Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Constrained Multi-Agent Optimal Control for Football Tactical Decision Support: Analysis, Solver Verification, and Reproducible Scenario Experiments

Abbes Kadyri

kadyri.a@ucd.ac.ma

Chouaib Doukkali University

Abdelfettah Belkahla

Chouaib Doukkali University

Otmane Bencheihb

Chouaib Doukkali University

Research Article

Keywords: Optimal control, Football analytics, Multi-agent systems, Tactical decision support, Projected gradient, Spatial control, Pontryagin minimum principle

Posted Date: September 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10816173/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Constrained Multi-Agent Optimal Control for Football Tactical Decision Support: Analysis, Solver Verification, and Reproducible Scenario Experiments

Abbes Kadyri[a,*], Abdelfettah Belkahla[a], Otmane Bencheihb[a]

[a]*Department of Mathematics, Chouaib Doukkali University, El Jadida 24000, Morocco*

Abstract

This work develops a constrained multi-agent optimal-control framework for football tactical planning, where player positions are modeled as states and admissible velocities as controls. A dimensionless Bolza functional combines territorial control, defensive exposure, compactness, pressing coverage, movement effort, and terminal formation accuracy. Existence of an optimal control is established, first-order optimality conditions are derived using Pontryagin's Minimum Principle, and a time-discrete adjoint with an Armijo projected-gradient algorithm is employed for numerical solution. In a reproducible 11-player synthetic benchmark on a 105×68 m pitch, the optimized strategy reduces the objective from 0.20935 to 0.18511, corresponding to an 11.58% improvement over a feasible linear transition and a 38.26% improvement over the static compact baseline. All speed and pitch constraints are satisfied, while gradient verification and grid-refinement tests confirm numerical consistency. Priority experiments further reveal clear trade-offs between pressing intensity, defensive exposure, and movement effort. The proposed synthetic benchmark is intended to validate the mathematical and computational framework and provide a reproducible basis for future calibration with real tracking data.

Keywords: Optimal control, Football analytics, Multi-agent systems, Tactical decision support, Projected gradient, Spatial control, Pontryagin minimum principle

*Corresponding author.
Email address: kadyri.a@ucd.ac.ma (Abbes Kadyri)

1. Introduction

Football tactics are expressed through coordinated spatial and temporal adaptations rather than by a single static formation. A team may defend in a compact block, expand in possession, press in advanced zones, and rapidly reorganize after losing the ball. Positional and event data have therefore become central to quantitative analyses of compactness, spatial occupation, passing structure, pressing, and territorial dominance [1, 2, 3, 4]. Spatial-control models in particular estimate which team is better positioned to reach or influence regions of the pitch [5, 6, 7, 8, 9].

Much of the football-analytics literature is descriptive or predictive: it estimates what happened, detects tactical patterns, or predicts values from observed tracking and event data. Recent work has expanded this direction through counterpressing detection [10], space-control learning [11], graph representations of multi-agent tracking data [12], automatic tactical-periodization models [13], and tracking-based pressure indicators [14]. A recent systematic review confirms the rapid growth of artificial-intelligence methods for mapping collective tactical behavior [15]. These approaches are valuable, but they address a different question from the one considered here: given a tactical objective and explicit motion constraints, how should a collective configuration evolve?

Optimal-control theory provides a natural language for this prescriptive question because dynamics, admissible controls, state constraints, and competing objectives are treated within one mathematical problem [16, 17, 18]. Multi-agent coordination and differential games offer related perspectives when several decision-making agents or opponents are simultaneously strategic [19, 20, 21].

The present paper develops a transparent bridge between football tactical indicators and constrained optimal control. The opponent is prescribed rather than optimized, which isolates the single-team control problem and makes the computational structure explicit. The main contributions are:

(i) a constrained multi-agent formulation with normalized territorial, defensive, compactness, pressing, effort, and terminal losses;

(ii) an existence result, smoothness properties, Pontryagin necessary conditions, and an explicit interior control law;

(iii) a consistent time-discrete adjoint and a projected-gradient method with

Armijo backtracking, together with a first-order stationarity result for the discrete iteration;

(iv) an actual implementation of the stated optimization algorithm rather than a prescribed interpolation presented as an optimizer output;

(v) solver verification through finite-difference gradient checks, constraint audits, spatial-grid sensitivity, and multi-start experiments; and

(vi) fully documented synthetic coordinates, parameters, code, and numerical outputs, allowing independent reproduction without proprietary data.

The scope is deliberately methodological. The numerical benchmark is synthetic and does not establish predictive validity, causal coaching benefit, or performance in professional matches. Those claims require tracking-data calibration and independent empirical evaluation. This distinction is maintained throughout the paper.

2. Related work and positioning

Positional data support the analysis of collective shape, team centroids, occupied area, synchronization, compactness, and player roles [3, 4]. Spatial-control formulations extend these descriptions by constructing influence or arrival-time fields. Examples include dominant-region concepts, probabilistic movement zones, continuous influence fields, and location-dependent dangerousness [7, 8, 9].

Data-driven methods increasingly learn tactical representations directly from tracking or synchronized event data. Bauer and Anzer [10] formulate counterpressing detection as a supervised learning problem; Gu et al. [11] quantify in-game space-control efficiency; Raabe et al. [12] represent multi-agent spatiotemporal sports data using tactical graphs. Biermann et al. [13] combine tactical periodization with machine learning to create frame-by-frame match-phase annotations, while Li et al. [14] develop a tracking-based defensive pressure index. These studies emphasize the richness of real positional data and motivate the tactical quantities represented in the present objective.

The distinction of this work is not that territorial control, compactness, or pressing are new concepts. Rather, the contribution is to encode them simultaneously as differentiable losses in a constrained continuous-time control

problem, to derive the associated optimality system, and to connect that system to a verified numerical solver. This positioning also avoids claiming that the present synthetic benchmark replaces empirical football analytics. Instead, the framework is designed so that future tracking data can supply the initial state, prescribed opponent path, tactical anchor functions, and calibration targets.

