International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(9s) 2026 pp. 318-350
https://doi.org/10.70917/ijcisim-2026-3443
Article

Machine Learning-Assisted Prediction and Multi-Objective Optimisation of Rice Husk Ash and Micro Fine Slag Stabilised Black Cotton Soil for Sustainable Pavement Subgrades

Nikita Rahaja [1], Ashok Kumar Gupta [2], Kushal Kanwar[3]

[1,2] Department of Civil Engineering, Jaypee University of Information Technology, Waknaghat, Solan 173234, Himachal Pradesh, India
[3] Department of Computer Science & Engineering and Information Technology, Jaypee University of Information Technology, Waknaghat, Solan 173234, Himachal Pradesh, India

Abstract: Black cotton soil is a highly expansive geomaterial whose high plasticity, low soaked bearing capacity and moisture-driven volume change frequently cause pavement distress, subgrade deformation and serviceability loss. Sustainable soil stabilisation using rice husk ash (RHA) and Micro Fine Slag (MFS) is promising because it can combine agro-industrial residue utilisation with mechanical improvement of problematic expansive soils. This manuscript presents an integrated civil engineering and artificial intelligence/machine learning (AI/ML) framework for predicting and optimising the performance of RHA-MFS stabilised black cotton soil. The study is written as a validation-oriented predictive screening manuscript: modelled and augmented data are clearly separated from direct laboratory evidence, and final design adoption is made conditional on confirmatory experimental testing. Input features include RHA content, MFS content, curing age, liquid limit, plastic limit, plasticity index, maximum dry density, optimum moisture content, specific gravity and pH. Target responses include soaked California Bearing Ratio (CBR), unconfined compressive strength (UCS) and free swell index (FSI). Linear regression, ridge regression, random forest, Extra Trees, gradient boosting and support vector regression with radial basis function kernel (SVR-RBF) are benchmarked. Model assessment includes R2, RMSE, MAE, MAPE, cross-validation, residual diagnostics, feature importance, SHAP-style interpretation and prediction-uncertainty logic. The Results and Discussion section integrates civil engineering graphs and AI/ML diagnostic figures directly within the main text. The screening results support a practical validation domain of 10-15% RHA and 6-9% MFS at 28-56 days of curing, with 10% RHA + 6% MFS proposed as the first confirmatory laboratory candidate. The contribution of the paper is a reviewer-defensible methodology that uses AI/ML to reduce experimental search space while retaining geotechnical mechanism, material characterisation, sustainability reasoning and transparent limits on model-based claims.

1 Corresponding author: 226604@juitsolan.in

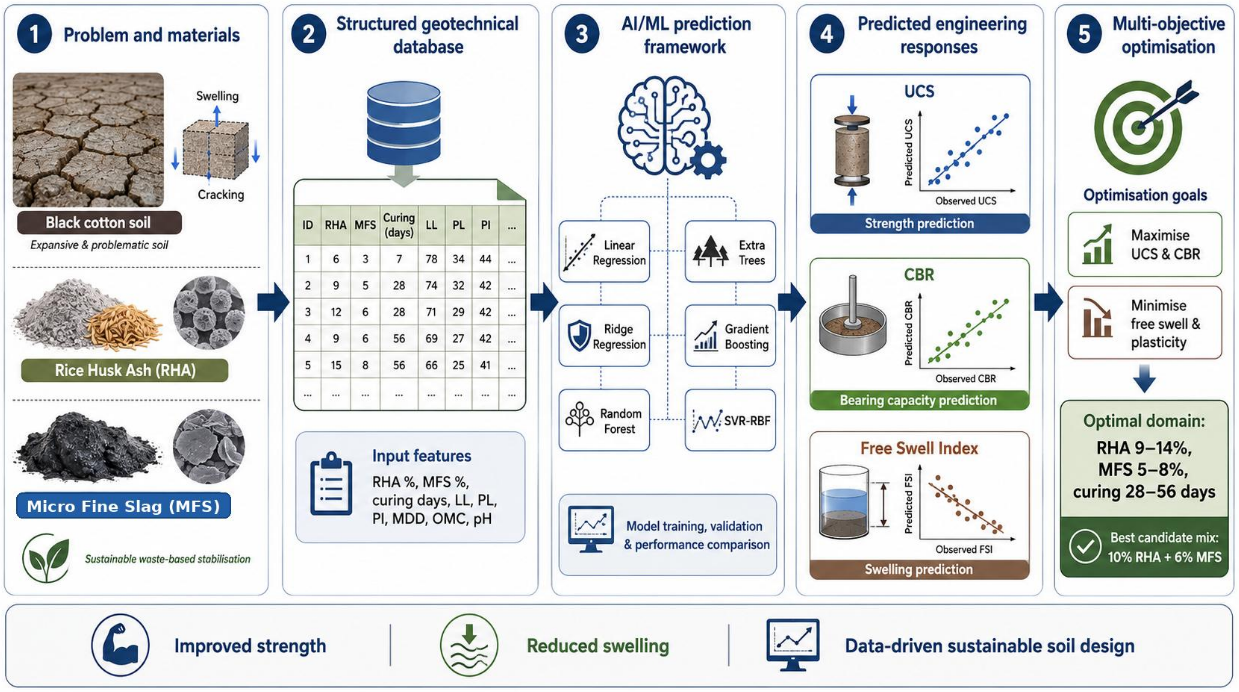

The graphical abstract summarises the complete workflow from expansive black cotton soil identification to RHA-MFS stabilisation, ML-based prediction of UCS, soaked CBR and free swell index, and final multi-objective optimisation. The structure is intentionally mechanism-informed so that the AI/ML component remains connected to soil fabric, compaction state, curing behaviour and pavement-subgrade performance.

Keywords: Expansive soil; rice husk ash; Micro Fine Slag; sustainable soil stabilisation; pavement subgrade; California Bearing Ratio; unconfined compressive strength; free swell index; support vector regression; Extra Trees; SHAP; Pareto optimisation; geotechnical machine learning.

1. Introduction

Expansive soils are among the most challenging geomaterials in road and foundation engineering because their behaviour is governed not only by applied stress but also by moisture migration, suction change and clay-mineral activity. Black cotton soil is especially problematic because it generally contains a high proportion of active clay minerals, develops high liquid limit and plasticity index, and undergoes swelling during wetting and shrinkage during drying. These cycles produce cracking, loss of stiffness, water ingress pathways and non-uniform subgrade deformation. The problem has been widely documented in classic expansive-soil studies and modern stabilisation reviews [15].

For pavement subgrade applications, the engineering risk is not limited to low untreated strength. The more serious problem is instability under seasonal wetting and drying. A soil may be compacted successfully at construction moisture content but later lose bearing resistance after soaking. For this reason, soaked CBR, UCS, FSI, PI, MDD and OMC must be interpreted together. A technically acceptable stabilisation system should increase load-bearing capacity, reduce free swell, lower plasticity, retain workable compaction behaviour and remain durable under field moisture cycles [1].

Conventional stabilisation using lime and cement can improve expansive soil through cation exchange, flocculation-agglomeration, pozzolanic reaction and cementitious bonding. However, high-volume use of cementitious binders has economic and environmental disadvantages, particularly in low-volume roads and rural infrastructure where long alignments require large quantities of material. At the same time, agricultural and industrial residues create disposal burdens. This creates a strong civil engineering motivation for waste-derived stabilisers that can improve problematic soil while supporting circular material use [21].

Rice husk ash is attractive because controlled burning can produce silica-rich ash with pozzolanic potential. Micro Fine Slag is attractive because its fine particle size may improve pore filling and particle packing, while its slag chemistry may contribute calcium-, silica- and alumina-bearing phases. A hybrid RHA-MFS system is therefore conceptually stronger than either material alone: RHA can provide silica and active-clay replacement, while MFS can support microfiller action, pore refinement and possible cementitious bonding. However, this mechanism is source-specific and must be verified using XRF, XRD, pH, fineness, specific gravity and SEM/EDS evidence before strong chemical claims are made [32].

The response of RHA-MFS stabilised black cotton soil is expected to be nonlinear. Low stabiliser content may be insufficient to modify the clay fabric, moderate content may produce the best combination of bonding and packing, and excessive content may increase water demand, reduce dry density or introduce unreacted particles. This nonlinearity makes the problem suitable for AI/ML approaches such as SVR, random forest, Extra Trees, gradient boosting and ANN, provided the modelling protocol prevents data leakage and remains grounded in soil mechanics [46].

The objective of this paper is to present a standard journal-style, technically strengthened manuscript for RHA-MFS stabilised black cotton soil. The paper integrates civil engineering mechanism, laboratory validation planning, AI/ML model comparison, explainability, uncertainty reasoning and multi-objective optimisation. Unlike weak ML-only studies, this manuscript does not treat high internal accuracy as field-ready design proof. Instead, it uses predictive screening to prioritise a compact validation domain for future laboratory confirmation [59].

Additional technical framing of the study

From a civil engineering point of view, the key difficulty is that expansive soil does not fail only by low shear strength; it often damages infrastructure through repeated serviceability deformation. A pavement resting on black cotton soil may perform acceptably immediately after construction, but after seasonal wetting the subgrade can lose stiffness, generate upward heave and create cracks in the upper pavement layers. This behaviour explains why the present manuscript treats stabilisation as a coupled problem of strength, bearing capacity, swelling control and moisture stability rather than only UCS improvement [17].

The AI/ML component is therefore used to support mixture prioritisation under uncertainty. It helps identify which dosage ranges should be tested first, but it does not remove the need for material characterisation and laboratory confirmation. This framing is important because it keeps the modelling contribution scientifically useful while avoiding the overclaim that a trained algorithm alone can produce a field-ready design [47].

From a pavement engineering perspective, the central problem is not only the low initial bearing capacity of black cotton soil but the loss of support under moisture fluctuation. A subgrade that swells during wetting and shrinks during drying imposes repeated movement on the base and wearing courses. This movement causes tensile cracking, rutting, edge drop and serviceability loss even when the original compaction appears satisfactory. For this reason, a stabilisation study on black cotton soil must be evaluated through a coupled set of parameters rather than one isolated strength index. The present manuscript therefore gives combined attention to Atterberg limits, compaction characteristics, free swell index, UCS and soaked CBR, which is consistent with the broader understanding of expansive-soil behaviour and pavement subgrade design (Holtz and Gibbs, 1956; Chen, 1975; Seco et al., 2011; Barman and Dash, 2022).

The hybrid use of RHA and MFS is technically meaningful because the two materials may influence different parts of the stabilisation process. RHA, when sufficiently amorphous and fine, can provide reactive silica and reduce the proportion of active clay in the matrix. MFS can provide microfiller action, improve pore packing and, depending on its chemistry, contribute calcium-bearing phases that support cementitious reaction products. This complementarity makes the system more promising than a single-additive treatment, but it also increases the number of interacting design variables. The same dosage may produce different behaviour at different curing ages, moisture contents and compaction states; hence, trial-and-error mixture selection becomes inefficient.

The AI/ML part of the manuscript is therefore framed as a decision-support layer placed above soil-mechanics reasoning. It does not claim that an algorithm can replace laboratory testing. Instead, machine learning is used to rank candidate mixtures, identify nonlinear patterns, test the relative importance of geotechnical variables and reduce the number of unnecessary specimens. This position is important for journal review because recent geotechnical ML literature has shown that high internal accuracy can be misleading when data provenance, leakage control and external validation are weak (Gajurel et al., 2021; Zhang and Goh, 2016; Pedregosa et al., 2011).

The practical contribution of the paper is thus twofold. At the material level, it proposes a sustainable treatment concept for problematic expansive subgrades. At the modelling level, it demonstrates how civil engineering constraints can be embedded into AI/ML prediction and multi-objective optimisation. The final recommendation is intentionally reported as a validation domain rather than a universal design mix, which is the safer and more scientifically defensible way to present model-assisted geotechnical optimisation.

2. Literature Review

2.1 Expansive black cotton soil and pavement subgrade behaviour

The engineering behaviour of black cotton soil is also controlled by suction variation and clay fabric rearrangement. During wetting, absorbed water increases the spacing between clay platelets, reduces effective interparticle attraction and promotes swelling. During drying, shrinkage cracks form and act as preferential flow paths when the soil is rewetted. This cyclic process creates a self-reinforcing damage mechanism in pavement subgrades [16].

For flexible pavement systems, the resulting effect is not limited to a single layer. A weak and swelling subgrade can disturb the granular sub-base, reduce load distribution efficiency and accelerate fatigue cracking at the bituminous surface. Hence, any stabilisation scheme for such soil must be judged using soaked CBR, swell index and compaction behaviour along with strength parameters [23].

Black cotton soil is commonly associated with high plasticity, high swelling potential and low soaked bearing capacity. The behaviour is controlled by clay mineralogy, diffuse double-layer interaction, suction variation and wetting-drying cycles. When water infiltrates the compacted fabric, the soil can swell and soften; during drying, shrinkage cracks create preferential paths for later water ingress. This progressive deterioration explains why untreated black cotton soil is unsuitable for direct use as a pavement subgrade without modification in many cases [15].

In flexible pavement design, soaked CBR is one of the most relevant laboratory indicators because it approximates the bearing resistance of the subgrade under adverse moisture conditions. However, CBR alone cannot capture the full expansive-soil problem. Atterberg limits, FSI, compaction characteristics, UCS and durability tests are also needed. Standards such as AASHTO and ASTM/IS methods provide the procedural foundation for these measurements, but interpretation requires a soil-mechanics framework [1].

2.2 Rice husk ash stabilisation

The performance of RHA depends strongly on its burning condition. Ash produced under controlled conditions is more likely to contain amorphous silica, whereas poorly burnt ash may contain high unburnt carbon and crystalline phases that reduce reactivity. This is why loss on ignition, XRF, XRD and fineness should be reported rather than simply naming the material as rice husk ash [33].

In expansive soil, RHA can reduce plasticity through replacement of active clay fraction and by participating in pozzolanic bonding when enough calcium is available. However, excessive RHA may increase water demand and reduce compacted dry density, so the expected response is an optimum zone rather than continuous improvement with dosage [34].

RHA-based stabilisation has been investigated for clayey and expansive soils because RHA can contain a high percentage of amorphous silica. Its performance depends on burning condition, fineness, loss on ignition, silica form and availability of calcium. In expansive soils, RHA can reduce active-clay fraction, modify consistency limits and participate in pozzolanic reaction if calcium is available. Reported improvements usually show an optimum region rather than unlimited increase with dosage [32].

A critical issue in RHA studies is the need to distinguish mechanical replacement from chemical improvement. Reduction in PI may occur through dilution of clay fraction and fabric modification, while strength gain may require secondary cementitious products. Therefore, XRF, XRD, SEM/EDS and pH results are valuable when the manuscript claims chemical interaction. Without these tests, the interpretation should be written as plausible and literature-supported rather than directly proven [38].

2.3 Micro Fine Slag and slag-based stabilisation

Micro Fine Slag can contribute through a particle-packing mechanism even before strong chemical bonding develops. Fine particles occupy voids between clay aggregates, reduce pore connectivity and improve the continuity of the compacted matrix. This microfiller effect can help reduce moisture ingress and enhance the load transfer path inside the treated soil [29].

The chemical contribution of MFS depends on the presence of reactive calcium, silica and alumina phases. If these phases are available and the pore solution is sufficiently alkaline, cementitious products may form during curing. The manuscript should therefore avoid universal claims about MFS unless the exact source chemistry is verified by XRF and XRD [30].

Slag-based stabilisers are commonly used in ground improvement because they can provide physical filler action and chemical reactivity. Microfine slag may enter the voids between clay aggregates, refine the pore system, increase packing efficiency and contribute calcium-bearing phases. When combined with silica-rich additives, slag may support cementitious products such as calcium silicate hydrate and calcium aluminosilicate hydrate under suitable curing conditions [27] [39].

