AN AUDITABLE MOBILE DECISION-SUPPORT SYSTEM FOR EQUITABLE PLAYING TIME AND TEAM COHESION IN SCHOOL SOCCER

Ying Yee Yu[1] and Kyler Harris[2]

[1]Basis Independent Brooklyn, 556 Columbus St, NY 11231
[2]University of California, Santa Cruz, 1156 High St, Santa Cruz, CA 95064

ABSTRACT

Middle- and high-school soccer coaches decide who plays, for how long, and at what cost to a student's rest and schoolwork, yet they make those decisions from memory and a paper roster while validated monitoring tools remain confined to professional academies. This paper presents FairPlay Lab, a cross-platform mobile application that turns the records a volunteer coach can realistically keep — training attendance, a seven-day perceived-exertion average, match minutes, and tallied pitch events — into three transparent indicators: a squad fairness index, a hero-ball versus team-ball cohesion split, and a per-athlete academic crash-risk band. Every recommendation is deliberately rule-based rather than learned, and each substitution suggestion opens into an itemized explanation card naming the factors that produced it. An analysis of the shipped decision surfaces shows the rotation engine responds most strongly to perceived exertion, confines all recommendations to a 24–42 minute window, and contains a discontinuity at the high-load threshold that inverts the intended penalty.

KEYWORDS

youth soccer, playing-time equity, coaching decision support, explainable analytics, rating of perceived exertion, training load monitoring, passing networks, team cohesion, athlete wellbeing, mobile application, offline-first architecture, rule-based scoring

1. INTRODUCTION

Organized youth sport loses participants at a rate that dwarfs any other developmental activity, and the reasons are social rather than physical. A systematic review of forty-three studies of dropout among children and adolescents found that intrapersonal and interpersonal constraints — lack of enjoyment, eroded perceptions of competence, and social pressure — predict withdrawal far more strongly than structural factors such as cost or facility access [1]. Playing time sits at the intersection of all three: an athlete who watches from the bench receives fewer opportunities to demonstrate competence, reads that absence as a judgment, and leaves. The countervailing pressure is equally well documented. Overuse injuries account for between 45.9% and 54% of all youth sports injuries, and early intensive training coupled with competitive pressure raises the risk of both injury and burnout [2]. Sleep is the quiet casualty: adolescent athletes are a recognized at-risk group for disordered sleep, with measurable consequences for the cognitive tasks that schoolwork demands [3]. A school coach therefore optimizes three quantities at once — fairness, load, and academic capacity — using nothing but recall.

That recall is not neutral. Interviews with football coaches show selection decisions shaped by preconception and by pressure to win immediately, with relatively older and more physically

David C. Wyld et al. (Eds): BMLI, CACIT, EDUAIT, InWES – 2026
pp. 107-118, 2026. CS & IT - CSCP 2026 DOI: 10.5121/csit.2026.161809

mature players systematically favored [4]. The instruments that could discipline such judgment already exist and are inexpensive. Session rating of perceived exertion, introduced as a way to quantify internal training load from an athlete's own report [5], correlates strongly with heart-rate-derived measures in soccer specifically [6], and an international consensus statement establishes load monitoring as standard practice [7]. Almost none of this reaches a school team. A case study of practitioners at a category-one academy found that even well-resourced staffs were constrained by data complexity, interpretation time, and reluctance to let derived metrics displace contextual knowledge [8] — constraints that bind a volunteer coach far more tightly.

Meanwhile, sports analytics has moved toward deep learning and computer vision pipelines that presume tracking infrastructure a school does not own [9]. Opacity carries its own cost: people abandon algorithmic advice disproportionately after observing it err [10], and surveys of explainability research identify trust and understanding as the outcomes that determine whether an explanation is used at all [11].

FairPlay Lab responds with a deliberately modest proposition: a mobile application that computes fairness, cohesion, and wellbeing indicators from hand-entered data using transparent weighted arithmetic, and exposes the reasoning behind every substitution recommendation as an itemized card the coach can read and overrule. The design trades predictive sophistication for auditability. Because the inputs are quantities a coach already tracks informally, and because the arithmetic is inspectable rather than learned, a disputed recommendation can be traced to the specific attendance record or exertion average that produced it. This matters for the outcome the system targets — distributing minutes in a way an athlete perceives as fair — since cohesion and performance are meaningfully related in team sport [12], and a recommendation that cannot be explained cannot be defended to the player it disadvantages.

