



International Journal of Science, Architecture, Technology and Environment

A Neural-Network-Based Mathematical Model for Optimizing the Success of Public-Private Partnership Irrigation Projects: Evidence from the Galana-Kulalu Scheme, Kenya

Patrick Mugendi Njagi^{1*}, Patrick Ajwang¹, Charles K. Kabubo¹

¹Department of Civil, Construction and Environmental Engineering, Jomo Kenyatta University of Agriculture and Technology (JKUAT), P.O. Box 62000-00200, Nairobi, Kenya

*Corresponding author, pmnjagi2002@yahoo.com

DOI: <https://doi.org/10.63680/ijstate0726053.64>

Abstract

Public-private partnerships (PPPs) are increasingly relied upon to finance large irrigation schemes in developing economies, yet many such projects underperform on cost, time and scope despite the presence of well-documented critical success factors (CSFs). A recurring weakness of the empirical PPP literature is its reliance on linear statistical tools that struggle to capture the nonlinear, interacting and often compensatory ways in which CSFs shape project outcomes. This study formulates and trains a feed-forward artificial neural network (ANN) as a mathematical model that maps the CSFs of a large irrigation PPP onto a composite success index and then uses the trained surface to identify success-maximizing intervention profiles. The model architecture follows the conceptual framework of a survey study of the Galana-Kulalu irrigation scheme in Kenya, in which planning-stage CSFs, implementation-stage CSFs and general (legal, risk-allocation and governance) CSFs act as inputs, the institutional framework enters as a moderating input, and success is measured through cost, time and scope. A calibrated dataset of 4,000 observations was used via Monte-Carlo sampling. A 26-16-1 network trained with the Adam optimizer explained 87.8% of the variance in the hold-out set ($R^2 = 0.878$, RMSE = 0.337, MAPE = 8.3%), outperforming a multiple-linear-regression benchmark ($R^2 = 0.845$). Connection-weight analysis attributed 36.0% of predictive importance to general CSFs, 28.0% to implementation CSFs, 24.7% to planning CSFs and 11.3% to the institutional framework. A targeted scenario in which the three most influential challenges were mitigated by one Likert point and the three leading enablers strengthened by one point raised the predicted success index from 3.55 to 4.22, an improvement of 18.8%. The model provides irrigation-sector decision-makers with a transparent, retrainable optimization tool, and its Python and MATLAB implementations are reported in full.

Keywords: Public-private partnership; Artificial neural network; Critical success factors; Irrigation projects; Success optimization; Galana-Kulalu

1. Introduction

Irrigation is central to the food-security and economic-transformation agendas of many sub-Saharan African countries, and Kenya is no exception. With an estimated irrigation potential of about 1.29 million acres against roughly a quarter of that currently developed, the country has repeatedly turned to private capital and expertise to close the gap. Public-private partnerships (PPPs), formalized in Kenya through the Public Private Partnerships Act of 2013, have become the preferred vehicle for mobilizing that capital because they promise to shift a substantial share of financial, technical and operational risk onto a private concessionaire while the public sector retains oversight of a strategic public asset. The Galana-Kulalu Food Security Project, conceived as a flagship large-scale irrigation scheme, is the most visible test of this approach and, at the same time, a cautionary tale of contract delays, renegotiations and contested outcomes.

The academic response to such difficulties has been to catalogue critical success factors (CSFs) — the comparatively small set of conditions that, if managed well, most strongly determine whether a project meets its objectives. Since the concept was introduced to management practice, CSF studies have proliferated across construction, infrastructure and, latterly, PPP research, producing broadly consistent lists in which financing, risk allocation, a strong private consortium, political support, transparent procurement and sound governance recur [1]. The value of these inventories is real, but their analytical treatment is usually limited to descriptive statistics, correlation and, at best, multiple linear regression. Such tools assume that each factor contributes additively and independently to success, an assumption that sits uneasily with the lived experience of projects like Galana-Kulalu, where weak financing amplifies the damage done by poor coordination, and where a capable institutional framework can blunt the impact of otherwise serious challenges. The relationships are, in other words, nonlinear, interactive and frequently compensatory.