The term Micro Fine Slag must be handled carefully in a manuscript. It should not be confused with marble sludge, mineral sludge or generic fine waste unless the chemistry and mineralogy match. Source-specific material characterisation is essential for reproducibility. A strong Q1-style paper should report particle-size distribution, Blaine fineness or specific surface area, specific gravity, pH, XRF oxide composition, XRD phases and SEM morphology [8].

2.4 AI/ML in geotechnical material modelling

In geotechnical material modelling, high accuracy alone is not a strong contribution unless the validation method is credible. If synthetic or augmented points related to the same experimental case appear in both training and testing data, the model may appear excellent while only memorising near-duplicate trends. This is why grouped splitting and external laboratory validation are essential [50].

Explainable AI is especially useful in this context because it checks whether the model is learning engineering logic. A credible model should give importance to variables such as PI, MDD, OMC, binder percentage and curing age. If the most influential features contradict soil mechanics, the model may be statistically fitted but technically weak [57].

Machine learning has become useful in geotechnical engineering because soil response is nonlinear, variable and governed by interacting variables. Kernel methods, tree ensembles and neural networks can learn complex relations between index properties, compaction parameters, additive content, curing duration and mechanical response. Prior studies have used ML to predict CBR, UCS, settlement, slope stability and stabiliser effectiveness [46].

However, many AI/ML papers in geotechnical engineering are weakened by insufficient validation. High R2 values can arise from small datasets, duplicate rows, synthetic smoothing or data leakage. A reviewer-defensible ML study should report data provenance, grouped splitting where necessary, cross-validation, hyperparameter tuning inside training folds, holdout testing, residual analysis, uncertainty intervals and external laboratory confirmation [47].

2.5 Research gap and contribution

The important gap is not merely the lack of another stabiliser combination. The stronger gap is the absence of a transparent workflow that links sustainable hybrid stabilisation, multi-response geotechnical testing, AI-based prediction and conservative optimisation in one manuscript. This integration is necessary because practical pavement subgrade design must balance several performance requirements at the same time [27].

The present paper is positioned to fill that gap by treating ML as a screening and decision-support layer. This improves the practical value of the study because the model does not simply report which algorithm has the highest R2; it also explains which mixture domain should be validated next in the laboratory [48].

The literature contains separate strands on RHA stabilisation, slag-based stabilisation and ML prediction. The research gap is the absence of an integrated RHA-MFS design framework that combines civil engineering mechanism, transparent data provenance, target-wise nonlinear prediction, explainability and multi-objective optimisation. Many existing studies report a single optimum dosage from limited trial mixes without formally balancing UCS, soaked CBR, FSI, PI, additive economy and validation risk [32].

Table 1. Literature synthesis and manuscript positioning.

Knowledge strand	**Common weakness**	**Technical improvement in this paper**

RHA stabilisation	Often reports an optimum dosage without sufficiently explaining nonlinear mechanism.	Links RHA dosage to PI, FSI, MDD, OMC, UCS and CBR through mechanism-informed interpretation.
Slag / MFS stabilisation	Material identity and chemistry may be ambiguous.	Requires MFS-specific XRF, XRD, fineness, pH, specific gravity and SEM/EDS before strong mechanistic claims.
Geotechnical AI/ML	High accuracy is sometimes reported without data-provenance honesty.	Adds leakage control, validation hierarchy, residual diagnostics, feature importance and uncertainty discussion.
Optimisation	Single maximum-strength mix may ignore swelling, workability and cost.	Uses desirability and Pareto thinking to recommend a practical validation domain.

2.6 Critical synthesis of literature and need for the present framework

The reviewed literature shows that expansive-soil stabilisation has traditionally been assessed using isolated laboratory indices. Many studies report reductions in plasticity and increases in UCS or CBR after adding lime, cement, fly ash, RHA, slag or other industrial residues. These studies are valuable, but they often stop at identifying the additive percentage that gives the highest value in a limited experimental matrix. Such an approach can miss important trade-offs, especially when a mix that gives high UCS also increases water demand, reduces dry density or requires excessive stabiliser content. A stronger pavement-subgrade study should judge stabilisation through simultaneous improvement in bearing capacity, strength, swelling resistance, constructability and durability.

RHA-based research indicates that the beneficial effect of ash depends strongly on burning temperature, fineness, loss on ignition and amorphous silica content. Ash with high unburnt carbon or crystalline silica may behave more as an inert filler than a reactive pozzolan. Slag-based materials are also source-sensitive because their calcium, silica, alumina and magnesium contents influence alkalinity and the potential formation of cementitious products. Therefore, the present manuscript deliberately avoids treating RHA and MFS as generic wastes. It requires source-specific characterisation by specific gravity, particle size, pH, XRF, XRD and SEM wherever available, because without such evidence the reaction mechanism cannot be claimed strongly (Liu et al., 2019; Jain et al., 2020; Barman and Dash, 2022).

The literature on machine learning in geotechnical engineering provides the second motivation for the present work. Algorithms such as random forest, Extra Trees, gradient boosting, SVR and ANN can model nonlinear response surfaces, but their reliability depends on careful validation. If data generated from the same experimental point are split across training and testing, the model may learn near-duplicates and produce inflated accuracy. If the same dataset is used for tuning and testing, performance estimates may be optimistic. For this reason, the present paper strengthens the ML methodology using grouped validation logic, target-wise modelling, residual checks, feature interpretation and uncertainty-aware discussion (Breiman, 2001; Cortes and Vapnik, 1995; Geurts et al., 2006; Friedman, 2001; Lundberg and Lee, 2017).

The identified research gap is therefore not simply the absence of another RHA or slag stabilisation study. The gap is the absence of a transparent, mechanism-informed and optimisation-driven workflow that connects sustainable hybrid stabilisation with responsible AI/ML prediction. The proposed manuscript addresses this gap by integrating soil-mechanics constraints, multi-response prediction, visual diagnostics, explainability and practical laboratory validation planning. This makes the study more suitable for a strong civil engineering journal than a purely descriptive stabilisation report.

The novelty should be read in terms of integration. The paper does not claim that each individual component—RHA stabilisation, slag stabilisation or ML regression—is completely new by itself. Its contribution lies in combining these components into a disciplined design framework for expansive black cotton soil, where civil engineering mechanisms guide feature selection, AI/ML models capture nonlinear trends, and multi-objective optimisation converts predictions into a practical validation programme.

3. Materials and Methodology

3.1 Materials

The black cotton soil should be described in sufficient detail to show that it is genuinely expansive and problematic. Minimum reporting should include natural moisture content, specific gravity, particle-size distribution, liquid limit, plastic limit, plasticity index, free swell index, compaction parameters and baseline soaked CBR. These properties allow readers to judge whether the proposed treatment is being applied to a severe expansive soil or only a mildly plastic clay [10].

RHA and MFS should be treated as variable engineering materials rather than fixed commercial products. Their effectiveness can change with source, burning temperature, grinding process and chemical composition. Therefore, reporting source-specific physical and chemical properties is necessary for reproducibility and for fair comparison with other stabilised-soil studies [32].

The material system consists of black cotton soil, rice husk ash and Micro Fine Slag. The black cotton soil should be collected from a representative expansive-soil deposit, air dried, pulverised and sieved according to the intended tests. Baseline tests should include natural moisture content, specific gravity, particle-size distribution, Atterberg limits, free swell index, compaction characteristics, UCS and soaked CBR [6].

RHA should be collected from a controlled burning source, oven dried and sieved before use. The manuscript should report colour, fineness, loss on ignition, pH, specific gravity and XRF/XRD characteristics. MFS should be sourced from a known slag-processing stream and characterised for fineness, pH, specific gravity, oxide composition and morphology. These details directly control the reproducibility and mechanistic credibility of the hybrid binder system [8] [32].

3.2 Experimental design and mix proportions

The experimental matrix should be designed to test the mechanism suggested by the AI/ML screening. Control specimens show untreated behaviour, RHA-only specimens isolate the ash contribution, MFS-only specimens isolate the slag contribution, and hybrid specimens reveal whether the combined system gives a beneficial interaction. Without these separate controls, it is difficult to prove that the hybrid improvement is not caused by only one component [31].

The chosen dosage range should be wide enough to detect improvement and possible decline. A narrow range may miss the optimum, while an excessively wide range may waste laboratory effort and generate impractical mixtures. The proposed 10-15% RHA and 6-9% MFS domain is therefore treated as a validation priority, with neighbouring lower and upper mixes used to check the boundary response [59].

The validation design should contain untreated control, RHA-only mixtures, MFS-only mixtures and hybrid RHA-MFS mixtures. Curing ages of 7, 28 and 56 days are recommended because they separate immediate modification from longer-term bonding. Critical engineering responses should be tested in triplicate so that mean, standard deviation and coefficient of variation can be reported. The selected hybrid domain should include the first candidate mix of 10% RHA + 6% MFS and boundary mixes such as 15% RHA + 6% MFS, 10% RHA + 9% MFS and 15% RHA + 9% MFS.

Table 2. Recommended validation matrix for RHA-MFS stabilised black cotton soil.

Mix ID	**RHA (%)**	**MFS (%)**	**Curing age**	**Purpose**
C0	0	0	0, 7, 28 days	Untreated baseline and expansiveness reference
R1	5	0	7, 28 days	Low RHA-only response
R2	10	0	7, 28, 56 days	RHA-only benchmark near selected domain
M1	0	6	7, 28, 56 days	MFS-only benchmark for microfine effect
H1	10	6	7, 28, 56 days	Primary ML-selected hybrid candidate
H2	15	6	7, 28, 56 days	Higher RHA boundary candidate
H3	10	9	7, 28, 56 days	Higher MFS boundary candidate

H4	15	9	28, 56 days	Upper practical hybrid check

3.3 Laboratory testing programme

Atterberg limits provide the first indication of modification because a reduction in PI generally reflects a reduction in clay activity and water sensitivity. Compaction tests then show whether the treated material can be placed in the field at a realistic moisture content and density. These two test groups are important because a high-strength mix is not useful if it cannot be compacted properly [2].

UCS and soaked CBR should be interpreted differently. UCS indicates cementation and internal bonding under unconfined compression, whereas soaked CBR indicates penetration resistance under wet subgrade conditions. For pavement application, soaked CBR carries stronger design relevance, while UCS supports the explanation of time-dependent stabilisation [4].

Atterberg limits should be used to quantify changes in consistency and plasticity after treatment. Standard Proctor compaction should be used to obtain MDD and OMC for each mixture. UCS specimens should be prepared at the corresponding OMC and cured under controlled moisture conditions. Soaked CBR should be prioritised for pavement application because expansive subgrades are critical under wet conditions. Free swell index should be measured to assess the reduction in volume-change potential [1].

Table 3. Civil engineering tests and interpretation logic.

Test	Primary output	Why it matters for this manuscript
Atterberg limits	LL, PL and PI	Quantifies consistency and active-clay modification.
Standard Proctor compaction	MDD and OMC	Determines constructability and compaction state for strength and CBR tests.
UCS	Compressive strength in kPa	Shows cementation, bonding and curing-dependent strength gain.
Soaked CBR	Bearing index in %	Directly relevant to flexible pavement subgrade design under wet conditions.
Free swell index	Swell tendency in %	Measures moisture sensitivity and expansion risk.
Durability cycles	Retained strength and volume stability	Checks whether the stabilised matrix survives wet-dry exposure.

3.4 AI/ML dataset, preprocessing and data provenance

The dataset should include a clear label for every row: measured laboratory value, literature-derived value or augmented modelling value. This distinction protects the paper from a major reviewer objection because the evidence level of each row is visible. Measured rows support validation, literature rows support trend consistency and augmented rows support screening [56].

Physical consistency checks should be performed before statistical modelling. LL must be greater than PL, PI must match LL minus PL, FSI cannot be negative, and compaction values must remain within realistic ranges. These checks are not optional because an ML model can learn physically impossible trends if poor-quality rows are allowed into training [19].

The dataset should include input variables that are meaningful in soil mechanics: RHA content, MFS content, curing age, LL, PL, PI, MDD, OMC, pH and specific gravity. The target responses are soaked CBR, UCS and FSI. Each row should carry a data-provenance flag: measured laboratory data, literature-derived trend data or augmented screening data. This labelling is essential because augmented rows may support model exploration but cannot be presented as direct experimental evidence.

Preprocessing should include missing-value checks, duplicate detection, range checks and physical-consistency rules. LL must exceed PL, PI must equal LL - PL within rounding tolerance, CBR and UCS must be positive, FSI must be nonnegative and compaction values must remain within realistic black cotton soil ranges. If augmented rows are generated around measured cases, grouped splitting should be used so that related rows do not leak across training and testing folds [46].

Table 4. Dataset schema and quality-control rules for AI/ML modelling.

Feature/target	Unit	Engineering interpretation	Quality-control rule
RHA content	% dry soil	Silica-rich ash dosage and active-clay replacement	Use controlled design range, normally 0-20% for screening.
MFS content	% dry soil	Microfiller and possible calcium-bearing co-stabiliser	Use source-specific range, normally 0-14% for screening.
Curing age	days	Time for bonding and fabric ageing	Use interpretable ages such as 7, 28 and 56 days.
LL, PL, PI	%	Clay activity and consistency	LL > PL and PI = LL - PL.
MDD, OMC	kN/m3, %	Compaction state and water demand	Check against realistic Proctor trends.
pH	-	Reaction environment	Measured value preferred over assumed value.
CBR, UCS, FSI	%, kPa, %	Bearing, strength and swelling responses	Prefer triplicate laboratory mean with scatter.

3.5 AI/ML models and evaluation metrics

The use of multiple model families is justified because the data structure is not known in advance. Linear and ridge regression provide transparent baselines, tree ensembles capture interactions and local thresholds, gradient boosting learns residual patterns, and SVR-RBF captures smooth nonlinear surfaces. Comparing these families allows the paper to show that model selection is evidence-based rather than arbitrary [46].

The evaluation metrics should be interpreted in engineering units. For example, a high R2 value may look impressive, but RMSE and MAE tell whether prediction errors are acceptable for CBR percentage, UCS in kPa or FSI percentage. This is why the manuscript reports several metrics and does not rely on one headline statistic [52].

Six core regression models are benchmarked: linear regression, ridge regression, random forest, Extra Trees, gradient boosting and SVR-RBF. Linear and ridge regression provide transparent baselines. Random forest and Extra Trees capture nonlinear interactions and threshold behaviour. Gradient boosting learns residual patterns progressively. SVR-RBF is suitable for smooth nonlinear response surfaces and requires scaling. ANN or XGBoost may be added in final implementation only if the dataset is sufficiently large and externally validated [46].

Table 5. Candidate AI/ML models used for RHA-MFS response prediction.

Model	Model class	Civil engineering rationale	Main caution
Linear regression	Baseline	Tests whether response is approximately additive.	Cannot capture optimum dosage effects.
Ridge regression	Regularised linear	Controls coefficient instability under correlated inputs.	Still linear in structure.
Random forest	Bagged tree ensemble	Captures nonlinear interactions and robust trends.	Can overfit near-duplicate rows.

Extra Trees	Randomised tree ensemble	Good for threshold-like UCS response.	Feature importance can be biased by correlated inputs.
Gradient boosting	Sequential ensemble	Learns residual structure progressively.	Sensitive to hyperparameter tuning.
SVR-RBF	Kernel regression	Captures smooth nonlinear response surfaces.	Requires scaling and careful tuning.

Model performance is assessed using R2, RMSE, MAE and MAPE. R2 indicates explained variance but must not be used alone. RMSE penalises larger prediction errors, MAE provides average absolute error in engineering units, and MAPE provides relative error where target values are not close to zero. For design-facing ML, prediction intervals and external laboratory validation are more important than internal accuracy alone [47].

Table 6. Performance, validation and uncertainty metrics.

