PERFORMANCE INTELLIGENCE
WHITEPAPER SERIES

PAPER SIX OF SEVEN · ELITE FOOTBALL

The Augmented Touchline

How responsible AI can transform elite football performance, coaching and player development

Paul Forrest
Founder and Chief AI Officer, Athleet.AI
February 2026

About this series

This is the sixth of seven papers examining how artificial intelligence can be used responsibly in performance sport. Each paper stands alone. Together they set out a single argument, which is that the durable advantage in sport comes from better decisions rather than from more data, and that better decisions require systems designed around the person who is accountable for them.

Football is the best-instrumented sport in the world and among the worst at retaining why it decided anything. That gap is the subject of this paper. The series shares one framework, described in Chapter 1, and one management thesis, described in Chapter 5.

Editorial balance. Roughly seven parts in ten of this document is evidence and sector analysis, two parts in ten is original framework, and one part in ten concerns the Athleet.AI platform itself.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence supports it.

PEER-REVIEWED A primary study in a peer-reviewed journal.

REVIEW A systematic review or meta-analysis.

CONSENSUS A consensus statement, governing-body position or regulator guidance.

COMMENTARY Methodological criticism, which carries weight without being primary data.

PRACTICE Practitioner opinion, including the author's own.

ATHLEET.AI The author's own framework, proposed rather than proven.

Where the evidence is thin, the paper says so. Where it points the other way, the paper says that too.

Legal and regulatory notice. Parts of this paper describe legislation, regulatory guidance and governing-body requirements, in particular in Chapter 9. Nothing in this document constitutes legal advice. The provision of legal advice sits outside the terms of any engagement with the author or with Athleet.AI, and the material is presented to support discussion and further review by qualified advisers.

Statements about what a source says are attributed to that source. Where the author draws a practical implication the source does not itself state, the text says so. Legislative timetables in this area have moved more than once, and the position described is the position the author was able to verify in February 2026.

ILLUSTRATIVE MATERIAL

Every club, player, scenario and figure of speech in this paper is invented. Cases are marked **illustrative scenario** wherever they appear, and none describes a real club, player or member of staff. Published research is cited and numbered; invented material is never presented as evidence.

CITATION

Forrest, P. (2026). *The Augmented Touchline: How Responsible AI Can Transform Elite Football Performance, Coaching and Player Development.* Athleet.AI Performance Intelligence Whitepaper Series, Paper Six. Version 1.0, February 2026.

ALSO IN THIS SERIES

- **Paper one.** *The Responsible Performance System.* Governance and trust in sports AI.
- **Paper two.** *The Intelligent Track.* AI and event-specific expertise in elite athletics.
- **Paper three.** *Closing the Coaching Gap.* Responsible AI in grassroots athletics.
- **Paper four.** *Casting the Net Wider.* AI as a force multiplier in football pathways.
- **Paper five.** *The Connected Endurance Athlete.* Integrated preparation in triathlon.
- **Paper seven.** *The Augmented Golfer.* AI across technique, practice and strategy.

Contents

A NOTE FROM THE AUTHOR

The most measured game in the world

Football measures more than any sport in history and forgets faster than most. Every touch is logged, every metre is tracked, every match is filmed from eight angles, and then a manager leaves in November and the reasoning behind a season of decisions leaves with him. The spreadsheets stay. The judgement does not.

I have spent my working life on how organisations turn what they know into what they do, first through business process work in the 1990s and since 2014 through machine learning and language models. I also coach, in a different sport and at a very different level, which is why the thing that strikes me about elite football is not the sophistication of the analysis. It is how little of that analysis reaches the person standing on the grass on a Thursday morning deciding whether a particular player should train.

What prompted this paper was a pattern I kept meeting in performance departments. Enormous capability, genuine expertise, and a recurring inability to answer a simple question after the fact, which is why did we do that. The evidence had been in the room. The argument had been had. Nobody had written down what won it.

This paper argues that football's next advantage is an explainable decision layer, and it argues equally firmly that several things being sold to clubs at present should be declined. Both halves matter. I have graded the evidence throughout, and where the research contradicts something the industry likes to say, I have quoted the research.

TWO PAGES THAT CIRCULATE ON THEIR OWN

Executive summary

Elite football has more data than any sport and less shared understanding of it than its investment implies. Tracking, event feeds, video, medical systems and scouting platforms all improved. The decision they exist to inform, which is what this player should do this week and why, is still made in a meeting, from memory, and is still very rarely recorded in a form that survives a change of manager.

This paper argues that the opportunity is an explainable decision layer that connects tactical, technical, physical, medical and contextual information, states its own uncertainty, and leaves a named human being accountable. It argues against three things the market currently sells, which are single-score readiness indices, individual injury prediction, and models trained on a club's own past selections that then recommend the same selections back with a decimal point attached.

1

Availability is a performance variable, and the evidence is unusually good. Across 11 seasons and 7,792 injuries in the UEFA Champions League injury study, lower injury burden was associated with better league ranking, more points per match and a better European ranking.

2

The injury picture is improving overall and worsening in one place. Across 18 seasons, overall injury incidence fell by around three per cent a year. Over 21 seasons, hamstring injuries roughly doubled as a share of the total, from about 12 to 24 per cent, with 18 per cent of them recurring.

3

Congestion is dangerous in a way the running data will not show you. A meta-analysis of 16 studies found fixture congestion of 96 hours or less had no significant effect on total distance covered. Muscle injury rates rise with congestion regardless. A department watching only distance will conclude the squad is fine.

4

Tracking systems agree about the easy variables and disagree about the useful ones. Total distance is robust. Accelerations, decelerations and derived metabolic measures are considerably less so, and they are the variables load models lean on hardest. A club that changes provider has changed its own history.

5

The strongest AI result in football is a decision-support result. In the TacticAI study, expert raters at a professional club favoured the system's corner-kick suggestions over existing tactics in about 90 per cent of cases. The system proposed. People chose. That is the shape of the opportunity.

6

The barriers clubs report are organisational, not algorithmic. In a survey of 29 national federations and 32 professional clubs, the leading obstacles were resourcing, in-house expertise and the absence of a shared metric vocabulary, with only around 30 per cent finding provider positional data sufficiently transparent.

What follows for a football club

The recommendation is less exciting than the technology and more useful than the sales cycle. Clubs should join their records under one player identity before buying any intelligence, because a system reasoning across five incompatible platforms will reason confidently about the wrong player. They should insist that every recommendation names its sources, states its confidence and offers alternatives, and they should place a named person between any recommendation and any player.

Above all they should record decisions. A decision record costs nothing at the moment of the decision and cannot be reconstructed afterwards at any price. It is the only asset in this field that appreciates, and it is the one thing a change of manager currently destroys.

Chapter 12 sets out a 90-day proof of value on one squad group, with availability, completed against planned work and decision-record completeness as the measures. Chapter 13 describes what a club looks like three years later, which is much as it looks now, except that it can answer the question a poor result always provokes.

CHAPTER ONE

Why football's next advantage is decision intelligence

Every club in the top tier now has the same data from the same vendors. The advantage has moved to what happens between the dashboard and the decision.

For 15 years football's analytics arms race was an acquisition race. Better tracking, richer event data, more cameras, more wearables, more staff to read them. That race is largely over, in the sense that the leading clubs in every major league now buy comparable data from a small number of comparable suppliers. Whatever advantage existed in having the numbers has been competed away.

What has not been competed away is the quality of the decision that follows. Two clubs holding identical data will reach different conclusions about the same player, and the difference lies in how they assemble evidence, how they handle disagreement between departments, and whether anybody records why they chose what they chose. Those are organisational properties rather than analytical ones, which is precisely why they remain unevenly distributed.

THE CORE CLAIM

Football's remaining advantage is in decision quality and institutional memory, neither of which can be bought from a data vendor.

A survey of 29 national federations and 32 professional clubs associated with a recent World Cup found the leading barriers to data use were resourcing, a lack of in-house analytics expertise and the absence of shared performance-metric definitions, with only around 30 per cent of federation respondents finding provider positional data sufficiently transparent.[32] PRACTICE Those are self-reported perceptions rather than measured outcomes. They also describe a set of problems no additional data feed will solve.

1.1 The cycle this series uses

Every paper in this series describes the same five-stage cycle, because it is the sequence a performance decision actually follows whether or not anybody writes it down. Information is gathered, converted into insight, weighed as a decision, delivered as work, and compared afterwards against what happened. Football does the first and fourth stages superbly. The second, third and fifth are where institutional knowledge either accumulates or evaporates.

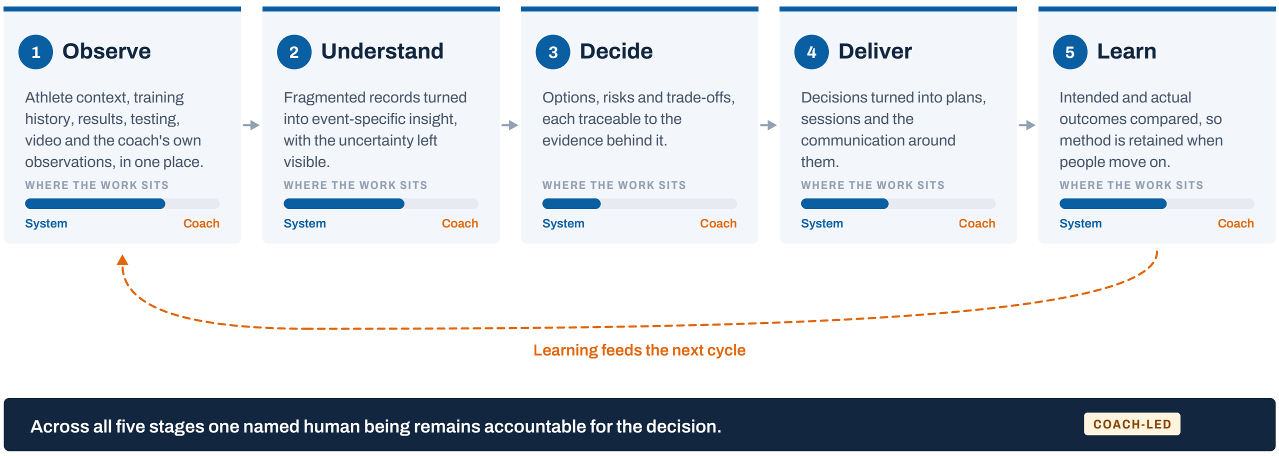

FIGURE 1 The performance intelligence cycle. The bar beneath each stage shows the balance of work between system and staff, which shifts deliberately towards the human at the point of decision.

