The Systemic Architecture of High Performance: An Integrated Philosophy

This document is the philosophical foundation of Performance Systems Science: the principles that shape how I think about, lead, and build performance systems. Where a principle is drawn from a named source, that source is cited so the idea is never presented as originating here when it didn't.

Chris Cothern · chriscothern.com · November 2026

1. The Foundational Mindset: Sport as a Complex Adaptive System (CAS)

In elite sport, the ultimate competitive advantage is no longer the mere acquisition of data, but the strategic architecture of the mental model through which that data is processed. To succeed in the modern era, organizations must abandon the mechanistic "bicycle" analogy: the reductionist fallacy that a system can be taken apart, optimized in isolation, and reassembled for a predictable outcome. As Scott McLean observes, sport is a **Complex Adaptive System (CAS)** defined by non-linearity; the components cannot be understood outside of their interactions. Pushing a pedal on a bike results in a linear chain of movement, but in a CAS, behavioral properties are **emergent**, arising from the constant, shifting interactions between human actors (athletes, coaches) and non-human artifacts (GPS, technology, environments).

This complexity necessitates an **"Up-and-Out" methodology**. Rather than focusing solely "down and in" on the athlete's internal physiology, we adopt the perspective of Bittencourt et al. (2016), viewing injury and performance as systemic outcomes. For example, a cluster of hamstring injuries is rarely a failure of individual strength training; it is more often a manifestation of higher-level system pressures such as league-wide fixture congestion, travel density, and organizational scheduling. When we ignore these broader determinants, we succumb to the "root cause" fallacy. True resilience is found by "dancing with complexity": staying adaptable to what the system tells us rather than clinging to rigid, data-backed certainty that may arrive too late.

2. The Coder-Practitioner Symbiosis: Reclaiming Cognitive Capacity

Elite performance departments frequently fall into the "Data Janitor" trap, accumulating significant **technical debt** as high-level practitioners spend their hours cleaning spreadsheets and managing ETL (Extract, Transform, Load) pipelines. To scale, we must separate "the complicated" from "the complex." As Ian Kadish notes, the modern "arms race" for shiny new technology often leads to decreased efficiency because the system lacks the architecture to house it.

The **Coder-Practitioner Symbiosis** is the mechanical solution to this inefficiency:

- **The Complicated (The Coder):** Automated via systems engineering (Anand) and data architecture (Kerth). This includes ETL pipelines, data cleaning, and AMS database maintenance.
- **The Complex (The Practitioner):** Irreplaceable human elements including context, clinical intuition, and the "human-to-human" interface.

The split is lopsided. Mapping Premier League and NBA departments task by task, Clancy and Peric (2026) found that in one medical department they counted 94 recurring tasks, of which only 5 were true judgment calls (the pick, the offer, return to play, resting a key player, in-game decisions); the other 89 assembled what those calls rest on. The goal is to hand those hours back to the five calls. This is a practitioner mapping, not published research.

A useful illustration, one Scott McLean uses, is the **Challenger Disaster (O-rings)**: engineers' concerns about launching in the cold were overruled because they could not prove the launch was unsafe (Rogers Commission, 1986; Vaughan, 1996). In our architecture, a coach's or clinician's observed "drift" in an athlete triggers a check immediately; it does not have to wait for data to confirm a catastrophe that has already occurred. This symbiosis is the prerequisite for the cultural emergence of **Leadership Density**, returning the coach to the floor and maximizing the system's focus on actionable insights.

Intuition as trigger, not verdict. Human judgment plays two different roles here, and they belong at different points. As a *trigger*, intuition comes first: "something looks off" is reason enough to look closer. As a *verdict*, it comes last: in a structured decision such as return to play, the separate assessments are completed before anyone asks what their gut says, so instinct informs the call without anchoring it (Kahneman, Sibony & Sunstein, 2021). Skilled intuition is trustworthy only where the environment is regular enough to learn and the practitioner gets timely feedback on whether they were right (Kahneman & Klein, 2009). The closed loop in Section 4 supplies that feedback, but outcomes in sport are noisy feedback: a sound call can be followed by a bad result, and a weak one by a good result. So review asks two separate questions (Section 4, *Continuous Reassessment*).

This is how I treat my own judgment too. I see patterns across departments; that is what a background in physical therapy, strength and conditioning, and sports science gives me. But I don't trust my pattern recognition on its own; I build systems that check it. That is what Performance Systems Science is.