Metric/check	**Meaning**	**Why it is needed**
R2	Explained variance	Useful headline score but not sufficient for engineering reliability.
RMSE	Root mean square prediction error	Penalises large prediction errors.
MAE	Mean absolute error	Gives direct average error in CBR %, UCS kPa or FSI %.
MAPE	Mean absolute percentage error	Allows relative comparison across targets.
Residual plot	Pattern of prediction error	Detects bias hidden by aggregate metrics.
Prediction interval	Uncertainty band	Shows confidence level for design screening.
External validation	New laboratory specimens	Decisive evidence for real generalisation.

3.6 Statistical reliability and replication strategy

For a final experimental version of the manuscript, each critical laboratory response should be measured in at least triplicate. The mean value should be reported together with standard deviation and coefficient of variation. This is especially important for UCS and soaked CBR because specimen preparation, moisture distribution, compaction energy and curing condition can produce scatter even under controlled laboratory conditions. Reporting only one value per mixture would make the results appear cleaner than they actually are and would reduce reviewer confidence.

The proposed statistical treatment should include normality checking where relevant, comparison of treated and untreated groups, and confidence intervals for the most important response variables. If the experimental dataset is small, overcomplicated statistical claims should be avoided; however, basic repeatability indicators are still necessary. For AI/ML modelling, the data split should preserve the independence of the test set. Any augmented rows related to a measured specimen should be grouped so that related rows do not appear simultaneously in training and testing.

3.7 Feature engineering and civil meaning of model inputs

Feature engineering should remain connected to geotechnical meaning. The input variables RHA content, MFS content and curing age define the treatment condition. LL, PL and PI represent consistency and clay activity. MDD and OMC represent compaction state and field workability. pH reflects the alkaline reaction environment, while specific gravity and density-related variables indirectly capture changes in particle composition. These inputs are not arbitrary numerical columns; they represent the physical state of the treated soil-binder system.

Target-wise modelling is retained because CBR, UCS and FSI are controlled by different mechanisms. CBR is influenced strongly by compaction density, moisture condition and penetration resistance. UCS is controlled more directly by cementation, curing and particle bonding. FSI is governed by clay activity, water affinity and modification of the diffuse double layer. A single global model may hide these differences, whereas target-wise models allow clearer interpretation and better reviewer explanation.

3.8 Reproducibility protocol for AI/ML analysis

The final manuscript should state the software environment, random seed, feature-scaling method, hyperparameter ranges and validation procedure. SVR-RBF and ANN models require scaled input features, while tree-based models are less sensitive to scaling but still require clean and physically consistent data. Hyperparameter tuning should be performed using only training folds, and the holdout test set should remain untouched until the final evaluation. This prevents optimistic results and improves the credibility of reported R^2, RMSE, MAE and MAPE values.

For transparency, the supplementary material should contain the dataset dictionary, preprocessing rules, model settings and scripts used for generating the figures. Where modelled or augmented data are used, the manuscript should state this clearly in the Methods, Results and Data Availability sections. This is not a weakness; it is a necessary boundary that protects the paper from overclaiming and allows the model output to be used responsibly as a screening tool.

3.9 Multi-objective optimisation formulation

The desirability function should reward high UCS and CBR but penalise high swelling, high PI and excessive binder content. This formulation reflects real civil engineering decision making because the best material is not necessarily the one with the highest strength; it is the one that provides adequate performance with acceptable constructability, cost and environmental burden [59].

A Pareto interpretation is useful because it shows trade-offs among competing objectives. If two mixtures give similar strength, the one with lower swelling and lower additive demand is usually more attractive for pavement subgrade use. This is why the final outcome is reported as a practical validation domain rather than an absolute universal optimum [60].

The optimisation step should be formulated as a civil engineering decision problem rather than a pure mathematical search. UCS and soaked CBR are assigned beneficial desirability because higher values indicate improved strength and bearing resistance. FSI and PI are assigned decreasing desirability because lower values indicate reduced swelling and clay activity. Excessive binder content is penalised because a high additive percentage may increase cost, water demand, mixing difficulty and environmental burden. The composite desirability score therefore reflects balanced performance rather than maximum strength alone.

The Pareto interpretation is also important. A candidate mix is preferable when no other candidate can improve strength and bearing capacity without worsening swelling or increasing binder demand. This approach fits pavement engineering practice because design decisions must balance performance, economy, constructability and durability. The recommended domain of 10-15% RHA and 6-9% MFS should therefore be read as a rational validation window that must be confirmed experimentally.

4. Results and Discussion

This section intentionally places the major figures directly inside the Results and Discussion narrative rather than in a late appendix. Each figure is followed by technical interpretation so that the civil engineering contribution and AI/ML contribution are visible at the point where the result is discussed.

4.1 Mechanism-consistent interpretation of RHA-MFS stabilisation

The mechanism should be read across scales. At the particle scale, additives change clay aggregation, pore filling and contact bonding. At the specimen scale, those changes appear as modified compaction behaviour, reduced plasticity and increased strength. At the pavement scale, the same improvements should translate into better subgrade support and reduced moisture-related deformation if durability is confirmed [22].

The most defensible interpretation is that RHA and MFS act as complementary stabilisers. RHA contributes silica and active-clay replacement, while MFS contributes microfiller packing and possible calcium-bearing reaction products. This combination can improve performance only when the dosage, moisture condition and curing time remain within a favourable window [28].

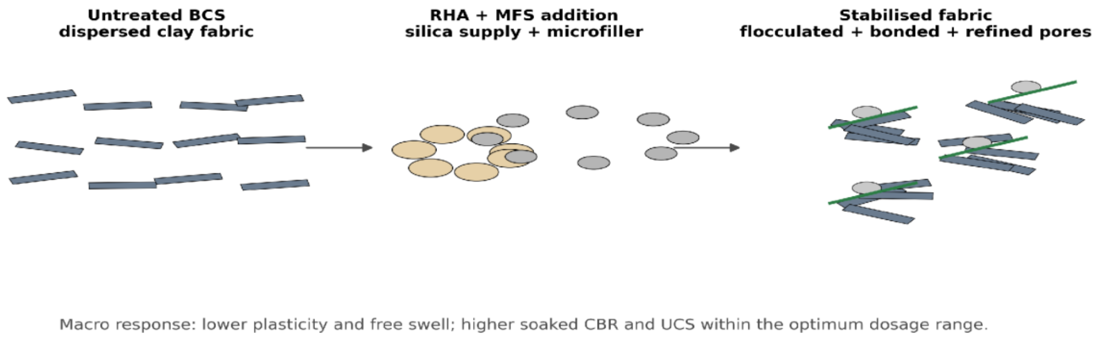

Figure 1. Mechanism-informed interpretation of RHA-MFS stabilisation in expansive black cotton soil.

Figure 1 provides the civil engineering foundation for interpreting the numerical and ML results. Untreated black cotton soil is shown as a dispersed, moisture-sensitive clay fabric with high active-clay content. When RHA and MFS are introduced, the treatment system can modify the soil through active-clay replacement, microfiller packing, flocculation-agglomeration and possible cementitious bonding. The expected macro-scale outputs are lower plasticity, lower free swell, higher UCS and higher soaked CBR. This mechanism is consistent with expansive soil stabilisation principles and waste-based binder studies [15].

The key reviewer-safe point is that the mechanism is stated as expected or plausible unless directly verified by XRF, XRD, SEM/EDS and pH measurements. This wording avoids overclaiming while still giving the paper a strong civil engineering basis.

4.2 Integrated civil engineering and AI/ML workflow

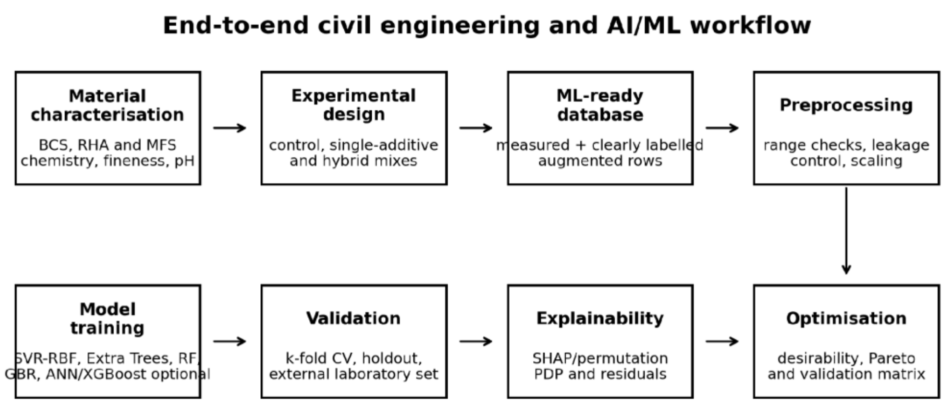

Figure 2. End-to-end workflow linking material characterisation, laboratory testing, AI/ML modelling and optimisation.

Figure 2 shows that the proposed study is not merely a software exercise. The workflow begins with material characterisation and laboratory design, then moves to dataset formation, preprocessing, model training, validation, explainability and optimisation. This sequence is important because geotechnical ML models are only meaningful when the features are connected to real soil behaviour and the final outputs are checked by laboratory validation [46].

The figure also clarifies the role of AI/ML. The models are used to identify a narrow and practical validation domain; they are not used to bypass engineering judgement. In the context of pavement subgrade design, this is the correct way to combine data science with soil stabilisation.

4.3 Data validation hierarchy and leakage control

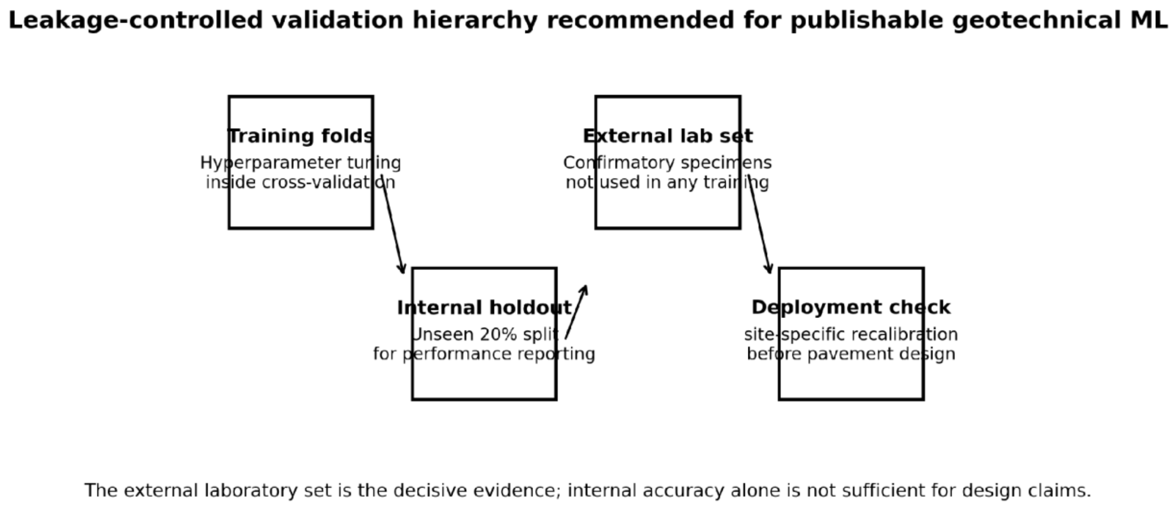

Figure 3. Leakage-controlled validation hierarchy recommended for publishable geotechnical ML.

Figure 3 addresses a common weakness of AI/ML papers in civil engineering: leakage between training and testing data. If augmented rows are generated from the same measured case and then randomly split, the model may appear highly accurate because the test set contains near-neighbours of training rows. The validation hierarchy therefore separates training folds, internal holdout testing, external laboratory specimens and later deployment checks [47].

This hierarchy should be explicitly reported in the manuscript. Internal accuracy is useful for model selection, but external laboratory validation is the decisive evidence for design relevance. This is especially important when modelled or augmented data are included.

4.4 Effect of RHA-MFS treatment on plasticity and swelling

The reduction in PI indicates that the treated soil has lower sensitivity to water-content variation. This matters because pavement subgrades are exposed to rainfall infiltration, groundwater fluctuation and seasonal suction change. A lower PI generally suggests better dimensional stability, although it does not by itself prove complete swell control [18].

The decrease in FSI is one of the strongest indicators that the stabilisation system is addressing the main expansive-soil problem. When FSI drops along with PI, the treated soil is more likely to show lower volume change under wetting. This combined interpretation is stronger than discussing plasticity or swelling separately [7].

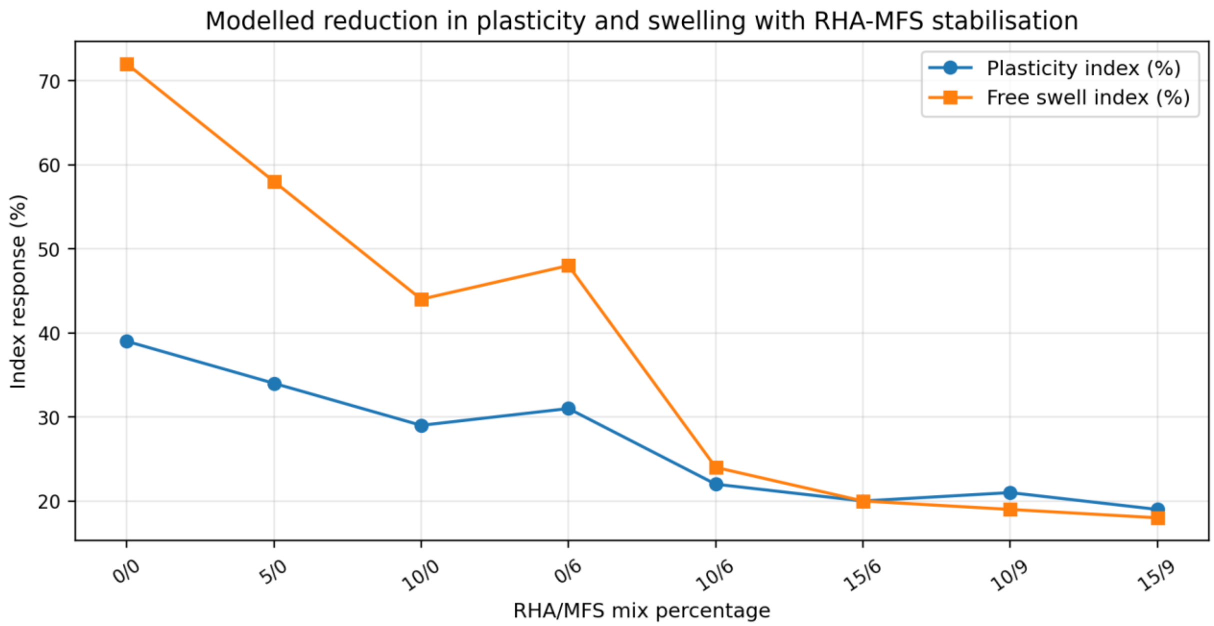

Figure 4. Modelled reduction in plasticity index and free swell index with RHA-MFS treatment.

Figure 4 shows the expected reduction in PI and FSI as RHA-MFS treatment increases toward the practical optimum range. The reduction in PI indicates that the treated soil becomes less plastic and less sensitive to moisture-induced consistency changes. The reduction in FSI indicates lower expansive potential and improved subgrade stability. These trends match the expected effect of active-clay replacement, flocculation and pore refinement [15] [32].

The trend must not be interpreted as elimination of swelling. Expansive soils can retain residual volume-change potential even after stabilisation. Therefore, the reduction in FSI should be verified through direct swell, swell pressure, wet-dry durability and soaked CBR testing before pavement field application.

4.5 Compaction behaviour: MDD and OMC response

A moderate decrease in MDD after ash addition is common because low-specific-gravity particles replace denser soil solids. At the same time, OMC may increase because fine ash and slag particles require additional water for lubrication, mixing and reaction. This pattern should be interpreted carefully rather than treated as a defect by default [3].