Athleet.AI framework, proposed rather than validated. ATHLEET.AI

The proportions carry the argument. A system earns its place in observation, where it holds more than a person can, and in learning, where it compares intention against outcome without getting bored or defensive. At the moment of decision its contribution falls, and it should. Any product whose contribution rises at that point is selling something the evidence does not support and football should not buy.

THE ARGUMENT IN ONE PARAGRAPH

Football does not need a machine that picks the team. It needs one that remembers precisely, assembles evidence across departments that do not naturally talk, shows its working, and leaves the decision where the accountability already sits. Everything that follows is an elaboration of that sentence, including the parts explaining why several popular products fail it.

CHAPTER TWO

The data-rich, insight-poor club

A modern club stores an extraordinary quantity of information in five systems that cannot see one another, owned by four departments that do not entirely trust one another.

Ask a performance director where a player's information lives and the honest answer involves a list. Tracking sits with the sports science team. Event data sits with analysis. Medical records sit, correctly, behind a clinical wall. Scouting reports sit in a recruitment platform. Video sits in a third place, and the coach's own view of the player sits in the coach's head, which is the one source nobody has tried to capture.

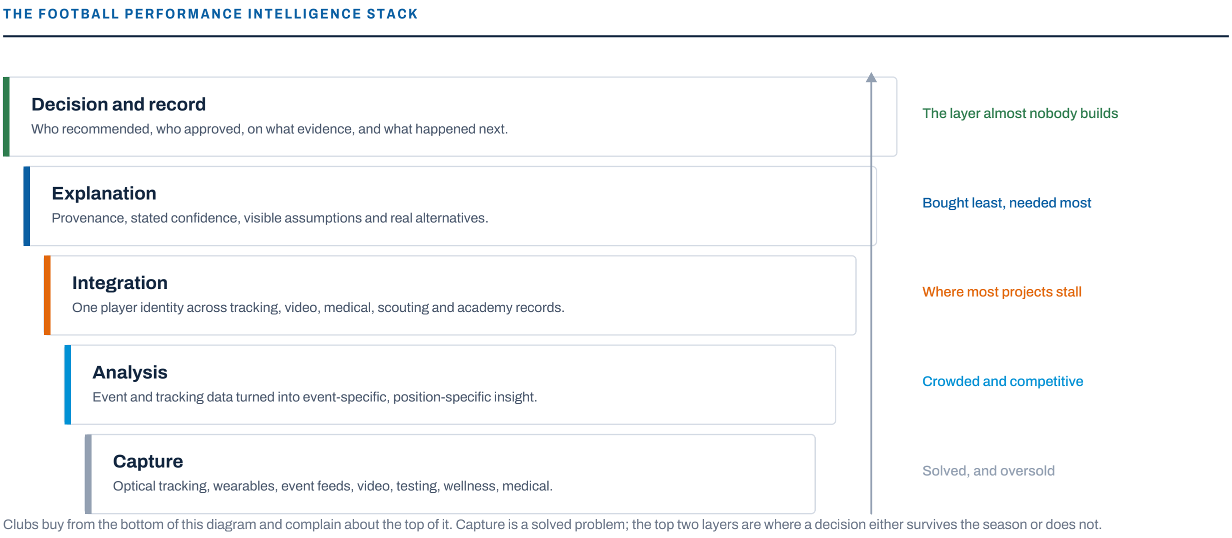

FIGURE 2 The football performance intelligence stack. Investment concentrates at the bottom, and the recurring complaints concentrate at the top. Athleet.AI framework. **ATHLEET.AI**

The departmental structure is not a fault. Medical confidentiality exists for good reasons and should not be dissolved for analytical convenience, a point Chapter 10 returns to. What is a fault is that the integration layer above those systems is usually a person, working from memory, under time pressure, on a Wednesday.

2.1 Volume is not the constraint

The instinct when a club feels short of insight is to add a data source. That instinct is almost always wrong, because the binding constraint is rarely how much the club records. It is how much of what the club already records reaches somebody making a decision, in a form they can act on, before the decision is made.

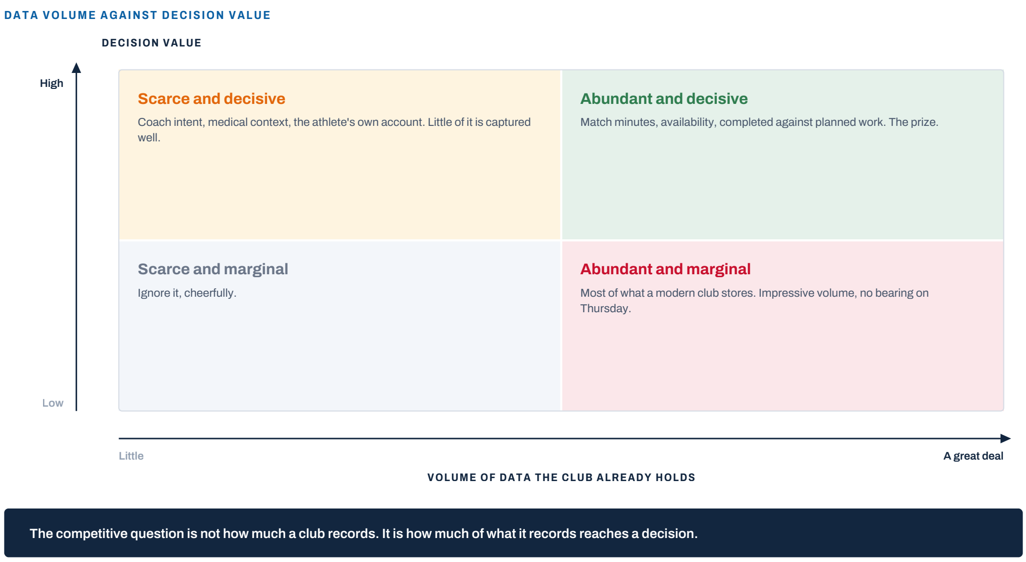

FIGURE 3 Data volume against decision value. The upper-left quadrant is where the scarce and decisive information sits, and it is the quadrant least served by the current market.

Athleet.AI framework. **ATHLEET.AI**

The upper-left quadrant deserves attention. Coach intent, the medical view of risk, and the player's own account of how they feel are all decisive and all poorly captured, largely because they are unstructured and arrive in conversation. A system that captures a coach's stated intention for the week and holds it against what actually happened is doing something no tracking vendor offers.

ILLUSTRATIVE SCENARIO

The invented Wednesday

An invented club is preparing for a Saturday fixture after a Wednesday cup tie. A midfielder played 90 minutes on Wednesday, has a history the medical department describes as sensitive, and is the only player in the squad comfortable in the role the manager wants on Saturday.

Four departments hold four pieces of that picture and each is confident in their own. The physical team can quantify Wednesday's exposure. The medical team can describe the history without disclosing the detail. Analysis can show what the role demands. The manager knows what the fixture is worth. The decision takes eight minutes in a corridor, and in September nobody can reconstruct it. The technology question is not how to make that decision automatically. It is how to make sure all four pieces are in the corridor at the same time, and how to keep the answer.

CHAPTER THREE

What the technology can actually do

Tracking systems agree closely about the variables that matter least and diverge about the ones a load model leans on hardest.

Football's measurement infrastructure is genuinely impressive and it is not uniformly reliable. The distinction matters commercially, because clubs sign multi-year contracts on the assumption that a number is a number, and then discover that changing supplier has altered their own historical record.

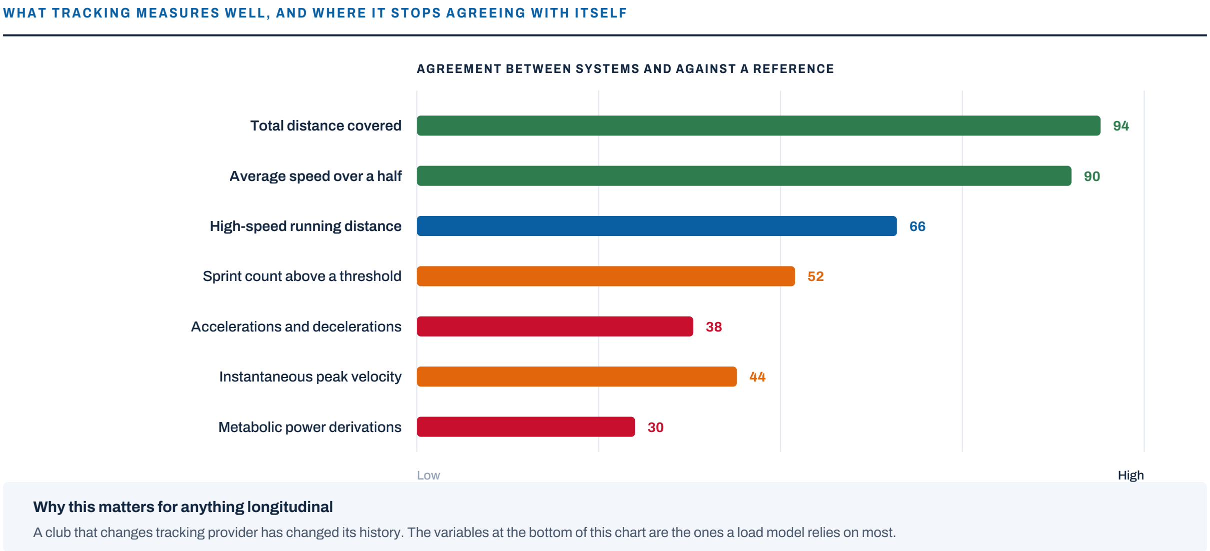

FIGURE 4 Where tracking systems agree, and where they stop. Positions reflect the author's reading of the validation and inter-system comparison literature cited at references 13 to 19, and are offered as a discussion instrument rather than a measurement.