3. Continuous optimal-control model

3.1. Pitch domain, state, and admissible controls

Let the pitch be the closed rectangle

$$\Omega = [0, L] \times [0, W] \subset \mathbb{R}^2, \tag{1}$$

where $L > 0$ and $W > 0$. For n_p modeled players, player i has position $x_i(t) \in \Omega$ and velocity control $u_i(t) \in \mathbb{R}^2$. The team state and control are

$$X(t) = (x_1(t), \dots, x_{n_p}(t)) \in \Omega^{n_p}, \qquad U(t) = (u_1(t), \dots, u_{n_p}(t)) \in \mathbb{R}^{2n_p}. \tag{2}$$

The first-order kinematics are

$$\dot{x}_i(t) = u_i(t), \qquad x_i(0) = x_{i,0}, \qquad i = 1, \dots, n_p. \tag{3}$$

For a maximum speed $V_{\max} > 0$, define $B_{V_{\max}} = \{v \in \mathbb{R}^2 : \|v\|_2 \le V_{\max}\}$. The admissible set is

$$\mathcal{U}_{\mathrm{ad}} = \left\{ U \in L^2(0, T; \mathbb{R}^{2n_p}) : u_i(t) \in B_{V_{\max}} \text{ a.e., } X_U(t) \in \Omega^{n_p} \ \forall t \in [0, T] \right\}. \tag{4}$$

Thus pitch feasibility is part of the optimization problem rather than a post-processing correction.

3.2. Prescribed opponent and smooth spatial control

The opponent is represented by a prescribed path

$$Y(t) = (y_1(t), \dots, y_{n_o}(t)) \in \Omega^{n_o}. \tag{5}$$

For a regularization parameter $\kappa > 0$, define

$$d_\kappa(z, x) = \sqrt{\|z - x\|_2^2 + \kappa^2}. \tag{6}$$

With arrival-time sensitivity $\alpha > 0$, the influence contributions are

$$q_i^A(z;X) = \exp\left[-\alpha \frac{d_\kappa(z,x_i)}{V_{\max}}\right], \qquad q_j^B(z;Y) = \exp\left[-\alpha \frac{d_\kappa(z,y_j)}{V_{\max}}\right]. \quad (7)$$

Let $I_A = \sum_i q_i^A$ and $I_B = \sum_j q_j^B$. The analyzed team's normalized spatial-control share is

$$P_A(z;X,Y) = \frac{I_A(z;X)}{I_A(z;X) + I_B(z;Y)}. \quad (8)$$

Because all influence terms are positive, $0 < P_A < 1$. The quantity is a model-based normalized share; it is not interpreted as an empirically calibrated probability of possession or ball recovery.

Let $\omega(z,t) \geq 0$ be a strategic spatial weight normalized by

$$\int_\Omega \omega(z,t)\, dz = 1. \quad (9)$$

The territorial loss is

$$\ell_{\text{terr}}(X,t) = 1 - \int_\Omega \omega(z,t) P_A(z;X,Y(t))\, dz. \quad (10)$$

3.3. Normalized tactical losses

Let $D_\Omega = \sqrt{L^2 + W^2}$ denote the pitch diameter. The collective centroid and compactness loss are

$$c(X) = \frac{1}{n_p} \sum_{i=1}^{n_p} x_i, \qquad \ell_{\text{comp}}(X) = \frac{1}{n_p D_\Omega^2} \sum_{i=1}^{n_p} \|x_i - c(X)\|_2^2. \quad (11)$$

For a nonempty defensive set $\mathcal{D}$ and adjacency set E_D, define the smooth positive-part approximation

$$h_\delta(s) = \frac{1}{2}\left(s + \sqrt{s^2 + \delta^2}\right), \qquad \delta > 0. \quad (12)$$

Let $c_D(X) = |\mathcal{D}|^{-1} \sum_{i \in \mathcal{D}} x_i$ and let $a_D(t)$ be a prescribed defensive anchor. Then

$$\ell_{\text{def}}(X,t) = \frac{1}{|E_D| D_\Omega^2} \sum_{(i,j) \in E_D} h_\delta(\|x_i - x_j\|_2 - d_{\text{safe}})^2 + \frac{\gamma_D}{D_\Omega^2} \|c_D(X) - a_D(t)\|_2^2. \quad (13)$$

The first term penalizes excessive separation of linked defenders while remaining differentiable.

Let $\mathcal{P}$ be the pressing unit, let $q_r(t) \in \Omega$, $r = 1, \dots, n_q$, be prescribed pressure points, and let $\eta_r \geq 0$ with $\sum_r \eta_r = 1$. For $\rho > 0$, define

$$\pi_r(X,t) = 1 - \prod_{i \in \mathcal{P}} \left[1 - \exp\left(-\frac{\|x_i - q_r(t)\|_2^2}{2\rho^2} \right) \right], \tag{14}$$

$$\ell_{\text{press}}(X,t) = 1 - \sum_{r=1}^{n_q} \eta_r \pi_r(X,t). \tag{15}$$

The movement-effort loss is

$$\ell_{\text{eff}}(U) = \frac{1}{n_p V_{\max}^2} \sum_{i=1}^{n_p} \|u_i\|_2^2 . \tag{16}$$

For a target formation X_{tar} and $\beta \geq 0$, define the terminal loss

$$\Phi(X(T)) = \frac{\beta}{n_p D_\Omega^2} \sum_{i=1}^{n_p} \|x_i(T) - x_{i,\text{tar}}\|_2^2 . \tag{17}$$

3.4. Bolza functional

Let $\theta_r \geq 0$ denote raw tactical priorities and normalize them as

$$w_r = \frac{\theta_r}{\sum_s \theta_s}, \qquad \sum_r w_r = 1. \tag{18}$$

With $r \in \{\text{terr}, \text{def}, \text{comp}, \text{press}, \text{eff}\}$, define

$$\ell_0(X,U,t) = w_{\text{terr}}\ell_{\text{terr}} + w_{\text{def}}\ell_{\text{def}} + w_{\text{comp}}\ell_{\text{comp}} + w_{\text{press}}\ell_{\text{press}} + w_{\text{eff}}\ell_{\text{eff}} \tag{19}$$

and the time-normalized Bolza functional

$$J(U) = \frac{1}{T} \int_0^T \ell_0(X_U(t), U(t), t)\, dt + \Phi(X_U(T)). \tag{20}$$