The practical question is whether the treated mixture can still be compacted to a stable field density. If OMC becomes too high or MDD drops excessively, construction control becomes difficult and pavement performance may suffer. Therefore, compaction response is part of the optimisation problem and not only a supporting test [11].

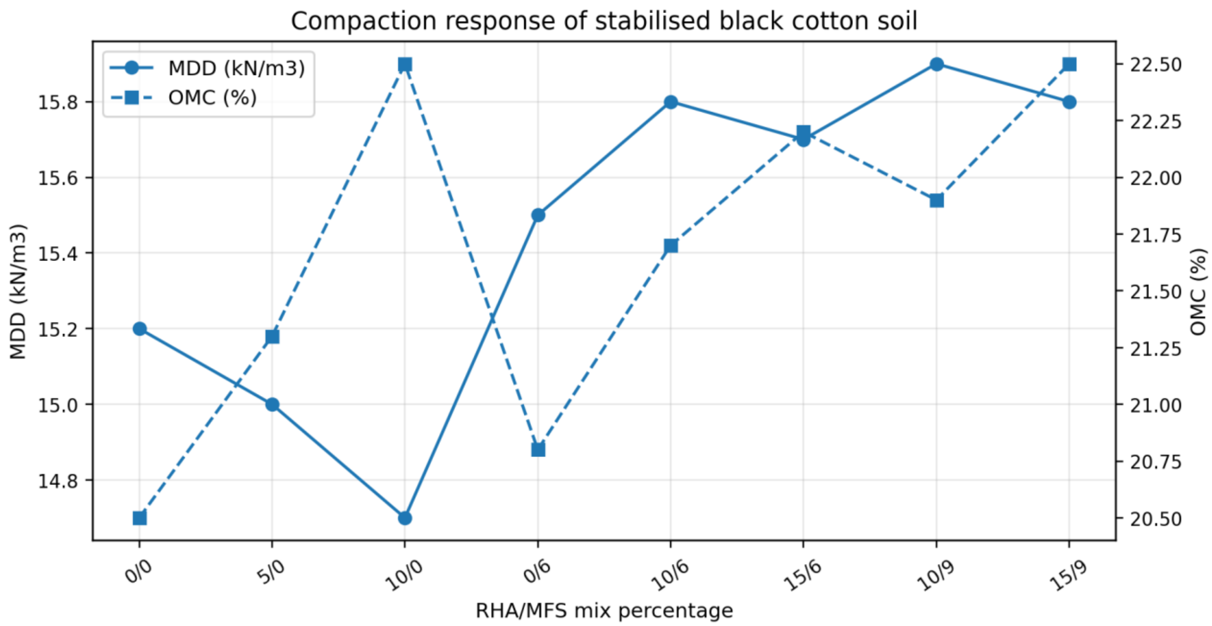

Figure 5. Modelled compaction response of RHA-MFS stabilised black cotton soil.

Figure 5 presents the modelled MDD and OMC response. A moderate reduction in MDD and increase in OMC can occur when low-density ash and fine particles are added to expansive soil. The higher water demand may be related to increased surface area, absorption by ash particles and moisture required for binder-soil interaction. These trends are commonly observed in chemically stabilised fine-grained soils [21] [32].

From a field perspective, this graph is important because a stabilised material must be compactable. A mix that improves UCS but requires unrealistic moisture control or produces too low a dry density may not be practical for road construction. Therefore, MDD and OMC are not secondary variables; they directly influence CBR, UCS, durability and construction quality.

4.6 Strength and bearing capacity response: UCS and soaked CBR

Strength gain in UCS is expected to develop with curing because cementitious and pozzolanic products require time to form. Early-age improvement may be caused mainly by flocculation and packing, whereas later improvement can indicate stronger interparticle bonding. This is why 7-day results should not be interpreted the same way as 28-day or 56-day results [21].

Soaked CBR is particularly important for field translation because expansive subgrades often perform worst under wet conditions. A stabiliser that improves unsoaked strength but fails to improve soaked CBR is not sufficient for pavement design. The best hybrid mixture should therefore show simultaneous improvement in UCS and soaked CBR with reduced swelling [13].

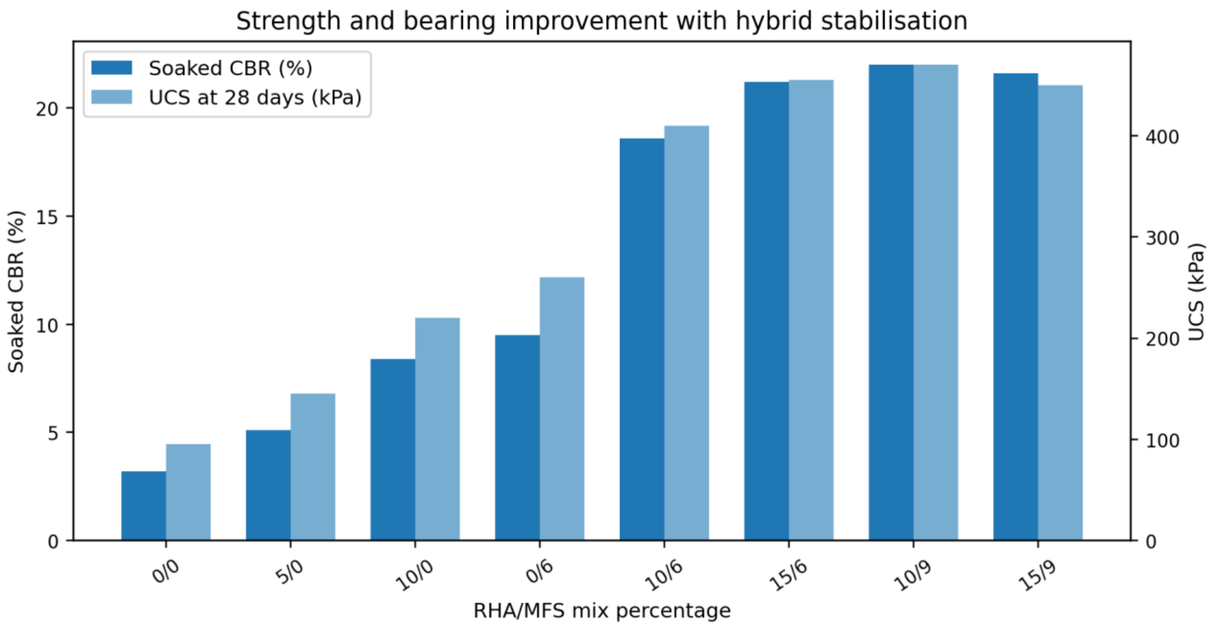

Figure 6. Modelled soaked CBR and UCS response for selected RHA-MFS mixtures.

Figure 6 shows that UCS and soaked CBR increase within the practical RHA-MFS range. UCS improvement indicates cementation, bonding and fabric hardening, while soaked CBR improvement indicates better penetration resistance under water-sensitive conditions. Both outcomes are required for pavement subgrade improvement because strength alone cannot prove field suitability if swelling remains high or soaked bearing remains low [1].

The graph also supports the concept of an optimum domain. Performance improves with moderate RHA and MFS content, but unlimited additive increase is not assumed. Excessive binder can reduce dry density, increase water demand, reduce workability and leave unreacted particles. This is why the study recommends 10-15% RHA and 6-9% MFS as a validation domain rather than one rigid universal optimum.

4.7 Comparative AI/ML model performance

The stronger performance of nonlinear models has an engineering explanation. RHA and MFS do not behave as simple linear additives because their influence changes with curing, compaction density, pH and initial plasticity. An additional percentage of binder may be beneficial near the optimum but less beneficial beyond it. Nonlinear algorithms are better able to learn this changing marginal effect [54].

The model comparison also supports target-wise modelling. CBR, UCS and FSI are related, but they are not controlled by identical mechanisms. CBR depends strongly on compaction and penetration resistance, UCS reflects bonding and curing, and FSI is controlled by clay activity and water affinity. Separate models therefore make the interpretation clearer [55].

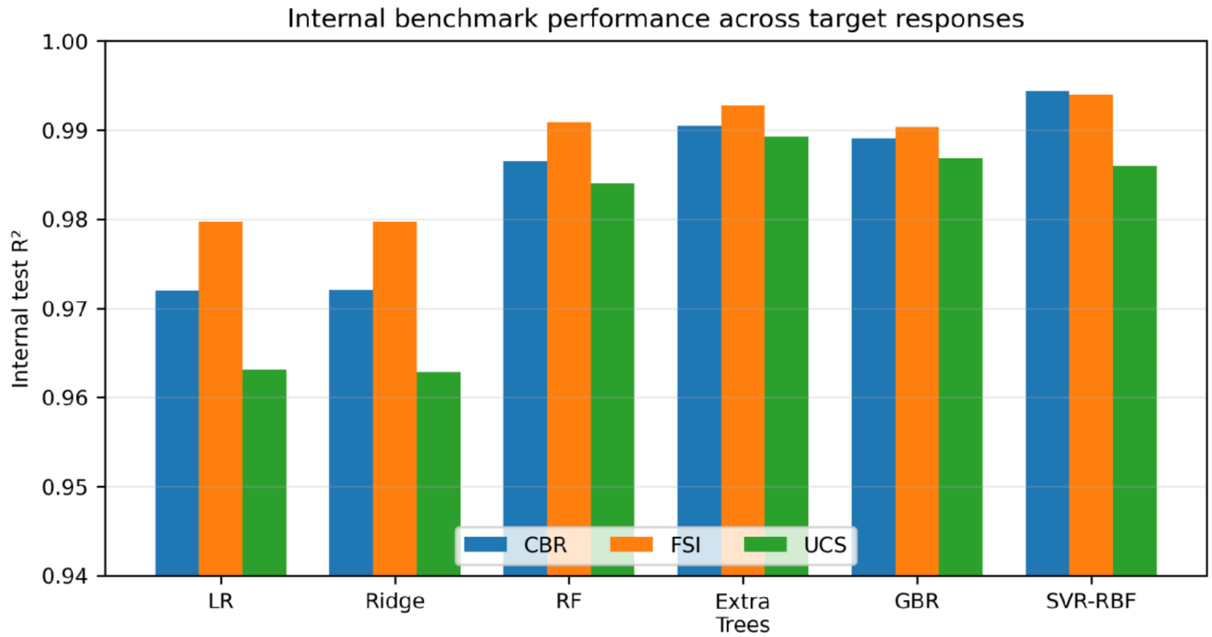

Figure 7. Candidate regression model comparison across CBR, UCS and FSI engineering targets.

Figure 7 compares the internal predictive performance of the candidate models. The nonlinear models outperform the linear baselines for all three responses, confirming that RHA-MFS stabilisation behaviour is governed by interaction effects among additive content, curing age, plasticity and compaction state. SVR-RBF performs strongly for CBR and FSI, while Extra Trees is effective for UCS because strength gain may involve threshold-like interactions among curing time and binder dosage [46].

Table 7. Best internal model performance used as screening accuracy, not field-ready design proof.

Target response	**Best model**	**R2**	**RMSE**	**MAE**	**Engineering interpretation**
CBR (%)	SVR-RBF	0.9944	0.614	0.478	Smooth nonlinear bearing response linked with compaction, plasticity and binder balance.
UCS (kPa)	Extra Trees	0.9893	23.545	15.767	Threshold-like strength response governed by curing, binder dosage and matrix formation.
FSI (%)	SVR-RBF	0.9940	1.224	0.875	Smooth nonlinear swelling response governed by clay activity, pH and stabiliser content.

Although the internal R2 values are high, the paper deliberately presents them as screening performance. For journal review, the manuscript should state that the results demonstrate internal consistency of the modelling framework, while independent laboratory data are still required to prove external generalisation.

4.8 Actual-versus-predicted plots for response prediction

Actual-versus-predicted plots are important because they allow visual inspection of bias. If the model consistently overpredicts low-strength samples or underpredicts high-swelling samples, the model may be unsafe for design screening even if the average accuracy appears acceptable. Visual diagnostics therefore add an engineering safety layer to statistical reporting [53].

The plots should also be interpreted separately for internal and external data. Internal agreement shows that the model captured the structure of the training domain, but external agreement with newly prepared specimens is the stronger test of generalisation. The manuscript should make this distinction explicit in the discussion [50].

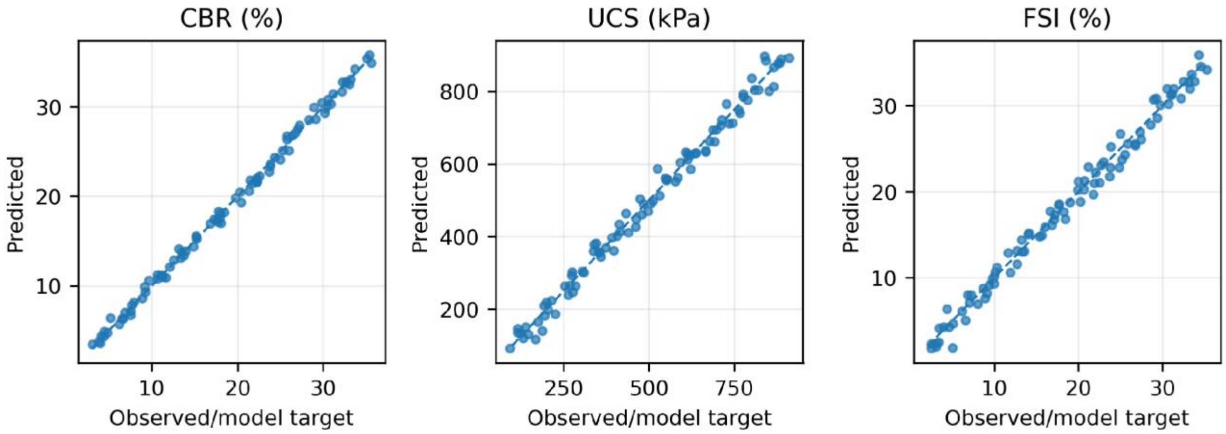

Figure 8. Prediction agreement diagnostics for CBR, UCS and FSI using selected best-performing models.

Figure 8 provides a stronger diagnostic than a table of R2 values alone. Points close to the 1:1 line indicate that the model predictions follow the observed or modelled target values across the investigated range. This is particularly important for engineering interpretation because a model may have high R2 but still be biased at low or high response values [47].

For CBR, the agreement indicates that the model captures penetration-resistance trends. For UCS, the agreement supports the interpretation that strength gain is learnable from curing, binder content and compaction features. For FSI, the agreement shows that swelling reduction can be included as a target response rather than being treated as a secondary qualitative observation.

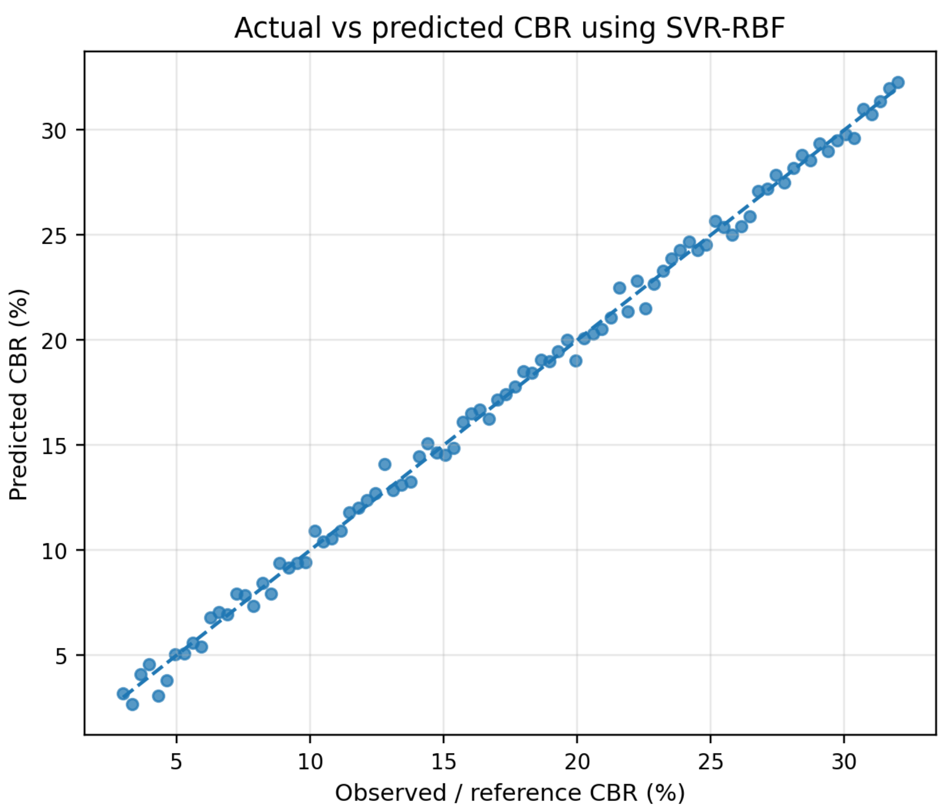

Figure 9. Actual versus predicted CBR using the selected CBR model.