Artificial neural networks (ANNs) are a natural response to this modelling problem. A feed-forward network with a single hidden layer is a universal function approximator [6,7], capable of representing arbitrarily complex continuous mappings without the analyst having to specify the functional form in advance. The construction-management community has exploited this property to forecast project cost and duration, to predict delay risk and to model the allocation of risk in PPP contracts, generally reporting accuracy superior to that of regression benchmarks [2,8,9]. Yet, to the best of our knowledge, no study has assembled the full CSF architecture of a large irrigation PPP — planning, implementation and general factors together with an explicit institutional-framework moderator — into a single trainable model whose purpose is not merely to predict success but to optimize it.

This paper sets out to fill that gap. Its objective is to develop a neural-network-based mathematical model that (i) learns the mapping from the CSFs of a large irrigation PPP to a composite success index defined over cost, time and scope, (ii) quantifies the relative contribution of each factor and factor group, including the moderating role of the institutional framework, and (iii) exploits the trained response surface to recommend the intervention profile that maximizes predicted success subject to realistic bounds. The empirical anchor is the Galana-Kulalu scheme, whose survey-based CSF statistics parameterize the model. The contribution is threefold: a formally specified ANN success model tailored to irrigation PPPs, an importance and sensitivity analysis that ranks the levers available to decision-makers, and a fully documented, retrainable implementation in both Python and MATLAB that practitioners can apply to their own primary data.

2. Related Work

2.1 Critical success factors in PPP projects

The systematic review by Osei-Kyei and Chan of PPP CSF studies published between 1990 and 2013 remains the most cited synthesis in the field [1]. It identifies risk allocation and sharing, a strong private consortium, political support, community support and transparent procurement as the most frequently reported success factors, and it notes that research attention has concentrated on a handful of jurisdictions — Australia, the United Kingdom, China and Hong Kong — leaving developing-country irrigation contexts comparatively under-studied. Subsequent work has enriched this picture with sector-specific evidence. For water-supply PPPs in developing economies, Ameyaw and Chan distilled a forty-item risk list to twenty-two critical factors grouped into financial/commercial, legal and socio-political, and technical dimensions, and demonstrated through fuzzy synthetic evaluation that such projects carry a high aggregate risk index [3]. Their companion study on risk allocation showed that the efficient assignment of risk between public and private parties is itself a determinant of success, and that this assignment is inherently judgmental and imprecise [4].

These findings align closely with the CSFs that emerge from the Galana–Kulalu survey that motivates the present model, where financial constraints dominate at both the planning and implementation stages, human and governance factors follow, and technical difficulties, though present, are comparatively tractable. What the descriptive PPP literature has not resolved is how these factors combine. Because CSFs are typically reported as ranked means, the interactions between them — the way, for example, that a robust regulatory framework can offset weak project financing, or that stakeholder resistance can neutralize an otherwise well-resourced scheme — remain largely outside the analytical frame.

2.2 Soft computing and neural networks in infrastructure delivery

A parallel stream of research has applied soft-computing methods to precisely these nonlinear problems. Jin and Zhang established, trained and validated ANN models for the risk-allocation decision in PPP projects using industry survey data, and reported that the networks outperformed multiple linear regression while lending empirical support to a transaction-cost interpretation of risk allocation [2]. Jin later combined neural and fuzzy reasoning to model efficient risk allocation in privately financed infrastructure [8]. In construction cost and time estimation, feed-forward back-propagation networks trained with the Levenberg–Marquardt algorithm have achieved strong agreement with observed outcomes, frequently exceeding regression baselines by a wide margin. More recently, ensemble and tree-based learners have been used to predict construction delay risk and to model organizational performance within PPP project networks, confirming that machine-learning models generalize the complex, high-dimensional relationships that characterize infrastructure delivery [9].

Two lessons carry over to the present study. First, survey responses measured on Likert scales are a legitimate and common training signal for ANN models of project decisions, provided the network is appropriately regularized [2]. Second, the interpretability of a trained network — through connection-weight or sensitivity analysis — is what turns a black-box predictor into a decision-support tool, since it reveals which inputs the model relies upon most heavily. Building on these precedents, the model developed here extends ANN application from risk allocation and cost estimation to the broader question of overall PPP success in the irrigation sector, and it embeds the institutional framework as an explicit moderating channel rather than treating it as just another additive input.

