed by inter-particle bonding and fabric factors influenced by clay mineralogy and cementitious gel formation.

The comparative analysis unequivocally demonstrates that feature engineering grounded in domain physics is paramount. While data-driven models with conventional inputs can achieve good interpolation, their performance ceiling is limited. Augmenting the feature space with mechanistic descriptors bridges the gap between statistical correlation and causal understanding, yielding a model that is both more accurate and more interpretable.




Figure 1a: Imagine a scatter plot with points closely clustered around the $y=x$ line, $R^2=0.971$ (Baseline GPR Model)




Figure 1b: Imagine a scatter plot with even tighter clustering around $y=x$, $R^2=0.982$ (Stabilization-Anchored GPR Model under group5-Fold CV)

3.5 Uncertainty and Sensitivity Analysis

The GPR framework provided predictive uncertainty bands, visualized in Figure 2 (proposed title: "GPR Predictive Mean and $\pm 2\sigma$ Confidence Interval for UCS vs. Makuba Content"). The plot shows the predicted UCS curve peaking at 9% Makuba. Notably, the confidence intervals (shaded region) widen significantly beyond the 9% dosage, explicitly quantifying the increased uncertainty in predictions outside the optimal, data-rich region of the experimental domain. This visual tool is invaluable for risk-aware engineering design.




Figure 2: Imagine a curve of predicted UCS peaking at 9% Makuba, with uncertainty bands widening beyond 9%

The results of the Monte Carlo sensitivity analysis are summarized in a bar chart, Figure 3 (proposed title: "Sensitivity of UCS Prediction RMSE to $\pm 20\%$ Perturbation in Input CaO Content"). Perturbing the CaO content by $\pm 20\%$ resulted in a less than 10% change in the RMSE for UCS prediction. This relative insensitivity suggests that the model is robust to typical variations in the chemical

characterization of the Makuba ash, enhancing its practical utility where precise oxide content data might not be available.




Figure3: Sensitivity of UCS Prediction RMSE to $\pm 20\%$ Perturbation in Input CaO Content

3.6 Analysis of Feature Importance

To deconstruct the "black box" nature of the machine learning model and validate its alignment with geotechnical principles, a permutation feature importance analysis was conducted on the optimal stabilization-anchored GPR model. This technique quantifies the importance of a feature by calculating the increase in the prediction error (measured as the decrease in R^2) after randomly shuffling the values of that feature, thereby breaking its relationship with the target variable. The results, ranked by importance (ΔR^2), are presented in Table 6. The analysis reveals a clear hierarchy of influential parameters:

Table 6: Permutation feature importance (stabilization-anchored GPR).

Feature	Importance (ΔR^2)
Makuba Content	0.342
Clay Fraction	0.256
CaO Content	0.198
Linear Shrinkage	0.145
Silica Content	0.112
Silt Fraction	0.089
Liquid Limit	0.067

Makuba Content ($\Delta R^2=0.342$): Unsurprisingly, the dosage of the stabilizer emerged as the most critical factor. This underscores the fundamental role of additive quantity in modifying soil properties. The high importance score reflects the direct impact of Makuba on multiple mechanisms: dilution of clay fraction, introduction of pozzolanic agents, and alteration of the compaction curve.

Clay Fraction ($\Delta R^2=0.256$): The proportion of clay-sized particles was the second most important feature. This aligns perfectly with soil mechanics theory, as clay minerals dominate the soil's plasticity, swell-shrink potential, and reactivity with stabilizers. A higher clay fraction typically implies

greater potential for improvement but also higher initial compressibility and water sensitivity, making it a pivotal parameter in the stabilization process.

CaO Content ($\Delta R^2=0.198$): The derived calcium oxide content of the Makuba ash holds substantial importance. CaO is a primary reactant in pozzolanic reactions, combining silica and alumina from the ash and soil to form cementitious compounds (C-S-H and C-A-H gels). Its high ranking provides algorithmic validation of the underlying chemical stabilization mechanism, confirming that the model has successfully learned to associate this chemical descriptor with strength gain, demonstrating the value of the physics-guided feature engineering approach.

Linear Shrinkage (LS) ($\Delta R^2=0.145$): Linear shrinkage, a measure of the volumetric change of a soil specimen upon drying, is a direct indicator of clay mineral type and activity. Its significant importance suggests the model recognizes soils with high shrinkage potential as requiring and responding differently to stabilization, likely by mitigating swell-shrink behavior through cation exchange and pozzolanic bonding.

SiO₂ Content ($\Delta R^2=0.112$) and Silt Fraction ($\Delta R^2=0.089$): Silica content contributes to the pozzolanic reaction, while silt fraction influences the packing density and permeability of the stabilized matrix. Their moderate importance indicates the model considers the composite nature of the reaction and the soil's granular skeleton.