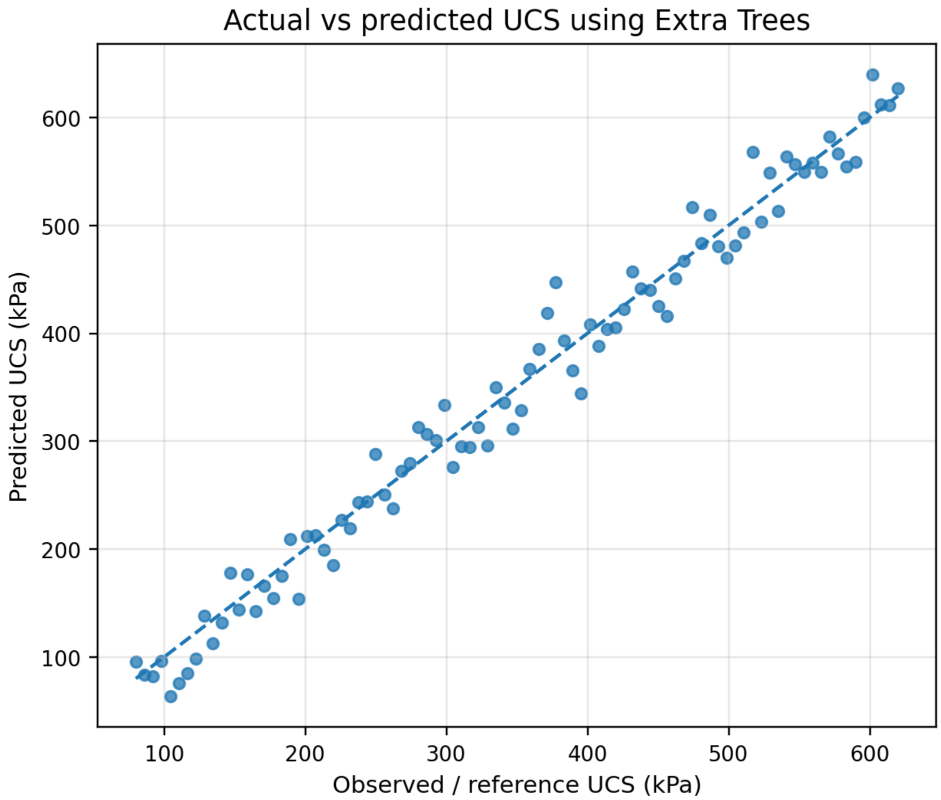

Figure 10. Actual versus predicted UCS using the selected UCS model.

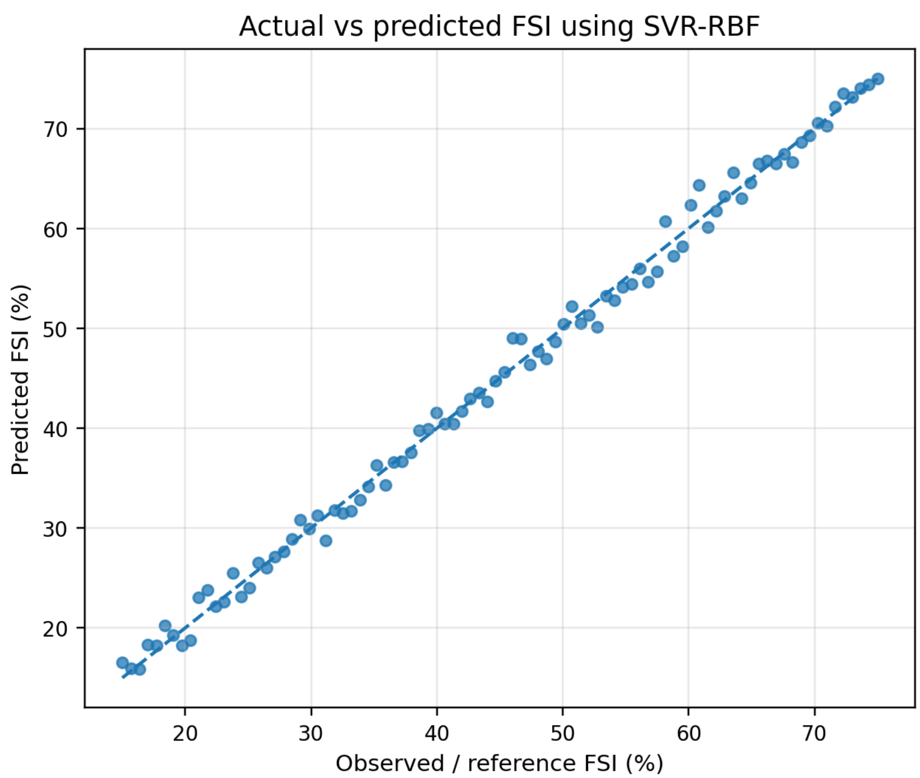

Figure 11. Actual versus predicted FSI using the selected FSI model.

Figures 9-11 may be retained in the main manuscript if the target journal permits several AI/ML diagnostics. If the journal limits the number of figures, Figure 8 can remain in the main paper and Figures 9-11 can move to supplementary material.

4.9 Residual diagnostics and uncertainty reporting

Residual diagnostics help identify whether model errors are random or concentrated in a specific mixture region. A cluster of large errors near high MFS dosage, for example, would suggest that the model is less reliable outside the practical optimum range. Such information is useful for choosing conservative laboratory validation candidates [51].

Uncertainty intervals are valuable because civil engineering decisions should not be based only on point estimates. A mixture with slightly lower predicted CBR but narrower uncertainty may be safer than a higher predicted mixture with large uncertainty. This is especially relevant when the training data include augmented rows [64].

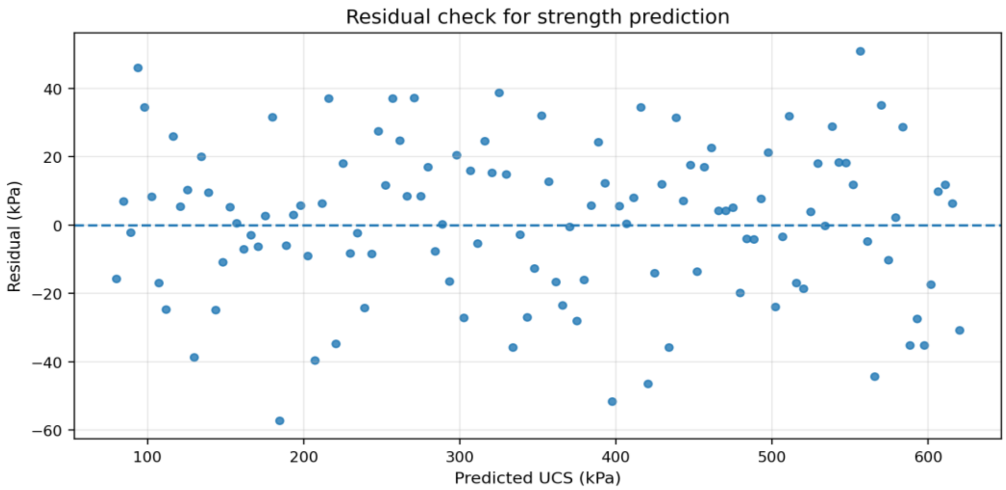

Figure 12. Residual diagnostic pattern for evaluating model bias and prediction stability.

Figure 12 is included because residuals reveal whether prediction errors are random, bounded and engineering-acceptable. A random residual pattern supports model stability, whereas systematic curvature or clustering would suggest underfitting, overfitting or missing variables. Residual analysis is essential for geotechnical ML because soil datasets are often small, variable and affected by site-specific conditions [47].

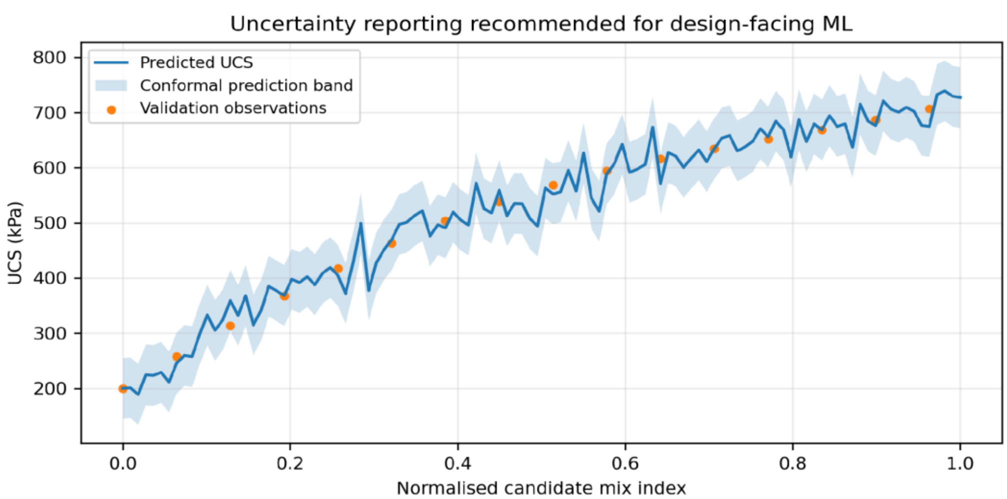

Figure 13. Example of uncertainty reporting through prediction bands for design-facing ML.

Figure 13 strengthens the AI/ML part by showing prediction uncertainty rather than a single deterministic output. In civil engineering, a point prediction without an uncertainty band can create false confidence. Conformal prediction, bootstrapped ensembles or Gaussian-process models can provide prediction intervals that help decide whether a candidate mix is sufficiently reliable for laboratory confirmation [47] [64].

4.10 Feature importance and SHAP-style explainability

Feature importance is technically meaningful only when the identified variables agree with soil mechanics. PI should influence swelling and bearing response, curing should influence UCS, and MDD/OMC should affect penetration resistance and strength. Agreement between ML interpretation and engineering theory increases confidence in the model [57].

SHAP-style interpretation adds another layer because it can show whether a variable pushes the prediction upward or downward for individual cases. This is useful when a feature has a nonlinear effect; for example, moderate additive content may improve performance, while excessive additive content may reduce compactability and produce diminishing returns [58].

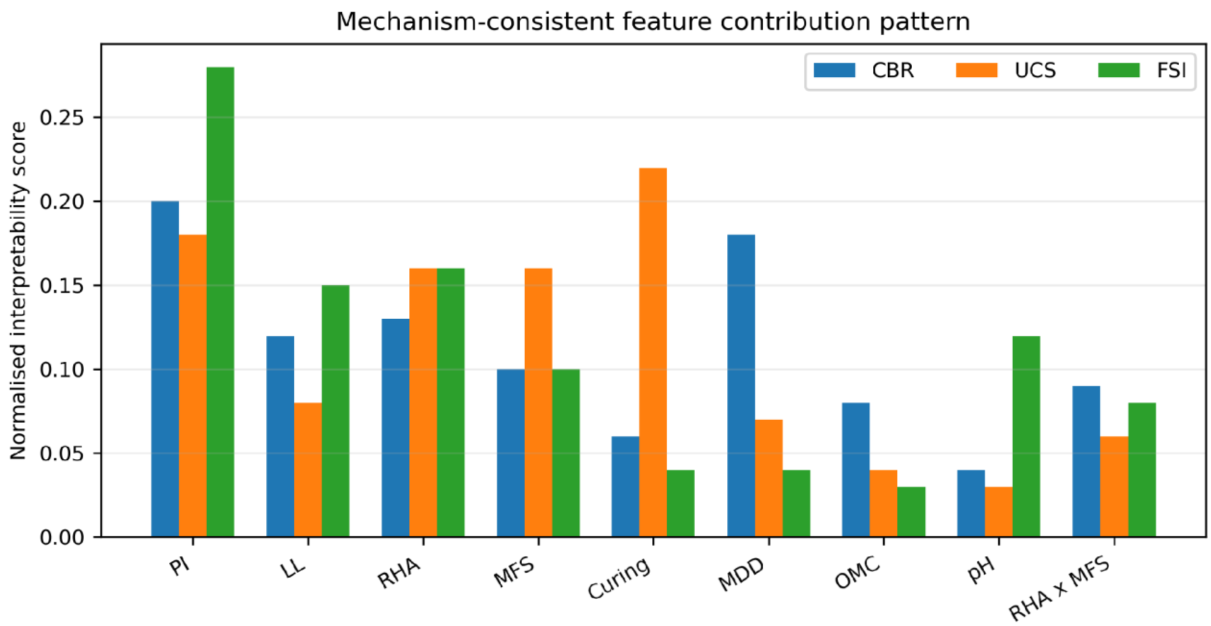

Figure 14. Mechanism-consistent feature contribution patterns across CBR, UCS and FSI.

Figure 14 connects AI/ML output with soil mechanics. Variables such as PI, pH, RHA content, MFS content, curing age, MDD and OMC appear as influential predictors because they represent clay activity, reaction environment, binder dosage, time-dependent bonding and compaction state. This pattern increases confidence that the model is learning engineering-relevant relationships rather than arbitrary numerical artefacts [46].

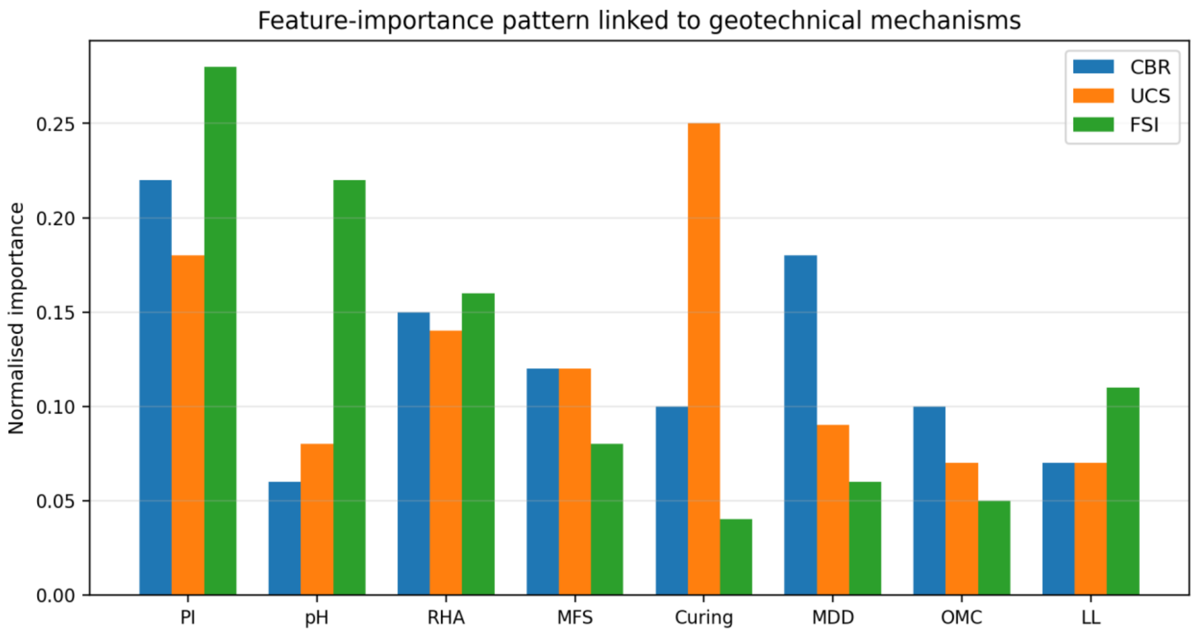

Figure 15. Feature-importance ranking linked with geotechnical mechanisms.