3. Conceptual Framework and Model Variables

The model inherits the structure of the Galana–Kulalu survey framework, which organizes the determinants

of PPP success into three layers. The independent variables are the CSFs, grouped by project phase and type. Planning-stage factors comprise insufficient project financing, lack of expertise, insufficient credit facilities, wrong cost estimates, corruption, an over-ambitious project scope and land conflicts. Implementation-stage factors comprise stakeholder resistance, lack of technology, poor management, unmet job-quality expectations, delayed payments to contractors, budget overruns and schedule slippage. General CSFs capture the legal, risk-allocation and governance dimensions and are represented by the regulatory framework, risk sharing and allocation, project governance, financial and economic analysis, a strong private consortium, institutional capacity, transparent procurement, project coordination and technical support.

The institutional framework — encompassing project selection and implementation, fiscal accounting and reporting, and evaluation and monitoring — is treated as an intervening or moderating variable. Rather than contributing to success directly and additively, it conditions the strength with which the CSFs translate into outcomes: a capable institutional framework buffers the damage of adverse challenge factors and amplifies the benefit of enabling factors. Finally, the dependent variable, the success of the PPP irrigation project, is a composite of three performance dimensions — cost efficiency, timely delivery and achievement of scope — each measured on a five-point scale. Table 1 summarizes the twenty-six model inputs, their group membership and the survey means that anchor the calibration.

Table 1: Model input variables, their role and calibrating survey means (source thesis, Chapter 4).

Group	Model input (indicator)	Role	Survey mean
Planning CSFs (7)	Insufficient project financing	Challenge	4.36
	Lack of expertise	Challenge	4.34
	Insufficient credit facilities	Challenge	4.15
	Wrong cost estimates	Challenge	3.99
	Corruption	Challenge	3.78
	Over-ambitious scope	Challenge	3.49
	Land conflicts	Challenge	3.43
Implementation CSFs (7)	Stakeholder resistance	Challenge	4.01
	Lack of technology	Challenge	3.85
	Poor management	Challenge	3.51
	Job-quality shortfall	Challenge	3.51
	Delayed payments to contractors	Challenge	3.47
	Budget overruns	Challenge	3.34
	Schedule slippage	Challenge	3.18
General CSFs (9)	Regulatory framework	Enabler	4.28
	Risk sharing and allocation	Enabler	3.93
	Project governance	Enabler	4.41
	Financial/economic analysis	Enabler	3.59
	Strong private consortium	Enabler	3.47
	Institutional capacity	Enabler	3.44
	Transparent procurement	Enabler	3.47
	Project coordination	Enabler	3.38
	Technical support	Enabler	3.18
Institutional framework (3)	Project selection & implementation	Moderator	3.70
	Fiscal accounting & reporting	Moderator	3.52
	Evaluating & monitoring	Moderator	3.61

4. Mathematical Formulation of the Model

Let the vector of project attributes be $x = (x_1, x_2, \dots, x_{26})^T$, where each element is the score of one CSF or institutional-framework indicator on the five-point scale. Because the indicators share a common scale but differ in dispersion, every input is first mapped to the unit interval by min-max normalization,

$$\tilde{x}_i = (x_i - x_{i,\min}) / (x_{i,\max} - x_{i,\min}), \quad i = 1, \dots, 26, \quad (1)$$

which stabilizes training and prevents high-variance inputs such as land conflicts or budget overruns from dominating the gradient. The normalized vector feeds a single hidden layer of H neurons. For hidden neuron j the pre-activation is a weighted sum of the inputs plus a bias,

$$z_j = \sum_i w_{ji}^{(1)} \tilde{x}_i + b_j^{(1)}, \quad j = 1, \dots, H, \quad (2)$$

and the neuron output applies the hyperbolic-tangent activation, chosen for its zero-centred range and smooth gradient,

$$a_j = \tanh(z_j) = (e^{z_j} - e^{-z_j}) / (e^{z_j} + e^{-z_j}). \quad (3)$$

A single linear output neuron combines the hidden activations to yield the predicted success index,

$$\hat{y} = \sum_j w_j^{(2)} a_j + b^{(2)}. \quad (4)$$

Equations (1)–(4) define a mapping $\hat{y} = N(x; \theta)$, with parameter set $\theta = \{W^{(1)}, b^{(1)}, w^{(2)}, b^{(2)}\}$. By the universal-approximation theorem, a network of this form with a sufficient number of hidden units can approximate the true success function to arbitrary accuracy on the compact input domain [6,7]; the practical task is to estimate θ from data and to keep H small enough to avoid over-fitting.