ATHLEET.AI

The underlying literature supports the shape of that picture rather than its precise values. A review of global positioning systems in team sport concluded that accuracy improves with sampling frequency and remains imprecise for short, high-intensity and curvilinear movement, which is most of what a footballer does.[13] COMMENTARY A systematic review of local positioning systems found them generally valid for load monitoring and often more accurate than satellite or video alternatives, with caution still required at high speed and on complex trajectories.[14] REVIEW

Cross-system work is more pointed still. A study integrating multi-camera semi-automatic tracking, local position measurement and satellite systems within official and training football settings established that the outputs are not straightforwardly comparable,[16] PEER-REVIEWED and a subsequent paper asked directly whether position tracking data from different technologies can be merged.[18] PEER-REVIEWED A validation study of a widely used optical system under football-specific conditions provides the reference point for what good looks like.[17] PEER-REVIEWED

A PROCUREMENT QUESTION

Ask a tracking supplier to state, in the contract, the agreement between their system and the one it replaces, variable by variable. The answer determines whether the club keeps its history.

My own position, which goes beyond what any single paper states, is that a club should treat any change of tracking provider as a break in its longitudinal record, and should say so in writing to everyone who uses that record. The alternative is a load model quietly trained across a discontinuity nobody flagged.

3.1 What artificial intelligence has actually achieved here

The most substantial published result in football AI is instructive precisely because it is a decision-support result. In the TacticAI study, a system trained on corner-kick data produced tactical suggestions that expert raters at a professional club favoured over existing tactics in around 90 per cent of cases, alongside a receiver-prediction accuracy of 0.782 in its top three candidates.[20] PEER-REVIEWED The rating panel was small and drawn from one club, so the 90 per cent figure describes expert preference in one setting rather than a general truth.

Around that sit several other genuine contributions. A framework for valuing every on-ball action rather than only goals and assists has been widely adopted.[23] PEER-REVIEWED A broad position paper from the same research group set out what artificial intelligence might do for football and what football might do for artificial intelligence.[21] COMMENTARY Work on multi-agent reinforcement learning has produced coordinated team behaviour in simulation, which is a striking result about learning and not a result about coaching.[22] PEER-REVIEWED

The contrary evidence deserves equal space. A review of machine learning applied to match-outcome prediction across sports found that models frequently fail to beat simple bookmaker-odds baselines.[24] REVIEW A systematic review of big data in tactical performance analysis found most studies remained descriptive rather than predictive or prescriptive, and identified a persistent gap between the science and the practice.[25] REVIEW Expected goals models, meanwhile, are demonstrably sensitive to feature choice and data quality, which is a useful reminder that even the sport's most familiar derived metric is a model rather than a measurement.[26] PEER-REVIEWED

CHAPTER FOUR

From match analysis to coaching action

The distance between a good analysis and a changed training session is measured in relationships, not in processing.

Analysis departments in elite football produce more good work than gets used. The reason is rarely the quality of the analysis and usually the shape of it, since a 40-slide review delivered on Monday competes for a coach's attention against a squad, a fixture, an owner and a press conference. What survives that competition is short, specific and framed as a decision.

This is where a well-designed system earns its keep, and the work it does is unglamorous. Retrieval comes first, meaning the ability to find every instance of a defined situation across four seasons in seconds rather than an evening of an analyst's time. Preparation comes second, meaning a draft of the week's evidence assembled before the meeting rather than during it. Translation comes third, meaning the same finding expressed differently for a coach, a player and a director.

DESIGN CONSEQUENCE

Evidence that arrives as a document gets filed. Evidence that arrives inside the decision it bears on gets used. The difference is a product design choice, not an analytical one.

4.1 Turning an observation into a training priority

The step football does least well is the one from pattern to session. An analysis showing that the team concedes disproportionately from a particular transition is a finding. Deciding that Thursday's session will therefore rehearse a specific defensive rotation, at a specific intensity, with specific players, is a coaching decision that draws on the finding and on a great deal the finding does not contain.

TABLE 1 Four analytical outputs and the decision each can honestly support

OUTPUT	WHAT IT ESTABLISHES	THE DECISION IT SUPPORTS	WHAT IT CANNOT DO
Expected goals	The quality of chances created and conceded, given a model's assumptions	Whether a run of results reflects process or variance	Explain why the chances occurred, or settle selection
Action values	The contribution of on-ball actions beyond goals and assists	Comparison of players in similar roles within one data source	Compare across providers, or value off-ball work well
Tracking exposure	What a player physically did, with the caveats of Chapter 3	Planning the week's distribution of work	Predict who will be injured
Set-piece modelling	Likely receivers and outcomes from a defined situation	Rehearsing a routine and preparing for an opponent's	Substitute for the coach's judgement about the squad

Compiled by the author from references 13 to 26. The final column is the column that gets skipped in a sales meeting.

Two habits distinguish clubs that convert analysis into action. They state the coaching intention for a week in writing before the week starts, which makes it possible to ask afterwards whether the week did what it was for. And they record which analytical outputs actually changed a session, which after a season tells them what to stop producing.

PART TWO

From data to Thursday

Four chapters on the week, the player, the pathway and the point at which a great deal of analysis has to become one decision about one person.

CHAPTER FIVE

Smarter training and availability decisions

Football's load literature contains one robust finding, one dangerous null result and a great deal of overclaiming built on both.

Every paper in this series examines the same management problem through the demands of a different sport. What the competition requires, what the staff plan, how the player responds, what is then decided, whether those decisions are building capacity, and what the resulting continuity is worth. Football presents that problem at unusual density, because the squad is large, the fixtures are relentless and the decisions are made weekly for 25 people at once.

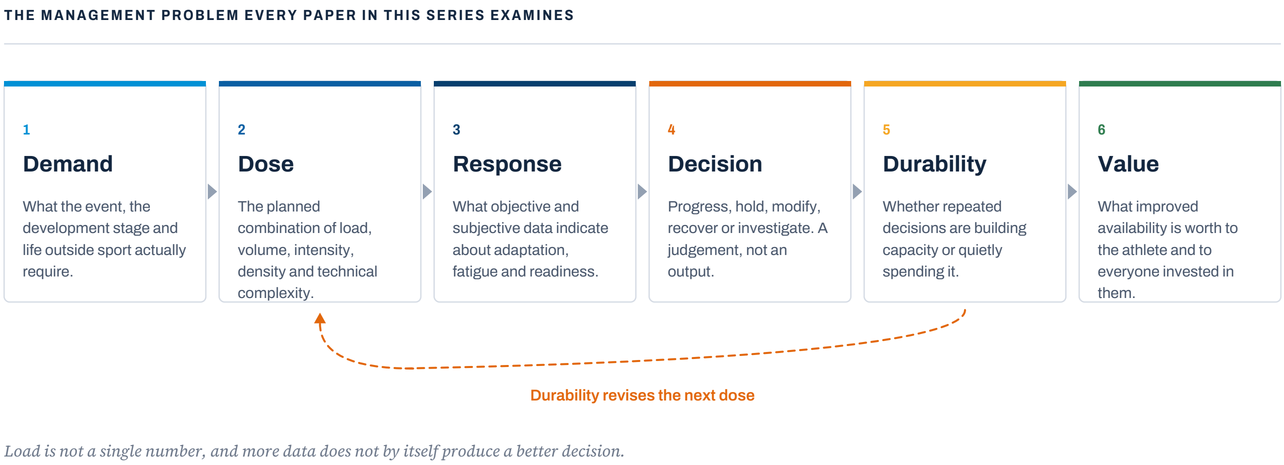

FIGURE 5 Demand, dose, response, decision, durability and value. The dotted return path is the part that distinguishes a training record from a training system.

Athleet.AI framework, used across the whole series. ATHLEET.AI

5.1 There is no squad, only players

The most common analytical error in football is the squad average. A squad-level load report describes a fictional player who started, came on, sat on the bench and flew to an international fixture, all in the same week.

ONE MICROCYCLE, FOUR PLAYERS, FOUR DIFFERENT WEEKS

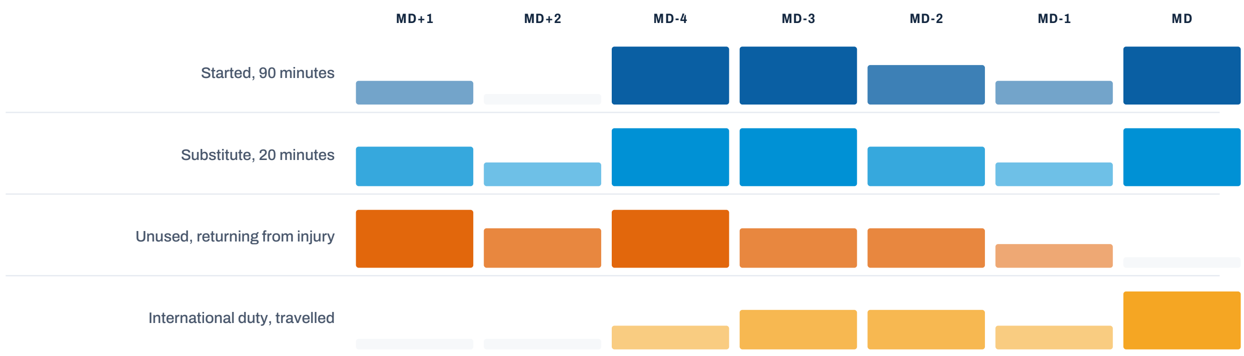

The point of the diagram

Four players in the same squad, in the same week, on four incompatible schedules. Illustrative rather than measured. A squad-level load average describes none of them, and the coach who asks for one has asked the wrong question.

FIGURE 6 One microcycle, four players. The pattern is illustrative rather than measured, and the point is that no single squad number describes any of these four weeks.

Athleet.AI framework. ATHLEET.AI

That is a problem a system solves well and a person solves slowly. Producing 25 individual weekly pictures takes an afternoon of a sports scientist's time; producing 25 drafts for that scientist to correct takes minutes, and the correcting is where the expertise lives.