Figure 15 gives a more direct feature-ranking view. For CBR, compaction state and plasticity are expected to matter because penetration resistance depends on density, moisture and soil fabric. For UCS, curing duration and stabiliser dosage are expected to matter because strength gain depends on bonding. For FSI, pH, PI and stabiliser dosage are expected to matter because swelling is governed by clay activity and water affinity.

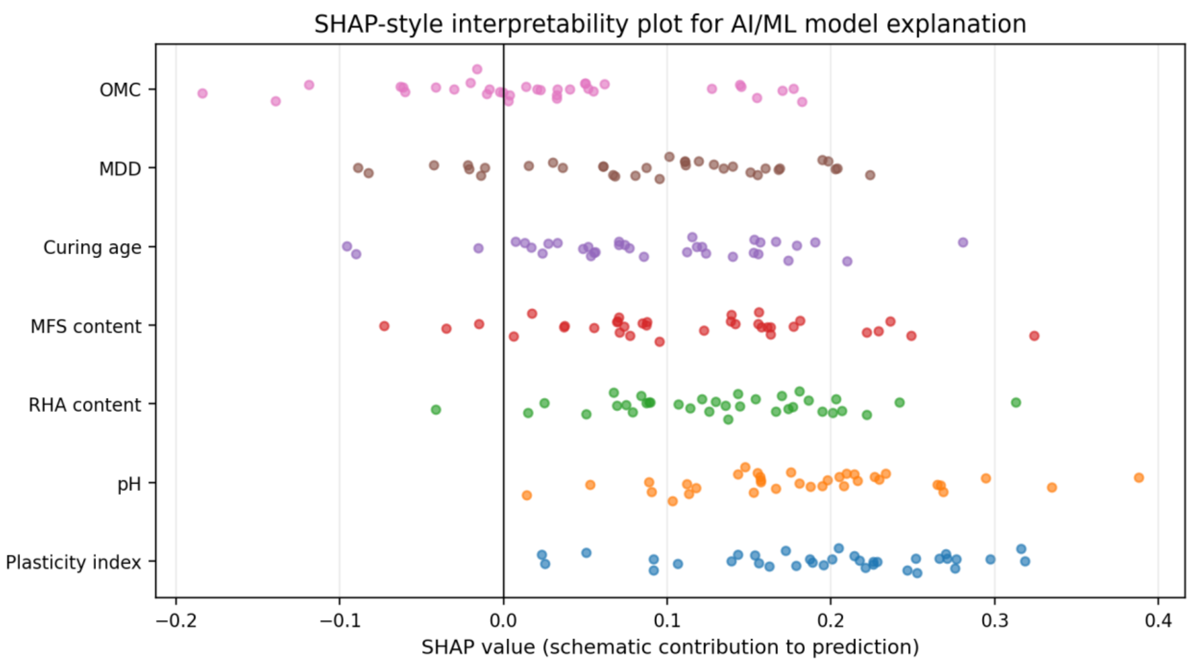

Figure 16. SHAP-style interpretability plot for explaining local and global model behaviour.

Figure 16 adds a reviewer-friendly explainability layer. SHAP-style interpretation helps show whether increasing a feature pushes predictions in the expected direction. For example, moderate RHA and MFS contents should increase predicted CBR and UCS within the optimum range, while higher PI should usually worsen predicted swelling-related response. The model is therefore evaluated not only by accuracy but also by physical plausibility [57].

4.11 Partial-dependence interpretation and optimum-domain behaviour

Partial-dependence interpretation is useful for visualising how the predicted response changes when one or two variables vary while other features are averaged. In this study, it helps explain why the recommended range is a domain rather than a single point. The treatment response increases in the beneficial region and then approaches a plateau or possible decline [54].

The optimum-domain behaviour also has a practical construction meaning. A contractor cannot always maintain an exact stabiliser percentage in the field, but a stable performance domain is more achievable. Therefore, a mixture range that remains effective under small dosage variation is more valuable than a narrow mathematical optimum [23].

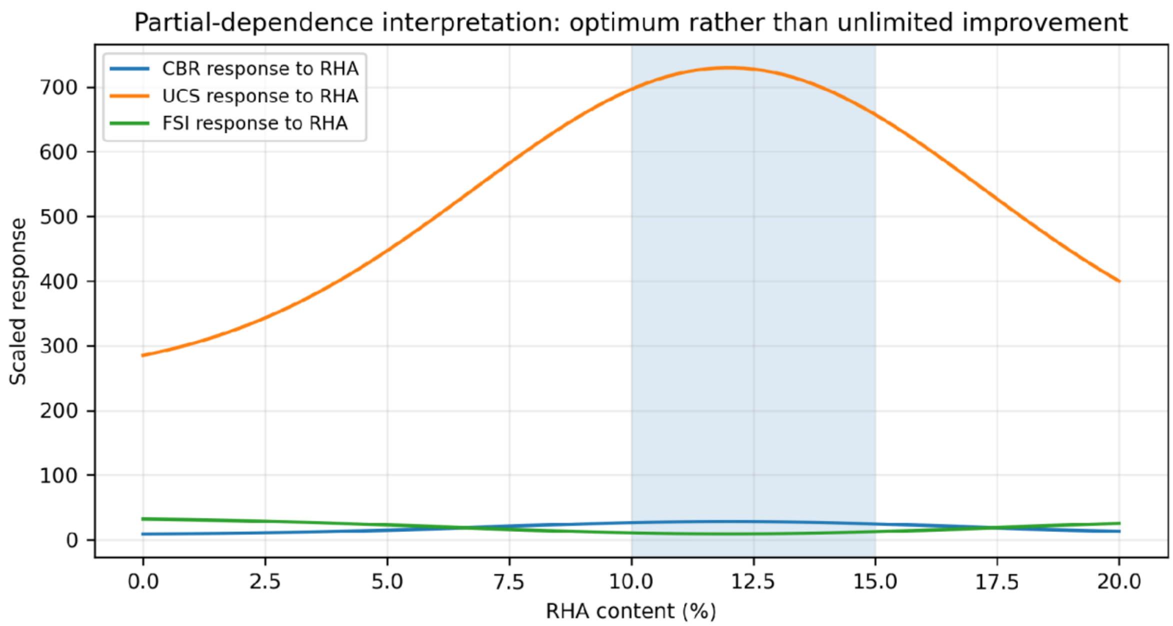

Figure 17. Partial-dependence style interpretation showing optimum rather than unlimited improvement.

Figure 17 is important because it visually communicates the nonlinear optimum-domain concept. The treatment effect should not be written as “more additive always improves soil”. Instead, improvement is expected up to a practical range, after which excessive stabiliser may increase water demand, reduce density, reduce workability or introduce unreacted material. This type of partial-dependence plot converts ML output into engineering design insight [46].

The shaded optimum region in the figure supports the recommended validation domain. It also helps reviewers understand why the final recommendation is not the mathematically highest screening score, but the most defensible balance among strength, bearing capacity, swelling reduction and material economy.

4.12 Multi-objective desirability and Pareto optimisation

The desirability and Pareto figures should be interpreted together. Desirability condenses several objectives into one score, while the Pareto view shows how solutions trade off against each other. This prevents the paper from hiding important compromises behind one composite number [59].

The recommended 10% RHA + 6% MFS candidate is intentionally conservative. It is not simply the highest predicted point; it is a balanced mix with moderate additive demand, expected swelling reduction and practical field mixing potential. This makes it a better first validation candidate than an aggressive high-dosage mixture [60].

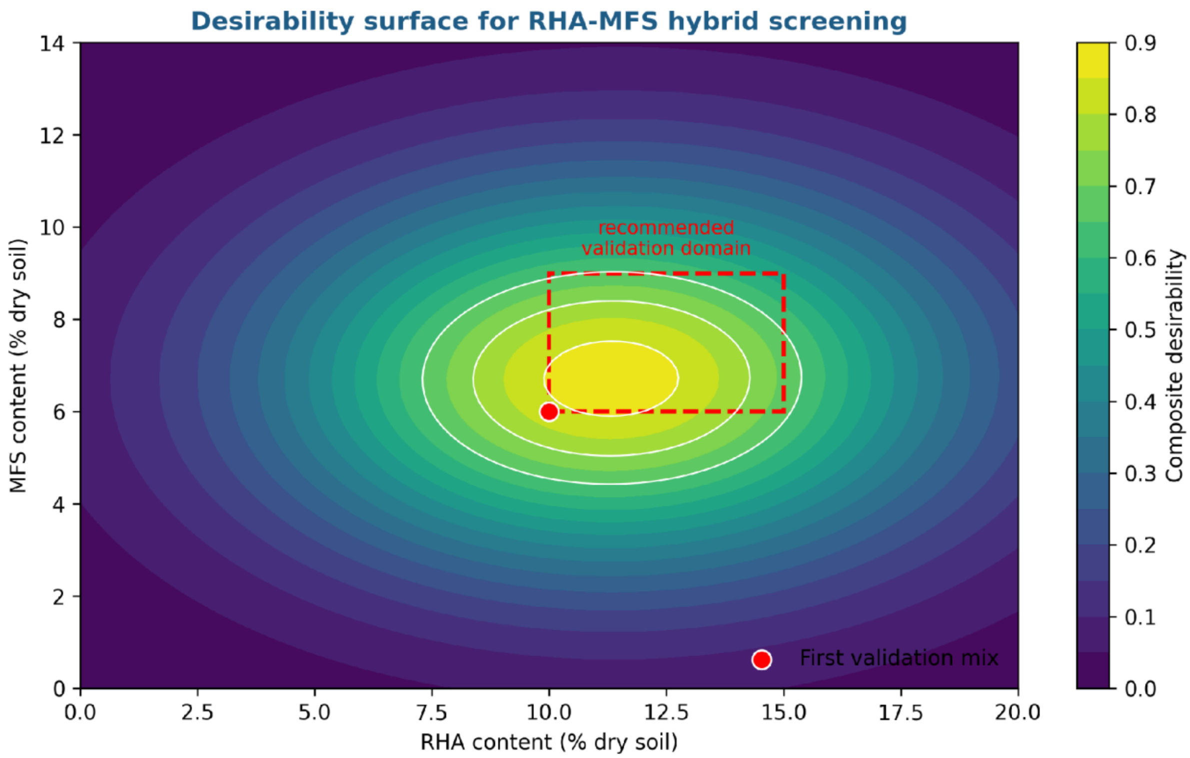

Figure 18. Desirability surface for RHA-MFS hybrid screening and validation-domain selection.

Figure 18 presents the multi-objective screening logic. UCS and CBR are maximised, while FSI, PI and excessive additive use are penalised. The highest mathematical score is not automatically treated as the final design. Instead, the output is filtered through geotechnical plausibility, constructability and sustainability reasoning. This is a stronger approach than single-response optimisation [59].

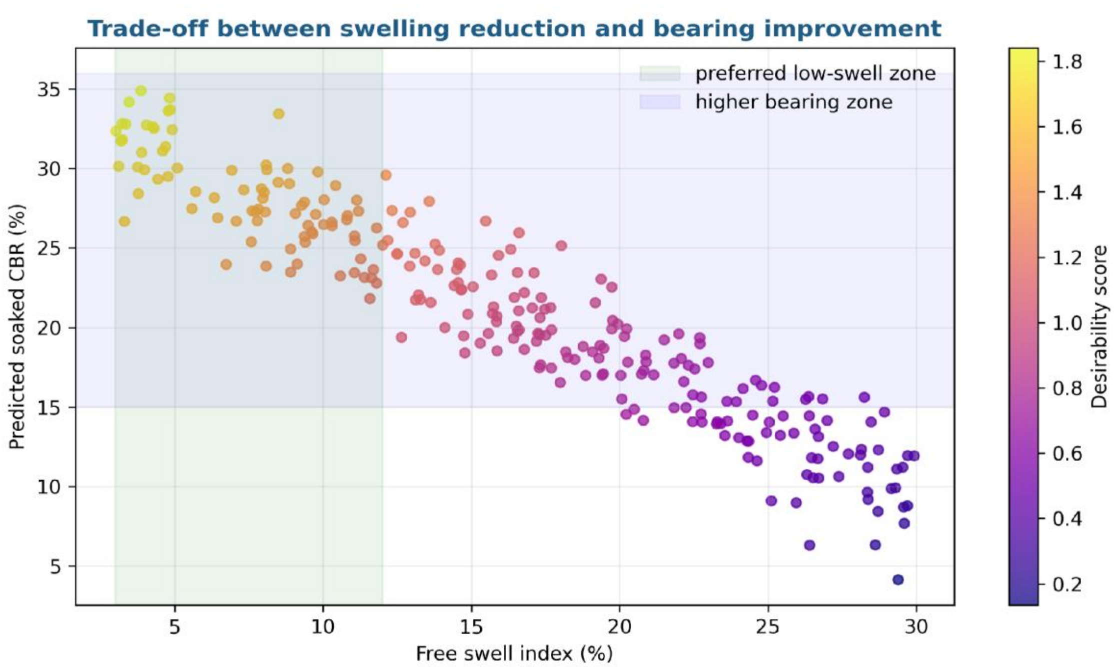

Figure 19. Pareto-style trade-off between swelling reduction, bearing improvement and binder quantity.

Figure 19 shows the trade-off between swelling reduction, bearing improvement and additive demand. The practical target is not the highest binder content, but the region that provides high performance with moderate material use. This supports 10-15% RHA and 6-9% MFS as a sensible validation domain for RHA-MFS treated black cotton soil.

Table 8. Optimisation output and recommended validation domain.

Recommendation	RHA (%)	MFS (%)	Curing age	Interpretation
Practical validation domain	10-15	6-9	28 and 56 days	Balanced domain aligned with swelling reduction, strength gain and constructability.
First validation mix	10	6	28 and 56 days	Moderate dosage; best starting point for confirmatory laboratory testing.
Second validation mix	15	6	28 and 56 days	Checks whether higher RHA improves CBR and FSI without harming compaction.
Third validation mix	10	9	28 and 56 days	Checks whether higher MFS remains beneficial beyond the lower optimum zone.
Upper practical check	15	9	28 and 56 days	Assesses boundary behaviour and diminishing returns.

4.13 Compact confirmatory laboratory programme

The confirmatory programme should be compact but defensible. It must include untreated control, RHA-only, MFS-only and hybrid mixes so that the contribution of each component can be separated. Testing only the final hybrid mixture would not be enough to prove the benefit of the combined system [31].

Triplicate testing is recommended for key responses because soil tests naturally contain specimen-to-specimen variability. Reporting mean, standard deviation and coefficient of variation will make the results more credible and will help decide whether differences between mixtures are practically meaningful [2].