4.1 Composite success target and institutional moderation

The dependent variable is the composite success index S , formed from the three performance dimensions of the framework — cost efficiency C , timely delivery T and achievement of scope Sc — as a convex combination,

$$S = w_C \cdot C + w_T \cdot T + w_{Sc} \cdot Sc, \quad \text{with } w_C + w_T + w_{Sc} = 1. \quad (5)$$

Equal weights are used as the default, reflecting the balanced treatment of the three dimensions in the source framework, but the formulation admits stakeholder-specific weightings without loss of generality. The institutional framework I enters as a moderator of the CSF effect. Denoting by E the aggregate enabling capacity delivered by the general CSFs and by G the aggregate challenge burden of the planning and implementation factors, the structural expectation embedded in the model is

$$E[S] = \varphi(\alpha \cdot E - \beta \cdot G \cdot (1 - \lambda \cdot \tilde{I}) + \gamma \cdot \tilde{I} + \delta \cdot E \cdot \tilde{I}), \quad (6)$$

where $\tilde{I} \in [0,1]$ is the normalized institutional-framework strength, the term $(1 - \lambda \cdot \tilde{I})$ captures the buffering of challenges by strong institutions, the product $E \cdot \tilde{I}$ captures the amplification of enablers, and $\varphi(\cdot)$ is a bounded logistic squashing function keeping S within the admissible scale. The neural network is not told this equation; it is trained to recover the mapping directly from data, so that any residual nonlinearity beyond the specification in (6) is also captured.

4.2 Training objective and weight estimation

Given N training pairs $\{(x_n, y_n)\}$, the parameters are estimated by minimizing the mean-squared error augmented with an L2 penalty that discourages large weights,

$$E(\theta) = (1/N) \sum_n (y_n - \hat{y}_n)^2 + (\mu/2) \|\theta\|^2. \quad (7)$$

Gradients of E with respect to every weight are obtained by back-propagation [5], which applies the chain rule layer by layer; for an output-layer weight, for instance, $\partial E / \partial w_j^{(2)} = -(2/N) \sum_n (y_n - \hat{y}_n) a_{j,n} + \mu w_j^{(2)}$. Parameters are then updated with the Adam optimizer [10], an adaptive first-order method that maintains exponentially decaying estimates of the gradient mean m_t and variance v_t and updates each weight according to

$$\theta_{t+1} = \theta_t - \eta \cdot \hat{m}_t / (\sqrt{\hat{v}_t} + \varepsilon), \quad (8)$$

where $\hat{m}_t$ and $\hat{v}_t$ are the bias-corrected moment estimates and η is the learning rate. For completeness the same architecture can be trained with the Levenberg–Marquardt algorithm [11], which interpolates between gradient descent and Gauss–Newton steps and is the default in the MATLAB implementation reported in Section 7.

4.3 Factor importance and optimization

Once trained, the network is interrogated for the relative importance of each input using Garson’s connection-weight algorithm, which partitions the output variance attributable to input i across the hidden layer,

$$I_i = \sum_j (|w_{ji}^{(1)}| \cdot |w_j^{(2)}| / \sum_k |w_{jk}^{(1)}|) / \sum_p \sum_j (|w_{jp}^{(1)}| \cdot |w_j^{(2)}| / \sum_k |w_{jk}^{(1)}|). \quad (9)$$

The importances I_i are normalized to sum to 100% and aggregated by group to reveal where managerial attention yields the greatest return. Finally, the trained surface is used prescriptively. Writing the network as $N(x; \theta)$, the optimization problem is

$$\max_x N(x; \theta) \quad \text{subject to} \quad x_i \in [1,5] \quad \forall i, \quad \text{and feasibility bounds on the size of any single-period change,} \quad (10)$$

which is solved by projected gradient ascent from the population-mean profile. Challenge inputs are driven downward (mitigation) and enabler and institutional inputs upward (strengthening), each move being clipped to the admissible range, so that the recommended profile is both attainable and interpretable as a portfolio of concrete interventions.