The continuous tactical-control problem is

$$\min_{U \in \mathcal{U}_{\text{ad}}} J(U). \tag{21}$$

4. Analytical properties

4.1. Existence

Assumption 1. *Ω is nonempty, compact, and convex; $T > 0$; $X_0 \in \Omega^{n_p}$; and $\mathcal{U}_{\mathrm{ad}} \neq \varnothing$. The prescribed functions Y, ω, a_D, and q_r are continuous and bounded on their domains. The parameters $V_{\max}, \alpha, \kappa, \rho, \delta, d_{\mathrm{safe}}$ are strictly positive and all remaining coefficients are finite and nonnegative.*

Theorem 1 (Existence)**.** *Under Assumption 1, problem* (21) *admits at least one minimizer $U^\star \in \mathcal{U}_{\mathrm{ad}}$.*

Proof. Let (U^m) be a minimizing sequence. The pointwise speed bound yields boundedness in $L^2(0,T;\mathbb{R}^{2n_p})$, so a subsequence converges weakly to some $U^\star$. The associated states satisfy

$$X^m(t) = X_0 + \int_0^t U^m(s)\, ds.$$

They are bounded in $W^{1,2}(0,T;\mathbb{R}^{2n_p})$. In one time dimension, the compact embedding into $C([0,T];\mathbb{R}^{2n_p})$ gives a uniformly convergent subsequence $X^m \to X^\star$. Passing to the limit in the integral state equation identifies $X^\star$ as the trajectory generated by $U^\star$. The pointwise control constraint is weakly closed because $B^{n_p}_{V_{\max}}$ is closed and convex, and the state constraint is preserved by uniform convergence and closedness of Ω^{n_p}. Hence $U^\star \in \mathcal{U}_{\mathrm{ad}}$.

All state-dependent integrands are continuous on a compact configuration set and therefore converge under uniform state convergence. The effort term is convex in U and weakly lower semicontinuous; the terminal term is continuous. Thus

$$J(U^\star) \leq \liminf_{m\to\infty} J(U^m) = \inf_{U\in\mathcal{U}_{\mathrm{ad}}} J(U),$$

which proves existence; see also the standard direct-method arguments in [16, 17]. □

4.2. Differentiability

Proposition 1 (Smooth state losses)**.** *Under Assumption 1, with $\kappa, \delta, \rho > 0$, the mappings $X \mapsto \ell_{\mathrm{terr}}$, ℓ_{def}, ℓ_{comp}, and ℓ_{press} are continuously differentiable on Ω^{n_p}. Their gradients are Lipschitz on Ω^{n_p}.*

Proof. The regularized distance d_κ is smooth because its argument is bounded away from zero. The influence fields are therefore smooth, and the denominator in (8) is strictly positive. The compactness term is polynomial in X. The map h_δ is smooth for $\delta > 0$, while the pressing term is a finite product of smooth Gaussian factors. Continuity of the Hessians on the compact set Ω^{n_p} implies bounded Hessians and therefore Lipschitz gradients. □

4.3. Pontryagin conditions

Assume temporarily that the optimal trajectory remains in the interior of the pitch. Define

$$L_0(X, U, t) = \frac{1}{T}\ell_0(X, U, t), \qquad \mathcal{H}(X, U, \Lambda, t) = L_0(X, U, t) + \Lambda^\top U. \tag{22}$$

Theorem 2 (First-order necessary conditions)**.** *Let $(X^\star, U^\star)$ be an optimal pair for which the pitch constraints are inactive. Then there exists an absolutely continuous adjoint Λ such that*

$$\dot{X}^\star(t) = U^\star(t), \tag{23}$$

$$\dot{\Lambda}(t) = -\frac{1}{T}\nabla_X \ell_0(X^\star, U^\star, t), \tag{24}$$

$$\Lambda(T) = \nabla_X \Phi(X^\star(T)), \tag{25}$$

and, almost everywhere,

$$U^\star(t) \in \arg\min_{V \in B^{n_p}_{V_{\max}}} \mathcal{H}(X^\star(t), V, \Lambda(t), t). \tag{26}$$

If $w_{\text{eff}} > 0$, each player control satisfies

$$u_i^\star(t) = \mathrm{Proj}_{B_{V_{\max}}}\left(-\frac{T n_p V_{\max}^2}{2 w_{\text{eff}}}\lambda_i(t)\right). \tag{27}$$

Proof. The standard finite-dimensional Pontryagin Minimum Principle gives the state, adjoint, transversality, and Hamiltonian-minimization relations [22, 23, 18]. Since the only control-dependent running term is quadratic, minimizing the Hamiltonian over the Euclidean ball gives (27) by projection. □

Remark 1. *When a pitch boundary is active, state-constrained necessary conditions include additional normal-cone or measure-valued multipliers. Equation* (27) *is therefore an interior formula. The discrete method below enforces pitch feasibility directly and does not require assuming interiority.*

5. Time discretization and projected-gradient solver

5.1. Discrete problem

Let $t_k = k\Delta t$, $\Delta t = T/K$, and use piecewise-constant controls. The discrete dynamics are exact for (3):

$$X^{k+1} = X^k + \Delta t\, U^k, \qquad X^0 = X_0. \tag{28}$$

The rectangle-rule approximation is

$$J_h(U) = \frac{\Delta t}{T} \sum_{k=0}^{K-1} \ell_0(X^k, U^k, t_k) + \Phi(X^K). \tag{29}$$

The feasible set is

$$\mathcal{A}_h = \{U : \left\|u_i^k\right\|_2 \le V_{\max},\ X^{k+1} = X^k + \Delta t U^k,\ X^k \in \Omega^{n_p}\}. \tag{30}$$

Because the dynamics are affine, the pitch is convex, and the velocity balls are convex, $\mathcal{A}_h$ is closed and convex. It is also bounded, hence compact in the finite-dimensional discrete space.

