Evolving Football Formations Across Phases of Play from Tracking Data

Sam Nejati[1][0009-0005-3421-1695], Daniele Notarangelo[1][0009-0004-0990-2075], Jasnoor Singh[1][0009-0001-2824-976X], Erik Nielsen[1][0009-0007-8176-5396], Stefano Genetti[1][0009-0004-2417-0319], and Giovanni Iacca[1][0000-0001-9723-1830]

Department of Information Engineering and Computer Science
University of Trento, Italy
{s.nejatizadehdaryaei, daniele.notarangelo, jasnoor.singh}@studenti.unitn.it
{erik.nielsen, stefano.genetti, giovanni.iacca}@unitn.it

Abstract. Synthesizing realistic football team formations is a high-dimensional coordination problem that balances rigid physical constraints with fluid tactical objectives. We frame this problem as a 22-dimensional continuous black-box optimization task, where the goal is to optimize 11-player (x, y) coordinates tailored to specific phases of play: attacking possession, defensive possession, and defensive. We introduce the first application of Covariance Matrix Adaptation Evolution Strategies (CMA-ES) to phase-conditioned tactical pattern synthesis, leveraging its covariance adaptation to capture the intricate spatial dependencies between players. Benchmarked against a Differential Evolution (DE) baseline and validated using empirical tracking data, our results show that CMA-ES (in a static scenario where the opponent configuration is fixed) consistently converges to professional-grade tactical structures. Furthermore, while introducing reactive opponent behavior creates a challenging non-stationary landscape, the CMA-ES approach significantly outperforms the DE baseline in generating coherent, coordinated team formations. This work establishes a new intersection between evolutionary search and automated sports analytics, providing a robust framework for objective tactical simulation.

Keywords: Pattern Recognition · Evolutionary Computation · Spatiotemporal Pattern Analysis · Evolution Strategies · Sports Analytics

1 Introduction

Recognizing and generating structured spatiotemporal patterns from high-dimensional data is a central challenge in modern pattern recognition. Team sports provide a particularly demanding setting, where collective behaviors emerge from

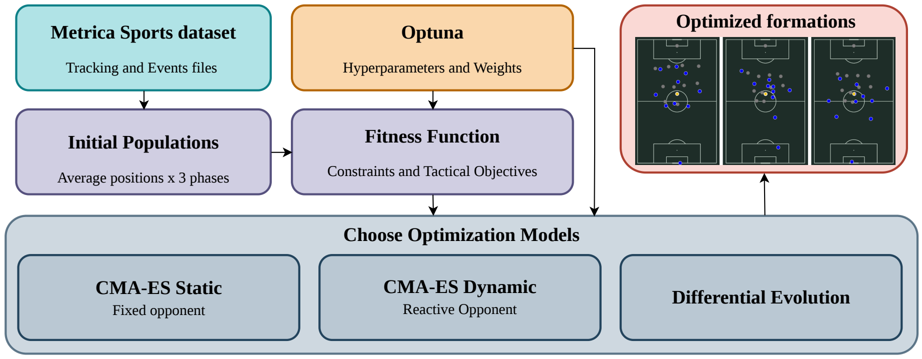

Fig. 1: Overview of the proposed framework.

strongly coupled agent interactions that continuously adapt to dynamic environmental conditions. In football, tactical formations can be interpreted as evolving spatial patterns whose structure depends on contextual factors such as possession, opponent pressure, and phase of play.

In this context, the increasing availability of high-fidelity spatiotemporal tracking data has opened new opportunities for bio-inspired pattern recognition systems capable of modeling, adapting, and synthesizing collective tactical behaviors, enabling, for instance, quantitative models for pitch control [10, 15] and learned tactical assistants [28]. However, turning high-level tactical principles (e.g., compact defending, spacing in possession, offside awareness) into an actionable team formation remains difficult: the space of feasible team spatial patterns is continuous, high-dimensional, and dominated by inter-player dependencies.

In this work, we cast the synthesis of phase-specific team formations as a continuous black-box optimization problem over the 11 on-pitch player locations. The objective combines (i) hard physical constraints to avoid implausible configurations (pitch boundaries, goalkeeper region, minimum distances) and (ii) phase-dependent tactical terms that reward desirable behaviors (coverage and passing support in possession; compactness, marking, and pressure out of possession). We optimize this objective with Covariance Matrix Adaptation Evolution Strategies (CMA-ES) [13] and compare against Differential Evolution (DE) [26] as a simpler baseline. An overview of the overall framework is shown in Figure 1.