2. CHALLENGES

2.1. Turning sparse sideline records into a defensible minutes recommendation

A school coach records attendance intermittently and exertion scores rarely, so any rotation engine must produce a usable recommendation from partial data without silently inventing the missing parts. Weighting attendance, minute history, and perceived exertion into a single score risks letting one well-populated field dominate two empty ones. A workable approach would gate the recommendation on a minimum set of populated inputs, refusing to advise at all until both attendance and exertion exist, and would surface each contributing factor separately so a coach can see which input is actually driving a number.

2.2. Separating individual brilliance from collective build-up in a sparse event log

Distinguishing a team that constructs chances from one that relies on a single forward requires knowing how a shot was reached, not merely that it occurred. Full passing-network reconstruction needs tracking data no school possesses. A tractable substitute would ask the coach to tag each shot with the length of the passing sequence preceding it, then combine shot concentration, sequence length, and involvement dispersion into a single contrast. The hazard is that such a composite invites over-reading: a percentage split computed from four tallies can appear far more precise than the underlying counts justify.

2.3. Keeping a wellbeing indicator useful without overstating what it can predict

Combining sleep, homework hours, and training load into an “academic risk” label is useful only if the label is honest about its provenance. Thresholds chosen by hand encode assumptions, not evidence, and a three-level band can imply a validated instrument where none exists. Any responsible design would keep the rubric small enough to read in full, expose it to the coach, and describe its output as a prompt for conversation rather than a prediction — while remaining alert to whether the intermediate band is wide enough to carry meaning.

3. SOLUTION

FairPlay Lab is a Flutter application organized into four strictly separated layers, shown in Figure 1. A presentation layer of five tabs holds no analytical logic; a single TeamController owns all mutable squad state; a deterministic analytics layer recomputes every indicator whenever that state changes; and a repository interface abstracts storage so the same controller runs against Firestore or on-device storage without modification.

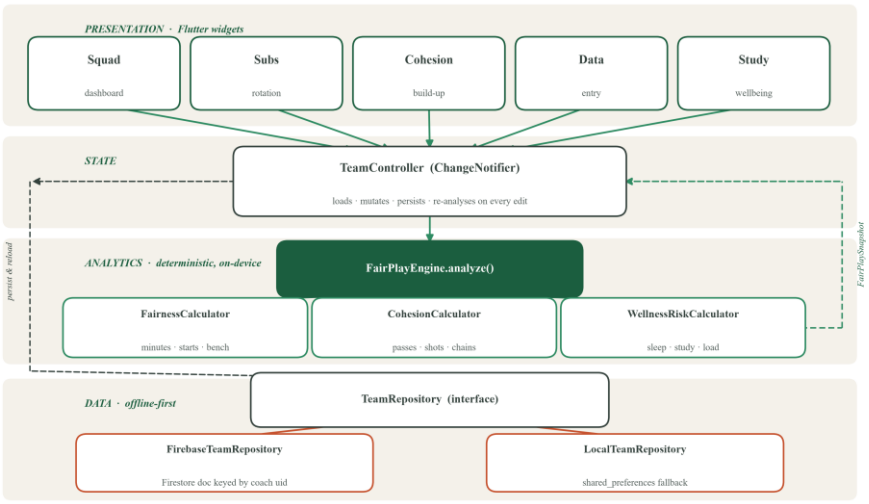

Figure 1. System architecture. User edits travel down through the controller into the analytics layer; a FairPlaySnapshot returns up the right-hand path, while persistence loops down the left.

System overview. The end-to-end flow begins on the Data tab, where the coach enters a squad, attendance, seven-day exertion values, match minutes, and tagged pitch events. Each edit mutates the TeamInput held by the controller, which persists the change and immediately re-invokes FairPlayEngine.analyze(). That method delegates to three independent calculators and assembles their outputs into a single immutable FairPlaySnapshot consumed by every screen, so the dashboard in Figure 2 can never disagree with the rotation or cohesion tabs — all three read the same computed object. The dashboard headline figures are therefore derived quantities, not stored ones.