5.2. Discrete adjoint and reduced gradient

For a feasible control sequence, compute the state forward from (28). The discrete adjoint satisfies

$$\Lambda^K = \nabla_X \Phi(X^K), \tag{31}$$

$$\Lambda^k = \Lambda^{k+1} + \frac{\Delta t}{T} \nabla_X \ell_0(X^k, U^k, t_k), \qquad k = K-1, \dots, 0. \tag{32}$$

The exact reduced gradient is

$$G^k = \nabla_{U^k} J_h = \Delta t \left[\frac{1}{T} \nabla_U \ell_0(X^k, U^k, t_k) + \Lambda^{k+1} \right]. \tag{33}$$

For the effort term,

$$\nabla_{u_i} \ell_0 = \frac{2 w_{\text{eff}}}{n_p V_{\max}^2} u_i. \tag{34}$$

5.3. Projection and Armijo update

Given $s_m > 0$, define

$$U^{m+1} = \mathrm{Proj}_{\mathcal{A}_h}\big(U^m - s_m G^m\big). \tag{35}$$

The projection is the unique solution of the strongly convex problem

$$\min_{V \in \mathcal{A}_h} \frac{1}{2} \left\| V - (U^m - s_m G^m) \right\|_2^2. \tag{36}$$

Because the speed constraints are Euclidean balls and pitch constraints become linear cumulative-state inequalities, the general projection is a convex quadratically constrained problem rather than a plain unconstrained quadratic program.

The Armijo test accepts a step when

$$J_h(U^{m+1}) \le J_h(U^m) - \frac{\sigma}{s_m} \left\| U^{m+1} - U^m \right\|_2^2, \qquad \sigma \in (0,1). \tag{37}$$

Define the projected-gradient residual

$$R_m = \frac{1}{s_m} \left\| U^m - \mathrm{Proj}_{\mathcal{A}_h}(U^m - s_m \nabla J_h(U^m)) \right\|_2. \tag{38}$$

Proposition 2. *Under Assumption 1, J_h is continuously differentiable and has a Lipschitz gradient on $\mathcal{A}_h$.*

Proof. The discrete state is an affine function of the control sequence. Proposition 1 gives Lipschitz gradients for the state-dependent losses on the compact state set, and the effort term is quadratic. Composition with the affine state map preserves continuous differentiability and Lipschitz continuity of the gradient on compact $\mathcal{A}_h$. □

Theorem 3 (Stationarity of accumulation points)**.** *Assume that Armijo backtracking starts from a fixed maximum step $\bar{s} > 0$ and reduces the trial step by a fixed factor $\tau \in (0,1)$ until (37) holds. Then the generated objective values are nonincreasing, $R_m \to 0$, and every accumulation point $\bar{U} \in \mathcal{A}_h$ satisfies the first-order stationarity condition*

$$0 \in \nabla J_h(\bar{U}) + N_{\mathcal{A}_h}(\bar{U}), \tag{39}$$

where $N_{\mathcal{A}_h}$ is the convex normal cone.

Proof. By Proposition 2, ∇J_h is Lipschitz with some constant L_h on $\mathcal{A}_h$. The projection variational inequality implies, with $D^m = U^{m+1} - U^m$,

$$\langle \nabla J_h(U^m), D^m \rangle \leq -\frac{1}{s_m} \|D^m\|_2^2 .$$

The descent lemma therefore gives

$$J_h(U^{m+1}) \leq J_h(U^m) - \left(\frac{1}{s_m} - \frac{L_h}{2} \right) \|D^m\|_2^2 .$$

Hence the Armijo condition is guaranteed whenever $s_m \leq 2(1-\sigma)/L_h$. Fixed-factor backtracking consequently produces accepted steps bounded below by a positive constant depending on $\bar{s}$, τ, σ, and L_h. Summing (37) shows $\sum_m \|D^m\|_2^2 < \infty$, so $\|D^m\|_2 \to 0$. The positive lower bound on s_m gives $R_m \to 0$. Compactness of $\mathcal{A}_h$ supplies accumulation points, and continuity of the projected-gradient mapping yields $\bar{U} = \mathrm{Proj}_{\mathcal{A}_h}(\bar{U} - s\nabla J_h(\bar{U}))$ for an admissible limit step $s > 0$, which is equivalent to $0 \in \nabla J_h(\bar{U}) + N_{\mathcal{A}_h}(\bar{U})$; see standard projected-gradient analysis in [24, 25]. □

6. Reproducible numerical benchmark

6.1. Setup

The benchmark uses the same 11-player compact and proactive configurations listed in Appendix Appendix A. The compact configuration is the initial state and the proactive configuration is the target. The opponent is fixed over the 12 s horizon. These coordinates are synthetic model states and do not correspond to official tracking records.

The strategic spatial weight is fixed in time and evaluated on a rectangular quadrature grid:

$$\widetilde{\omega}(z) = \exp\left[-\frac{1}{2}\left(\frac{z_1 - 75}{22} \right)^2 - \frac{1}{2}\left(\frac{z_2 - 34}{18} \right)^2 \right] + 0.15, \tag{40}$$

with numerical normalization so that the discrete weights sum to one. The defensive unit is players 2–5, with adjacency edges $(2,3)$, $(3,4)$, and $(4,5)$. The pressing unit is players 6–11 and the three pressure points are opponent players 6, 9, and 11, with equal weights. The defensive anchor is the centroid of target defenders 2–5.

Table 1: Numerical parameters for the reproducible benchmark.