5.2 The congestion trap

TWO FINDINGS A PERFORMANCE DEPARTMENT SHOULD HOLD TOGETHER

HAMSTRING INJURIES AS A SHARE OF ALL INJURIES

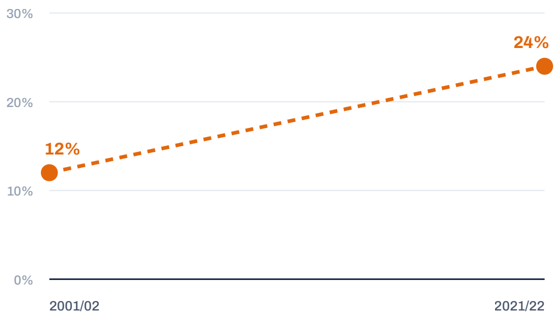

Across 21 seasons of the UEFA Elite Club Injury Study, hamstring injuries roughly doubled as a share of the total. The rise concentrated from 2014/15 onwards. Only the two endpoints are plotted; the dashed line marks direction rather than measured intermediate values.

No significant effect

A meta-analysis of 16 studies found that fixture congestion of 96 hours or less had no significant effect on total distance covered in a match, with a pooled effect of 0.12 and a p value of 0.134.

Congestion is associated with higher muscle injury rates and yet does not reliably reduce running output. A department reading only the running data will conclude the squad is fine.

JULIAN AND COLLEAGUES, 2021

Sources at references 5 and 11 in the reference list.

FIGURE 7 Two findings that need to be held together. Only the endpoints of the hamstring series are plotted; the dashed line marks direction rather than measured intermediate values.

References 5 and 11. PEER-REVIEWED

Muscle injury rates in professional football rise with fixture congestion, established across 11 seasons of the UEFA Champions League injury study.[3] PEER-REVIEWED A meta-analysis of 16 studies then found that congestion of 96 hours or less had no significant effect on total distance covered, with a pooled effect of 0.12 and a p value of 0.134.[11] REVIEW A narrative review reached a compatible position, noting that markers of muscle damage persist for around 72 hours while distance-based running output remains largely unchanged.[12] COMMENTARY

Put those together and the practical implication is uncomfortable. The variable most performance departments watch most closely is the variable least sensitive to the risk that congestion creates. A club reassured by stable distance numbers through a congested December has been reassured by the wrong instrument.

READ THIS TWICE

Higher injury risk with unchanged running output means the most-watched dashboard in football is the one least able to see the problem.

5.3 What the injury evidence does and does not license

Football's injury epidemiology is among the best in sport. A meta-analysis pooled overall incidence at around 8.1 injuries per 1,000 hours, with match incidence roughly ten times training incidence, at about 36 against 3.7 per 1,000 hours.[9] REVIEW Across 18 seasons, 49 teams and 19 countries, overall injury incidence fell by around three per cent a season and reinjury rates by around five per cent.[6] PEER-REVIEWED Surveillance outside Europe found a compatible pattern, with match injury burden far exceeding training burden.[8] PEER-REVIEWED

Against that improving picture sits one deteriorating strand. Hamstring injuries rose by around four per cent annually over 13 seasons,[4] PEER-REVIEWED and across 21 seasons they grew from roughly 12 to 24 per cent of all injuries, with 18 per cent recurring and 69 per cent of those recurrences arriving within two months.[5] PEER-REVIEWED The recurrence figure is the one a decision-support system can act on, because a recurrence within two months is a decision problem rather than a prediction problem.

WHAT THIS RULES OUT

Three claims a supplier should be asked to withdraw

- That a single readiness score indicates whether an individual player should train today.
- That a model can identify which players will be injured in the coming weeks at a useful individual accuracy.
- That a workload ratio has an optimal zone applicable to a specific player in a specific role.

None of these is a question of tuning. A systematic review of 204 musculoskeletal injury-prediction models across 30 studies found 98 per cent at high or unclear risk of bias, with external validation inadequate or absent in most, and concluded that none was ready for practical use.[44] REVIEW The underlying mathematics is unkind too, since injuries are low-incidence multifactorial events for which even a statistically significant group-level marker will have poor individual predictive value.[45] COMMENTARY The dominant load ratio has been shown to produce its familiar association even when fed meaningless inputs.[46] COMMENTARY

What survives is substantial. Exposure records describe what happened accurately. The separation of external dose from internal response remains the right conceptual frame.[47] COMMENTARY Injury is better understood as emerging from an interacting system than as the output of a single variable.[48] COMMENTARY And availability, which Chapter 11 takes up, is measurable within a season and connected to results.

CHAPTER SIX

Technical and tactical learning

Video is football's richest asset and its most wasteful process. Most of the value is lost in the search.

An analyst preparing individual feedback for a squad of 25 spends the majority of that time finding clips rather than thinking about them. That ratio is the single clearest efficiency argument in this paper, and it requires no model to make a decision about anybody.

Retrieval at scale changes what is possible. When every instance of a defined situation across four seasons is available in seconds, a coach can ask questions they would not previously have bothered to ask, because the cost of asking has collapsed. Whether the player has ever done this successfully. Whether the pattern predates the current manager. Whether it appears against one shape of opponent and not another.

6.1 Holding the club's own idea of the game

Two respected coaches will describe the same defensive situation in ways that partly contradict each other, and both win matches. A system that imposes a third view is worse than useless, since it introduces an authority the club never agreed to. A system that records what this coaching staff believe, applies it consistently across 25 players and an academy, and shows where an individual departs from it, is doing something no member of staff can do alone.

That is also the mechanism by which a club survives a change of manager without losing everything. A written, applied, versioned model of how this club wants to play is an asset. It is currently held almost entirely in conversation.

WORTH HIGHLIGHTING

The most valuable thing a system can hold in a technical setting is the coach's own model, applied consistently. A third opinion helps nobody.

ILLUSTRATIVE SCENARIO

The invented full-back and the tidy conclusion

An invented full-back has produced fewer progressive passes in each of five matches. An automated report flags the decline and, if poorly designed, presents it as a technical fault.

The coaching staff know three further things. The team changed shape in October, the winger ahead of him is a different player with different movement, and two of those five matches were against opponents who deliberately blocked that pass. The measurement is real. The conclusion drawn from it alone is worthless, and a system offering it as an instruction has made the analyst's job harder rather than easier.

6.2 Feedback the player will actually use

Personalised feedback is one of the clearest wins available, and it is limited by production time rather than by insight. A player who receives two of their own clips with one question attached, on the morning of a session about that exact thing, learns faster than one who receives a 20-clip compilation on a Monday. Producing the first at squad scale is a retrieval and drafting problem, which is to say a solved one.

CHAPTER SEVEN

Academy and pathway intelligence

The academy is where a club's information problem is worst and where the consequences of getting it wrong last longest.

A first-team squad is 25 people watched constantly. An academy is a far larger group of children observed intermittently by staff who turn over, in age groups that reset annually, with the record of each year's judgement largely evaporating into the next. It is the part of the club where continuity of information would be worth most and where it is weakest.

THE PLAYER CONTEXT GRAPH

The medical node is deliberately drawn as availability rather than diagnosis. A performance system has no business holding the second.

FIGURE 8 The player context graph. The medical node is deliberately expressed as availability rather than diagnosis, and the decision-history node is the one clubs most often lack.

Athleet.AI framework. ATHLEET.AI

Two evidential points frame what a club should expect of itself here, and the companion paper in this series treats both at length. The first is that early selection in football is heavily contaminated by birth month, with players born in the first quarter of the selection year making up 43 to 46 per cent of European national youth squads against 9 to 18 per cent born in the last.[27] PEER-REVIEWED The second is that the effect is a selection bias rather than a performance advantage, since a large analysis of European professionals found early-born players over-represented and no better valued in the transfer market.[28] PEER-REVIEWED

Maturity-matched competition, in which players are grouped by biological rather than chronological development, has been examined empirically. In a maturity-banded tournament involving 115 academy players, earlier maturers reported greater physical and technical challenge and later maturers reported more opportunity to show technical and tactical ability.[29] PEER-REVIEWED That is perception data from a single tournament, so it establishes that the format changes the experience rather than that it changes outcomes.

COMPANION PAPER

Casting the Net Wider, paper four in this series, treats talent discovery, deselection and re-entry in full. This chapter covers only what an elite club needs from its own academy record.

7.1 What the progression evidence actually says

Clubs and journalists both quote conversion rates that turn out to be difficult to source. A verified ten-year follow-up of 198 academy prospects at two clubs found 6.1 per cent became professional, of whom seven reached the top division, meaning about 3.5 per cent of the original cohort.[30] PEER-REVIEWED That is a Spanish sample and should not be read as an English figure. An English study of 382 academy graduates across 33 clubs measured how much those graduates actually played rather than how many turned professional, finding wide variation between clubs and only a weak relationship between the number produced and the minutes they received.[31] PEER-REVIEWED

No verified England-wide academy conversion rate appears in the peer-reviewed literature. That absence is stated here as a fact about the evidence rather than a gap in the search, and any supplier quoting a precise national figure should be asked for its source.

7.2 The record that should outlive the age group

What an academy needs from a system is unglamorous and specific. Every observation dated and attributed. Birth month and maturity status recorded as context alongside the observation rather than buried. Individual development objectives that persist across a change of age-group coach. And a decision record for every retain and release judgement, stating what was decided, by whom, on what evidence, and when it will be revisited.

My own view, which goes beyond anything the literature establishes, is that the release decision is the single most consequential judgement any club makes about a person, and it is the one least well documented. A club that cannot explain a release two years later has failed a duty it should recognise as its own.

CHAPTER EIGHT

The multidisciplinary performance meeting

The most valuable hour in a football club is the one where five departments disagree, and it is almost never recorded.

Every club runs some version of the same meeting. Coaching, physical performance, medical, analysis and sometimes recruitment assemble, review the squad, and reach positions on availability and selection. It is where the club's real intelligence lives, and it typically produces no durable output beyond a shared verbal understanding that decays within a fortnight.