From a pattern-recognition perspective, football formations represent structured multi-agent spatial patterns whose organization changes over time in response to contextual game dynamics. Unlike conventional supervised recognition tasks, these patterns emerge through decentralized interactions between agents and therefore require adaptive, population-based search strategies capable of capturing coordination and self-organization. This motivates the use of evolutionary algorithms, such as CMA-ES or DE, as bio-inspired optimizers able to learn correlated spatial structures directly from tracking-derived tactical distributions.

This paper makes four main contributions: (1) a compact, phase-conditioned formation objective grounded in standard tactical heuristics and tracking data; (2) a data-driven initialization scheme that seeds the search from empirical phase averages; (3) an empirical comparison of static vs. simple dynamic-opponent settings, illustrating when non-stationarity can harm convergence; and (4) a bio-inspired interpretation of football formations as emergent collective patterns that can be optimized through evolutionary search.

The remainder of the paper is organized as follows. First, Section 2 summarizes the related works and defines the problem formulation. Section 3 describes the main aspects of the methodology. Section 4 presents the experimental setup. Section 5 reports the empirical results, while Section 6 discusses implications and limitations. Section 7 outlines the conclusions and directions for future work.

2 Related Work

Our work falls into three fields of prior research: spatial analysis of football tracking data, learning-based models of tactical and multi-agent behavior, and evolutionary black-box optimization. We review each in turn and position our contribution with respect to them.

2.1 Spatial value models from tracking data.

The availability of high-fidelity tracking data has driven substantial work on quantifying spatial value in football. Pitch-control models estimate, for each location on the pitch, which team is most likely to control the ball, using geometric or physics-based formulations [10, 24, 15]. Fernández and Bornn [10] introduced space-creation and pitch-value layers to reason about off-ball movement, later extended into a fine-grained expected-possession-value framework [11], while Spearman [24] connected pitch control to off-ball scoring opportunities. These models inform several of our tactical terms (e.g., coverage and passing-lane quality), but they are largely *descriptive*: they evaluate a given configuration rather than synthesize one.

2.2 Formations, roles, and defensive positioning.

A complementary line of work represents and detects team formation directly from tracking data. Bialkowski et al. [3] proposed a role-based representation that assigns each player a latent role per frame via entropy-minimizing assignment, enabling large-scale formation detection. Closer to the defensive phase, recent work evaluates positioning through counterfactual reasoning [27] and off-ball defensive role modeling [12]. Our phase-conditioned initialization is related in spirit, aggregating role-consistent empirical positions per phase to seed the optimizer; however, we then *generate* new formations under explicit tactical objectives rather than only clustering or evaluating observed ones.

2.3 Learning-based and multi-agent tactical models.

Beyond descriptive analysis, machine learning has been used to model and recommend tactical behavior. Le et al. [21, 22] learn league-average defensive trajectories and coordinated multi-agent policies via (deep) imitation learning, producing counterfactual movements. Decroos et al. [9] value individual on-ball actions in context, and TacticAI [28] applies geometric deep learning to corner kicks to predict and generate alternative player setups. Simulation environments such as Google Research Football [20] further enable reinforcement learning of football behavior. These approaches learn implicit behavioral models and typically require substantial trajectory data; in contrast, we adopt an explicit, inspectable objective with a black-box optimizer, making tactical assumptions transparent and easy to modify, at the cost of relying on hand-specified terms.

2.4 Evolutionary black-box optimization.

We cast formation synthesis as continuous black-box optimization and tackle it through evolutionary search. CMA-ES, introduced in [14] and summarized in a widely used tutorial [13], is a state-of-the-art evolutionary algorithm for non-convex, ill-conditioned problems that adapts both a global step size and a full covariance matrix, which suits the strongly coupled coordinates of a team formation. Among others, it has been extensively used in memetic approaches or hybrid metaheuristics with restarts [4, 23, 5] or for constrained optimization [8]. DE [26, 7], on the other hand, is a simple yet effective population-based algorithm that perturbs candidates with scaled difference vectors according to various strategies [29]; as such, it has been widely used as a component of adaptive metaheuristics [6], e.g., with ensembles of hyperparameters [18, 16], or in conjunction with more complex algorithmic schemes [30, 19, 17], with various applications such as in neuroevolution [31].

Of note, sport itself has also inspired metaheuristic design [2]. To our knowledge, however, combining evolutionary search with phase-conditioned tactical objectives for formation *synthesis* from tracking data has not been previously explored.

3 Methods

We use the public Metrica Sports tracking and event datasets [25], containing (i) event annotations (including possession changes) and (ii) (x, y) trajectories for the ball and 22 players. Events are used to segment the match into the three phases of interest based on possession and ball location, while tracking data provide the empirical player distributions used both for evaluation and for initialization.