Quantity	Symbol	Value
Pitch size	$L \times W$	105×68 m
Horizon	T	12 s
Time steps	K	60
Time step	Δt	0.20 s
Spatial grid	$N_x \times N_y$	40×26
Maximum speed	$V_{\max}$	7.0 m s^{-1}
Influence sensitivity	α	0.75 s^{-1}
Distance regularization	κ	10^{-3} m
Pressing radius	ρ	10 m
Smooth gap parameter	δ	0.5 m
Safe defender separation	d_{safe}	19 m
Defensive anchor weight	γ_D	0.6
Terminal coefficient	β	2.0
Armijo maximum step	$\bar{s}$	800
Armijo parameter	σ	10^{-4}
Stopping residual		2×10^{-4}

The balanced weights are

$$(w_{\text{terr}}, w_{\text{def}}, w_{\text{comp}}, w_{\text{press}}, w_{\text{eff}}) = (0.30, 0.20, 0.10, 0.25, 0.15). \tag{41}$$

Two additional vectors are used later: pressing-first $(0.25, 0.10, 0.05, 0.45, 0.15)$ and defense-first $(0.20, 0.45, 0.15, 0.05, 0.15)$.

The main solver is initialized with the feasible constant control

$$U_{\text{lin}}^k = \frac{X_{\text{tar}} - X_0}{T}, \tag{42}$$

which produces a linear formation transition. During the reported optimization all iterates remain strictly inside the pitch; consequently, the full projection onto $\mathcal{A}_h$ reduces in this benchmark to independent Euclidean projections onto the player speed balls. Pitch feasibility is checked after every trial step.

AI-assisted software development.. OpenAI ChatGPT was used as a code-drafting aid for portions of the Python implementation. The numerical

algorithm is specified explicitly in Sections 5–6 and the complete reproducibility script is supplied. Finite-difference gradient checks, constraint audits, spatial-grid sensitivity, and multi-start tests are reported so that the computational outputs can be independently audited.

6.2. Solver verification

A finite-difference test was performed at $0.8U_{\mathrm{lin}}$ on 12 deterministically selected control components. With a centered perturbation of 10^{-6}, the maximum absolute difference between finite differences and the adjoint gradient was 9.74×10^{-11}. Excluding one component whose derivative magnitude is below 10^{-7}, the maximum relative difference was 3.05×10^{-5}. The full component-by-component audit is supplied in `gradient_check.csv`.

Spatial quadrature sensitivity is summarized in Table 2. The territorial loss changes by only 0.135% between the 40×26 grid used by the solver and a 120×78 evaluation grid. This verifies that the reported territorial value is not driven by a coarse spatial discretization.

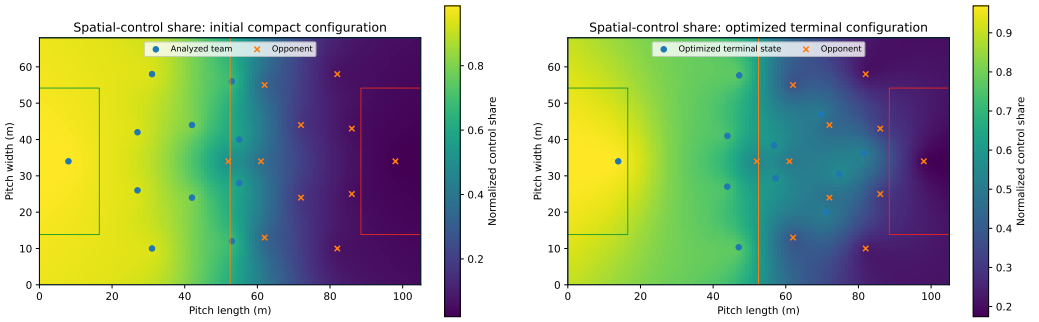
Table 2: Spatial-grid sensitivity of the time-averaged territorial loss for the balanced optimized trajectory.

N_x	N_y	Grid points	Territorial loss
40	26	1040	0.543152
60	40	2400	0.542801
80	52	4160	0.542613
120	78	9360	0.542422

A multi-start test under the balanced weights uses static, linear, and randomized feasible initializations. The final objectives are 0.185159, 0.185113, and 0.185193, respectively, giving a relative spread of only 0.043%. This does not prove global optimality of the nonconvex problem, but it provides a useful numerical robustness check for the reported benchmark.

6.3. Balanced optimization versus feasible baselines

Figures 1a and 1b show the model-based spatial-control field at the compact initial state and at the optimized terminal state. Figure 2 displays the trajectories returned by the projected-gradient solver. These curves are therefore optimizer outputs, not a prescribed linear interpolation.



(a) Model-based spatial-control share for the initial compact configuration against the fixed opponent.

(b) Model-based spatial-control share at the optimized terminal configuration under balanced priorities.

Figure 1: Model-based spatial-control comparison between the initial compact and optimized terminal configurations.

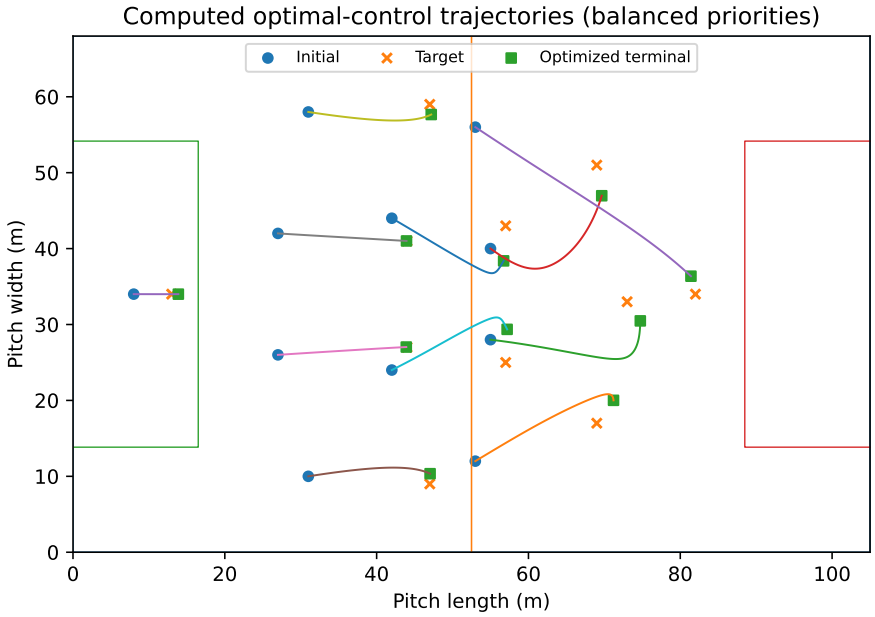


Figure 2: Computed player trajectories for the balanced-priority optimal-control problem. Initial, target, and optimized terminal locations are shown explicitly.