THE MULTIDISCIPLINARY DECISION LOOP

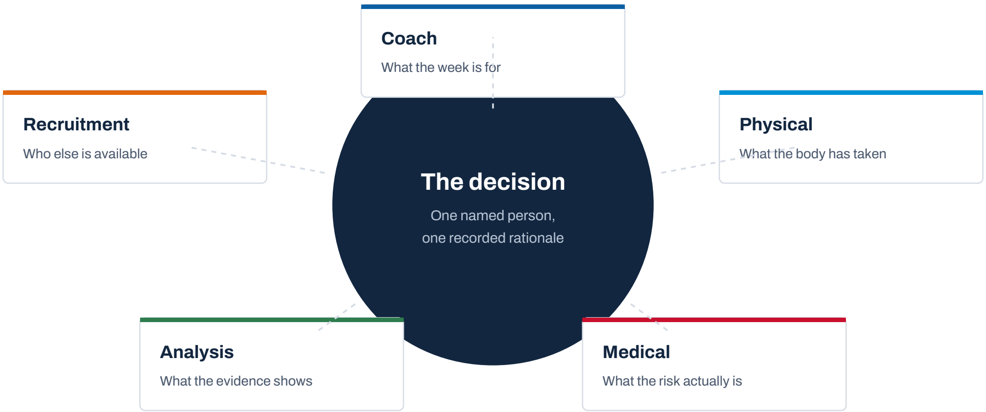

FIGURE 9 The multidisciplinary decision loop. Each department holds a legitimate and partial view, and the disagreement between them carries more information than any single view.

Athleet.AI framework. ATHLEET.AI

There is direct evidence that how these groups communicate matters to outcomes. A substudy of 14 teams within the UEFA Elite Club Injury Study found that a higher level of communication between medical and performance staff was associated with a lower hamstring injury burden.[10] PEER-REVIEWED A sample of 14 teams is modest, and communication quality was rated rather than measured directly, so the finding is suggestive rather than settled. It is nonetheless the clearest published signal that the organisational variable is a performance variable.

8.1 What a system should do before the meeting

The useful contribution is preparation rather than conclusion. Before the meeting, a system can assemble each player's exposure, availability status as stated by the medical department, the coaching intention recorded for the week, what was decided about this player last time a comparable situation arose, and what happened afterwards. That is four departments' worth of retrieval, done in advance, arriving as a single page.

What it must not do is arrive with a recommendation attached. A prepared recommendation anchors a meeting, and anchoring a meeting of experts is the specific failure mode this design is meant to avoid. The system's job is to ensure everybody is arguing about the same facts.

THE CHEAPEST ASSET

A decision record costs nothing at the moment of decision and cannot be reconstructed afterwards at any price.

8.2 What the meeting should produce

SPECIFICATION

The five fields of a decision record

- What was decided, in one sentence a player could understand.
- Who decided it, by name and role.
- What evidence was in front of the meeting, listed and dated.
- What the dissenting position was, and who held it.

- What would cause the decision to be revisited, and when it will be.

The fourth field is the one clubs resist and the one that matters most. A record showing that the medical department dissented, and that the decision went the other way for stated reasons, protects everybody in the room and creates the only honest basis for learning afterwards.

Two objections arise reliably. The first is that recording dissent creates legal exposure, which I suspect is close to the opposite of the truth, since a documented, reasoned disagreement is a far better record than a reconstructed memory. That is a view rather than legal advice, and a club should take its own. The second is that nobody has time, which is answered by the observation that the record is a by-product of a meeting the club is already holding.

CHAPTER NINE

Trust, explainability and accountability

A recommendation a coach cannot interrogate will be ignored by the good ones and followed by the inexperienced, and in football the second group is holding somebody's career.

Explainability in football is often treated as a procurement checkbox. It is a performance requirement, because the first response of any competent coach to a suggestion is to ask why, and a system that cannot answer has failed at the opening exchange. The more worrying case is the junior staff member who accepts the suggestion without asking, at which point the system has substituted its confidence for their judgement and nobody has noticed.

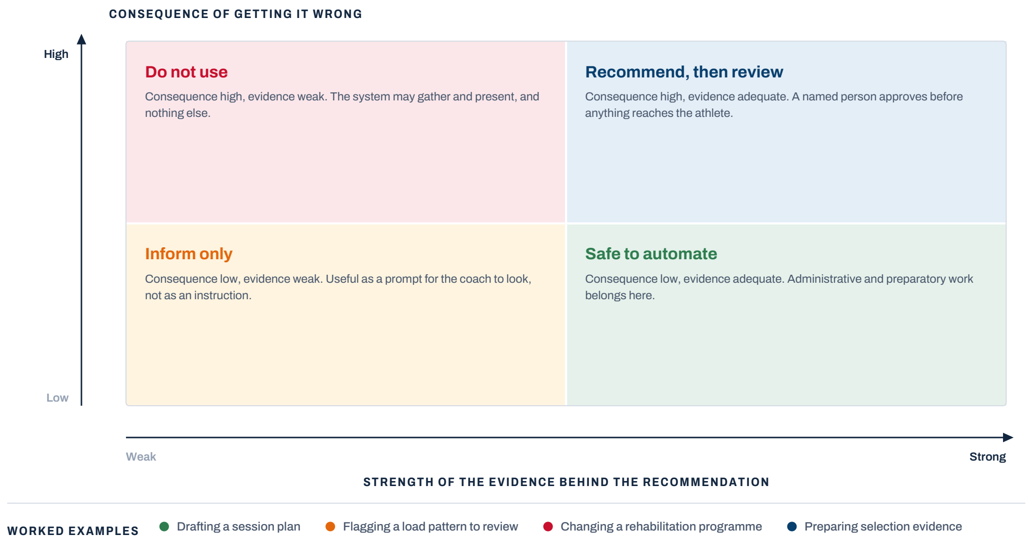

FIGURE 10 Consequence of error against strength of evidence. Most of what is currently sold into football sits in the upper-left quadrant and is marketed as though it sat in the upper-right.

Athleet.AI framework, used across the whole series. ATHLEET.AI

9.1 Interpretable models and post-hoc explanations

An influential argument in machine learning holds that post-hoc explanations of opaque models can be unreliable or misleading in high-stakes settings, and that inherently interpretable models should be preferred wherever comparable accuracy is achievable.[37] PEER-REVIEWED Applied to football, that points somewhere unfashionable. For most decisions a club actually faces, a transparent rule, a clear comparison against the player's own history and a stated confidence will be more useful than a sophisticated model with an explanation bolted on afterwards.

Bias arrives in football the same way it arrives everywhere, through a proxy rather than through intent. The best-documented case in the literature concerns a healthcare algorithm that used cost as a stand-in for need and systematically underestimated the needs of Black patients, because less had historically been spent on their care.[38] PEER-REVIEWED Football is full of cost-like proxies. Minutes played, squad membership, transfer fee and academy category are all records of past decisions, and a model trained on them will reproduce the club's history with new authority.

The sport-science literature carries its own gap. An analysis of 5,261 publications covering roughly 12.5 million participants found female participants made up around 34 per cent of the total.[50] PEER-REVIEWED A model trained on that literature will be more confident about men than women and will not volunteer the fact, which matters increasingly as clubs run women's teams on the same platforms.

A USEFUL TEST

Ask a supplier what their model would have said about a player the club got wrong. If the answer is that it would have agreed with the club, the model has learned the club rather than the game.

9.2 What the regulators expect

Guidance produced jointly by the Information Commissioner's Office and the Alan Turing Institute sets out the types of explanation an organisation should be able to give about AI-assisted decisions, covering rationale, responsibility, data, fairness, safety and impact.[34] CONSENSUS The guidance is not binding, and my own view is that it remains the most practical specification available to anyone buying in this area.

Two instruments bear directly on clubs operating in the United Kingdom. Article 9 of the UK General Data Protection Regulation treats data concerning health, and biometric data used to identify a person, as special category data requiring a further condition beyond a lawful basis.[35] CONSENSUS Routine performance metrics are not automatically special category data on that definition, a distinction commonly blurred in the industry, and on my own reading the boundary sits at the point where a metric reveals health status. A club should take its own advice rather than rely on that reading. Separately, section 80 of the Data (Use and Access) Act 2025 replaces Article 22 of the UK GDPR with new Articles 22A to 22D on automated decision-making, defining meaningful human involvement as the test of whether a decision was taken solely by machine, and retaining stricter conditions where special category data drives a solely automated significant decision. It comes into force on 5 February 2026, by commencement regulations made on 29 January 2026, and it does not apply to decisions taken before that date.[36] CONSENSUS The same instrument commences section 81, which is the provision that carries the child-specific duty, requiring a service likely to be accessed by children to take what the Act calls children's higher protection matters into account. Academies should read the two together and not assume that section 80 does the work for children.

The European Union's Artificial Intelligence Act entered into force on 1 August 2024 and applies in stages.[33] CONSENSUS Football performance systems are not named in the high-risk annexes, and on my own reading would ordinarily sit outside that tier, though a system used to inform employment decisions about a contracted player would need separate consideration. As the Regulation stands, prohibitions and the AI literacy duty have applied since 2 February 2025, general-purpose model obligations since 2 August 2025, and general application follows on 2 August 2026, with product-embedded high-risk systems a year later. A simplification package proposed by the Commission in November 2025 would defer the high-risk obligations, but it is a proposal and not the law, so any club relying on a particular deadline should verify it against the current official position rather than against this paper.

9.3 Whose data it is

The question of player agency is being settled by practice faster than by law. A charter of player data rights developed in professional football sets out eight rights over performance and personal data, including access, portability, rectification and erasure, citing survey evidence that 80 per cent of professional players surveyed wanted access to their own performance data.[42] CONSENSUS That charter addresses professional adults and not academy minors, which is a gap clubs should close in their own policies rather than wait for.

CHAPTER TEN

Reference architecture

The hard part of the architecture is not the connection. It is deciding what must stay separate, and enforcing it.

Most integration projects in football fail for one of three reasons. There is no single player identity, so records cannot be joined reliably. There is no owner for each dataset, so nobody is accountable for its quality. Or the medical boundary is handled by informal trust rather than by design, and the clinicians quite properly withdraw.