3.1 Problem Formulation

We optimize the spatial configuration of a given football team for a given phase of play. A candidate formation is a real-valued vector:

$$\mathbf{x} = (x_1, y_1, \ldots, x_{11}, y_{11}) \in \mathbb{R}^{22}, \tag{1}$$

where (x_i, y_i) are the pitch coordinates of player i (goalkeeper included). We consider three phases: *attacking possession* (team in possession in the opponent's half), *defensive possession* (team in possession in its own half), and *defending* (team out of possession and mainly in its own half). Each phase induces different tactical weights and, in our implementation, a different data-driven initialization.

3.2 Phase-Conditioned Initialization

Purely random initial formations almost always violate constraints or produce tactically meaningless formations. We therefore seed each phase with empirical phase averages, similarly in spirit to data-driven initialization used in counterfactual "ghosting" models [21]. Concretely, for each phase, we aggregate the tracking frames assigned to that phase and compute the mean position of each player role. We then sample an initial population by adding small perturbations around these means. This provides (i) a realistic starting point and (ii) a role-consistent reference scale for the optimizer.

3.3 Objective Function

We minimize a single scalar objective that combines phase-conditioned tactical terms with validity penalties:

$$F(\mathbf{x}) = \sum_{k=1}^{K} \lambda_k \, T_k(\mathbf{x}; \text{phase}) + \sum_{j=1}^{J} \mu_j \, P_j(\mathbf{x}), \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^{22}$ is the candidate formation (the stacked (x, y) coordinates of the 11 players), and phase $\in$ {attacking possession, defensive possession, defending} selects the active tactical configuration. The first sum aggregates the K tactical terms T_k, each rewarding a desirable tactical behavior (e.g., coverage, marking, pressure) and weighted by a non-negative, phase-dependent coefficient λ_k that is calibrated per phase (see Section 4.2). The second sum aggregates the J penalty terms P_j, each measuring a violation of a physical or structural constraint (e.g., pitch boundaries, goalkeeper region, minimum inter-player distance) and weighted by a fixed coefficient μ_j. Lower values of F therefore correspond to formations that are simultaneously tactically coherent and physically valid.

Tactical Terms (Phase-Conditioned) The tactical component models the desired team behavior and changes by phase:

1. Field coverage (offensive phase): maximizes the occupation of central and advanced areas of the pitch while penalizing redundancy in less relevant regions, aligning with modern pitch control principles [10].
2. Offside avoidance (offensive possession): penalizes players positioned beyond the offside line, dynamically calculated based on opponent positioning.
3. Marking and compactness (defensive phase): ensures optimal distances between defenders and opponents [12], while minimizing player dispersion relative to the team centroid to favor a compact defensive block.
4. Defensive line height: dynamically regulates the height of the defensive line relative to the ball, penalizing excessively deep lines or lines positioned ahead of the ball.
5. Preventive marking: penalizes unmarked opponents located in the defending team's half during offensive phases to mitigate counterattack risks.
6. Ball pressure and passing lanes: minimizes the distance between the closest defender and the ball to model pressing [28], while simultaneously evaluating and maximizing the availability and quality of passing options for the attacking team [15].

Penalty Terms (Validity) The penalty component enforces structural validity and prevents physically implausible formations:

1. Boundary constraint: penalizes solutions where one or more players are located outside the dimensions of the playing field.
2. Goalkeeper constraint: restricts the goalkeeper to a predefined penalty area, with penalties scaling proportionally to the distance outside this region.
3. Collision and proximity constraint: penalizes pairs of players whose distance falls below a minimum physical threshold, preventing spatial overlaps.
4. Transition smoothness: penalizes excessive positional shifts between configurations belonging to consecutive phases of play, promoting tactical coherence.

3.4 Optimizers

We optimize Eq. (2) with CMA-ES [13] and compare against DE [26] as a simpler baseline. CMA-ES maintains a multivariate Gaussian search distribution and adapts both the global step size and the covariance matrix, which is particularly suitable here because player coordinates are strongly coupled (e.g., line height depends on multiple defenders). The covariance adaptation mechanism of CMA-ES naturally captures emergent dependencies between player roles, enabling the optimizer to discover self-organized collective structures rather than independent player movements. DE, in contrast, relies on difference-vector mutations and crossover, which can struggle to preserve such coordinated structures.