Table 3 compares the static compact state, the feasible linear transition, and the optimized control. The optimized objective is 0.185113, compared with 0.209347 for the linear transition and 0.299825 for the static baseline.

This is an 11.58% reduction relative to the linear transition and a 38.26% reduction relative to remaining static. The optimized solution accepts greater movement effort than the linear baseline but substantially reduces territorial and especially pressing losses. The terminal RMS formation error is 2.87 m because the target is a soft terminal objective rather than a hard terminal constraint.

Table 3: Balanced-priority baseline comparison. All component entries are time-averaged normalized losses except the terminal loss.

Case	J_h	Terr.	Def.	Comp.	Press	Effort	Terminal	Final RMS error (m)
Static compact	0.299825	0.628260	0.010438	0.027982	0.247245	0.000000	0.044650	18.691
Linear transition	0.209347	0.565815	0.003567	0.029502	0.114047	0.049513	≈ 0	≈ 0
Optimized	0.185113	0.543152	0.002736	0.029215	0.026944	0.072718	0.001055	2.874

The accepted iteration reaches the stopping conditions after 88 iterations with $R_m = 1.57 \times 10^{-4}$. The maximum observed player speed is 7.0 m s^{-1}, exactly the imposed bound. The optimized states remain within $x \in [8.0, 81.41]$ m and $y \in [10.0, 58.0]$ m, so no pitch-boundary constraint becomes active. Figures 3 and 4 report the objective and stationarity histories.

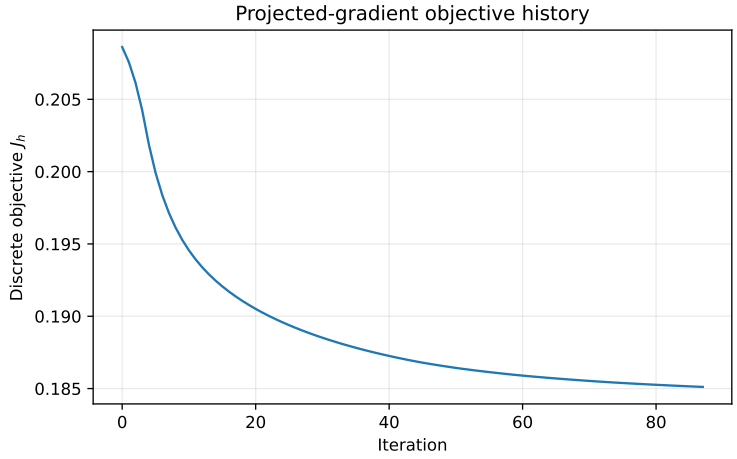


Figure 3: Monotone objective history of the balanced projected-gradient run.

Figure 5 isolates the time-averaged loss components and makes the optimization trade-off visible. The effort increase is not an artifact hidden inside a composite score: it is displayed explicitly beside the reductions in pressing and territorial losses.

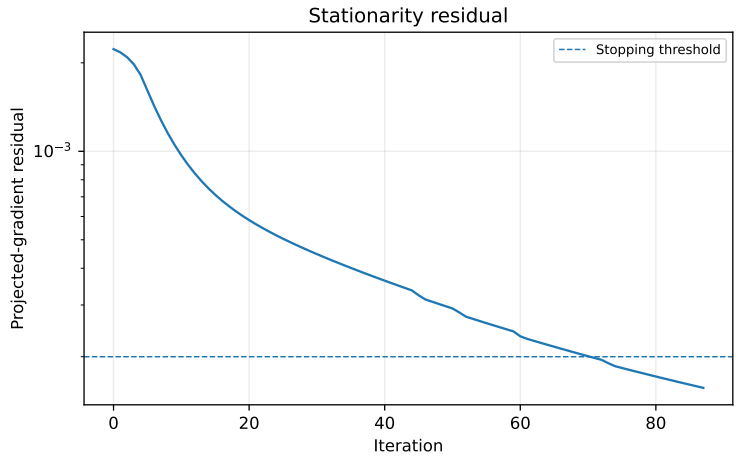


Figure 4: Projected-gradient residual for the balanced run. The dashed line is the stopping threshold.

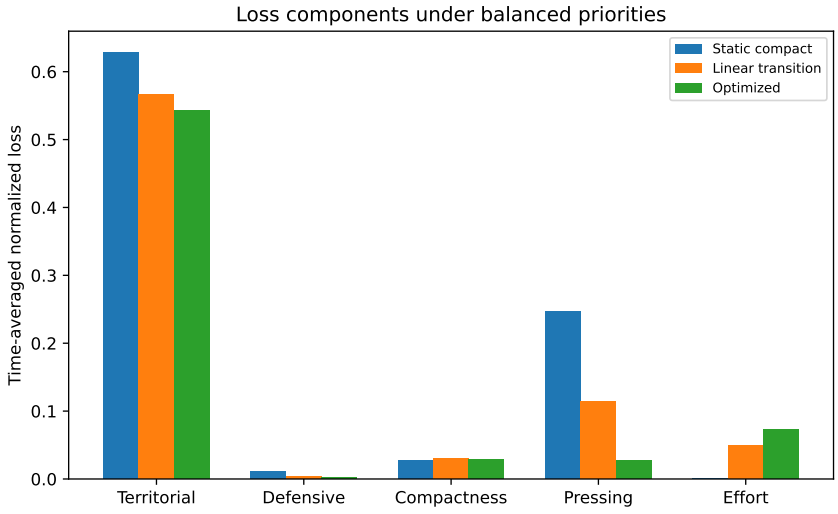


Figure 5: Balanced-priority loss components for the static, linear-transition, and optimized controls.