10.1 Five design rules

SPECIFICATION

- **One player identity, owned by the club.** Every source keyed to it, including academy and loan records, and portable if a supplier is replaced.
- **Medical shares availability, not diagnosis.** The performance layer receives a status and a set of constraints. The clinical detail stays behind the clinical wall, permanently and by architecture rather than by policy.
- **Planned and completed work stored separately.** The difference between them is the most informative quantity a club holds and it disappears the moment the two are merged.
- **Every modification carries its reason.** A session shortened for a calf complaint and one shortened because the pitch was unusable look identical in a database and mean entirely different things.
- **Provenance on every derived number.** Which source, which version, which provider, which date. Chapter 3 explains what happens to a club that cannot answer this after changing supplier.

Two documentation practices from machine learning transfer directly and cost almost nothing. Model cards record a model's intended use, its performance across subgroups and its known limitations.[39] PEER-REVIEWED Datasheets record a dataset's motivation, composition, collection and maintenance.[40] PEER-REVIEWED Almost nobody in football asks for either, and either would improve most procurement conversations immediately.

At organisational level, the NIST AI Risk Management Framework organises the work into governing, mapping, measuring and managing risk across a system's life,[41] CONSENSUS and ISO/IEC 42001 provides a certifiable management-system standard.[43] CONSENSUS The International Olympic Committee has published principles covering privacy, integrity, fairness, robustness and human accountability across the Olympic Movement.[49] CONSENSUS None binds a football club. All three give a board a vocabulary for asking a supplier sensible questions.

10.2 Six questions to put to any supplier

TABLE 2 Procurement questions, and what a good answer looks like

QUESTION	WHAT A GOOD ANSWER LOOKS LIKE
What was this trained on?	A named dataset with its population, leagues, sexes and limitations stated. A refusal is itself an answer.
Where does it perform worst?	A specific answer naming positions, populations or conditions. Uniform performance claimed is uniform performance unmeasured.
What does it do when uncertain?	It says so and declines to recommend. A system that always produces an answer is not measuring its own confidence.
Who owns the player record?	The club and the player, with export in a usable format and a written position on what happens at transfer or release.
Does our data train your models?	A clear yes or no, in the contract rather than the meeting.
What has been independently validated?	A published or auditable evaluation. Internal benchmarks are marketing until somebody outside the company has checked them.

Compiled by the author, drawing on references 39 to 43 and 49. ATHLEET.AI

CHAPTER ELEVEN

Availability, durability and asset value

Availability is measurable within a season, connected to results, and the only claim in this whole field that the football evidence actually supports.

Chapter 5 argued that individual injury prediction is beyond the current evidence. This chapter argues that football should pursue something less exciting and far better supported, which is availability, meaning the proportion of the season in which a player is fit to be selected.

WHAT THE AVAILABILITY EVIDENCE ACTUALLY SHOWS

27 CLUBS, 11 SEASONS

7,792

injuries recorded in the UEFA Champions League injury study between 2001 and 2012.

Lower injury burden was associated with better league ranking, more points per match and a better European ranking.

18 SEASONS, 19 COUNTRIES

3%

the annual fall in overall injury incidence across 1.78 million hours of exposure.

Reinjury rates fell faster still, at around five per cent a season, and availability rose slowly.

21 SEASONS

24%

of all injuries in men's professional football are now hamstring injuries, up from around 12 per cent.

Of those, 18 per cent recur, and 69 per cent of recurrences happen within two months.

The limit of the claim

Every one of these findings is observational. Squad quality, wealth and medical resourcing plausibly drive both fewer injuries and better results, so the association between availability and success cannot be read as proof that reducing injuries causes points. It does establish availability as a performance variable worth managing, which is a smaller and more useful claim.

FIGURE 11 Three findings from the UEFA Elite Club Injury Study series, and the limit of what they license.

References 2, 5 and 6. PEER-REVIEWED

The central finding comes from a prospective cohort of 27 clubs across 11 seasons, covering 7,792 injuries, which reported that teams with a lower injury burden achieved significantly better league ranking, more points per match and a better European ranking.[2] PEER-REVIEWED A separate analysis in a different league and continent found the same direction of association between low injury rates and final league position.[7] PEER-REVIEWED

Both are observational. Squad quality, wealth and medical resourcing plausibly drive both fewer injuries and better results, so the association cannot be read as proof that reducing injuries causes points. What it does establish is that availability behaves as a performance variable, which is a smaller claim and a sufficient one for anybody deciding where to spend attention.

THE HONEST CAVEAT

Wealthier clubs have better squads, better medical departments and better results. That confound cannot be ruled out, and it should be stated whenever this evidence is used commercially.

11.1 The value chain, and where to stop believing it

THE PLAYER AVAILABILITY AND CAREER VALUE CHAIN

Better decisions	More available days	More minutes	A longer career	Value to others
Evidence assembled, disagreement visible, reasoning recorded.	Fewer avoidable interruptions to training and selection.	Selection based on readiness rather than guesswork.	More seasons in which to develop and to earn.	Squad stability, protected development spend, sponsor continuity.

Read this chain from left to right, and doubt it from right to left

Each link is plausible and only the first two have supporting evidence in football. No published study shows that a software system lengthens a career. A club running a pilot should measure the left of the chain and treat the right of it as a hypothesis.

FIGURE 12 The availability and career value chain. Each link is plausible; only the first two have supporting evidence in football.

Athleet.AI framework. ATHLEET.AI

Commercial conversations about this technology tend to run the whole chain in one breath, from better decisions to longer careers to protected transfer value. The first two links have evidence behind them in football. The rest are hypotheses, and a club should treat any supplier presenting career extension as a deliverable as describing a hope.

11.2 What to measure in a season

MEASURES AVAILABLE WITHIN ONE SEASON

- Availability rate, meaning the proportion of matches for which a player was fit to be selected.
- Completed against planned work at session level, recorded with the reason for every modification.
- Time-loss episodes, their duration, their category and their recurrence.
- Decision-record completeness, meaning the proportion of availability and selection decisions with a written rationale.
- The proportion of prepared evidence that staff recorded as useful, alongside the proportion they recorded as noise.

The fourth measure is the unusual one and the one I would weight most heavily in a first season. It is entirely within the club's control, it requires no supplier to achieve, and it is the foundation everything else in this paper rests on.

Career length is not on that list, because a season cannot measure it. Neither is transfer value, which is determined by a market rather than by a performance department. Both belong in a three-year review and neither belongs in a pilot.

CHAPTER TWELVE

Implementation roadmap

A pilot small enough to finish, specific enough to fail honestly, and designed so that the club learns something even if the technology does not survive.

The instinct in football is to run a technology trial across the whole first team in a season the club needs to win. That design cannot produce a clean answer and cannot be stopped without embarrassment. A better one is narrower.

PROPOSED STRUCTURE

- **Days 1 to 25.** Records only. One player identity across tracking, training, availability and academy systems for one squad group. Planned and completed work stored separately. Measures agreed in writing, along with a written definition of a disappointing result.
- **Days 26 to 60.** Preparation and retrieval only. The system assembles the weekly evidence pack and the pre-meeting page described in Chapter 8. It produces no recommendations at all. Staff log what was useful and what was noise, and the noise count is reported.

- **Days 61 to 85.** Explainable prompts based on departures from a player's own established pattern. Each names its sources, states its confidence and offers alternatives. Nothing reaches a player without a named approver, and every approval is recorded.
- **Days 86 to 90.** Review against the agreed measures with somebody outside the performance department in the room. Decide to stop, continue or extend, and write down the reasoning.

Three features of that design carry the weight. The first phase deliberately contains no artificial intelligence, so that any improvement from simply joining the records can be separated from any improvement attributable to the technology. Declaring the disappointing result in advance prevents the near-universal habit of redefining success at the end. And counting useless prompts alongside useful ones is the only honest measure of whether a tool is helping, since a system producing 40 alerts a week is switched off by February whatever its accuracy.

One further point about sequencing. The academy is usually the better place to start than the first team, because the time pressure is lower, the staff turnover the system is meant to survive is higher, and the value of continuity is greater. It is also, as Chapter 9 sets out, where the governance requirements are strictest, so a club that builds it correctly there has built it correctly everywhere.

START WHERE THE RISK IS LOWEST

Preparation and retrieval carry almost no risk of harm and deliver most of the available time saving. They are the right place to begin and, for some clubs, the right place to stop.

CHAPTER THIRTEEN

The augmented club

What a club looks like when it has done this well, which is much as it looks now, with one important difference.

Bringing the argument together produces a sequence rather than a purchase. A club that joins its records, insists on explainable output, keeps the medical boundary intact by design, places a named person between every recommendation and every player, and records decisions will have done almost everything this paper recommends, and can begin the first and last of those tomorrow without buying anything.

FIGURE 13 The football AI maturity model. The test beneath each stage is the question a club at that stage can answer about its own decisions. Athleet.AI framework. ATHLEET.AI

Elite football is unusual in being extremely well instrumented and weak at the fourth stage. Most top-tier clubs sit between the second and third, holding excellent data in systems that partially connect, with the reasoning behind decisions held in people who will not all be there next season.

13.1 Three years on

A club that has done this properly will not look dramatically different from outside. The coaches will still coach, the arguments in the performance meeting will still happen and will still be worth having, and the players will still do the work. What will have changed is that a decision made in September can be found in April with its reasoning intact, that a new head of performance inherits method rather than folklore, and that nobody is being asked to trust a number whose provenance they cannot establish.

That is a modest description of success and a demanding one to achieve. It is also, on the evidence assembled in this paper, considerably more valuable than the prediction the sport keeps being offered instead.

THE THREE REFUSALS

A club that adopts nothing else from this paper should adopt these. Decline any product that scores a player's readiness in a single number. Decline any product that claims to predict individual injury. And decline any product trained on the club's own past selections that then recommends those selections back, because that is not intelligence, it is an expensive mirror.

END MATTER

Implementation checklist

Fourteen things to settle before a single recommendation reaches a coach. Most cost nothing and none requires a supplier.

