INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN APPLIED SCIENCE (IJRIAS)
ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

A Projected Steepest-Descent Algorithm for the Nonlinear Constrained Optimization of Football Team Performance

Chikwendu, C.R.; *Ndibe, O.E.; Unaegbu, E.N.; AGU, C.M.

Department of Mathematics, Nnamdi Azikiwe University, Awka, Nigeria

***Correspondence Author**

DOI: https://doi.org/10.51584/IJRIAS.2026.11080121

Received: 28 August 2026; Accepted: 2 September 2026; Published: 17 September 2026

ABSTRACT

Coaching staff must repeatedly decide how to divide a finite pool of training and match-day effort across tactical roles, and evidence from load-monitoring research indicates that the resulting gains in synergy and the resulting build-up of fatigue do not scale in proportion to that effort (Mandorino et al., 2025; Harris et al., 2025). Yet the great majority of published squad-composition and lineup-selection models still rely on binary, integer or purely linear-programming formulations (Zhao et al., 2021; Deb and Das, 2023), which cannot represent such curvature. This paper instead formulates team-performance maximisation as a continuously differentiable nonlinear programme whose objective combines power-law direct contributions, pairwise synergistic interactions between roles, and quadratic fatigue penalties, subject to individual workload ceilings, a global squad-load budget, a nonlinear fatigue-interaction inequality and an effort-sum equality constraint. A projected steepest-descent algorithm equipped with an Armijo backtracking line search and an exact Euclidean projection is developed, analysed and implemented in Python using NumPy, SciPy and automatic differentiation via JAX. On calibrated professional-season data the algorithm converges in 47 iterations to a stationary allocation that improves performance by 9.9% relative to a uniform allocation and by 6.2% relative to the club's historical average. A Karush–Kuhn–Tucker multiplier analysis shows that only two constraints bind at the optimum — the global load budget and the pairwise fatigue-interaction inequality — while all five individual role ceilings remain slack. Scenario and parametric sensitivity experiments confirm the robustness of this qualitative structure and show that a linearised counterpart of the same problem underperforms the nonlinear optimum by 3.6% while misidentifying the operative bottleneck. The resulting Lagrange multipliers translate directly into prioritised, quantitatively grounded recommendations for rotation policy and midfield-freshness preservation, illustrating the practical value of a transparent first-order method as decision support for coaching staff.

Keywords: nonlinear constrained optimisation; projected steepest descent; Karush–Kuhn–Tucker conditions; binding constraints; sports analytics; workload management; football performance modeling

INTRODUCTION

Coaches managing association football squads face a continuous resource-allocation problem: how to distribute practice intensity, match minutes and rotation across tactical roles so as to raise collective output without accumulating excessive physical risk. Longitudinal load-monitoring studies indicate that this trade-off is neither additive nor monotone — workload and recovery interact with on-field output and injury probability in ways that shift with fixture congestion and playing history (Bache-Mathiesen et al., 2021; Mandorino et al., 2025). A 40-week generalised-estimating-equation analysis of an English Championship squad similarly identified eighteen significant acute-to-chronic load interactions governing neuromuscular

and immune readiness, underscoring that treating workload components as independently additive misrepresents how fatigue actually accumulates (Harris et al., 2025). Notwithstanding this body of evidence, published quantitative models for lineup selection, squad composition and transfer-market planning have overwhelmingly relied on binary or mixed-integer linear programmes (Mahrudinda et al., 2020; Zhao et al., 2021), multi-criteria ranking procedures, or heuristic search over a linear scoring function (Deb and Das, 2023), each of which by

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

construction excludes the curved, interactive dose-response effects documented in the sports-science literature above.

This paper closes that gap in three linked steps. It first specifies a continuously differentiable objective that combines power-law direct contributions, pairwise synergy terms between tactical roles, and quadratic fatigue penalties, so that the curvature documented empirically above is built into the model rather than approximated after the fact. It then defines a constraint set encoding individual physiological ceilings, a global squad-load budget, the requirement that effort allocations sum to unity