6.4. Priority-dependent optimized behavior

To test whether the computed solution responds coherently to tactical preferences, the optimization is repeated with the pressing-first and defense-first weight vectors. Table 4 and Figure 6 report the resulting components.

Raw objective values are not compared across weight vectors because each profile defines a different scalarized objective.

The pressing-first solution decreases pressing loss from 0.026944 to 0.023432 while maintaining a similar movement effort. The defense-first solution reduces defensive loss and movement effort, with RMS speed decreasing from 1.888 to 1.630 m s^{-1} and maximum speed from 7.0 to 4.08 m s^{-1}, but its pressing loss increases to 0.054828. These changes are consistent with the intended meaning of the weights and provide an internal behavioral check on the model.

Table 4: Optimized component losses under three tactical-priority vectors. Objectives are not compared because the scalarization weights differ.

Profile	Terr.	Def.	Comp.	Press	Effort	RMS speed (m/s)
Balanced	0.543152	0.002736	0.029215	0.026944	0.072718	1.888
Pressing-first	0.545947	0.003054	0.029501	0.023432	0.072980	1.891
Defense-first	0.552377	0.002804	0.028207	0.054828	0.054247	1.630

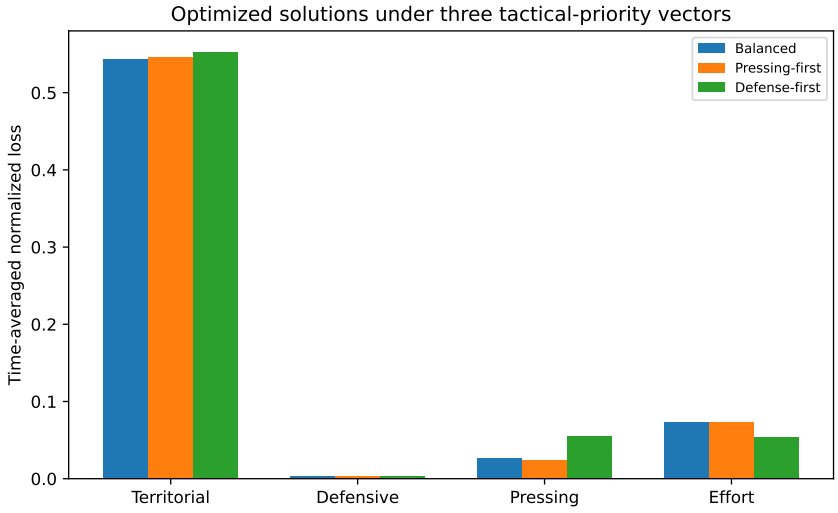


Figure 6: Optimized tactical-loss components under balanced, pressing-first, and defense-first priorities.

7. Discussion

The strengthened computational results close an important methodological gap between formulating an optimal-control problem and merely drawing

a tactical transition. The proposed state equations, adjoint recursion, and projected-gradient update are now connected to trajectories that are actually returned by the solver. The finite-difference audit verifies the adjoint implementation, while the Armijo history and residual provide explicit convergence diagnostics. Multi-start and grid-sensitivity experiments further reduce the risk that the reported behavior is a plotting or discretization artifact.

The optimized balanced solution illustrates the intended multi-objective interpretation. Compared with a direct formation interpolation, the algorithm expends more movement energy but moves influential players so that the pressing loss falls sharply and territorial control improves, while the terminal formation remains close to the proactive target. The priority experiments show that this behavior is not hard-coded: changing weights alters the compromise between pressure, defensive structure, and movement cost.

Several limitations remain. First, the motion model is first order and does not represent acceleration, orientation, individual speed heterogeneity, fatigue, or biomechanical limits. A second-order model $\dot{x}_i = v_i$, $\dot{v}_i = a_i$ with acceleration bounds is a natural extension. Second, the opponent is prescribed rather than adaptive. If both teams optimize strategically, the appropriate extension is a differential game rather than a single-team optimal-control problem [21]. Third, the ball, pass feasibility, score state, substitutions, and stochastic events are omitted. Fourth, the tactical loss weights are illustrative rather than learned from coaches or tracking data. Finally, and most importantly, the present experiments are synthetic. They verify the internal mathematics and computation but do not establish empirical football validity.

A data-calibrated study should extract short tactical phases from open or authorized tracking data, use observed positions for X_0 and $Y(t)$, define pressure points from the ball carrier and nearby passing options, estimate or elicit tactical weights, and compare optimized counterfactual trajectories with held-out match sequences. Recent tracking-based work on tactical periodization and defensive pressure provides suitable empirical concepts for such a stage [13, 14].

8. Conclusions

A constrained multi-agent optimal-control framework for football tactical decision support has been developed and computationally verified. The continuous problem admits a minimizer under standard compactness and

continuity assumptions, and Pontryagin conditions provide an interpretable continuous first-order characterization. The discrete formulation admits an exact adjoint gradient, and projected gradient with Armijo backtracking approaches first-order stationarity under a Lipschitz-gradient assumption.

The reproducible 11-player benchmark demonstrates that the actual solver improves the chosen balanced objective by 11.58% relative to a feasible linear transition and by 38.26% relative to a static compact baseline. Constraint satisfaction, gradient verification, grid sensitivity, and multi-start consistency are reported explicitly. The priority experiments reveal interpretable trade-offs rather than a universally superior tactical profile.

The study should therefore be read as a solver-verified mathematical-computational framework, not as a validated match-prediction system. The next step is direct calibration and out-of-sample testing with player-and-ball tracking data, followed by second-order motion and adaptive-opponent extensions.