Static vs. Dynamic CMA-ES Configurations. We evaluate two settings for CMA-ES. In the *static* setting, the opponent configuration is held fixed (taken from data-derived phase averages). In the *dynamic* setting, the opponent changes deterministically at each generation via simple heuristic rules inspired by marking/pressing. While this is not full co-evolution, it makes the objective non-stationary and provides a first stress-test of the optimizer stability in dynamic scenarios and possible real-time usages. This behavior highlights the broader challenge of adaptive pattern recognition in non-stationary environments, where the target spatial distribution evolves continuously during optimization.

4 Experimental Setup

4.1 Dataset and Phase Definition.

From the Metrica Sports data, given a match, we extract the three above-mentioned phase-specific scenarios (*attacking possession*, *defensive possession*, and *defending*), using event-based possession changes together with ball location. For each phase, we build (i) a reference formation from empirical phase averages and (ii) a phase-conditioned initialization distribution around that reference.

4.2 Hyperparameter and Weight Tuning with Optuna

We use the hyperparameter optimization framework Optuna [1] in two distinct and independent roles, both employing its default Tree-structured Parzen Estimator (TPE) sampler with the study direction set to minimization.

Stage 1: Tactical weight calibration. For each phase, we calibrate the tactical weights λ_k in Eq. (3) so that the resulting formation reproduces a reference target formation. Each trial runs a full static CMA-ES optimization with the suggested weights and is scored by the mean squared positional error between the generated formation and the reference, $\frac{1}{11}\sum_i \|\mathbf{x}_i - \mathbf{x}_i^{ref}\|^2$.

Stage 2: Algorithmic hyperparameter tuning. Independently, we tune the optimizer hyperparameters by minimizing the final fitness $F(\mathbf{x})$ over 30 trials. For CMA-ES, we search the initial step size $\sigma_{init} \in [0.02, 0.2]$ and the population size in $[10, 40]$; for DE, we search the population size in $[10, 40]$ and the crossover rate in $[0.5, 0.9]$. Trials whose solutions violate the physical penalty terms beyond a fixed threshold are rejected and assigned a high constant cost, so that hyperparameters are never selected on the basis of physically invalid formations. The hyperparameter configurations finally adopted are those reported in Table 1.

4.3 Evaluation Protocol.

Let $\mathbf{x}$ be the optimized player positions and $\mathbf{x}_i^{ref}$ the empirical reference coordinates for a given phase. We quantitatively assess performance using three metrics (all to be minimized).

Table 1: Hyperparameter configuration.

Parameter	CMA-ES	DE
Population Size	20	20
Max. Iterations	160	160
Init. Step-Size	0.15	–
Mutation Factor	–	(0.5, 1.0)
Crossover Rate	–	0.7

1. **Total Fitness**: it is the weighted sum of the above-mentioned $K = 6$ tactical objectives T_k and $J = 4$ penalty terms P_j for constraint violations. Expanding Eq. (2) we have:

$$F(\mathbf{x}) = \sum_{k=1}^{K=6} \lambda_k T_k(\mathbf{x}; \text{phase}) + \sum_{j=1}^{J=4} \lambda_j P_j(\mathbf{x}). \tag{3}$$

   All λ weights are non-negative and, in the case of the tactical objectives, they are tuned per phase. It should be noted that here we focus on analyzing the qualitative impact of these terms, rather than claiming a definitive optimal parameter configuration yielding the "best" possible tactical model.
2. **RMSE**: it quantifies the average positional deviation from the target:

$$RMSE = \sqrt{\frac{1}{11} \sum_{i=1}^{11} \|\mathbf{x}_i - \mathbf{x}_i^{ref}\|^2}. \tag{4}$$

3. **Centroid Shift**: it measures the displacement of the team's geometric center $\bar{\mathbf{x}} = \frac{1}{11} \sum \mathbf{x}_i$ to evaluate global structural stability relative to the geometric center of the reference $\bar{\mathbf{x}}^{ref}$:

$$\Delta C = \|\bar{\mathbf{x}} - \bar{\mathbf{x}}^{ref}\|. \tag{5}$$

4.4 Computational Setup.

All experiments were conducted on a Linux workstation equipped with an Intel Core i9-7940X CPU @ 3.10 GHz and 64 GB of RAM. We make our code publicly available on GitHub[1].

5 Experimental Results

We report qualitative formation examples and quantitative evaluation metrics for the compared methods under the evaluation protocol described in Section 4. Specifically, we assess performance across Total Fitness, spatial error (RMSE), and structural consistency (Centroid Shift), as summarized in Figure 5.