METHODS

Mathematical Formulation

Decision Variables and Objective

Let $x = (x_1, ..., x_n)^T \in \mathbb{R}^n$ denote the vector of normalised effort allocations across n tactical roles — central defence, full-back overlapping, central midfield retention, wide attacking transition and central striking ($n = 5$ in the calibrated application). Following the interaction structure reported for fitness-fatigue coupling in elite squads (Mandorino et al., 2025), team performance is modelled as a sum of three components: a concave, diminishing-returns term for each role's direct contribution, a bilinear term capturing synergy between pairs of roles, and a quadratic term penalising excessive individual load. This yields the continuously differentiable performance function

$$f(x) = \alpha\, \Sigma_i\, \beta_i x_i^{\gamma_i} + \delta\, \Sigma_i \Sigma_{j>i}\, \theta_{ij} x_i x_j - \eta\, \Sigma_i\, \varphi_i x_i^2 \qquad (2.1)$$

where $\gamma_i \in (0,1)$ encode diminishing returns on each role's direct contribution, the double sum captures pairwise tactical synergies between roles i and j, and the quadratic term penalises excessive individual load. The scalars α, δ and η weight the direct-contribution, synergy and fatigue components respectively, while β_i, θ_{ij} and φ_i are role-specific coefficients calibrated from historical event and player-tracking data. The concavity induced by $\gamma_i < 1$ and the negative quadratic term ensures that f is bounded above on the feasible region defined below, which is essential for the existence of a maximiser.

Constraint Set

Individual workload ceilings, a global load budget and a nonlinear fatigue-interaction inequality constraint are written as:

$$g_i(x) = x_i - u_i \le 0, \qquad i = 1, ..., n, \qquad (2.2)$$

$$g_{n+1}(x) = \Sigma_i\, w_i x_i - L \le 0, \qquad (2.3)$$

$$g_{n+2}(x) = \Sigma_i \Sigma_{j>i}\, \kappa_{ij} x_i x_j - F \le 0, \qquad (2.4)$$

while the effort-sum equality reads

$$h_1(x) = \Sigma_i x_i - 1 = 0. \qquad (2.5)$$

Here u_i is the physiological ceiling entering role i, w_i is the marginal load coefficient entering the club's global budget L, and κ_{ij} is the pairwise fatigue-interaction coefficient bounded by the collective threshold F. The feasible region Ω defined jointly by equations (2.2)–(2.5) is compact, and because equation (2.3) is linear and the quadratic form in equation (2.4) is taken with non-negative interaction coefficients κ_{ij} on the non-negative orthant, Ω is convex. The Lagrangian of the maximisation problem satisfies standard Karush–Kuhn–Tucker (KKT) conditions; complementary slackness identifies the binding constraints whose multipliers quantify the marginal performance value of relaxing each resource limit.

INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN APPLIED SCIENCE (IJRIAS)
ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

This framing is consistent with Agu, Unaegbu & Chikwendu (2024), who characterise Karush–Kuhn–Tucker conditions as a first-derivatives test for optimality in nonlinear programming, extending the classical Lagrange-multiplier approach so that inequality, as well as equality, constraints can be accommodated and a maximum or minimum can occur at an interior point of the feasible region rather than only at its boundary.

Projected Steepest-Descent Algorithm

Given a feasible iterate $x_k \in \Omega$, the search direction is the gradient, $d_k = \nabla f(x_k)$ (an ascent direction, since the problem is a maximisation). A trial point $y_k = x_k + \alpha_k d_k$ is formed and projected onto the feasible set:

$$x_{k+1} = P_\Omega(y_k) := \arg\min_{z\in\Omega} \|z - y_k\|_2. \qquad (2.6)$$

Step lengths α_k are selected by an Armijo backtracking rule that guarantees a sufficient increase in the objective while preserving feasibility of the trial point after projection. The algorithm terminates when the projected-gradient norm falls below a prescribed tolerance $\varepsilon = 10^{-6}$. Under continuous differentiability of f and compactness of Ω, every accumulation point of the generated sequence satisfies the first-order necessary condition for a constrained local maximum; the argument follows the standard convergence theory for projected-gradient methods on compact convex sets (Bertsekas, 2016; Nocedal and Wright, 2006).

The procedure outlined below summarises one outer iteration of the method:

Step 1. (Gradient). Evaluate $d_k = \nabla f(x_k)$ by reverse-mode automatic differentiation.

Step 2. (Trial step). Set $y_k = x_k + \alpha_k d_k$ with an initial trial step length.

Step 3. (Projection). Compute $x_{k+1} = P_\Omega(y_k)$ by solving the sequential-quadratic-programming subproblem in (2.6).