Figure 2. The Squad tab, showing the four headline indicators recomputed from the sample squad on every edit.

Data entry is deliberately gated rather than free-form. The Data tab tracks six required steps and reports completion explicitly, as shown in Figure 3, because several calculators return empty results until their prerequisites exist — FairnessCalculator refuses to compute any rotation until both attendance and exertion records are present, so an incomplete squad yields no advice rather than misleading advice.

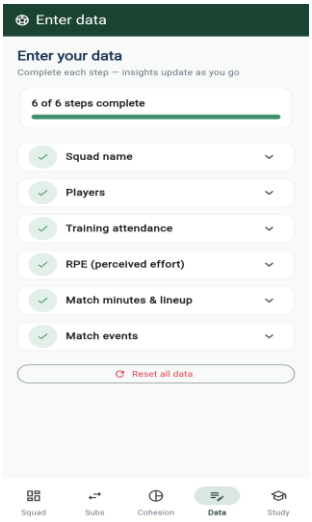

Figure 3. The Data tab gates analysis behind six explicit setup steps.

Persistence is offline-first by necessity, since sideline edits happen on fields with poor signal. Firestore is configured with unlimited local cache and persistence enabled, and if Firebase fails to initialize the application falls back to shared_preferences and runs entirely on-device with no sign-in. Squad data is scoped per coach by naming each Firestore document after the coach's authentication identifier, enforced server-side by the rules in Figure 4 rather than by client code alone.

```
rules_version = '2';

service cloud.firestore {
  match /databases/{database}/documents {
    // A squad document is named after the coach's uid, so a signed-in coach can
    // only read and write their own team. Everything else is denied.
    match /teams/{teamId} {
      allow read, write: if request.auth != null && request.auth.uid == teamId;
    }
  }
}
```

Figure 4. Firestore security rules restricting each squad document to the coach who owns it.

Rotation and fairness engine. The first component converts attendance, minute history, and perceived exertion into a recommended minutes range for every player, together with a status label and an explanation. It relies on no external service; the entire computation is local arithmetic. Figure 5 shows the resulting guidance for the sample squad.

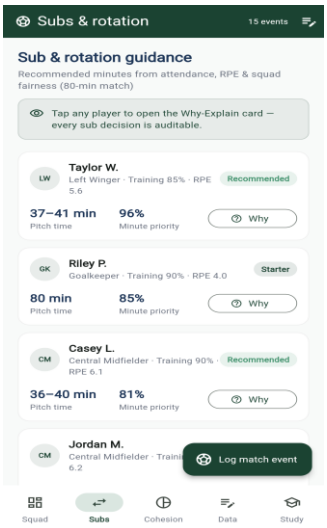

Figure 5. The Subs tab, listing each player's recommended minutes range, minute-priority percentage, and status.

```
TeamFairnessMetrics computeTeamMetrics(TeamInput team) {
  final fieldPlayers = team.players.where((p) => !p.isGoalkeeper).toList();
  final minutesByPlayer = _averageMinutes(fieldPlayers, team.games);

  final minutesEquity = _minutesEquity(minutesByPlayer.values.toList());
  final starterRotationBalance = _starterRotationBalance(fieldPlayers, team.games);
  final benchOpportunityRate = _benchOpportunityRate(fieldPlayers, team.games);

  final fairnessIndex = _clamp01(
    0.45 * minutesEquity +
        0.30 * starterRotationBalance +
        0.25 * benchOpportunityRate,
  );
```

Figure 6. The squad-level fairness index, a fixed-weight blend of three independently computed sub-measures.