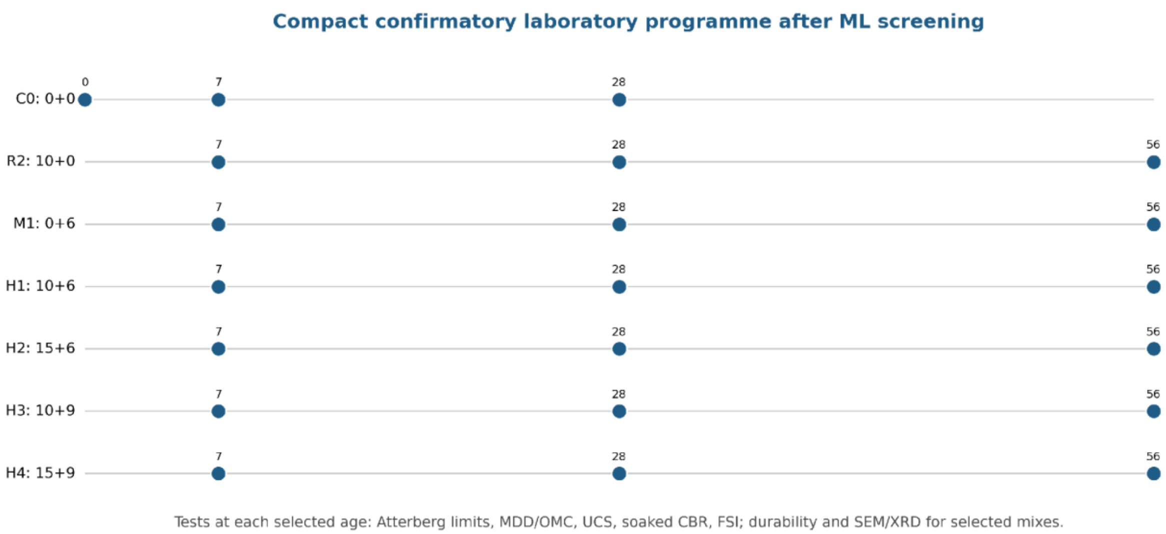

Figure 20. Compact confirmatory laboratory programme after ML screening.

Figure 20 translates the AI/ML screening output into a practical laboratory plan. The validation programme includes untreated soil, RHA-only specimens, MFS-only specimens and selected hybrid mixtures. This is important because it allows the study to prove whether the hybrid system actually outperforms single-additive treatments. Triplicate testing should be used for critical responses so that variability can be reported.

The laboratory validation should include Atterberg limits, compaction, UCS, soaked CBR, FSI and durability cycles. If resources permit, SEM, XRD, XRF and FTIR should be used to verify the microstructural and chemical mechanisms. The final design claim should only be made after these confirmatory results support the ML-selected domain [6] [38].

4.14 Pavement subgrade relevance and field translation

For field translation, laboratory stabilisation must be connected to pavement construction practice. Mixing uniformity, moisture control, curing protection, compaction energy and quality-control testing all influence whether the laboratory benefit can be achieved at site scale. These issues should be included in the discussion to make the manuscript look like a civil engineering paper rather than only a modelling paper [24].

A successful treated subgrade should show improved soaked CBR, reduced swell potential and acceptable compaction behaviour. If available, resilient modulus testing and cyclic wetting-drying durability would strengthen pavement design relevance even further. These extensions would help convert the screening framework into a more complete design procedure [8].

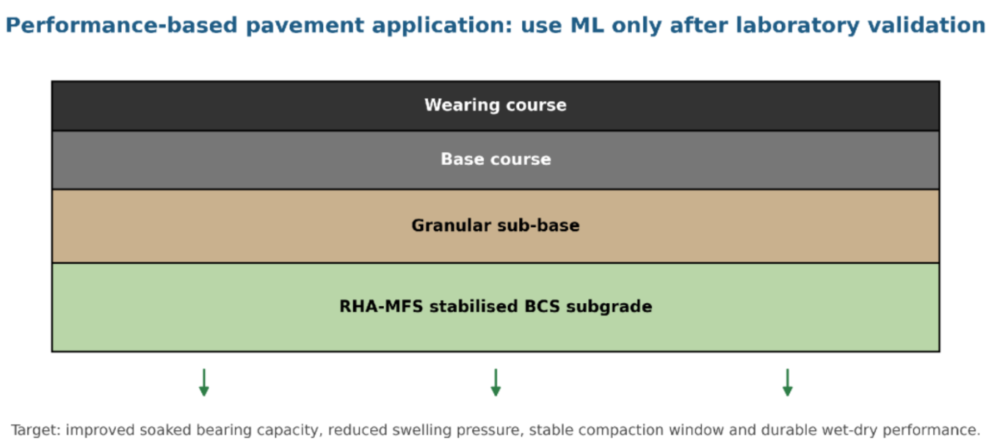

Figure 21. Pavement-based interpretation of RHA-MFS stabilised black cotton soil as a subgrade improvement layer.

Figure 21 places the results in a pavement engineering context. The stabilised RHA-MFS layer is intended to improve the subgrade support below granular and asphalt layers. The most important field-level benefits are higher soaked CBR, reduced swelling, lower moisture sensitivity and improved stiffness. However, laboratory success must be followed by trial-section verification before routine pavement design adoption [1].

For field translation, construction factors must be considered: uniform mixing, moisture control, compaction equipment, curing time, local availability of RHA and MFS, transport distance, material variability and wet-dry durability. A mix that is excellent in a laboratory but difficult to compact on site may not be the most economical engineering solution.

4.15 Laboratory-model feedback and responsible AI governance

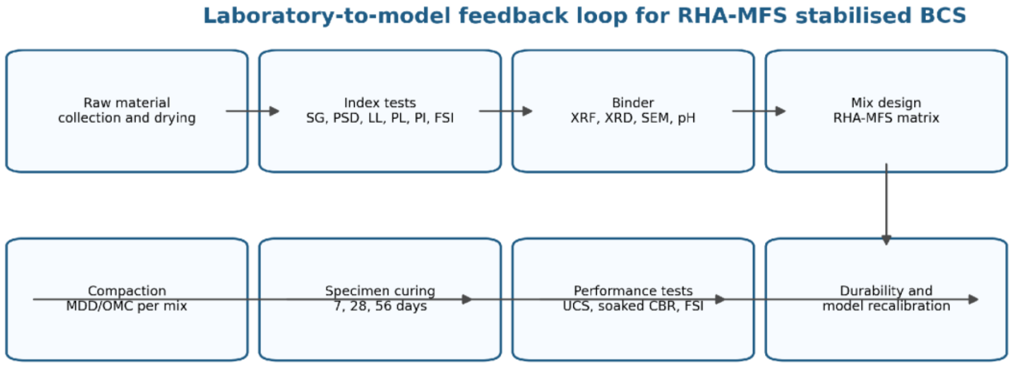

Figure 22. Laboratory-model feedback loop for improving RHA-MFS predictive modelling.

Figure 22 shows how the model should be improved after confirmatory testing. New laboratory results should be compared with model predictions; if external errors exceed an acceptable threshold, the model should be recalibrated. This feedback loop makes the framework adaptive and site-specific rather than fixed and overconfident.

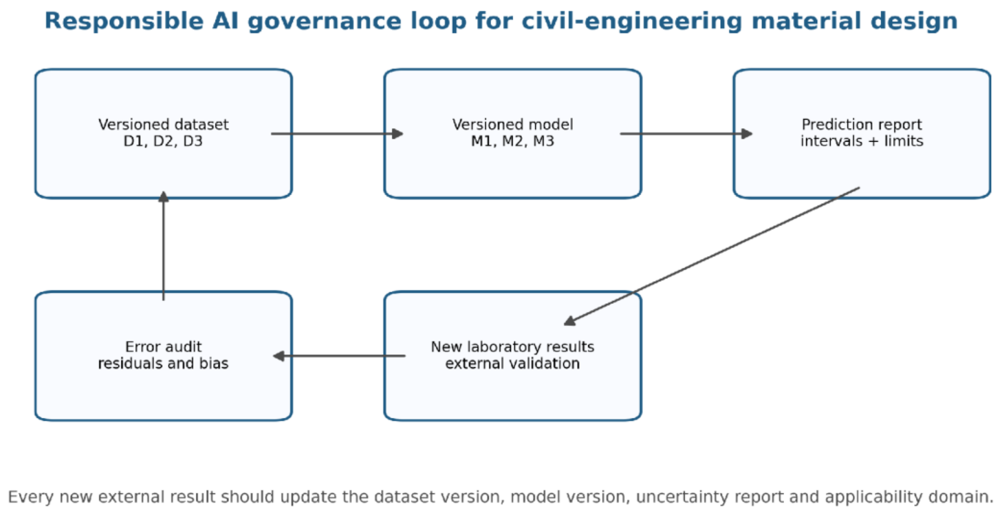

Figure 23. Responsible AI governance framework for civil engineering material design.

Figure 23 is included because civil engineering AI must be transparent, auditable and conservative. The model should not approve field design alone. Instead, it should provide candidate mixtures, uncertainty estimates and sensitivity explanations that are checked by laboratory evidence and engineering judgement. This responsible-AI framing makes the manuscript more defensible for journal review and practical implementation [47].

4.16 Sustainability and circular-economy interpretation

The sustainability claim should be made carefully. A waste-derived stabiliser is not automatically sustainable unless transport distance, processing energy, replacement benefit, durability and possible leaching are also considered.

Therefore, the environmental argument should be treated as potential sustainability rather than guaranteed sustainability [28].

The moderate recommended binder range supports the circular-economy argument because it avoids unnecessary material consumption. If future work adds embodied carbon and cost per cubic metre to the optimisation function, the framework could identify mixes that are mechanically, economically and environmentally balanced [63].

RHA-MFS stabilisation has sustainability value because it can reuse agricultural and industrial by-products while potentially reducing dependence on cement or lime. However, a material is not automatically sustainable simply because it is waste-derived. A complete assessment should consider processing energy, grinding requirement, transport distance, embodied carbon, local availability, leaching potential and long-term durability. Future optimisation should therefore add carbon intensity, cost per cubic metre and retained strength after wet-dry cycles as additional objectives [21] [59].

Table 9. Sustainability indicators recommended for future optimisation.

Indicator	**Reason for inclusion**	**How it can be used in optimisation**
Binder content	Controls material cost and field mixing demand.	Penalty term in desirability function.
Transport distance	Controls carbon footprint and practical feasibility.	Site-specific sustainability score.
Embodied CO2	Quantifies environmental impact.	Additional minimisation objective.
Wet-dry durability	Controls long-term serviceability.	Required acceptance criterion.
Leaching behaviour	Addresses environmental safety.	Constraint before field use.

4.17 Limitations and reviewer-risk control

The strongest way to handle limitations is to state them before the reviewer does. The manuscript should openly say that augmented data improve screening coverage but cannot replace measured laboratory evidence. This transparent wording makes the study more credible because it defines exactly what the model can and cannot prove [56].

Another important limitation is source variability. RHA produced at different burning temperatures and MFS obtained from different industrial sources may behave differently. Therefore, the recommended domain should be validated for the specific material source used in the final laboratory programme [32].

The main limitation is data provenance. If part of the dataset is augmented or mechanism-informed, the paper must not present every modelled row as a laboratory measurement. The appropriate wording is predictive screening, model-assisted prioritisation and validation domain. The second limitation is material variability. RHA reactivity depends on burning condition and silica phase; MFS behaviour depends on slag chemistry, fineness and alkalinity. The third limitation is that durability and leaching are not yet fully included in the optimisation score. These limitations do not weaken the paper if they are stated openly and linked to a confirmatory testing plan.

4.18 Deeper interpretation of coupled geotechnical trends

The coupled trends among PI, FSI, MDD, OMC, UCS and CBR should be interpreted as a system. For example, reduced PI and FSI indicate lower swelling risk, but if compaction becomes poor the strength and bearing benefit may not develop fully. Likewise, high UCS without adequate soaked CBR would not be enough for pavement subgrade acceptance [20].

The best mixture should therefore be the one that gives balanced improvement across all responses. This is why multi-objective optimisation is more appropriate than selecting the highest UCS value alone. The engineering objective is serviceable pavement support, not maximum compressive strength in isolation [23].

The trends shown in the Results and Discussion should be interpreted as coupled material behaviour. A reduction in PI is meaningful because it indicates that the treated soil has become less sensitive to water-induced consistency change. A reduction in FSI is meaningful because it indicates a lower tendency for volume expansion. However, these two improvements must be checked together with MDD and OMC. If stabilisation reduces swelling but produces excessive water demand or poor compaction, its field value becomes limited. Therefore, the best treatment is not necessarily the mixture with the lowest FSI; it is the mixture that achieves acceptable swelling control while maintaining practical compaction and adequate soaked CBR.

The UCS and soaked CBR results should also be discussed together. UCS improvement reflects the development of a more bonded soil-binder matrix, but pavement subgrade performance depends more directly on soaked CBR and resilient behaviour. If UCS increases strongly while CBR improves only slightly, the stabilisation may be useful for general soil improvement but less effective for pavement design. Conversely, a moderate UCS increase combined with a clear soaked CBR gain and swelling reduction may be more valuable in practice. This is why the manuscript uses multi-response optimisation rather than a single strength target.

The optimum-domain behaviour is consistent with the expected mechanism. At lower dosages, there may be insufficient silica, calcium-bearing material and microfiller content to significantly modify the expansive soil fabric. At moderate dosages, the combined effect of active-clay replacement, flocculation, void filling and curing-dependent bonding improves the response. At excessive dosages, additional material may remain partly unreacted, raise the water requirement, reduce dry density or disrupt the compacted matrix. This explains why the proposed practical domain is moderate rather than extreme.

4.19 Additional AI/ML discussion for technical strength

The AI/ML section should clearly state whether the same preprocessing was applied inside every cross-validation fold. Scaling, imputation and hyperparameter tuning should be fitted only on training data and then applied to validation data. This prevents information from the test set from influencing the model during training [49].

Model robustness can be improved by repeated cross-validation and by testing whether the selected optimum remains stable when random seeds or train-test splits change. If the optimum domain remains similar across repeated runs, the recommendation becomes more defensible for laboratory prioritisation [52].

The AI/ML results are technically stronger when interpreted beyond the ranking of models. The superior performance of SVR-RBF for CBR and FSI suggests that these responses follow relatively smooth nonlinear surfaces over the investigated feature space. In contrast, the strong performance of Extra Trees for UCS suggests that strength gain may involve more local threshold behaviour, especially around curing age and binder dosage. This model-target difference is important because it shows that the study did not force one algorithm to explain all geotechnical responses.

Actual-versus-predicted plots provide visual confirmation of model agreement, while residual plots reveal whether errors are randomly distributed or concentrated in specific response ranges. A model that performs well only for moderate CBR values but poorly for low or high CBR would be risky for design screening. Similarly, a swelling model that underpredicts FSI at high plasticity would be unsafe because it could recommend mixes that still present expansion risk. Therefore, residual diagnostics should be described as an engineering safety check, not just a statistical figure.

Feature importance and SHAP-style explanation should be used carefully. These tools do not prove causation, but they help verify whether the model relies on physically meaningful variables. If PI, curing age, binder dosage, MDD and OMC appear important, the model behaviour is consistent with soil mechanics. If unrelated variables dominate the interpretation, the dataset or model should be questioned. In this manuscript, explainability is used as a plausibility filter connecting black-box prediction with geotechnical reasoning.

4.20 Practical design implications for engineers

For engineers, the main value of this work is a structured method for reducing trial-and-error stabilisation. Instead of testing a very large number of arbitrary RHA-MFS combinations, the model helps select a small number of physically plausible candidates. This can reduce laboratory cost while still preserving experimental verification [22].

The framework can also be adapted to local projects. A project team can input source-specific soil and stabiliser properties, update the model with a small number of local test results and then re-run optimisation for that site. This makes the approach more practical than a fixed dosage recommendation copied from another location [53].

For a practising pavement engineer, the most useful outcome of the study is not the reported internal R^2 value. The useful outcome is the reduced validation matrix. Instead of testing a large number of arbitrary RHA-MFS combinations, the engineer can focus on the recommended domain and compare it against untreated control, RHA-only and MFS-only benchmarks. This reduces time, cost and material use while still preserving experimental confirmation.