Step 4. (Line search). If the Armijo sufficient-increase condition fails, halve α_k and return to Step 2; otherwise accept x_{k+1}.

Step 5. (Stopping test). Terminate if the projected-gradient norm is below ε; otherwise set $k \leftarrow k + 1$ and return to Step 1.

The procedure is implemented in Python 3.11. Gradients are obtained by reverse-mode automatic differentiation (JAX); the Euclidean projection in equation (2.6) is realised by a sequential-quadratic-programming subproblem that calls SciPy's constrained least-squares solver. Multi-start initialisation from Dirichlet-distributed feasible seeds mitigates the risk of convergence to a poor local maximum, and the reported solution is the best objective value obtained across all restarts.

NUMERICAL RESULTS AND CONSTRAINT DIAGNOSTICS

Baseline Solution and Convergence

On calibrated data from a professional season the algorithm converges in 47 iterations to the stationary point

$$x^* = (0.214, 0.187, 0.241, 0.198, 0.160)^T, \qquad f(x^*) = 1.4183.$$

The largest share is allocated to central midfield retention, reflecting its high direct coefficient and its strong synergistic links to both defence and attack. Relative to a uniform allocation the performance gain is 9.9%; relative to the club's historical average allocation the gain is 6.2%. Table 1 summarises the convergence history for the best multi-start seed.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

Table 1: Convergence history (baseline, best multi-start seed)

Iter. k	$f(x_k)$	‖proj. grad‖	α_k	$\max\lvert g_i\rvert$
0	1.1247	0.3124	—	0.000
10	1.3519	0.0712	0.125	0.000
20	1.3984	0.0248	0.0625	0.000
30	1.4126	0.0083	0.03125	0.000
47	1.4183	8.7×10^{-7}	0.0078	0.000

Convergence is monotone in the objective and exhibits the expected diminishing step-length pattern of the Armijo rule; the projected-gradient norm decreases by more than five orders of magnitude over the 47 iterations, and the maximum constraint residual remains at machine-zero throughout, confirming feasibility is preserved at every iterate by the exact projection step.

Binding Constraints

Complementary-slackness residuals (Table 2) show that the global load budget and the nonlinear fatigue-interaction constraint are the only active inequalities; all five individual role ceilings remain inactive at the optimum. The associated Lagrange multipliers are 0.312 and 0.187 respectively, indicating that a marginal expansion of the squad-load allowance yields a larger performance increment than an equivalent relaxation of the pairwise fatigue threshold.

Scenario and Sensitivity Analysis

Three alternative scenarios were examined: a congested fixture period (global load budget reduced by 12%), an injury-affected midfield, and an attacking tactical preference. Fixture congestion activates an additional individual ceiling and reduces the objective by 4.0%; the midfield-injury scenario produces the largest loss,

Table 2: Constraint residuals and multipliers at the optimum

Constraint	Residual	Multiplier	Status
Individual ceilings (all five)	< 0	0.000	Inactive
Global load budget	0.000	0.312	Binding
Fatigue interaction	0.000	0.187	Binding
Effort-sum equality	0.000	1.094	Binding (eq.)

at 4.8%, because it removes capacity from the role with the highest direct coefficient and the strongest synergistic linkages. Across a parametric grid spanning plausible ranges of L and F, the global load budget remains binding in every instance, identifying it as the structurally critical constraint of the model. A linearised counterpart of the same problem — obtained by suppressing the synergy and fatigue terms — attains an objective 3.6% lower than the nonlinear optimum and fails to recognise the pairwise fatigue bottleneck altogether, underscoring the practical cost of imposing linearity on an intrinsically nonlinear system.

PRACTICAL IMPLICATIONS AND LIMITATIONS

Because the binding constraints are collective rather than position-specific, rotation policy should prioritise expansion of the high-intensity rotation pool and tactical adjustments that reduce simultaneous activation of the full-back, wide-attacking and central-striking roles. The persistently elevated optimal share allocated to central midfield underscores the need to preserve the freshness of that unit across a congested calendar. The dual multipliers supply quantitative guidance for resource-allocation trade-offs; a unit increase in the global load budget yields a marginal return that is equivalent to relaxation of the fatigue-interaction threshold, and investment in squad-load capacity — additional rotation-eligible players, recovery infrastructure, or schedule management — should be prioritised over attempts to relax individual fatigue tolerances.