The team-level index combines three quantities. _minutesEquity computes the coefficient of variation of average minutes across field players and returns its complement, so a squad in which everyone plays similar minutes approaches 1.0. _starterRotationBalance returns one minus the proportion of matches started by the most frequent starter, penalizing a fixed eleven. _benchOpportunityRate measures what fraction of substitutes received at least ten minutes — the constant that encodes "meaningful" bench time. A parallel per-player score replaces these with a minute-deficit term, attendance, and a load-fit term at weights of 0.45, 0.30, and 0.25. Two properties deserve emphasis because the interface does not state them. The weights are hand-chosen constants, not fitted parameters, so the index expresses a coaching philosophy rather than an empirical finding. Goalkeepers bypass the computation entirely: they receive a hardcoded score of 0.85 and the full match, so the displayed index describes outfield rotation only.

The interface claims on the Subs tab that "every sub decision is auditable," and this claim is met. Each recommendation carries a list of WhyFactor records naming the attendance ratio, the exertion average, the minute differential, and the resulting range, rendered as the card in Figure 7.

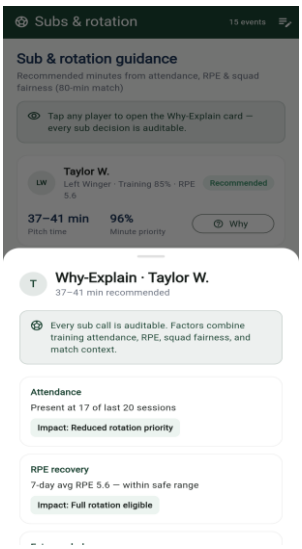

Figure 7. The Why-Explain card, itemizing the factors behind one player's recommended range.

Cohesion analyzer. The second component reads tagged pitch events and estimates whether chances arise from collective build-up or individual action. Each shot is logged with the length of the passing sequence preceding it, which supplies the structural information a passing network would otherwise provide.

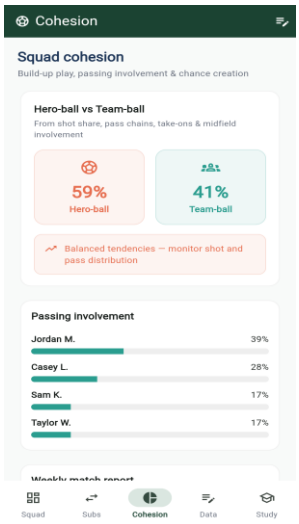

Figure 8. The Cohesion tab, reporting the hero-ball / team-ball split and passing involvement.

```
final heroSignal = _clamp01(
  0.35 * shotConcentration +
      0.25 * shortChainRate +
      0.25 * soloRate +
      0.15 * (1 - involvementBalance),
);
final teamSignal = _clamp01(
  0.40 * involvementBalance +
      0.35 * linkedShotRate +
      0.25 * (1 - shotConcentration),
);

final heroBallPercent = heroSignal + teamSignal == 0
    ? 50.0
    : (heroSignal / (heroSignal + teamSignal)) * 100;
final teamBallPercent = 100 - heroBallPercent;
final cohesionScore = teamBallPercent;
```

Figure 9. Competing hero and team signals, normalized against one another to produce the displayed split.

Five sub-measures feed the two composites: the share of shots taken by the most frequent shooter, the fraction of shots preceded by fewer than three passes, the ratio of take-ons to all events, one minus the Gini coefficient of pass counts, and the fraction of shots following sequences of three or more passes. The two signals are then normalized against each other. This construction has a consequence the interface does not disclose: because the split is a ratio of one signal to their sum, it always totals 100% and can never report that both tendencies are weak — a team logging very few events still produces a confident-looking percentage. The dashboard's "Cohesion score" is also not an independent metric; as the final line shows, it is simply the team-ball percentage under another name.

Academic wellbeing monitor. The third component estimates an academic crash-risk band from weekly training intensity, average sleep, and after-school study hours, entered per athlete on the Study tab shown in Figure 10.