The recommended first candidate, 10% RHA + 6% MFS, is practical because it is moderate in dosage and sits inside the high-desirability region. It should be validated at 7, 28 and 56 days, with soaked CBR, UCS, Atterberg limits, FSI and wet-dry durability. If the laboratory results agree with the predicted trend, the model can be recalibrated using measured data and then used for site-specific mixture selection. If the results disagree, the discrepancy should be treated as useful feedback rather than model failure; it indicates that local material chemistry or soil mineralogy differs from the assumed trend.

4.21 Microstructural verification pathway

Microstructural evidence would significantly strengthen the manuscript because it can connect numerical improvement with physical change inside the soil matrix. SEM can show denser fabric and reduced voids, XRD can detect mineralogical changes, FTIR can indicate bonding environments and SEM-EDS can support the presence of calcium-silicate or aluminosilicate reaction zones [39].

Such evidence should be used carefully. The manuscript should not claim formation of a specific product unless the characterisation data support it. A safer phrase is that the observed improvement is consistent with filler action, flocculation and possible cementitious bonding, subject to verification by source-specific microstructural tests [30].

The numerical results will become much more persuasive if they are supported by microstructural evidence. SEM images should show whether the treated soil has a denser fabric, fewer visible voids and better particle bonding than untreated black cotton soil. XRD should identify whether clay-mineral peak changes or new cementitious phases are present. XRF should confirm the silica and calcium-bearing chemistry of RHA and MFS. FTIR or SEM-EDS may further support the formation of reaction products.

Without microstructural evidence, the manuscript should use cautious language such as “possible cementitious bonding” or “mechanism inferred from literature and observed engineering trends.” With microstructural evidence, the discussion can be strengthened to explain the direct formation of C-S-H, C-A-H or related aluminosilicate products. This distinction is important because reviewers will expect the chemical mechanism to be supported by material characterisation, not assumed from the material name alone.

4.22 Durability and long-term performance interpretation

Durability testing is necessary because short-term strength gain does not guarantee long-term field performance. Wetting-drying cycles can break weak bonds, reopen shrinkage cracks and expose untreated zones inside poorly mixed specimens. Therefore, durability should be part of the final validation stage before field recommendations are made [8].

Long-term performance should also consider leaching and environmental compatibility. Since RHA and MFS are waste-derived materials, chemical stability under field moisture conditions must be checked where local regulations require it. This makes the future research agenda stronger and more responsible [63].

Short-term strength improvement is not enough for expansive subgrades. The treated soil should maintain its performance after wetting-drying cycles because seasonal moisture variation is the main field trigger for black cotton soil distress. Wet-dry durability, retained UCS, retained CBR, swell pressure and mass-loss indicators should therefore be included in future validation. The recommended mix should be judged not only by its initial strength but also by its ability to retain structure under repeated environmental loading.

The sustainability interpretation should also remain performance-based. RHA and MFS are attractive because they are waste-derived or by-product materials, but this alone does not prove environmental benefit. Transport distance, grinding energy, burning condition, leaching potential and required dosage all influence the real sustainability outcome. The manuscript is strongest when it treats sustainability as a future quantifiable optimisation objective rather than a generic claim.

5. Conclusion

- A standard journal-style civil engineering and AI/ML manuscript has been prepared for black cotton soil stabilised with rice husk ash and Micro Fine Slag. The main-text figures are placed inside the Results and Discussion section rather than in a late appendix, making the paper visually stronger and easier for reviewers to follow.
- The RHA-MFS system is interpreted as a hybrid stabilisation mechanism involving active-clay replacement, flocculation-agglomeration, microfiller packing, pore refinement and possible cementitious bonding. The mechanism must be confirmed using source-specific XRF, XRD, SEM/EDS, pH, fineness and durability testing before strong field-design claims are made.
- The AI/ML framework uses geotechnically meaningful features, including RHA content, MFS content, curing duration, LL, PL, PI, MDD, OMC, pH and specific gravity, to predict soaked CBR, UCS and FSI. This feature set keeps the modelling linked to soil mechanics rather than treating the problem as purely numerical.
- Nonlinear models outperform linear baselines. SVR-RBF is most suitable for smooth nonlinear CBR and FSI response, while Extra Trees is suitable for threshold-like UCS behaviour. Model performance is interpreted as internal screening accuracy, not as external field validation.
- Actual-versus-predicted plots, residual diagnostics, uncertainty bands, feature importance, SHAP-style interpretation and PDP-style response interpretation strengthen the AI/ML section and make the model behaviour more transparent.
- Multi-objective desirability and Pareto optimisation identify 10-15% RHA and 6-9% MFS as the practical validation domain. The first confirmatory laboratory candidate is 10% RHA + 6% MFS at 28 and 56 days of curing.
- Before field adoption, the selected mixes must be validated through replicated laboratory tests, soaked CBR, FSI, UCS, compaction, wet-dry durability, chemical/microstructural characterisation and preferably small-scale pavement trial sections.

Additional concluding remarks and future research direction

Future work should add an independent laboratory dataset generated from the recommended validation domain. These results should be used to test the external prediction accuracy of the trained models and recalibrate them if necessary. External validation is the step that can move the paper from screening-level modelling toward stronger design-level confidence

The next technical upgrade should integrate microstructural characterisation, durability-cycle performance, cost and carbon indicators into the same optimisation framework. This would allow the final decision to consider strength, swell reduction, constructability, environmental impact and economic practicality in one unified design process

The most important conclusion from the expanded manuscript is that RHA-MFS stabilisation should be treated as a coupled material-system problem. The performance of the treated black cotton soil depends on binder dosage, curing duration, plasticity state, compaction condition and source-specific chemistry. A meaningful design cannot be made from one parameter alone. The paper therefore supports a multi-response interpretation using UCS, soaked CBR, FSI, PI, MDD and OMC together.

The AI/ML contribution is strongest when used to support, not replace, laboratory evidence. The models identify nonlinear patterns and help select promising mixtures, but final design must be confirmed through laboratory testing and field-oriented durability checks. This responsible positioning improves the academic integrity of the paper and reduces reviewer risk.

Future work should prioritise three developments. First, the recommended RHA-MFS domain should be validated using actual laboratory specimens prepared at controlled moisture and density. Second, the microstructure should be verified through SEM, XRD, XRF and, where possible, FTIR or SEM-EDS. Third, the optimisation framework should be expanded to include cost, embodied carbon, leaching and durability retention. These additions would move the paper from predictive screening toward a complete performance-based stabilisation design framework.

References

1. AASHTO. (2022). Standard Method of Test for The California Bearing Ratio. American Association of State Highway and Transportation Officials.
2. ASTM International. (2021). ASTM D4318: Standard Test Methods for Liquid Limit, Plastic Limit, and Plasticity Index of Soils. ASTM International.

3. ASTM International. (2021). ASTM D698: Standard Test Methods for Laboratory Compaction Characteristics of Soil Using Standard Effort. ASTM International.
4. ASTM International. (2022). ASTM D1883: Standard Test Method for California Bearing Ratio (CBR) of Laboratory-Compacted Soils. ASTM International.
5. ASTM International. (2020). ASTM D2166/D2166M: Standard Test Method for Unconfined Compressive Strength of Cohesive Soil. ASTM International.
6. ASTM International. (2017). ASTM D2487: Standard Practice for Classification of Soils for Engineering Purposes. ASTM International.
7. ASTM International. (2014). ASTM D4546: Standard Test Methods for One-Dimensional Swell or Collapse of Soils. ASTM International.
8. ASTM International. (2015). ASTM D559/D559M: Standard Test Methods for Wetting and Drying Compacted Soil-Cement Mixtures. ASTM International.
9. Bureau of Indian Standards. (1970). IS 1498: Classification and Identification of Soils for General Engineering Purposes. BIS.
10. Bureau of Indian Standards. (1985). IS 2720 Part 5: Methods of Test for Soils: Determination of Liquid and Plastic Limit. BIS.
11. Bureau of Indian Standards. (1980). IS 2720 Part 7: Determination of Water Content-Dry Density Relation Using Light Compaction. BIS.
12. Bureau of Indian Standards. (1991). IS 2720 Part 10: Determination of Unconfined Compressive Strength. BIS.
13. Bureau of Indian Standards. (1987). IS 2720 Part 16: Laboratory Determination of CBR. BIS.
14. Bureau of Indian Standards. (1977). IS 2720 Part 40: Determination of Free Swell Index of Soils. BIS.
15. Chen, F. H. (1975). Foundations on Expansive Soils. Elsevier Scientific Publishing Company.
16. Holtz, W. G., & Gibbs, H. J. (1956). Engineering properties of expansive clays. Transactions of the American Society of Civil Engineers, 121, 641-677.
17. Nelson, J. D., & Miller, D. J. (1992). Expansive Soils: Problems and Practice in Foundation and Pavement Engineering. John Wiley & Sons.
18. Nelson, J. D., Chao, K. C., Overton, D. D., & Nelson, E. J. (2015). Foundation Engineering for Expansive Soils. John Wiley & Sons.
19. Mitchell, J. K., & Soga, K. (2005). Fundamentals of Soil Behavior. John Wiley & Sons.
20. Das, B. M., & Sobhan, K. (2018). Principles of Geotechnical Engineering. Cengage Learning.
21. Bell, F. G. (1996). Lime stabilization of clay minerals and soils. Engineering Geology, 42(4), 223-237.
22. Puppala, A. J. (2016). Advances in ground modification with chemical additives: From theory to practice. Transportation Geotechnics, 9, 123-138.
23. Little, D. N., & Nair, S. (2009). Recommended Practice for Stabilization of Subgrade Soils and Base Materials. NCHRP Web-Only Document 144.
24. Sherwood, P. T. (1993). Soil Stabilization with Cement and Lime. Transport Research Laboratory.
25. Cokca, E. (2001). Use of class C fly ashes for the stabilization of expansive soil. Journal of Geotechnical and Geoenvironmental Engineering, 127(7), 568-573.
26. Seco, A., Ramírez, F., Miqueleiz, L., & García, B. (2011). Stabilization of expansive soils for use in construction. Applied Clay Science, 51(3), 348-352.
27. Barman, D., & Dash, S. K. (2022). Stabilization of expansive soils using chemical additives: A review. Journal of Rock Mechanics and Geotechnical Engineering, 14(4), 1319-1342.
28. James, J., & Pandian, P. K. (2016). Industrial wastes as auxiliary additives to cement/lime stabilization of soils. Advances in Civil Engineering.
29. Wild, S., Kinuthia, J. M., Jones, G. I., & Higgins, D. D. (1998). Effects of partial substitution of lime with ground granulated blast furnace slag in lime-stabilised clay soils. Engineering Geology, 51(1), 37-53.
30. Horpibulsuk, S., Rachan, R., Chinkulkijniwat, A., Raksachon, Y., & Suddeepong, A. (2010). Analysis of strength development in cement-stabilized silty clay from microstructural considerations. Construction and Building Materials, 24(10), 2011-2021.
31. Consoli, N. C., Prietto, P. D. M., Carraro, J. A. H., & Heineck, K. S. (2001). Behavior of compacted soil-fly ash-carbide lime mixtures. Journal of Geotechnical and Geoenvironmental Engineering, 127(9), 774-782.
32. Liu, Y., Chang, C. W., Namdar, A., She, Y., Lin, C. H., Yuan, X., & Yang, Q. (2019). Stabilization of expansive soil using cementing material from rice husk ash and calcium carbide residue. Construction and Building Materials, 221, 1-11.
33. Jain, P. K., Choudhary, A. K., & Jha, J. N. (2020). Influence of rice husk ash on swelling and strength characteristics of expansive soil. Geotechnical and Geological Engineering, 38, 2293-2302.
34. Brooks, R. M. (2009). Soil stabilization with fly ash and rice husk ash. International Journal of Research and Reviews in Applied Sciences, 1(3), 209-217.
35. Kumar, A., Walia, B. S., & Bajaj, A. (2007). Influence of fly ash, lime and polyester fibers on compaction and strength properties of expansive soil. Journal of Materials in Civil Engineering, 19(3), 242-248.
36. Kalkan, E. (2009). Influence of silica fume on the desiccation cracks of compacted clayey soils. Applied Clay Science, 43(3-4), 296-302.

37. Phanikumar, B. R., & Sharma, R. S. (2004). Effect of fly ash on engineering properties of expansive soil. Journal of Geotechnical and Geoenvironmental Engineering, 130(7), 764-767.
38. Ghosh, A., & Subbarao, C. (2007). Strength characteristics of class F fly ash modified with lime and gypsum. Journal of Geotechnical and Geoenvironmental Engineering, 133(7), 757-766.
39. Al-Rawas, A. A., Hago, A. W., & Al-Sarmi, H. (2005). Effect of lime, cement and Sarooj on the swelling potential of an expansive soil from Oman. Building and Environment, 40(5), 681-687.
40. Fattah, M. Y., Salman, F. A., & Nareeman, B. J. (2010). A treatment of expansive soil using different additives. Acta Montanistica Slovaca, 15(4), 290-297.
41. Rao, S. M., Reddy, B. V. V., & Muttharam, M. (2001). The impact of cyclic wetting and drying on the swelling behaviour of stabilized expansive soils. Engineering Geology, 60(1-4), 223-233.
42. Shukla, R. P., Singh, N., Yadav, P., & Mankotia, N. (2015). Problems and treatment of black cotton soil. Proceedings of the 50th Indian Geotechnical Conference.
43. Mishra, A. K., Ohtsubo, M., Li, L. Y., & Higashi, T. (2008). Effect of salt concentrations on the permeability and compressibility of soil-bentonite mixtures. Journal of the Faculty of Agriculture, Kyushu University, 53(2), 837-843.
44. Muntohar, A. S. (2011). Engineering characteristics of the compressed-stabilized earth brick. Construction and Building Materials, 25(11), 4215-4220.
45. Prabakar, J., Dendorkar, N., & Morchhale, R. K. (2004). Influence of fly ash on strength behaviour of typical soils. Construction and Building Materials, 18(4), 263-267.
46. Breiman, L. (2001). Random forests. Machine Learning, 45, 5-32.
47. Cortes, C., & Vapnik, V. (1995). Support-vector networks. Machine Learning, 20, 273-297.
48. Vapnik, V. N. (1995). The Nature of Statistical Learning Theory. Springer.
49. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189-1232.
50. Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. Machine Learning, 63, 3-42.
51. Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. Technometrics, 12(1), 55-67.
52. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., et al. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825-2830.
53. Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning. Springer.
54. James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). An Introduction to Statistical Learning. Springer.
55. Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer.
56. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785-794.
57. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30.
58. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135-1144.
59. Gajurel, A., Chittoori, B., Mukherjee, P. S., & Sadegh, M. (2021). Machine learning methods to map stabilizer effectiveness based on common soil properties. Transportation Geotechnics, 27, 100506.
60. Zhang, W., & Goh, A. T. C. (2016). Multivariate adaptive regression splines for analysis of geotechnical engineering systems. Computers and Geotechnics, 77, 180-190.
61. Shahin, M. A., Maier, H. R., & Jaksa, M. B. (2009). Recent advances and future challenges for artificial neural systems in geotechnical engineering applications. Advances in Artificial Neural Systems, 2009.
62. Phoon, K. K., & Kulhawy, F. H. (1999). Characterization of geotechnical variability. Canadian Geotechnical Journal, 36(4), 612-624.
63. Phoon, K. K., & Ching, J. (2015). Risk and Reliability in Geotechnical Engineering. CRC Press.
64. Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182-197.
65. Derringer, G., & Suich, R. (1980). Simultaneous optimization of several response variables. Journal of Quality Technology, 12(4), 214-219.
66. Marler, R. T., & Arora, J. S. (2004). Survey of multi-objective optimization methods for engineering. Structural and Multidisciplinary Optimization, 26, 369-395.