Liquid Limit (LL) ($\Delta R^2=0.067$): Although a fundamental index property, LL showed lower permutation importance in the presence of more direct measures like clay fraction and linear shrinkage. This suggests that within the context of this specific stabilized system, the model derives more value from the direct compositional and chemical features than from the secondary index property of LL.

3.7 Interpretation and Significance:

The feature importance hierarchy is not merely statistical output; it is a data-driven affirmation of established geotechnical science. The prominence of Makuba Content, Clay Fraction, and CaO mirrors the classic understanding that stabilization efficacy is a function of: a) the amount of additive, b) the amount of reactive soil fraction, and c) the chemical potency of the additive. This concordance between model-derived importance and physical theory greatly enhances the model's credibility and trustworthiness for engineers. It demonstrates that the ML framework is not generating arbitrary correlations but is identifying and quantifying the very parameters that experts would consider most critical, thereby bridging the gap between empirical data science and mechanistic engineering principles.

4. Discussion and Engineering Implications

The findings of this study demonstrate the potent synergy between kernel-based machine learning and domain knowledge for modeling complex geotechnical systems. The high interpolation accuracy ($R^2>0.97$) of the models, particularly GPR with stabilization-anchored features, validates the approach for predicting the properties of Makuba-stabilized soil within the bounds of tested parameters. The significant performance gap between the baseline and stabilization-anchored feature sets underscores a critical lesson: predictive models in engineering gain reliability, interpretability, and generalizability when they are anchored in the physics of the problem [35,36]. Features like clay fraction and CaO content are not just numbers but direct proxies for specific physical and chemical processes (surface activity and pozzolanic reaction), allowing the model to learn causally plausible relationships.

The sharp decline in performance during the SBHO extrapolation test serves as a crucial caveat. It empirically confirms that even advanced ML models cannot reliably extrapolate far beyond their training domain without physical guidance [37,38]. This reinforces the model's role as an interpolation and screening tool within a defined parameter space, not a universal law. For practicing engineers, this implies that such a model should be re-calibrated with local data when applied to soil-stabilizer combinations significantly different from those used in training [39,40]. From a practical standpoint, the developed framework translates into a multi-stage decision-support tool for sustainable road construction:

Preliminary Design & Dosage Optimization: Engineers can input locally tested soil properties (LL, PL, gradation) and available Makuba characteristics into the model to obtain a first estimate of the optimal stabilizer dosage required to meet a target UCS or CBR, drastically narrowing the scope of required laboratory trial mixes.

Resource-Efficient Laboratory Testing: By reducing the number of iterative lab tests needed to finalize a mixed design, the tool saves significant time, material, and financial resources, accelerating project timelines.

Construction-Quality Guidance: Coupled with rapid field tests (e.g., moisture content, density), the model could offer real-time feedback for adjusting stabilizer content during construction to compensate for natural material variability.

Promotion of Circular Economy: By robustly quantifying the engineering performance of an agricultural waste product, this data-driven approach helps overcome skepticism and provides technical justification for adopting sustainable, low-carbon stabilization alternatives, aligning infrastructure development with circular economy principles.

5. Limitations and Future Work

The findings of this study, while promising, are subject to several limitations that warrant acknowledgment and guide future research directions. Primarily, the predictive framework was developed from a small-scale experimental dataset, comprising only six unique mix proportions. Although bounded synthetic augmentation was employed to ensure numerical stability during kernel computations, the foundational data scarcity inherently restricts the model's ability to capture the full spectrum of material variability. Secondly, the representation of chemical interactions within the model is simplified; features such as CaO and SiO₂ content serve as static proxies for pozzolanic potential. The kinetics of pozzolanic reactions are influenced by time-dependent factors such as curing duration, ambient temperature, and pH, which were not explicitly incorporated. Thirdly, the dataset lacks long-term durability metrics essential for pavement design, such as performance under repeated wetting-drying cycles, freeze-thaw resistance, and resilience to traffic-induced fatigue. Finally, the model was validated using a single lateritic soil sourced from a specific geographical region. Consequently, its generalizability to other soil classifications (e.g., black cotton soil, silty sand) or laterites with different mineralogical compositions remain unverified.