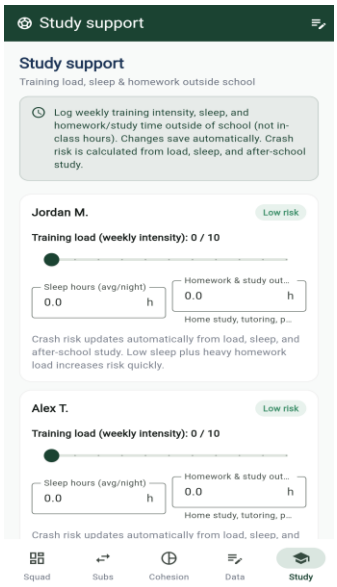

Figure 10. The Study tab, collecting sleep, homework, and training-intensity values per athlete.

```
// Compound strain: poor sleep + heavy study/training stacks quickly.
if (sleepHours > 0 && sleepHours < 7 && studyHours >= 4) {
  score += 2;
}
if (sleepHours > 0 && sleepHours < 6 && studyHours >= 6) {
  score += 2;
}
if (sleepHours > 0 && sleepHours < 7 && trainingLoad >= 6) {
  score += 1.5;
}
if (studyHours >= 5 && trainingLoad >= 6) {
  score += 1;
}

if (score >= 5) return CrashRisk.high;
if (score >= 2.5) return CrashRisk.moderate;
return CrashRisk.low;
```

Figure 11. The interaction terms and cut points of the crash-risk rubric.

The rubric assigns tiered points for short sleep, long homework hours, and heavy training, then adds four interaction terms for combinations judged to compound, before applying two cut points. The design is legible in full, which is its principal virtue. It is equally important to state what it is not: the thresholds are authored judgments, and no validation against actual academic outcomes has been performed. The word “risk” in the interface consequently carries more epistemic weight than the implementation earns, and the indicator is best read as a structured prompt for a conversation with an athlete rather than a prediction about one.

4. EXPERIMENT

Because the scoring layer is deterministic and its parameters are fixed constants rather than fitted values, its behavior is fully characterized by its decision surfaces. The two analyses below evaluate those surfaces directly, reproducing the shipped Dart functions exactly. They describe

what the implemented system will do for any input, and are not claims about outcomes for real squads; no field study has been conducted.

4.1. Input sensitivity and reachable range of the rotation recommendation

Blind spot. The rotation engine presents three inputs as jointly determining a minutes range. If one input dominates, or if the reachable range is narrow, the interface implies a responsiveness the arithmetic does not deliver.

Design. The per-player score and _recommendedMinutes were reproduced from fairness_calculator.dart and evaluated over a grid of minute differentials from −24 to +24, attendance from 0.60 to 1.00, and exertion averages from 0 to 10, at the application's 80-minute match length. Sensitivity was measured one factor at a time from a reference athlete at the squad minute average, 90% attendance, and an exertion average of 6.0. The full grid was then swept to establish the reachable output range.

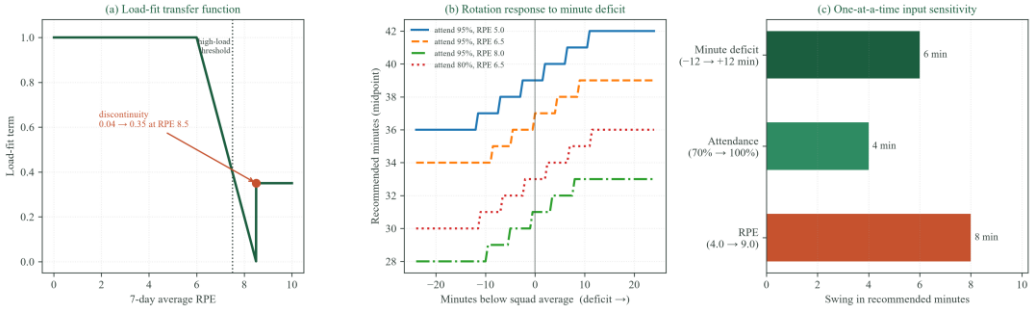

Figure 12. Load-fit transfer function (a), rotation response to minute deficit (b), and one-at-a-time input sensitivity (c).