Data and code availability

No proprietary or external match data are used in the numerical study. All synthetic coordinates are listed in Appendix Appendix A. The accompanying file `reproduce_results.py` regenerates the reported CSV tables, optimization histories, and figures from the stated parameters. Numerical arrays of the optimized controls and states are included in the compressed numerical archive supplied with the reproducibility files.

Appendix A. Synthetic scenario coordinates

Coordinates are in meters on a 105×68 m pitch. They are illustrative model configurations and do not correspond to official tracking records.

Table A.5: Compact initial state, proactive target, and fixed opponent coordinates.

	Compact X_0		Target X_{tar}		Opponent Y	
Player	x	y	x	y	x	y
1	8	34	13	34	98	34
2	31	10	47	9	82	10
3	27	26	44	27	86	25
4	27	42	44	41	86	43
5	31	58	47	59	82	58
6	42	24	57	25	72	24
7	42	44	57	43	72	44
8	53	12	69	17	62	13
9	55	28	73	33	61	34
10	55	40	69	51	62	55
11	53	56	82	34	52	34

References

[1] C. Carling, A. M. Williams, T. Reilly, Handbook of Soccer Match Analysis: A Systematic Approach to Improving Performance, Routledge, London, 2005.

[2] R. Rein, D. Memmert, Big data and tactical analysis in elite soccer: future challenges and opportunities for sports science, SpringerPlus 5 (2016) 1410. doi:10.1186/s40064-016-3108-2.

[3] J. Gudmundsson, M. Horton, Spatio-temporal analysis of team sports, ACM Computing Surveys 50 (2) (2017) 22. doi:10.1145/3054132.

[4] B. Low, D. Coutinho, B. Gonçalves, R. Rein, D. Memmert, J. Sampaio, A systematic review of collective tactical behaviours in football using positional data, Sports Medicine 50 (2020) 343–385. doi:10.1007/s40279-019-01194-7.

[5] W. Spearman, Beyond expected goals, in: Proceedings of the 12th MIT Sloan Sports Analytics Conference, Boston, 2018.

[6] J. Fernández, L. Bornn, Wide open spaces: A statistical technique for measuring space creation in professional soccer, in: Proceedings of the 12th MIT Sloan Sports Analytics Conference, Boston, 2018.

[7] D. Link, S. Lang, P. Seidenschwarz, Real-time quantification of dangerousity in football using spatiotemporal tracking data, PLOS ONE 11 (12) (2016) e0168768. doi:10.1371/journal.pone.0168768.

[8] U. Brefeld, J. Lasek, S. Mair, Probabilistic movement models and zones of control, Machine Learning 108 (1) (2019) 127–147. doi:10.1007/s10994-018-5725-1.

[9] F. Martens, U. Dick, U. Brefeld, Space and control in soccer, Frontiers in Sports and Active Living 3 (2021) 676179. doi:10.3389/fspor.2021.676179.

[10] P. Bauer, G. Anzer, Data-driven detection of counterpressing in professional football: A supervised machine learning task based on synchronized positional and event data with expert-based feature extraction, Data Mining and Knowledge Discovery 35 (2021) 2009–2049. doi:10.1007/s10618-021-00763-7.

[11] C. Gu, V. De Silva, M. Caine, A machine learning framework for quantifying in-game space-control efficiency in football, Knowledge-Based Systems 283 (2024) 111123. doi:10.1016/j.knosys.2023.111123.

[12] D. Raabe, R. Nabben, D. Memmert, Graph representations for the analysis of multi-agent spatiotemporal sports data, Applied Intelligence 53 (2023) 3783–3803. doi:10.1007/s10489-022-03631-z.

[13] H. Biermann, D. Memmert, N. Petersen, D. Raabe, Contextualization of soccer analysis with tactical periodization and machine learning, Data Mining and Knowledge Discovery 39 (2025) 23. doi:10.1007/s10618-025-01092-9.

[14] J. Li, F. Wu, A. Li, R. Martínez-Santos, Development and initial validation of a defensive pressure index using tracking data in the chinese super league, Frontiers in Sports and Active Living 8 (2026) 1833549. doi:10.3389/fspor.2026.1833549.

[15] J. E. Teixeira, E. Maio, P. Afonso, S. Encarnação, G. F. Machado, R. Morgans, T. M. Barbosa, A. M. Monteiro, P. Forte, R. Ferraz, L. Branquinho, Mapping football tactical behavior and collective dynamics with artificial intelligence: a systematic review, Frontiers in Sports and Active Living 7 (2025) 1569155. doi:10.3389/fspor.2025.1569155.

[16] L. Cesari, Optimization—Theory and Applications: Problems with Ordinary Differential Equations, Springer-Verlag, New York, 1983.

[17] W. H. Fleming, R. W. Rishel, Deterministic and Stochastic Optimal Control, Springer-Verlag, New York, 1975.

[18] D. Liberzon, Calculus of Variations and Optimal Control Theory: A Concise Introduction, Princeton University Press, Princeton, 2012.

[19] R. Olfati-Saber, J. A. Fax, R. M. Murray, Consensus and cooperation in networked multi-agent systems, Proceedings of the IEEE 95 (1) (2007) 215–233. doi:10.1109/JPROC.2006.887293.

[20] W. Ren, R. W. Beard, Distributed Consensus in Multi-Vehicle Cooperative Control, Springer, London, 2008. doi:10.1007/978-1-84800-015-5.

[21] T. Başar, G. J. Olsder, Dynamic Noncooperative Game Theory, 2nd Edition, SIAM, Philadelphia, 1999. doi:10.1137/1.9781611971132.

[22] E. B. Lee, L. Markus, Foundations of Optimal Control Theory, Wiley, New York, 1967.

[23] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, E. F. Mishchenko, The Mathematical Theory of Optimal Processes, Interscience, New York, 1962.

[24] J. Nocedal, S. J. Wright, Numerical Optimization, 2nd Edition, Springer, New York, 2006. doi:10.1007/978-0-387-40065-5.

[25] J. T. Betts, Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, 2nd Edition, SIAM, Philadelphia, 2010. doi:10.1137/1.9780898718577.