[1] https://github.com/samitos022/bio-ball

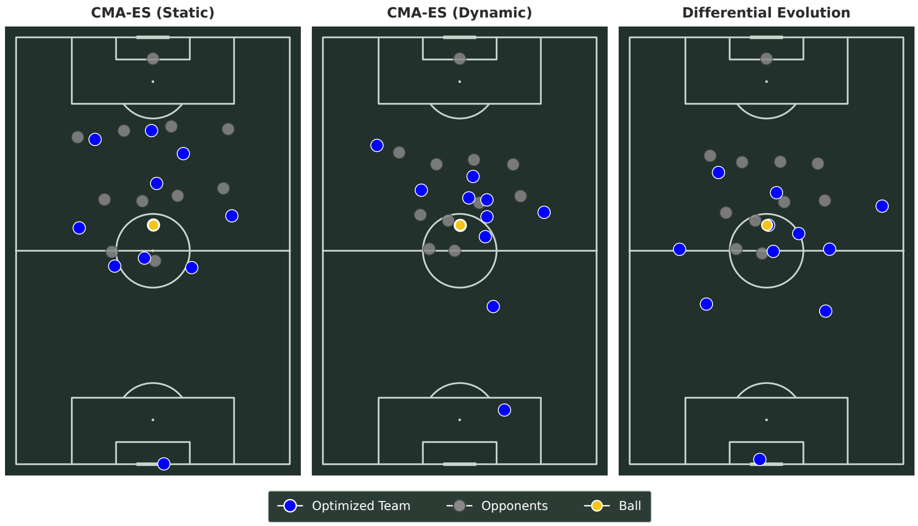

Fig. 2: Optimized formation obtained in the attacking possession phase.

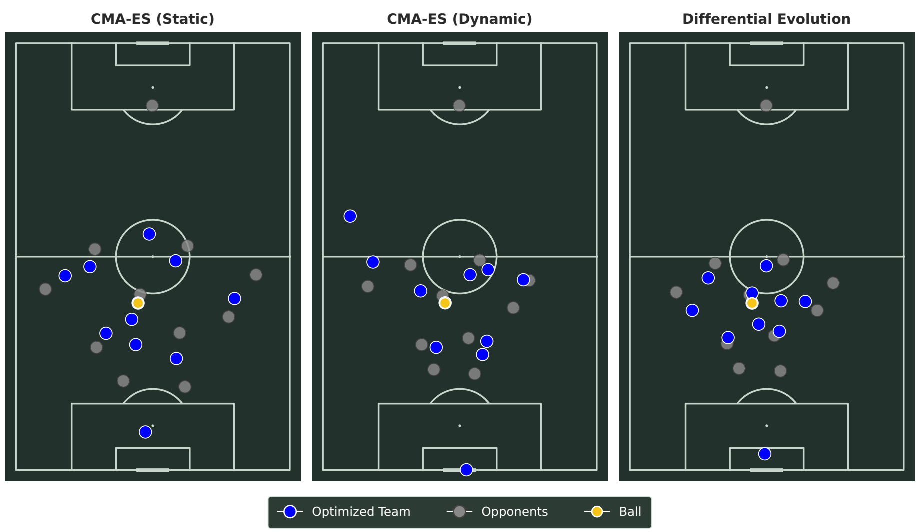

Fig. 3: Optimized formation obtained in the defensive phase.

Static CMA-ES. Across runs, the static CMA-ES variant is the most consistent and efficient approach. As shown in Figure 5, it achieves highly competitive fitness scores while maintaining the lowest centroid shift in possession phases, indicating strong team structure stability. Qualitatively, static CMA-ES produces tactically plausible formations with coordinated lines and role separation. This

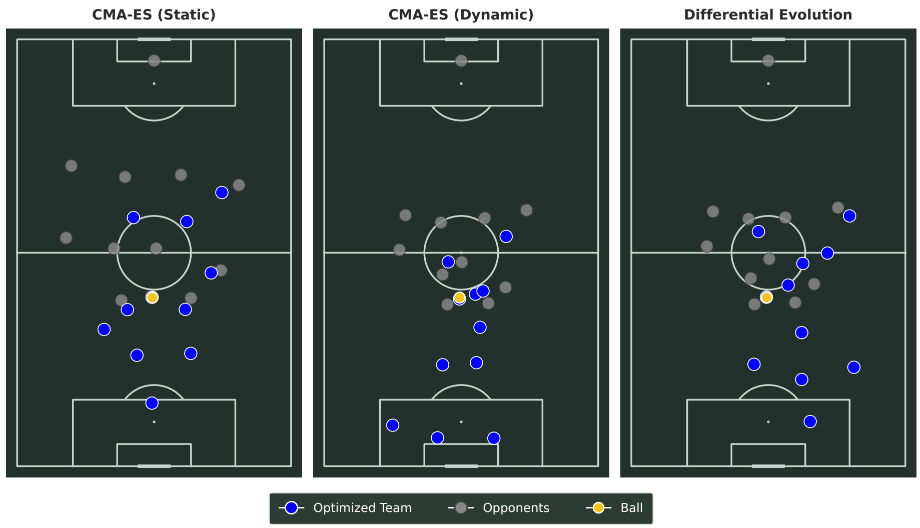

Fig. 4: Optimized formation obtained in the defensive possession phase.