Several factors that shape real match-day decisions are not represented explicitly: opponent-specific match-ups, the psychological state of individual players, and short-term tactical adjustments made in-game remain the province of coaching judgement. The aggregate, role-based representation of the squad — rather than an individual-player formulation — is a further methodological simplification that trades granularity for tractability and interpretability; extending the model to a player-indexed decision space, and nonlinear synergy and fatigue structure, is a natural direction for future work. Data quality, in particular the calibration of the synergy coefficients θ_{ij} and the fatigue-interaction coefficients κ_{ij} from a single club and season, also constrains the external validity of the reported coefficients and in the robustness of the identified bottlenecks.

CONCLUSION

A nonlinear constrained programme for football team-performance maximisation has been formulated and solved using a projected steepest-descent algorithm with an Armijo line search and an exact Euclidean projection. The algorithm converges reliably, produces a measurable performance gain over both a uniform allocation and the club's historically observed allocation, and identifies the global load budget and the pairwise fatigue interaction as the operative bottlenecks. Scenario and sensitivity experiments confirm the robustness of this structure to fixture congestion, injury and tactical-preference perturbations, while a linearised

counterpart demonstrates the quantitative and qualitative cost of imposing linearity on the problem. The binding-constraint diagnostics translate directly into prioritised recommendations for rotation policy and midfield preservation, illustrating the practical value of a transparent first-order method as evidence-based decision support for team management.

REFERENCES

1. Agu, C.M., Unaegbu, E.N., & Chikwendu, C.R. (2024). Harnessing Karush–Kuhn–Tucker (KKT) as optimality conditions for solving nonlinear constrained optimization problems. International Journal of Research and Innovation in Applied Science, 9(7), 503–508. https://doi.org/10.51584/IJRIAS.2024.907043
2. Andrei, N. (2013). Nonlinear Optimization Applications Using the GAMS Technology. Springer Optimization and Its Applications, Vol. 81. New York: Springer. https://doi.org/10.1007/978-1-4614-6797-7
3. Bache-Mathiesen, L.K., Andersen, T.E., Dalen-Lorentsen, T., Clarsen, B., & Fagerland, M.W. (2021). Not straightforward: modelling non-linearity in training load and injury research. BMJ Open Sport & Exercise Medicine, 7(3), e001119. https://doi.org/10.1136/bmjsem-2021-001119
4. Bertsekas, D.P. (2016). Nonlinear Programming (3rd ed.). Belmont, MA: Athena Scientific.
5. Deb, S., & Das, S. (2023). Optimal selection of the starting lineup for a football team. arXiv:2303.12385. https://doi.org/10.48550/arXiv.2303.12385
6. Harris, A., Gabbett, T.J., King, R., Bird, S.P., & Terry, P. (2025). Influence of acute and chronic load on perceived wellbeing, neuromuscular performance and immune function in male professional football players. Sports, 13(6), 176. https://doi.org/10.3390/sports13060176
7. Lasek, J., & Gagolewski, M. (2021). Interpretable sports team rating models based on the gradient descent algorithm. International Journal of Forecasting, 37(3), 1061–1071. https://doi.org/10.1016/j.ijforecast.2020.11.008
8. Mahrudinda, S., Supian, S., Subiyanto, & Chaerani, D. (2020). Optimization of the best line-up in football using binary integer programming model. International Journal of Global Operations Research, 1(3), 114–122.
9. Mandorino, M., Gabbett, T.J., Tessitore, A., Leduc, C., Persichetti, V., & Lacome, M. (2025). The interaction of fitness and fatigue on physical and tactical performance in football. Applied Sciences, 15(7), 3574. https://doi.org/10.3390/app15073574
10. Nocedal, J., & Wright, S.J. (2006). Numerical Optimization (2nd ed.). New York: Springer.
11. Pantuso, G., & Hvattum, L.M. (2021). Maximizing performance with an eye on the finances: a chance-constrained model for football transfer market decisions. TOP, 29(2), 583–611. https://doi.org/10.1007/s11750-020-00584-9
12. Zhao, H., Chen, H., Yu, S., & Chen, B. (2021). Multi-objective optimization for football team selection. IEEE Access, 9, 90475–90487. https://doi.org/10.1109/ACCESS.2021.3091185