- [ ] **One player identity, owned by the club**
  Across tracking, training, availability, academy, loan and scouting records, and portable if a supplier changes.
- [ ] **A named owner for every dataset**
  Accountable for its quality, its access rules and its retention.
- [ ] **The medical boundary enforced by architecture**
  Availability and constraints out, clinical detail in. Designed rather than trusted.
- [ ] **Planned and completed work stored separately**
  The difference between the two is the most informative quantity the club holds.
- [ ] **A reason captured with every modification**
  A session cut for a calf and one cut for a waterlogged pitch look identical in a database.
- [ ] **Provenance on every derived number**
  Source, version, provider and date, so a change of supplier is visible rather than silent.
- [ ] **A written statement of how this club wants to play**
  Versioned, applied consistently, and owned by the club rather than the current manager.
- [ ] **The coaching intention for each week, recorded before the week**
  So that afterwards the club can ask whether the week did what it was for.
- [ ] **A decision record from every performance meeting**
  Including the dissenting position and who held it.
- [ ] **A named approver between every recommendation and every player**
  Recorded with the recommendation, at the time, without exception.
- [ ] **A written rationale for every retain and release decision**
  Especially in the academy, where the consequence lasts longest and the record is weakest.
- [ ] **A contractual answer on third-party model training**
  Obtained in the contract rather than the sales meeting.
- [ ] **A written position on player data at transfer or release**
  Agreed while the record is small, rather than when it has become valuable.
- [ ] **An agreed definition of a disappointing pilot result**
  Written down before the pilot starts, by the people who will be tempted to redefine success at the end of it.

END MATTER

Glossary

Terms used in a specific sense in this paper, defined once and held to throughout.

Availability
The proportion of matches for which a player was fit to be selected. Measurable within a season and used in this paper in preference to injury prediction.

Coach-led
An arrangement in which a named person approves any recommendation before it reaches a player, and that approval is recorded alongside it.

Decision record

A written account of what was decided, by whom, on what evidence, against what dissent, and when it will be revisited.

Decision support

Assembling and presenting evidence, options and trade-offs so that a person can decide. Distinct from automation, in which the system acts.

Event data

A coded record of discrete on-ball actions in a match. Distinct from tracking data, which records position continuously.

Explainable

A recommendation whose sources, assumptions, confidence and alternatives are visible to the person deciding, in language they can interrogate.

Expected goals

A modelled estimate of the probability that a given chance results in a goal. A model rather than a measurement, and sensitive to feature choice.

Internal and external load

External load is the work as performed, in distance, sprints or minutes. Internal load is what that work demanded of a particular player on a particular day.

Microcycle

The training week between one match and the next, usually described by days relative to match day.

Provenance

The record of which inputs, which provider and which version produced a given output, so that a number can be traced to its source.

Responsible AI

Systems designed so that their limits are visible, their outputs are contestable, and accountability for consequences rests with an identified person.

Tracking data

Continuous positional data from optical, satellite or local positioning systems. Comparable within a provider and, as Chapter 3 sets out, not straightforwardly comparable between them.

END MATTER

Method and evidence grading

How the sources were selected and checked, what could not be verified, and what the reader should distrust.

Sources were gathered from the peer-reviewed sport science and machine learning literature, governing-body publications, regulator guidance and primary legislation, with priority given to systematic reviews, long-running prospective cohorts and methodological criticism over single primary studies. Every reference below was checked against the publishing journal, publisher or issuing body, and the bibliographic details were verified independently a second time before publication. Where a detail could not be confirmed, it was removed or explicitly flagged rather than approximated.

What could not be verified

Two things were searched for and not confirmed. No verified England-wide academy-to-professional conversion rate appears in the peer-reviewed literature, so the Spanish cohort is cited and labelled as Spanish. And the numeric incidence-rate ratio attached to short recovery periods in the congestion study could not be retrieved from the primary source, so this paper states the direction of that finding rather than a figure. A citable primary document for the governing-body electronic performance and tracking system quality programme was also not found, so a peer-reviewed account of it is cited in its place.

What to distrust in this paper

Three things. The tracking agreement chart in Chapter 3 is the author's reading of a literature others read differently, and it is published in a form designed to invite disagreement. The microcycle diagram in Chapter 5 and the value chain in Chapter 11 are models of a mechanism rather than measured data. And the regulatory position in Chapter 9 describes a timetable that has moved more than once already, and part of which is the subject of a live legislative proposal, so it should be verified against the current official position rather than relied on here.

A note on the availability evidence

The association between low injury burden and results is the strongest evidence in this paper and the most easily overstated. It is observational, drawn from clubs that differ in wealth and squad quality, and it cannot establish that reducing injuries causes points. Every commercial use of that finding, including by the author, should carry the caveat.

Conflict of interest

The author is the founder of Athleet.AI, which builds software in the area this paper examines. The paper recommends against several categories of product that would be commercially straightforward to build and sell into football. Readers should weigh the argument accordingly and are encouraged to test the citations rather than the author.

END MATTER

About the author

Paul Forrest is the founder and Chief AI Officer of Athleet.AI. He works with organisations on the deployment of artificial intelligence, having begun in business process reengineering in the 1990s and moved into natural language and machine learning work from 2014, building small, domain-specific models in preference to general-purpose ones.

He is a qualified athletics coach and coaches as a volunteer at club level. That combination of enterprise systems work and time on the grass is the reason this series exists, and it is also the reason the papers decline to recommend several things that would be straightforward to sell.

He is the author of *Tales from AI Development Hell* and *AI Made the Decision: Who Answers When Nobody Owns the Outcome?*, and holds an adjunct professorship teaching digital acumen and digital entrepreneurship.

END MATTER

References

Fifty sources, each verified twice against the publishing journal, publisher or issuing body.

Injury, availability and training load in football

1. **Hägglund, M., Waldén, M., Bahr, R., & Ekstrand, J.** (2005). Methods for epidemiological study of injuries to professional football players: developing the UEFA model. *British Journal of Sports Medicine*, 39(6), 340–346. doi:10.1136/bjsm.2005.018267
2. **Hägglund, M., Waldén, M., Magnusson, H., Kristenson, K., Bengtsson, H., & Ekstrand, J.** (2013). Injuries affect team performance negatively in professional football: an 11-year follow-up of the UEFA Champions League injury study. *British Journal of Sports Medicine*, 47(12), 738–742. doi:10.1136/bjsports-2013-092215
3. **Bengtsson, H., Ekstrand, J., & Hägglund, M.** (2013). Muscle injury rates in professional football increase with fixture congestion: an 11-year follow-up of the UEFA Champions League injury study. *British Journal of Sports Medicine*, 47(12), 743–747. doi:10.1136/bjsports-2013-092383
4. **Ekstrand, J., Waldén, M., & Hägglund, M.** (2016). Hamstring injuries have increased by 4% annually in men's professional football, since 2001: a 13-year longitudinal analysis of the UEFA Elite Club Injury Study. *British Journal of Sports Medicine*, 50(12), 731–737. doi:10.1136/bjsports-2015-095359
5. **Ekstrand, J., Bengtsson, H., Waldén, M., Davison, M., Khan, K. M., & Hägglund, M.** (2023). Hamstring injury rates have increased during recent seasons and now constitute 24% of all injuries in men's professional football: the UEFA Elite Club Injury Study from 2001/02 to 2021/22. *British Journal of Sports Medicine*, 57(5), 292–298. doi:10.1136/bjsports-2021-105407
6. **Ekstrand, J., Spreco, A., Bengtsson, H., & Bahr, R.** (2021). Injury rates decreased in men's professional football: an 18-year prospective cohort study of almost 12,000 injuries sustained during 1.8 million hours of play. *British Journal of Sports Medicine*, 55(19), 1084–1092. doi:10.1136/bjsports-2020-103159
7. **Eirale, C., Tol, J. L., Farooq, A., Smiley, F., & Chalabi, H.** (2013). Low injury rate strongly correlates with team success in Qatari professional football. *British Journal of Sports Medicine*, 47(12), 807–808. doi:10.1136/bjsports-2012-091040
8. **Chamari, K., Rekik, R. N., Chaabane, M., Chebbi, S., Daoud, R., Eirale, C., Schumacher, Y. O., Tabben, M., & Bahr, R.** (2025). Evolution of injury burden in Qatari professional football: eight-season data from the Aspetar injury and illness surveillance programme. *Biology of Sport*, 42(1), 201–209. doi:10.5114/biolsport.2025.139089
9. **López-Valenciano, A., Ruiz-Pérez, I., García-Gómez, A., Vera-García, F. J., De Ste Croix, M., Myer, G. D., & Ayala, F.** (2020). Epidemiology of injuries in professional football: a systematic review and meta-analysis. *British Journal of Sports Medicine*, 54(12), 711–718. doi:10.1136/bjsports-2018-099577
10. **Ekstrand, J., Braun, M., Waldén, M., Spreco, A., van Zoest, W., Hägglund, M., & Gauffin, H., et al.** (2025). Higher level of communication between the medical staff and the performance staff is associated with a lower hamstring injury burden: a substudy on 14 teams from the UEFA Elite Club Injury Study. *BMJ Open Sport & Exercise Medicine*, 11(1), e002182. doi:10.1136/bmjsem-2024-002182
11. **Julian, R., Page, R. M., & Harper, L. D.** (2021). The effect of fixture congestion on performance during professional male soccer match-play: a systematic critical review with meta-analysis. *Sports Medicine*, 51(2), 255–273. doi:10.1007/s40279-020-01359-9
12. **Carling, C., Gregson, W., McCall, A., Moreira, A., Wong, D. P., & Bradley, P. S.** (2015). Match running performance during fixture congestion in elite soccer: research issues and future directions. *Sports Medicine*, 45(5), 605–613. doi:10.1007/s40279-015-0313-z