Where AI Fits

AI belongs on the complicated side of this split. It is a tool for the coder's work, never a substitute for the practitioner's judgment, and it follows six rules:

- **Rule first, model second, language model last.** If a decision can be written as a fixed rule (a threshold, a lookup, a known calculation), it is plain code: cheaper, faster, and the same

answer every time. A statistical model is used where a pattern has to be learned from data, and a large language model only where the input is genuinely unstructured, such as free-text notes or a coach's question in plain English (Ajzenstadt, 2026b). The structured record is the source of truth: the model reads it and never rewrites it, it speaks only once quality checks show the data can support a conclusion, and if the model and the database disagree, the database wins (Pryor, 2026).

- **AI informs; a named person decides.** A model can surface a flag, draft a summary, or put a number into plain English. It does not clear an athlete, change a threshold, or prescribe a treatment; a named owner does (Section 5).
- **Grade the model in the open.** Report how often its flags were right, compare it against simple rules of thumb, and say "not enough data" rather than guess. Injury prediction is hard: injuries are rare, so even a good model raises many false alarms, and acting on a flag changes the outcome it was trying to predict (Carbone, 2026). No accuracy number is quoted that has not been checked. The same bar applies to a model that is bought: before purchase, ask for its sensitivity, false-alarm rate, injuries missed, how far ahead it flags, and independent prospective validation on athletes like ours, for the same injury and time window, plus evidence that it improves a defined decision (Blin, 2026, and Kerin in the comments). And the fair comparison is not perfection: does staff plus the model make fewer mistakes than staff alone, with a named person still answerable (Cox, in comments on Blin, 2026)?
- **People write the specification.** Before anything is built, the person accountable writes what it must never do and how everyone will know it is done, and tests check against that, not against whatever the code happens to do. An AI that writes its own specification is checking its own homework.
- **Agents get their own accountable identity.** An AI agent that acts on a system runs under its own account, with a named owner, limited permissions, keys that can be revoked, and logs that show which actions it took, never under a person's login (Ajzenstadt, 2026c).
- **Evaluate before adopting.** New AI tools go through *Choosing What to Build* (Section 4), and a health-facing assistant is held to the bar of published work such as Heydari et al. (2026): defined tasks, automated scoring, and human expert review before anyone calls it done.

The test of anyone building with AI, including me, is the system around the tool rather than the tool: which decision it changes, who owns it when it breaks, what it was deliberately not allowed to do on its own, what broke and how the system learned from it, and who still uses it without the builder (Cheng, 2026, and the practitioners responding to him).

3. Strategic Pillars and Leadership Density

High-performance alignment is the antidote to the "accidental promotion" bottleneck, where technical experts are elevated to leadership without a framework for managing volatility. We utilize the **"Pointing Arrows"** concept (Kadish), ensuring that Medical, S&C, Science, Psychology, and Nutrition operate under a transdisciplinary vision.

As staff numbers grow in the modern sports "arms race," we face a **quadratic explosion** of communication lines. If a department grows linearly, the number of possible communication lines grows quadratically: N staff have N(N-1)/2 possible pairwise lines, so 10 staff have 45 and 20 staff have 190. To prevent siloed expertise and "Fixes that Fail," we implement David Clancy's model of **Leadership Density**. This is not an increase in "bosses," but a distribution of the capacity to think critically under volatility across the entire staff. This culture is sustained by an **Idea Meritocracy** and a **Daily Operational Rhythm** that transforms communication from a buzzword into a performance infrastructure. By increasing the collective "Hive Mind," we minimize the decision-making bottlenecks traditional hierarchies create.

This is not only a practitioner observation. Stewart et al. (2024) treat elite performance support teams as **multiteam systems** (Zaccaro et al., 2012) and, from focus groups with 18 practitioners across Premier League football, Formula One, and Summer and Winter Olympic sport, identify four characteristics of effective teams: team structure, team member attributes, shared mental models, and social capital. Their participants describe the same size problem ("as capacity increases, you increase the chance of there being more noise"), prefer a flat "circle, not ladder" structure to top-down hierarchy, and name trust, and the psychological safety it creates, as a leading promoter of effectiveness: "without trust and respect, challenge becomes conflict." These are practitioner perceptions, not measured effects, but they are a peer-reviewed account of the human layer this section describes. Two of their findings sharpen this philosophy:

- **Shared at the objective level, distributed at the information level.** Stewart et al. frame effectiveness around *shared* mental models; Section 5 adopts *distributed* situation awareness, in which no one holds the whole picture. Performance Systems Science holds both, at different levels: objectives, roles, and vision are shared by everyone (the Pointing Arrows), while information is distributed, each role holding the piece it needs and the system, not any one person, connecting them. Their participants' wish for a flat structure that still routes information through one designated person is the same tension; the answer here is to move that filtering role from a person into the architecture (Section 4, *Reduce Noise*). One limit applies (King et al., 2024b): shared goals and distributed information produce coordination, not joint problem-solving. They are enough for simple problems; complex problems need the relevant disciplines deliberately convened (Section 4).
- **No agreed way to judge a support team.** Participants could name no common criterion for evaluating support team effectiveness. Decision Latency (Section 5) is this philosophy's answer: a measurable, system-level criterion that does not depend on any one department's outcomes.