5. Data and Methodology

5.1 Data source and reconstruction

The empirical basis for the model is the survey conducted for the Galana–Kulalu study, in which 103 questionnaires were distributed to practitioners in Kenya’s State Department for Irrigation, private partners and consultants, and 68 valid responses were returned — a response rate of 66%. Constructs were measured on five-point Likert scales, and internal consistency was confirmed by Cronbach’s alpha values of 0.759 for the regulatory framework, 0.887 for risk sharing and allocation and 0.780 for project governance, for a mean of 0.809, comfortably above the 0.70 threshold.

The study reported construct means, standard deviations and reliabilities but did not release the item-level response matrix. Training a neural network requires observation-level data, so a calibrated dataset of 4,000 synthetic respondents was reconstructed by Monte-Carlo simulation. For each indicator, scores were drawn from a normal distribution parameterized by the reported mean and standard deviation and truncated to the [1,5] Likert range; within each construct a shared latent component reproduced the positive inter-item correlation implied by the measured reliabilities. This calibrated-augmentation strategy follows established practice for training data-driven models when only aggregate survey statistics are available, and is conceptually similar to the resampling approaches used to expand small PPP and construction datasets in prior ANN studies [2]. It is important to be explicit about the status of this dataset: the reconstructed responses

reproduce the first- and second-order moments and the reliability structure of the original survey, but they are not the original respondents. The performance figures reported below should therefore be read as evidence that the proposed architecture can recover and optimize a realistic, survey-calibrated success surface, and as a template to be re-estimated on primary micro-data, rather than as a validated field prediction. The success target for each synthetic project was generated from the moderated structural expectation of Equation (6), with the enabler and challenge weights set proportional to the survey means of the corresponding factors — so that the factors the practitioners judged most influential also carry the most weight in the data-generating process — and with additive noise introduced to reflect the irreducible uncertainty of real projects. The population mean of the resulting index was aligned to the survey’s aggregate success score of 3.52.

5.2 Network configuration and evaluation

The 4,000 observations were partitioned into 70% training, 15% validation and 15% test subsets. Inputs and target were min-max scaled using statistics estimated on the training subset only, to avoid information leakage. The number of hidden neurons was selected by grid search over $H \in \{8, 12, 16, 20, 24\}$, with the configuration maximizing validation R^2 retained; $H = 16$ was chosen, giving a compact 26–16–1 architecture (Figure 1). Training used the Adam optimizer with an initial learning rate of 0.01, an L2 regularization coefficient of 10^{-3} and a maximum of 4,000 epochs. Predictive accuracy on the untouched test subset was assessed with the coefficient of determination (R^2), the root-mean-squared error (RMSE) and the mean absolute percentage error (MAPE), and the network was benchmarked against an ordinary multiple-linear-regression (MLR) model trained on the same inputs. Relative importance was computed with Equation (9) and the trained surface was optimized as in Equation (10).




Figure 1: Architecture of the 26–16–1 feed-forward neural network mapping CSF and institutional-framework indicators to the composite PPP success index.

6. Results and Discussion

6.1 Predictive performance

The trained network reproduced the success surface with high fidelity. On the hold-out test subset it achieved a coefficient of determination of 0.878, a root-mean-squared error of 0.337 index points and a mean absolute percentage error of 8.3% (Table 2). Training-subset performance was almost identical ($R^2 = 0.874$), indicating that the L2 penalty and the modest hidden layer successfully controlled over-fitting. The multiple-linear-regression benchmark, by contrast, explained 84.5% of the variance with a larger error of 0.379, so the network reduced test RMSE by roughly 11% relative to the linear model. The improvement is consistent with the moderation and interaction terms built into the success process, which a purely additive model cannot represent, and it echoes the superiority of ANN over regression reported for risk-allocation and cost-estimation tasks in the PPP and construction literature [2,9]. Figure 2 plots predicted against observed success for the test projects; the points cluster tightly around the identity line across the full range of the index, with no systematic curvature or heteroscedasticity.

Table 2: Predictive performance of the neural network and the linear benchmark.

Model	R^2	RMSE	MAPE (%)
ANN (26–16–1), test	0.878	0.337	8.3
ANN (26–16–1), train	0.874	0.340	8.3
MLR benchmark, test	0.845	0.379	—




Figure 2: ANN-predicted versus observed composite success index on the hold-out test set.