Tracking, measurement and validation

1. **Scott, M. T. U., Scott, T. J., & Kelly, V. G.** (2016). The validity and reliability of global positioning systems in team sport: a brief review. *Journal of Strength and Conditioning Research*, 30(5), 1470–1490. doi:10.1519/JSC.0000000000001221
2. **Pino-Ortega, J., Oliva-Lozano, J. M., Gantois, P., Nakamura, F. Y., & Rico-González, M.** (2022). Comparison of the validity and reliability of local positioning systems against other tracking technologies in team sport: a systematic review. *Proceedings of the Institution of Mechanical Engineers, Part P*, 236(2), 73–82. doi:10.1177/1754337120988236
3. **Rico-González, M., Pino-Ortega, J., Nakamura, F. Y., Moura, F. A., Rojas-Valverde, D., & Los Arcos, A.** (2020). Past, present and future of the technological tracking methods to assess tactical variables in team sports: a systematic review. *Proceedings of the Institution of Mechanical Engineers, Part P*, 234(4), 281–290. doi:10.1177/1754337120932023
4. **Buchheit, M., Allen, A., Poon, T. K., Modonutti, M., Gregson, W., & Di Salvo, V.** (2014). Integrating different tracking systems in football: multiple camera semi-automatic system, local position measurement and GPS technologies. *Journal of Sports Sciences*, 32(20), 1844–1857. doi:10.1080/02640414.2014.942687
5. **Linke, D., Link, D., & Lames, M.** (2020). Football-specific validity of TRACAB's optical video tracking systems. *PLOS ONE*, 15(3), e0230179. doi:10.1371/journal.pone.0230179
6. **Taberner, M., O'Keefe, J., Flower, D., Phillips, J., Close, G., Cohen, D. D., Richter, C., & Carling, C.** (2020). Interchangeability of position tracking technologies; can we merge the data? *Science and Medicine in Football*, 4(1), 76–81. doi:10.1080/24733938.2019.1634279
7. **Oliva-Lozano, J. M., & Muyor, J. M.** (2022). Understanding the FIFA quality performance reports for electronic performance and tracking systems: from science to practice. *Science and Medicine in Football*, 6(3), 398–403. doi:10.1080/24733938.2021.1984557

Artificial intelligence, analysis and modelling in football

1. **Wang, Z., Veličković, P., Hennes, D., Tomašev, N., Prince, L., Kaisers, M., & Bachrach, Y., et al.** (2024). TacticAI: an AI assistant for football tactics. *Nature Communications*, 15, 1906. doi:10.1038/s41467-024-45965-x
2. **Tuyls, K., Omidshafiei, S., Muller, P., Wang, Z., Connor, J., Hennes, D., Graham, I., & Spearman, W., et al.** (2021). Game plan: what AI can do for football, and what football can do for AI. *Journal of Artificial Intelligence Research*, 71, 41–88. doi:10.1613/jair.1.12505
3. **Liu, S., Lever, G., Wang, Z., Merel, J., Eslami, S. M. A., Hennes, D., Czarnecki, W. M., & Tassa, Y., et al.** (2022). From motor control to team play in simulated humanoid football. *Science Robotics*, 7(69), eabo0235. doi:10.1126/scirobotics.abo0235
4. **Decroos, T., Bransen, L., Van Haaren, J., & Davis, J.** (2019). Actions speak louder than goals: valuing player actions in soccer. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1851–1861. doi:10.1145/3292500.3330758
5. **Bunker, R., & Susnjak, T.** (2022). The application of machine learning techniques for predicting match results in team sport: a review. *Journal of Artificial Intelligence Research*, 73, 1285–1322. doi:10.1613/jair.1.13509
6. **Goes, F. R., Meerhoff, L. A., Bueno, M. J. O., Rodrigues, D. M., Moura, F. A., Brink, M. S., & Lemmink, K. A. P. M., et al.** (2021). Unlocking the potential of big data to support tactical performance analysis in professional soccer: a systematic review. *European Journal of Sport Science*, 21(4), 481–496. doi:10.1080/17461391.2020.1747552
7. **Mead, J., O'Hare, A., & McMenemy, P.** (2023). Expected goals in football: improving model performance and demonstrating value. *PLOS ONE*, 18(4), e0282295. doi:10.1371/journal.pone.0282295

Talent, academy and organisational practice

1. **Helsen, W. F., Van Winckel, J., & Williams, A. M.** (2005). The relative age effect in youth soccer across Europe. *Journal of Sports Sciences*, 23(6), 629–636. doi:10.1080/02640410400021310
2. **Doyle, J. R., & Bottomley, P. A.** (2018). Relative age effect in elite soccer: more early-born players, but no better valued, and no paragon clubs or countries. *PLOS ONE*, 13(2), e0192209. doi:10.1371/journal.pone.0192209
3. **Bradley, B., Johnson, D., Hill, M., McGee, D., Kana-ah, A., Sharpin, C., Sharp, P., Kelly, A., Cumming, S. P., & Malina, R. M.** (2019). Bio-banding in academy football: player's perceptions of a maturity-matched tournament. *Annals of Human Biology*, 46(5), 400–408. doi:10.1080/03014460.2019.1640284

4. **Moran, J., Cervera Raga, V., Jones, B., Tallent, J., Howe, L., Clemente, F. M., Arede, J., & Freeman, P.** (2024). The identification and development of young talent in Spanish soccer academies: a 10-year multi-study follow-up. *International Journal of Sports Science & Coaching*, 19(5), 1984–1994. doi: 10.1177/17479541241254767

5. **Balassiano, V., & Bullough, S.** (2021). Measures of academy productivity in English Championship clubs. *Team Performance Management*, 27(3/4), 332–350. doi:10.1108/TPM-09-2020-0083

6. **Lolli, L., Bauer, P., Irving, C., Bonanno, D., Höner, O., Gregson, W., & Di Salvo, V.** (2025). Data analytics in the football industry: a survey investigating operational frameworks and practices in professional clubs and national federations from around the world. *Science and Medicine in Football*, 9(2), 189–198. doi:10.1080/24733938.2024.2341837

Governance, explainability, regulation and assurance

1. **European Parliament and Council of the European Union** (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, OJ L, 2024/1689, 12 July 2024. In force 1 August 2024, applying in stages.

2. **Information Commissioner's Office & The Alan Turing Institute** (2020). *Explaining Decisions Made with AI.* Wilmslow, ICO. Guidance in three parts.

3. **United Kingdom General Data Protection Regulation**, Article 9. Processing of special categories of personal data, including data concerning health and biometric data used for unique identification.

4. **Data (Use and Access) Act 2025** (c. 18), sections 80 and 81. Section 80 replaces Article 22 of the UK GDPR with new Articles 22A to 22D on automated decision-making and contains no child-specific provisions; section 81 inserts the children's higher protection duty into Article 25. Royal Assent 19 June 2025; both commenced 5 February 2026 by the Data (Use and Access) Act 2025 (Commencement No. 6 and Transitional and Saving Provisions) Regulations 2026, S.I. 2026/82, made 29 January 2026.

5. **Rudin, C.** (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. doi:10.1038/s42256-019-0048-x

6. **Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S.** (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. doi:10.1126/science.aax2342

7. **Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T.** (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. doi:10.1145/3287560.3287596

8. **Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K.** (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. doi:10.1145/3458723

9. **National Institute of Standards and Technology** (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1. Gaithersburg, NIST. Published 26 January 2023.

10. **FIFPRO & FIFA** (2022). *Charter of Player Data Rights.* Hoofddorp, FIFPRO. Published 19 September 2022, setting out eight rights over performance and personal data.

11. **International Organization for Standardization** (2023). *ISO/IEC 42001:2023. Information technology, artificial intelligence, management system.* Geneva, ISO.

12. **Bullock, G. S., Mylott, J., Hughes, T., Nicholson, K. F., Riley, R. D., & Collins, G. S.** (2022). Just how confident can we be in predicting sports injuries? A systematic review of the methodological conduct and performance of existing musculoskeletal injury prediction models in sport. *Sports Medicine*, 52(10), 2469–2482. doi:10.1007/s40279-022-01698-9

13. **Bahr, R.** (2016). Why screening tests to predict injury do not work, and probably never will: a critical review. *British Journal of Sports Medicine*, 50(13), 776–780. doi:10.1136/bjsports-2016-096256

14. **Impellizzeri, F. M., Woodcock, S., Coutts, A. J., Fanchini, M., McCall, A., & Vigotsky, A. D.** (2021). What role do chronic workloads play in the acute to chronic workload ratio? Time to dismiss ACWR and its underlying theory. *Sports Medicine*, 51(3), 581–592. doi:10.1007/s40279-020-01378-6

15. **Impellizzeri, F. M., Marcora, S. M., & Coutts, A. J.** (2019). Internal and external training load: 15 years on. *International Journal of Sports Physiology and Performance*, 14(2), 270–273. doi:10.1123/ijspp.2018-0935

16. **Bittencourt, N. F. N., Meeuwisse, W. H., Mendonça, L. D., Nettel-Aguirre, A., Ocarino, J. M., & Fonseca, S. T.** (2016). Complex systems approach for sports injuries: moving from risk factor identification to injury pattern recognition, narrative review and new concept. *British Journal of Sports Medicine*, 50(21), 1309–1314. doi:10.1136/bjsports-2015-095850

17. **International Olympic Committee** (2024). *Olympic AI Agenda.* Lausanne, IOC. Launched 19 April 2024.

18. **Cowley, E. S., Olenick, A. A., McNulty, K. L., & Ross, E. J.** (2021). "Invisible sportswomen": the sex data gap in sport and exercise science research. *Women in Sport and Physical Activity Journal*, 29(2), 146–151. doi:10.1123/wspaj.2021-0028

WHERE THIS GOES NEXT

An elite football performance intelligence workshop

The argument in this paper is testable, and the fastest way to test it is on one squad group for 90 days with measures agreed in advance.

Athleet.AI works with clubs and national associations to run a structured workshop and a defined pilot using the design in Chapter 12. The platform is a configurable performance intelligence and coaching-assistant layer that ingests approved sources, retains the organisation's own methodology, and produces outputs that carry their sources, their confidence and their alternatives.

What the workshop covers

- Where the club sits on the maturity model in Chapter 13, assessed with the people who would have to change.
- The medical boundary, agreed with the clinicians before anything is connected.
- One squad group, one set of measures, and a written definition of a disappointing result.
- A first phase deliberately without artificial intelligence, so the gain from joining records is separated from the gain from the technology.
- An honest count of the prompts staff found useless alongside those they found useful.

Athleet.Ai

The Augmented Touchline · Paper Six of Seven · Athleet.AI Performance Intelligence Whitepaper Series
Paul Forrest · February 2026 · Version 1.0