Stewart et al. describe the human layer of the support team, and technology and data do not appear in their framework. Performance Systems Science designs the information layer that feeds that human layer (Section 1: non-human artifacts are actors in the system too).

The link between how departments communicate and what happens to athletes is not only a perception. Across 36 elite European football clubs, teams whose medical staff rated

communication with the head coach as low had a higher injury burden and lower availability than teams rating it moderate or high (match availability 82% vs 88%; Ekstrand et al., 2019). The design is observational, so it shows an association, not a cause, but it is the clearest outcome-linked evidence for this philosophy's starting claim that the hardest problems sit between departments.

Who sees what. King et al. (2024b), in focus groups with 28 practitioners, found that information inside support teams is often treated as a commodity: practitioners decide personally what to share, tend to under-share, and privileged information becomes a source of conflict, while medical confidentiality is a genuine limit on what can be shared at all. The information layer therefore needs an explicit access policy: routine performance information is shared by default, so it cannot be held as leverage, and clinical and mental-health information is protected by role. Sharing decisions are made once, in the design, rather than person by person, every day.

Vigor and rigor. Clancy (2026) reads an environment through two lenses. *Rigor* is the standard beneath the energy: clear roles and expectations, follow-through, honest challenge, decisions tied to a clear way of playing and working, and a team that learns when things go wrong. *Vigor* is the life of the place: engagement, curiosity, care, people noticing one another, and warmth without becoming comfortable at the expense of ambition. "Vigor without rigor can feel exciting but inconsistent. Rigor without vigor can feel organised but lifeless." The architecture in this document mostly supplies rigor. Vigor decides whether people actually use it: a well-designed information layer in a fearful or lifeless building still fails. Decision Latency measures how fast the system decides; it cannot tell you whether the people in it will speak up. For that, Clancy points to ordinary moments, where the signals are observable:

- **Transitions.** Between meetings, sessions, and shifts, do people share information, prepare for what's next, and close the loop? Handoffs are where distributed information is most often lost.
- **The quality of challenge.** Can someone disagree, ask for help, or raise a concern without being dismissed?
- **Respect across roles.** Is attention shared across the group, or reserved for the most senior or visible people?
- **The response to mistakes.** Does the team get curious about what happened, or move quickly to blame and protect?
- **Words against reality.** If the organization says it values collaboration, can you see people collaborating in the work?
- **What people are learning.** Can they name something the team changed because of feedback, reflection, or new information?

One generous gesture may be one person's character; a repeated pattern says something about the environment, the same logic as reading a cluster of injuries as a system pressure (Section 1). And the read is a starting point, not a verdict: like intuition (Section 2), it is tested against the

routines, decisions, and behaviors actually observed. This is a practitioner framework, not measured evidence.

Develop staff as deliberately as athletes. The athlete gets a baseline, a plan, and criteria for advancing; staff deserve the same. Thomas (2026), reading the four clubs left in the 2026 MLB Championship Series, finds pathways from intern to leadership, promotion from within, and responsibility given slightly ahead of readiness. His sharpest test is one question with a name and a date: what was the last decision that changed because someone less senior disagreed, and what happened to that person afterwards? Values that cannot be tested against a promotion, a departure, or a disagreement are slogans. This is observational reporting, which Thomas states plainly; it is consistent with this approach and does not prove it.

4. The Operational Framework: Closed Loops and the Michigan Model

Systems fail when they are "Open Loop": where data is ingested but never informs a structured feedback mechanism. We operationalize the **William Graham Closed-Loop Framework** to ensure every data point is tethered to action.

The Graham Framework: Strategic Contrast

	Reductionist Thinking (To Diagnose)	**Systems Thinking (To Contextualize)**
Objective	Isolate specific bottlenecks and track measurable KPIs.	Account for sleep, travel, culture, and sport demands.
Focus	Identifying the "broken component" for immediate repair.	Understanding how environment influences the component.
Process	Building objective and subjective readiness markers.	Acknowledging that athletes do not perform in a vacuum.