Analysis. Perceived exertion is the dominant input. Moving it from 4.0 to 9.0 shifts the recommendation from 39 to 31 minutes, a swing of 8 minutes, against 6 minutes for a ±12-minute change in minute differential and 4 minutes for attendance across 70% to 100%. Exertion therefore outweighs the minute-history term the feature is nominally built to correct, even though the equity term carries the largest nominal weight — a consequence of _playerMinutesEquity mapping an average player to 0.5 rather than 1.0, which halves that term's effective contribution. Sweeping the whole grid, every recommendation falls between 24 and 42 minutes: an 18-minute band inside an 80-minute match, with the clamp to matchLength - 10 never reached. Panel (a) exposes a defect. _loadFit decays linearly to 0.004 at an exertion average of 8.49, then jumps to 0.35 at 8.5, because the rpe >= 8.5 branch returns a floor that exceeds the value the preceding linear segment has already fallen to. The most fatigued athletes thus receive a higher load-fit score than moderately fatigued ones, inverting the intended penalty across that boundary.

4.2. Discrimination of the crash-risk rubric across the sleep–study–load space

Blind spot. A three-level risk band is useful only if all three levels are reachable across realistic inputs. If the intermediate level collapses, the indicator becomes an alarm rather than a gradient.

Design. WellnessRiskCalculator.calculate was reproduced exactly and evaluated over sleep from 4.0 to 10.0 hours and study from 0 to 10 hours in 0.25-hour steps, at three training-load settings spanning the scale. Class shares were computed over each grid.

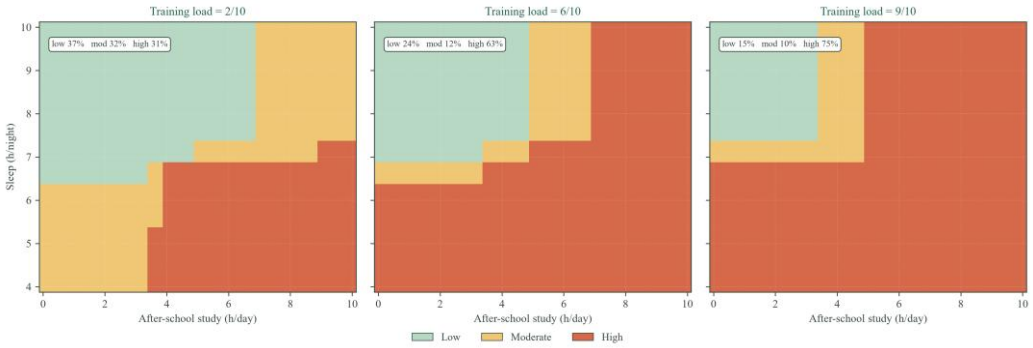

Figure 13. Crash-risk classification across the sleep–study plane at three training-load settings, with class shares.

Analysis. The rubric discriminates well at low training load and progressively less as load rises. At a load of 2 the three bands occupy 37%, 32%, and 31% of the plane, a usable gradient. At a load of 6 the moderate band collapses to 12.5% while high risk expands to 63%, and at a load of 9 moderate holds only 10% against 75% high. The cause is visible in the surface geometry: the four interaction terms each add 1 to 2 points, so once training load contributes its own 1 to 2 points, a single additional condition carries the total past the 5-point high-risk cut without ever resting in the 2.5-to-5 moderate window. For a competitive squad in season — precisely the population the tab targets — the indicator will report high risk for roughly three athletes in four, and an alert that common will be ignored. Widening the moderate window or scaling the interaction terms would restore the gradient; the present cut points were not chosen with the compounded totals in view.

5. Related Work

Training-load monitoring is the most mature thread and the closest prior art. Session rating of perceived exertion has been validated across sports, age groups, and expertise levels [13], and in soccer specifically it tracks heart-rate-derived internal load closely enough for practical use [6]. Building on it, the acute-to-chronic workload ratio has been widely adopted to flag injury risk, though systematic reviews report high methodological variability and caution against treating it as a settled predictor [14], [15]. Subjective self-report measures in fact outperform many objective ones for tracking training response [16]. These instruments quantify load well but say nothing about how minutes are distributed; they answer whether an athlete is overloaded, never whether the squad is being treated equitably. FairPlay Lab adopts the exertion average as an input and adds the distributional question on top of it.