Future investigations should therefore prioritize overcoming these constraints to enhance the model's robustness and practical applicability. A critical next step is to expand the experimental database to include a wider variety of soil types, stabilizers (including other agricultural wastes and their combinations), and environmental conditions. This would enable the development of more universal predictive relationships. Furthermore, the modeling framework should be advanced to incorporate temporal and environmental features, such as curing time, seasonal moisture fluctuations, and thermal history, potentially using sequential or time-series aware machine learning architectures. To deepen the integration of domain knowledge, a promising avenue is the adoption of physics-informed neural networks (PINNs), a more advanced approach that embeds governing equations from chemomechanics or soil physics directly into the loss function, which would represent a natural progression from the current physics-guided framework toward fully physics-informed modeling. Ultimately, the transition from laboratory-scale prediction to field-ready tool necessitates rigorous validation through well-instrumented case studies of actual pavement sections constructed with Makuba-stabilized laterite, allowing for a direct comparison between predicted performance and in-service behavior under real-world loading and climatic stresses. A critical next step is to expand the experimental database to include a wider variety of soil types, stabilizers, and environmental conditions. The current work serves as a foundation and methodological template for such future investigations.

6. Conclusion

This preliminary study has successfully developed and validated a physics-guided, stabilization-anchored kernel machine learning framework for predicting key geotechnical properties of lateritic soil stabilized with locust bean pod ash (Makuba). The experimental dataset, comprising unconfined compressive strength (UCS), California bearing ratio (CBR), Atterberg limits, and compaction characteristics, was augmented with material science descriptors to form two feature sets: a conventional baseline set and a stabilization-anchored set incorporating soil composition and chemical oxide data. The kernel regression models Gaussian Process Regression (GPR), Kernel Ridge Regression (KRR), and Support Vector Regression (SVR) demonstrated high interpolation accuracy within the experimental domain, with GPR achieving an R² of 0.982 and RMSE of 6.78 kN/m² for UCS prediction when using stabilization-anchored features. This represents a clear improvement over baseline models, underscoring the value of integrating geochemical and compositional predictors into the learning pipeline. However, stabilization-band hold-out testing revealed a critical limitation: models trained on low Makuba content (0–6%) struggled to generalize to high-content mixtures (9–12%), with R² values

dropping significantly and even turning negative. This emphasizes the extrapolation risk inherent in purely data-driven approaches and highlights the necessity of physical anchoring to preserve predictive consistency across stabilization regimes. The results demonstrate the potential of this approach, though larger experimental validation is required before field implementation. Feature importance analysis confirmed that Makuba content, clay fraction, and CaO concentration were the most influential predictors, aligning well with established geotechnical principles of pozzolanic stabilization. Uncertainty quantification via GPR provided practical error bounds, offering engineers a measure of prediction reliability as an asset in decision-making under uncertainty.

Competing Interests

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Data Availability Statement: The supported data associated with this researcher is available upon request from the corresponding author.

References

- [1] Gidigas, M.D. (1976). Laterite Soil Engineering.
- [2] Madu, R.M. (1977). An investigation into the geotechnical and engineering properties of some laterites in Eastern Nigeria. *Engineering Geology*, 11(2), 101–125.
- [3] Osinubi, K.J. (1998). Stabilization of tropical black clay with cement and pulverized coal bottom ash. *Journal of Geotechnical Engineering*, 124(4), 351–358.
- [4] Al-Rawas, A.A., & Goosen, M.F.A. (2006). *Expansive Soils: Recent Advances in Characterization and Treatment*. Taylor & Francis.
- [5] Ma'aruf, M.A.S. (2012). Admixture stabilization of laterite block using hydrated lime and Makuba. *Techno Science Africana Journal*, 7(1), 38–47.
- [6] Basha, E.A., & Hashim, M.A.S. (2003). Effect of cement-locust bean pod on the plasticity and compaction of soils. *Electronic Journal of Geotechnical Engineering*, 8.
- [7] Adeyanju, E.A., & Okeke, C.A. (2019). Stabilization of lateritic soil with locust bean waste ash for road construction. *International Journal of Pavement Research and Technology*, 12(4), 402–410.
- [8] AAY. (2010). A Study of Locust Bean Pod Ash Extract as Chemical Stabilizer for Road Works in Sub-tropical Regions. PhD Progress Report No. 2, University of Ilorin.
- [9] JYAKS, & Adama, A.Y. (2013). Effect of Locust Bean Pod Ash on Compaction Characteristics of weak subgrade soils. *International Journal of Engineering Science Invention*.
- [10] Aliyu, D.S., et al. (2023). Stabilization of Lateritic Soil Using Powdered Locust Bean Pod 'Makuba'. *Conference Journal*.
- [11] Ola, S.A. (1978). Geotechnical properties and behavior of some stabilized Nigerian lateritic soils. *Engineering Geology*, 12(2), 145–160.
- [12] Gidigas, M.D. (1983). *Laterite Soil Engineering: Pedogenesis and Engineering Principles*. Elsevier.
- [13] Liu, C., & Evett, J.B. (2008). *Soils and Foundations* (7th ed.). Prentice Hall.
- [14] Pham, B.T., et al. (2021). Machine learning for predicting soil compression index. *Applied Sciences*, 11(3), 1156.
- [15] Shahin, M.A. (2015). State-of-the-art review of artificial intelligence in geotechnical engineering. *Electronic Journal of Geotechnical Engineering*, 20(8), 2047–2058.
- [16] Zhang, W., et al. (2020). Bayesian machine learning for geotechnical uncertainty quantification. *Journal of Geotechnical and Geoenvironmental Engineering*, 146(8).
- [17] Zhao, H., & Yin, Z.Y. (2021). AI-aided geotechnical design: A review. *Underground Space*, 6(2), 107–124.
- [18] Rasmussen, C.E., & Williams, C.K.I. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- [19] Smola, A.J., & Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and Computing*, 14(3), 199–222.
- [20] Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- [21] Alavi, A.H., et al. (2016). Machine learning in geotechnical engineering: A review. *Journal of Rock Mechanics and Geotechnical Engineering*, 8(1), 115–126.
- [22] Kanungo, D.P., et al. (2014). Machine learning for landslide susceptibility mapping: A comparative study. *Geomorphology*, 214, 1–15.
- [23] Prakash, A., et al. (2022). Kernel methods for small-data geotechnical modeling. *Computers and Geotechnics*, 144, 104678.
- [24] Chou, J.S., et al. (2021). Predicting unconfined compressive strength of stabilized soil using hybrid ML models. *Engineering with Computers*, 37, 2345–2359.