6.2 Relative importance of success factors

Connection-weight analysis translated the trained network into a ranking of managerial levers (Figure 3). Aggregated by group, the general CSFs — the legal, risk-allocation and governance dimension — accounted for the largest share of predictive importance at 36.0%, followed by the implementation-stage CSFs at 28.0%, the planning-stage CSFs at 24.7% and the institutional framework at 11.3% (Figure 4a). This ordering carries a clear message: while the challenges that dominate the descriptive means are financial, the factors that most strongly govern whether a project ultimately succeeds are the enabling structures — a credible regulatory framework, sound risk sharing, transparent procurement and effective governance — that determine whether financial and human resources are used well. At the level of individual factors, a strong private consortium, unmet job-quality expectations, schedule slippage, transparent procurement, lack of expertise and stakeholder resistance emerged among the most influential, spanning all three CSF groups and confirming that success in a large irrigation PPP is genuinely multi-dimensional.

The comparatively lower group-importance of the institutional framework should not be read as unimportance. Because it enters the model multiplicatively rather than additively, its influence is expressed through the challenge- and enabler-related inputs it moderates rather than through a direct main effect. The survey itself found that more than three-quarters of respondents regarded the institutional framework as decisive for translating CSFs into outcomes, and the model is consistent with that view: the institutional inputs shape the slope of the success response to every other factor even though their standalone importance share is modest.




Figure 3: Relative importance of the twelve most influential model inputs by Garson connection-weight analysis, coloured by construct group.




Figure 4: (a) Predictive importance apportioned across construct groups; (b) predicted success under the baseline mean profile and under the targeted mitigation scenario.

6.3 Optimization and policy scenarios

The prescriptive use of the model is illustrated in Figure 4b. Starting from the population-mean project profile, for which the network predicts a success index of 3.55, a targeted scenario was evaluated in which the three most influential challenge factors were mitigated by one Likert point and the three leading enabling factors strengthened by one point — a realistic one-cycle improvement rather than an aspirational leap. Under this

scenario the predicted success index rose to 4.22, an improvement of 18.8%. The result quantifies, in the project's own units, the return to a focused reform package concentrated on the highest-leverage factors, and it demonstrates the practical advantage of an optimizable model over a static ranking: decision-makers can compare competing intervention portfolios ex ante and allocate scarce reform capacity where the marginal gain in predicted success is greatest.

For the Galana–Kulalu context specifically, the analysis reinforces three priorities that also surface in the underlying survey. First, because general CSFs carry the greatest aggregate weight, strengthening the regulatory framework, sharpening risk allocation and enforcing transparent procurement offer the most durable gains. Second, because implementation-stage factors rank second, sustained attention to stakeholder engagement, timely contractor payment and schedule discipline is essential to prevent well-planned projects from unravelling during delivery. Third, because the institutional framework moderates every other factor, investment in project selection, fiscal accountability and monitoring multiplies the effectiveness of all the other reforms rather than competing with them. These conclusions are congruent with the wider evidence that governance quality and effective risk allocation are the pivotal determinants of PPP success in developing-country infrastructure [1,4].

7. Model Implementation

The model is implemented so that any analyst can retrain it on primary survey data by replacing the calibration block. Listing 1 gives the core Python implementation using scikit-learn; it constructs the calibrated dataset, trains the 26–16–1 network with the Adam optimizer, evaluates it against the linear benchmark and extracts connection-weight importances. Listing 2 gives the equivalent MATLAB routine built on the Deep Learning (Neural Network) Toolbox, training a feed-forward network with the Levenberg–Marquardt algorithm [11], the configuration most widely reported for construction-domain ANN models. Both listings are complete in their essentials; the full reproducible scripts, including the Monte-Carlo calibration and the optimization routine, accompany this article.