A second thread analyzes team play structurally. Passing-network studies of youth football show that lower dependence on any single player and better-connected passing relations accompany better match outcomes [17], and a systematic review confirms that dense, low-centralization networks characterize successful teams [18]. This work is methodologically strong but presumes event or tracking data captured by professional providers, which places it out of reach of a school team, and it treats cohesion as a structural property of play rather than as something a coach acts on — even though the cohesion–performance relationship is well established in sport psychology [12] and in the wider teams literature [19]. The cohesion analyzer approximates the same contrast from four hand-tagged event types, accepting substantially coarser evidence in exchange for being collectible on a sideline.

A third thread concerns the delivery of analytics to practitioners. Reviews of artificial intelligence in sport document a shift toward deep learning and computer vision pipelines dependent on sensing infrastructure [9], while practitioner studies find adoption limited by interpretation time and by reluctance to let derived metrics displace contextual judgment [8]. Users abandon algorithmic advice sharply after observing errors [10], and explainability research identifies understanding and trust as the decisive outcomes [11]. Cross-platform frameworks now make a single-codebase deployment practical at near-native cost [20]. FairPlay Lab occupies the gap these findings define: transparent arithmetic over hand-entered data, with an explanation attached to every recommendation.

5. CONCLUSIONS

Two analyses of the shipped decision surfaces established what the system does across its whole input domain. The rotation engine responds most strongly to perceived exertion — an 8-minute swing, against 6 minutes for minute history and 4 for attendance — and confines every recommendation to a 24-to-42-minute band within an 80-minute match, narrower than the interface suggests. It also contains a boundary defect at which the most fatigued athletes receive a higher load-fit score than moderately fatigued ones. The crash-risk rubric discriminates across three usable bands at low training load, but its moderate band collapses to roughly a tenth of the input space once load rises, classifying three athletes in four as high risk for exactly the in-season squad it is meant to serve.

These findings define the immediate work. The load-fit floor should be removed so the function decays monotonically; the equity term should map an average player to a neutral rather than a half contribution; and the crash-risk cut points should be rescaled against the compounded totals the interaction terms actually produce. Three limitations bound the present claims more broadly. No field study has been conducted, so nothing here demonstrates that the indicators track outcomes coaches care about. The weights throughout are authored judgments encoding a coaching philosophy, not fitted parameters. And the cohesion split is structurally incapable of reporting that both tendencies are weak, so it will always look confident even on a thin event log.

The contribution is nonetheless a workable one. By restricting itself to records a volunteer coach can actually keep and to arithmetic a coach can actually read, the system makes the distribution of playing time inspectable and therefore contestable — which is the property a bench decision needs when it has to be defended to the athlete it disadvantages.

REFERENCES

[1] Crane, Jeff, and Viviene Temple. “A systematic review of dropout from organized sport among children and youth.” European Physical Education Review 21.1 (2015): 114-131.

[2] DiFiori, John P., Holly J. Benjamin, Joel S. Brenner, Andrew Gregory, Neeru Jayanthi, Greg L. Landry, and Anthony Luke. “Overuse injuries and burnout in youth sports: a position statement from the American Medical Society for Sports Medicine.” British Journal of Sports Medicine 48.4 (2014): 287-288.

[3] Coel, Rachel A., George G. A. Pujalte, Andres I. Applewhite, Tracy Zaslow, George Cooper, Angie N. Ton, and Holly J. Benjamin. “Sleep and the Young Athlete.” Sports Health 15.4 (2023): 537-546.

[4] Hill, Brad, and Popi Sotiriadou. “Coach decision-making and the relative age effect on talent selection in football.” European Sport Management Quarterly 16.3 (2016): 292-315.

[5] Foster, Carl, Jessica A. Florhaug, Jodi Franklin, Lori Gottschall, Lauri A. Hrovatin, Suzanne Parker, Pamela Doleshal, and Christopher Dodge. “A new approach to monitoring exercise training.” Journal of Strength and Conditioning Research 15.1 (2001): 109-115.

[6] Impellizzeri, Franco M., Ermanno Rampinini, Aaron J. Coutts, Aldo Sassi, and Samuele M. Marcora. “Use of RPE-based training load in soccer.” Medicine and Science in Sports and Exercise 36.6 (2004): 1042-1047.