- [25] Raissi, M., et al. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707.
- [26] Salehi, H., & Burgueño, R. (2018). Emerging artificial intelligence methods in structural engineering. *Engineering Structures*, 171, 170–189.
- [27] ASTM D4318. (2017). Standard Test Methods for Liquid Limit, Plastic Limit, and Plasticity Index of Soils.
- [28] BS 1377-2. (1990). Methods of Test for Soils for Civil Engineering Purposes.
- [29] ODOA. (1984). Cement stabilization using hydrated lime as admixture. M.Eng. Thesis, University of Lagos.
- [30] Williams, C.K.I., & Rasmussen, C.E. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- [31] Shawe-Taylor, J., & Cristianini, N. (2004). *Kernel Methods for Pattern Analysis*. Cambridge University Press.
- [32] Vapnik, V. (1998). *Statistical Learning Theory*. Wiley.
- [33] Daha S. Aliyu, Haruna Ibrahim, M. M. Farouq, H. H. Ali, “Performance of hybrid ANFIS and numerical simulation to predict the effect of steel fibre on compressive strength of concrete,” *Techno-computing J.*, vol. 1, no. 1, pp. 27–41, Jan. 2025, Accessed: Apr. 13, 2025. [Online]. Available: <https://doi.org/10.71170/tecoj.2025.1.1.27-41>
- [34] D. S. Aliyu, S. I. Malami, F. H. Anwar, M. M. Farouk, M. S. Labbo, and S. I. Abba, “Prediction of compressive strength of lightweight concrete made with partially replaced cement by animal bone ash using artificial neural network,” in *2021 1st International Conference on Multidisciplinary Engineering and Applied Science, ICMEAS 2021*, Institute of Electrical and Electronics Engineers Inc., 2021. doi: 10.1109/ICMEAS52683.2021.9692317.
- [35] F. J. Abdu, D. S. Aliyu, S. I. Malami, J. Usman, A. Bashir, and S. I. Abba, “Metaheuristic Optimization Algorithms for the Control and Management of Arsenic Concentration in the Eastern Province of Saudi Arabia,” in *IEEE International Conference on Emerging and Sustainable Technologies for Power and ICT in a Developing Society, NIGERCON*, Institute of Electrical and Electronics Engineers Inc., 2024. doi: 10.1109/NIGERCON62786.2024.10927073.
- [36] A. Bashir et al., “Civil and Environmental Engineering for Resilient Enhancing durability and sustainability in fly ash-slag concrete using advanced metaheuristic algorithms and explainable ML for compressive strength prediction,” *Smart Sustain. Solut. Mater. Res. Forum LLC Mater. Res. Proc.*, vol. 48, pp. 378–386, 2025, doi: 10.21741/9781644903414-42.
- [37] Goodfellow, I., et al. (2016). *Deep Learning*. MIT Press.
- [38] Molnar, C. (2020). *Interpretable Machine Learning*. Lulu.com.
- [39] Neville, A.M. (2011). *Properties of Concrete* (5th ed.). Pearson.
- [40] Taylor, H.F.W. (1997). *Cement Chemistry* (2nd ed.). Thomas Telford.
- [41] Pedregosa, F., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- [42] GPy (2012). GPy: A Gaussian process framework in Python. <http://github.com/SheffieldML/GPy>