matches the intuition that covariance adaptation captures correlations between players (e.g., shifting the back line and midfield line together). An example of the optimized formation obtained using static CMA-ES in the attacking possession phase is shown in Figure 2. The resulting configuration presents a coherent team structure, with players occupying tactically meaningful positions relative to both teammates and opponents. Similar structural coherence is observed in the defensive scenarios (Figure 3 and Figure 4).

Dynamic CMA-ES. The dynamic-opponent variant degrades performance. The optimized formations appear unrealistic, with players often converging toward locally optimal positions rather than contributing to a coherent collective structure. As reflected in the evaluation metrics (Figure 5), dynamic CMA-ES exhibits higher centroid shifts and generally worse fitness compared to the static variant. This instability can be attributed to the non-stationary fitness landscape introduced by the dynamic opponent behavior, which changes at every generation and prevents the algorithm from learning a stable search distribution.

Differential Evolution. The DE baseline generally produces less structured solutions, but it is able to achieve comparable RMSE values (Figure 5). Qualitatively, players tend to be more scattered, and line coordination is weaker, suggesting that standard DE operators do not preserve the coupled movement required by many tactical terms.

Why the methods differ. The three methods differ in how well they preserve the coupling between players. Static CMA-ES is the most stable because the

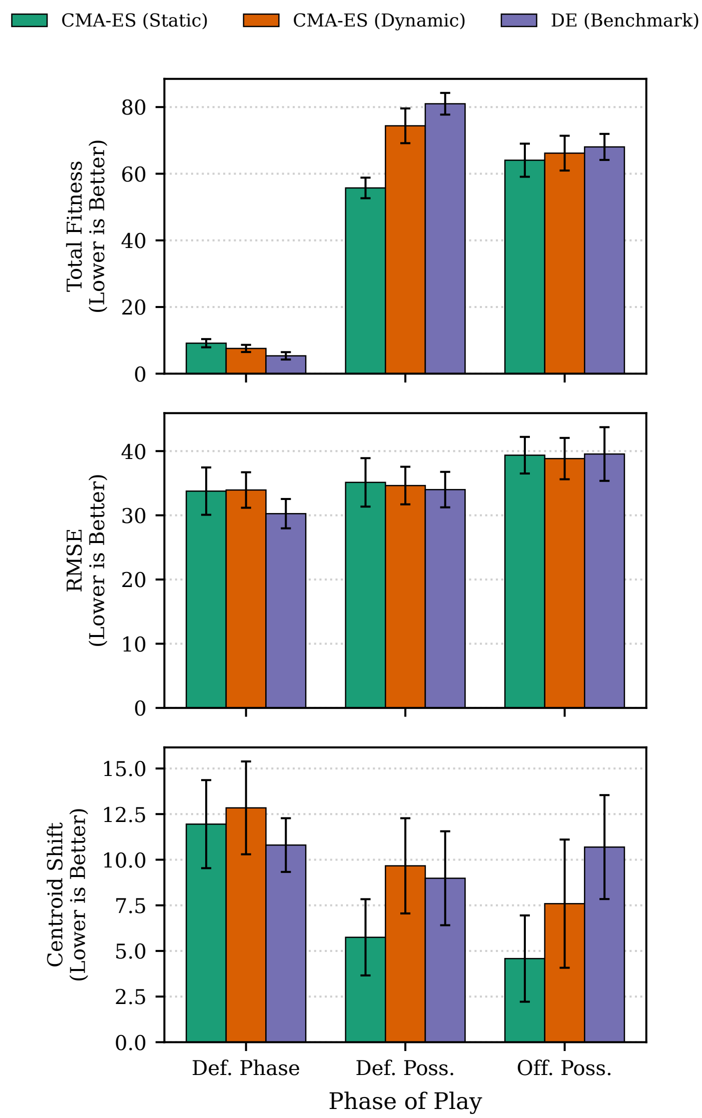

Fig. 5: Performance comparison of Static CMA-ES, Dynamic CMA-ES, and DE across three tactical phases. Metrics include Root Mean Square Error (RMSE), Centroid Shift, and Total Fitness, with lower values indicating better performance. Static CMA-ES consistently achieves superior structural stability (lower Centroid Shift) and competitive fitness in possession phases compared to the DE baseline. Error bars indicate standard deviation over 30 independent runs.