[7] Bourdon, Pitre C., Marco Cardinale, Andrew Murray, Paul Gastin, Michael Kellmann, Matthew C. Varley, Tim J. Gabbett, Aaron J. Coutts, Darren J. Burgess, Warren Gregson, and N. Timothy Cable. “Monitoring athlete training loads: consensus statement.” International Journal of Sports Physiology and Performance 12.s2 (2017): 161-170.

[8] Davidson, Tia-Kate, Steve Barrett, John Toner, and Chris Towlson. “Professional soccer practitioners’ perceptions of using performance analysis technology to monitor technical and tactical player characteristics within an academy environment: A category 1 club case study.” PLoS ONE 19.3 (2024): e0298346.

[9] Ghosh, Indrajeet, Sreenivasan Ramasamy Ramamurthy, Avijoy Chakma, and Nirmalya Roy. “Sports analytics review: Artificial intelligence applications, emerging technologies, and algorithmic perspective.” WIREs Data Mining and Knowledge Discovery 13.5 (2023): e1496.

[10] Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. “Algorithm aversion: People erroneously avoid algorithms after seeing them err.” Journal of Experimental Psychology: General 144.1 (2015): 114-126.

[11] Rong, Yao, Tobias Leemann, Thai-Trang Nguyen, Lisa Fiedler, Peizhu Qian, Vaibhav Unhelkar, Tina Seidel, Gjergji Kasneci, and Enkelejda Kasneci. “Towards human-centered explainable AI: A survey of user studies for model explanations.” IEEE Transactions on Pattern Analysis and Machine Intelligence 46.4 (2024): 2104-2122.

[12] Carron, Albert V., Michelle M. Colman, Jennifer Wheeler, and Diane Stevens. “Cohesion and performance in sport: A meta analysis.” Journal of Sport and Exercise Psychology 24.2 (2002): 168-188.

[13] Haddad, Monoem, Georgios Stylianides, Leo Djaoui, Alexandre Dellal, and Karim Chamari. “Session-RPE method for training load monitoring: Validity, ecological usefulness, and influencing factors.” Frontiers in Neuroscience 11 (2017): 612.

[14] Gabbett, Tim J. “The training-injury prevention paradox: should athletes be training smarter and harder?” British Journal of Sports Medicine 50.5 (2016): 273-280.

[15] Maupin, Danny, Ben Schram, Elisa Canetti, and Robin Orr. “The relationship between acute:chronic workload ratios and injury risk in sports: A systematic review.” Open Access Journal of Sports Medicine 11 (2020): 51-75.

[16] Saw, Anna E., Luana C. Main, and Paul B. Gastin. “Monitoring the athlete training response: subjective self-reported measures trump commonly used objective measures: a systematic review.” British Journal of Sports Medicine 50.5 (2016): 281-291.

[17] Gonçalves, Bruno, Diogo Coutinho, Sara Santos, Carlos Lago-Penas, Sergio Jiménez, and Jaime Sampaio. “Exploring team passing networks and player movement dynamics in youth association football.” PLoS ONE 12.1 (2017): e0171156.

[18] Caicedo-Parada, Sergio, Carlos Lago-Peñas, and Enrique Ortega-Toro. “Passing networks and tactical action in football: A systematic review.” International Journal of Environmental Research and Public Health 17.18 (2020): 6649.

[19] Grossman, Rebecca, Kevin Nolan, Zachary Rosch, David Mazer, and Eduardo Salas. “The team cohesion-performance relationship: A meta-analysis exploring measurement approaches and the changing team landscape.” Organizational Psychology Review 12.2 (2022): 181-238.

[20] Suri, Bhawna, Shweta Taneja, Isha Bhanot, Himanshi Sharma, and Aanchal Raj. “Cross-platform empirical analysis of mobile application development frameworks: Kotlin, React Native and Flutter.” Proceedings of the 4th International Conference on Information Management and Machine Intelligence (2022): 1-6.

©2026 By AIRCC Publishing Corporation. This article is published under the Creative Commons Attribution (CC BY) license.