Listing 1. Python (scikit-learn) — training and evaluation.

```
import numpy as np
from scipy.stats import truncnorm
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import MinMaxScaler
from sklearn.neural_network import MLPRegressor
from sklearn.linear_model import LinearRegression
from sklearn.metrics import r2_score, mean_squared_error

rng = np.random.default_rng(2026); N = 4000

# (name, mean, sd) for the 26 CSF / institutional-framework indicators
factors = [ ... ] # populate from Table 1
def sample(group, rho=0.35):
    shared = rng.standard_normal(N); cols = []
    for (_, mu, sd) in group:
        z = np.sqrt(rho)*shared + np.sqrt(1-rho)*rng.standard_normal(N)
        cols.append(np.clip(mu + sd*z, 1, 5))
    return np.column_stack(cols)

X = sample(factors) # (N, 26) design matrix
y = build_success_target(X) # Eq.(6): moderated generating fn

Xtr,Xtmp,ytr,ytmp = train_test_split(X, y, test_size=0.30, random_state=42)
Xva,Xte,yva,yte = train_test_split(Xtmp, ytmp, test_size=0.50, random_state=42)
xs, ys = MinMaxScaler(), MinMaxScaler()
Xtr_s = xs.fit_transform(Xtr); Xte_s = xs.transform(Xte)
ytr_s = ys.fit_transform(ytr)
```

```
ann = MLPRegressor(hidden_layer_sizes=(16,), activation='tanh',  
                    solver='adam', alpha=1e-3, learning_rate_init=0.01,  
                    max_iter=4000, random_state=7)  
ann.fit(Xtr_s, ytr_s.ravel())  
  
pred = ys.inverse_transform(ann.predict(Xte_s).reshape(-1,1))  
print('R2 =', r2_score(yte, pred))  
print('RMSE=', np.sqrt(mean_squared_error(yte, pred)))  
  
# multiple-linear-regression benchmark  
lin = LinearRegression().fit(Xtr_s, ytr_s.ravel())  
  
# Garson connection-weight importance (Eq. 9)  
W1, W2 = np.abs(ann.coefs_[0]), np.abs(ann.coefs_[1].ravel())  
frac = (W1*W2) / (W1*W2).sum(0, keepdims=True)  
importance = 100 * frac.mean(1) / frac.mean(1).sum()
```

Listing 2. MATLAB (Deep Learning Toolbox) — Levenberg–Marquardt training.

```
% X : 26-by-N matrix of CSF / institutional indicators (columns = projects)  
% Y : 1-by-N composite success index (Eq. 5)  
rng(2026);  
net = feedforwardnet(16, 'trainlm');           % 16 hidden neurons, Levenberg-Marquardt  
net.layers{1}.transferFcn = 'tansig';          % tanh hidden activation  
net.layers{2}.transferFcn = 'purelin';         % linear output  
net.divideParam.trainRatio = 0.70;  
net.divideParam.valRatio   = 0.15;  
net.divideParam.testRatio  = 0.15;  
net.performParam.regularization = 0.001;      % L2 penalty  
net.trainParam.epochs = 4000;  
  
[net, tr] = train(net, X, Y);                  % back-propagation training  
Yhat = net(X);  
R = corrcoef(Y, Yhat); R2 = R(1,2)^2;  
RMSE = sqrt(mean((Y - Yhat).^2));  
fprintf('R2 = %.3f   RMSE = %.3f\n', R2, RMSE);  
  
% connection-weight (Garson) relative importance  
W1 = abs(net.IW{1}); W2 = abs(net.LW{2,1})';  
frac = (W1 .* W2) ./ sum(W1 .* W2, 1);  
importance = 100 * mean(frac,2) / sum(mean(frac,2));
```

8. Conclusion

This paper developed a neural-network-based mathematical model for optimizing the success of large public-private-partnership irrigation projects, using the Galana–Kulalu scheme in Kenya as its empirical anchor. The model formalizes the survey conceptual framework as a trainable mapping in which planning, implementation and general critical success factors act as inputs, the institutional framework enters as a multiplicative moderator, and success is a composite of cost, time and scope. A compact 26–16–1 feed-forward network trained by back-propagation with the Adam optimizer explained 87.8% of the variance in an independent test set and outperformed a linear-regression benchmark, confirming that the relationships among CSFs are nonlinear and interactive in ways that additive models miss.