opponent configuration is fixed, so the fitness landscape is stationary and the covariance matrix can converge to a consistent set of inter-player correlations (e.g., back line and midfield line moving together). In contrast, dynamic CMA-ES degrades because the opponent reacts every generation, making the landscape non-stationary: the search distribution that was optimal at one generation is mismatched at the next, so the covariance estimate never stabilizes and the algorithm settles for locally good but globally incoherent positions. Finally, DE can reach a comparable RMSE because RMSE only measures per-player distance to the reference and is largely insensitive to whether the team formation is internally consistent; DE's difference-vector mutation moves players toward the reference individually but does not preserve the coordinated structure that the centroid-shift metric and the qualitative formations reveal.

6 Discussion and Limitations

Our framework automates the synthesis of phase-conditioned formations from tracking data, offering a decision-support tool for tactical "what-if" analysis without requiring hand-designed templates. However, several limitations remain.

Practical use. The model acts as an analytical aid rather than a replacement for coaching expertise, as it currently omits individual player attributes, fatigue, and broader strategic intent.

Objective sensitivity. The composite fitness function is sensitive to term weights (e.g., compactness vs. coverage). Although tuned via Optuna to approximate empirical formations, systematic ablations and cross-match validation are needed to ensure robust generalization.

Non-stationarity. The dynamic-opponent setting creates a shifting fitness landscape that destabilizes CMA-ES convergence, yielding noisier solutions. Future work should address opponent reactivity through explicit co-evolution or robust averaged evaluations.

Scalability. Despite a low-dimensional search space (22 variables), complex tactical evaluations (e.g., marking, passing-lane scoring) are computationally expensive. Scaling to larger datasets will necessitate careful profiling, surrogate models, or multi-fidelity optimization strategies.

7 Conclusions and Future Work

We presented an evolutionary framework that synthesizes phase-specific football formations by casting the 11-player formation as a 22-dimensional continuous optimization problem. Our results demonstrate that static CMA-ES reliably converges to coordinated, tactically plausible formations, outperforming DE by leveraging covariance adaptation to capture inter-player dependencies. However, introducing dynamic opponent reactions creates a non-stationary fitness landscape that can destabilize convergence, highlighting a trade-off between realism and optimization stability. More broadly, this work frames football analytics as a structured spatiotemporal pattern recognition problem in which collective tactical behaviors emerge through adaptive multi-agent interactions. We believe this provides a promising bridge between evolutionary computation and adaptive pattern recognition in complex collective systems.

Future work will transition from static snapshots to trajectory-level optimization and incorporate player-specific attributes (e.g., speed, roles) to enhance realism. We also aim to validate the generated formations using downstream metrics like pitch control and expected threat. Finally, we plan to replace heuristic opponent models with co-evolutionary or game-theoretic formulations to faithfully capture the interactive and adaptive nature of elite football.

References

1. Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 2623–2631 (2019)
2. Ben Ammar, R., Gharbi, A., Babai, M.Z.: Soccer Match Algorithm for Global Optimization: A Contender Metaheuristic. IEEE Access **12**, 93945 (2024)
3. Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Sridharan, S., Matthews, I.: Large-scale analysis of soccer matches using spatiotemporal tracking data. In: IEEE International Conference on Data Mining. pp. 725–730. IEEE (2014)
4. Caraffini, F., Iacca, G., Neri, F., Picinali, L., Mininno, E.: A CMA-ES super-fit scheme for the re-sampled inheritance search. In: IEEE Congress on Evolutionary Computation. pp. 1123–1130. IEEE (2013)
5. Caraffini, F., Iacca, G., Yaman, A.: Improving (1+1) Covariance Matrix Adaptation Evolution Strategy: a simple yet efficient approach. In: International Global Optimization Workshop (2019)
6. Caraffini, F., Neri, F., Cheng, J., Zhang, G., Picinali, L., Iacca, G., Mininno, E.: Super-fit multicriteria adaptive differential evolution. In: IEEE Congress on Evolutionary Computation. pp. 1678–1685. IEEE (2013)
7. Das, S., Suganthan, P.N.: Differential evolution: A survey of the state-of-the-art. IEEE Transactions on Evolutionary Computation **15**(1), 4–31 (2010)
8. De Melo, V.V., Iacca, G.: A modified covariance matrix adaptation evolution strategy with adaptive penalty function and restart for constrained optimization. Expert Systems with Applications **41**(16), 7077–7094 (2014)
9. Decroos, T., Bransen, L., Van Haaren, J., Davis, J.: Actions speak louder than goals: Valuing player actions in soccer. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 1851–1861 (2019)
10. Fernandez, J., Bornn, L.: Wide Open Spaces: A statistical technique for measuring space creation in professional soccer. In: Sloan sports analytics conference. vol. 2018 (2018)
11. Fernández, J., Bornn, L., Cervone, D.: A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. Machine Learning **110**(6), 1389–1427 (2021)