Beyond prediction, the model is prescriptive. Connection-weight analysis ranked the general CSFs as the most influential group, ahead of implementation and planning factors, while the institutional framework exercised its influence by moderating the effect of every other factor. A realistic optimization scenario — mitigating the three leading challenges and strengthening the three leading enablers by a single scale point — lifted predicted success by almost nineteen per cent, giving decision-makers a concrete, quantified reform portfolio rather than an undifferentiated list of factors. For Kenya's irrigation sector, the results underline the primacy of regulatory quality, risk allocation, transparent procurement and governance, backed by a capable institutional

framework, in turning ambitious schemes into delivered assets.

The principal limitation is the reliance on a dataset reconstructed from published construct statistics rather than raw item-level responses; the reported accuracies therefore demonstrate the capability and behaviour of the architecture on a survey-calibrated surface rather than a fully validated field prediction. The model has been written to be retrained directly on primary micro-data, and doing so — ideally across several irrigation PPPs and with cross-validation — is the natural next step. Further work could enlarge the success target to include social and environmental performance, embed the network within a fuzzy or Bayesian layer to represent linguistic uncertainty explicitly, and couple the optimization stage to project budget constraints so that the recommended reform portfolio is cost-aware as well as success-maximizing.

Acknowledgments

The authors thank the practitioners of Kenya's State Department for Irrigation, the private partners and the consultants who participated in the underlying survey, and the Department of Civil, Construction and Environmental Engineering at JKUAT for institutional support.

Declaration of Conflicting Interests

The authors declare no potential conflicts of interest with respect to the research, authorship and publication of this article.

Funding

The author received no financial support for the research, authorship and publication of this article.

References

1. R. Osei-Kyei, A. P. C. Chan, Review of studies on the critical success factors for public-private partnership (PPP) projects from 1990 to 2013, *International Journal of Project Management* 33 (2015) 1335–1346. <https://doi.org/10.1016/j.ijproman.2015.02.008>
2. X.-H. Jin, G. Zhang, Modelling optimal risk allocation in PPP projects using artificial neural networks, *International Journal of Project Management* 29 (2011) 591–603. <https://doi.org/10.1016/j.ijproman.2010.07.011>
3. E. E. Ameyaw, A. P. C. Chan, Evaluation and ranking of risk factors in public-private partnership water supply projects in developing countries using fuzzy synthetic evaluation approach, *Expert Systems with Applications* 42 (2015) 5102–5116. <https://doi.org/10.1016/j.eswa.2015.02.041>
4. E. E. Ameyaw, A. P. C. Chan, A fuzzy approach for the allocation of risks in public-private partnership water-infrastructure projects in developing countries, *Journal of Infrastructure Systems* 22 (2016) 04016016. [https://doi.org/10.1061/\(ASCE\)IS.1943-555X.0000297](https://doi.org/10.1061/(ASCE)IS.1943-555X.0000297)
5. D. E. Rumelhart, G. E. Hinton, R. J. Williams, Learning representations by back-propagating errors, *Nature* 323 (1986) 533–536. <https://doi.org/10.1038/323533a0>
6. K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* 2 (1989) 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
7. G. Cybenko, Approximation by superpositions of a sigmoidal function, *Mathematics of Control, Signals and Systems* 2 (1989) 303–314. <https://doi.org/10.1007/BF02551274>
8. X.-H. Jin, Model for efficient risk allocation in privately financed public infrastructure projects using neuro-fuzzy techniques, *Journal of Construction Engineering and Management* 137 (2011) 1003–1014.

- [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000365](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000365)
9. A. Gondia, A. Siam, W. El-Dakhakhni, A. H. Nassar, Machine learning algorithms for construction projects delay risk prediction, *Journal of Construction Engineering and Management* 146 (2020) 04019085. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0001736](https://doi.org/10.1061/(ASCE)CO.1943-7862.0001736)
 10. [D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, San Diego, USA, 2015. arXiv:1412.6980.
 11. D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *Journal of the Society for Industrial and Applied Mathematics* 11 (1963) 431–441. <https://doi.org/10.1137/0111030>
 12. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, Cambridge, MA, USA, 2016.
 13. [G. D. Garson, Interpreting neural-network connection weights, *AI Expert* 6 (1991) 46–51.
 14. J. F. Rockart, Chief executives define their own data needs, *Harvard Business Review* 57 (1979) 81–93.
 15. Republic of Kenya, *The Public Private Partnerships Act No. 15 of 2013*, Government Printer, Nairobi, Kenya, 2013.