12. Groom, S., Wang, S., Belo, F., Rice, A., Anderson, L.: A Machine Learning Framework for Off Ball Defensive Role and Performance Evaluation in Football (2026), arXiv preprint arXiv:2601.00748
13. Hansen, N.: The CMA evolution strategy: A tutorial (2016), arXiv preprint arXiv:1604.00772
14. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. Evolutionary computation **9**(2), 159–195 (2001)
15. Higgins, L., Galla, T., Prestidge, B., Wyatt, T.: Measuring the pitch control of professional football players using spatiotemporal tracking data. Journal of Physics: Complexity **4**(2), 025008 (2023)
16. Iacca, G., Caraffini, F., Neri, F.: Continuous Parameter Pools in Ensemble Self-Adaptive Differential Evolution. In: IEEE Symposium Series on Computational Intelligence. pp. 1529–1536 (2015)
17. Iacca, G., de Melo, V.V.: Cluster-centroid-based mutation strategies for Differential Evolution. Soft Computing **26**(4), 1889–1921 (2022)
18. Iacca, G., Neri, F., Caraffini, F., Suganthan, P.N.: A differential evolution framework with ensemble of parameters and strategies and pool of local search algorithms. In: European Conference on the Applications of Evolutionary Computation. pp. 615–626. Springer Berlin Heidelberg (2014)
19. Khalfi, S., Draa, A., Iacca, G.: A compact compound Sinusoidal Differential Evolution Algorithm for solving optimisation problems in memory-constrained environments. Expert Systems with Applications p. 115705 (2021)
20. Kurach, K., Raichuk, A., Stańczyk, P., Zając, M., Bachem, O., Espeholt, L., Riquelme, C., Vincent, D., Michalski, M., Bousquet, O., et al.: Google research football: A novel reinforcement learning environment. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 4501–4510 (2020)
21. Le, H.M., Carr, P., Yue, Y., Lucey, P.: Data-driven ghosting using deep imitation learning (2017)
22. Le, H.M., Yue, Y., Carr, P., Lucey, P.: Coordinated multi-agent imitation learning. In: International Conference on Machine Learning. pp. 1995–2003. PMLR (2017)
23. de Melo, V.V., Iacca, G.: A CMA-ES-based 2-stage memetic framework for solving constrained optimization problems. In: IEEE Symposium on Foundations of Computational Intelligence. pp. 143–150. IEEE (2014)
24. Spearman, W.: Beyond expected goals. In: Proceedings of the 12th MIT sloan sports analytics conference. pp. 1–17 (2018)
25. Sports, M.: Metrica Sports Sample Data. https://github.com/metrica-sports/sample-data (2020), accessed: 2026-03-23
26. Storn, R., Price, K.: Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. Journal of global optimization **11**(4), 341–359 (1997)
27. Umemoto, R., Fujii, K.: Evaluation of Team Defense Positioning by Computing Counterfactuals using StatsBomb 360 data. In: StatsBomb Conference (2023)
28. Wang, Zhe and Veličković, Petar and Hennes, Daniel and Tomašev, Nenad and Prince, Laurel and Kaisers, Michael and Bachrach, Yoram and Elie, Romuald and Wenliang, Li Kevin and Piccinini, Federico and others: TacticAI: an AI assistant for football tactics. Nature communications **15**(1), 1906 (2024)
29. Yaman, A., Iacca, G., Caraffini, F.: A comparison of three Differential Evolution strategies in terms of early convergence with different population sizes. In: International Global Optimization Workshop (2019)

30. Yaman, A., Iacca, G., Coler, M., Fletcher, G., Pechenizkiy, M.: Multi-strategy differential evolution. In: European Conference on the Applications of Evolutionary Computation. pp. 617–633. Springer International Publishing (2018)
31. Yaman, A., Mocanu, D.C., Iacca, G., Fletcher, G., Pechenizkiy, M.: Limited evaluation cooperative co-evolutionary differential evolution for large-scale neuroevolution. In: Genetic and Evolutionary Computation Conference. pp. 569–576. ACM (2018)