We map this 4-stage loop onto the **Michigan Sub-Departmental Model** to clarify functional ownership:

1. **Reduce Noise (Analytics, The Brain):** Identifying clear patterns and filtering distractions, including deciding how each issue should be handled (decision triage, below). This mirrors the first step of Ryan King's Blended Intelligence framework (Problem, Expertise, Team, Solution; King, 2026), and the system supports the triage because practitioners do not reliably make it themselves: in a survey of 111 practitioners, of those reporting complex problems, 37% still took a procedural approach and 42% a creative one, and the authors found no clear pattern overall (King et al., 2024a; descriptive survey data, not tested effects).
2. **Target Solutions (Performance Science, The Core):** Programming interventions based on the current context.
3. **Build in Environment (Performance Science, The Core):** Implementing solutions on the field/gym, not in the lab.

4. **Continuous Reassessment (R&D, The Future):** Validating methods and refining the loop via rigorous feedback. Review asks two questions, kept separate: *was it a sound call given what was known at the time?* and *what happened?* "A defensible process may precede an adverse result, while a weak process may precede success" (Rebelo et al., 2026), so judging decisions by outcome alone teaches the loop the wrong lessons. Review considers what was reasonably known, how uncertainty was handled, and whether the action followed from that information. That review needs four records for every decision: what we saw, what we decided, **what we expected to happen**, and what happened. The third is the one almost nobody writes down, and without it no outcome can score the call. A return-to-play clearance logged as "back in 9 weeks, no recurrence, full output by week 3" can be learned from; a bare "cleared" cannot (Clancy & Peric, 2026). The written expectation is also the timely feedback that turns intuition into trustworthy expertise (Section 2).

Decision Triage

Not every issue needs the whole team. Following Rebelo et al. (2026), the person first responsible for an issue triages it on six dimensions: **urgency** (how much time there is), **uncertainty** (whether more information or interpretation could change the decision), **risk and consequences**, **reversibility** (a decision that is hard to undo deserves more scrutiny before commitment), **interdependence** (whether contributions can stay separate), and **stakeholder differences** (where preferences, disagreement, or authority need explicit attention). These are weighed qualitatively, not scored: one serious condition, or several interacting ones, can justify broader input. The triage selects the least burdensome route that is still safe:

- **Discipline-led:** a bounded problem within one professional scope, with a clear pathway, limited consequences, and no material disagreement. The authorized practitioner or coach decides.
- **Consultation:** the problem and decision owner are clear, but specific advice is needed from another domain. The owner asks, considers, decides, and communicates, without convening everyone.
- **Integrated:** shared framing is needed, consequences are serious or hard to reverse, recommendations conflict, or implementation crosses roles. The team convenes, using the decision hygiene below.
- **Escalation, from any route:** when a decision exceeds the group's authority or expertise, or a health, safety, or professional obligation requires it. Immediate protective action never waits for deliberation.

Integration has costs as well as benefits: under-integration leaves complex decisions poorly informed, and over-integration adds delay, meeting burden, and unclear responsibility. Integrated work therefore uses **minimum sufficient integration**: the smallest relevant group, the decision-critical information, a named owner, existing channels, and a time limit or stopping rule. Information gathering stops when further input is unlikely to change the choice within the

time available. This is the operational form of the simple-versus-complex distinction (Cynefin: Snowden & Boone, 2007; Greenberg & Clubb, 2021).

Two Kinds of Noise

Reduce Noise applies to people as well as data. **Data noise** is variability in the measurements, filtered by personal baselines and Alert Yield Analysis. **Judgment noise** is variability in the calls: two clinicians looking at the same athlete and reaching different decisions, or the same clinician deciding differently on a Friday afternoon than on a Monday morning (Kahneman, Sibony & Sunstein, 2021). It is usually larger than anyone expects, and it stays invisible because calls are rarely placed side by side: in one insurance company the authors studied, leaders expected two underwriters pricing the same case to differ by about 10%; the median difference was 55%.

The remedy the authors call **decision hygiene** becomes the procedure for the integrated route, translated to return to play by Diamond (2026):

- **Start from the base rate.** How do athletes like this one usually do? Then adjust for the individual (Section 6), rather than starting from the individual alone.
- **Judge the parts separately.** The injury, the sport demand, and the risk are assessed independently before they are combined.
- **Decide alone, then compare.** Each discipline makes its own call before the group discusses, then the calls are combined. This counters the pattern King et al.'s participants described, meetings where "a lot of people just end up agreeing with one another" (2024b, p. 162), and reduces anchoring on the most senior voice in the room.
- **Gut last.** Intuition as verdict comes after the structured pieces (Section 2).
- **Compare, don't just score.** "Is he closer to his pre-injury baseline than last week?" is easier to answer well than "Is he a 7 out of 10?" WAIMS already works this way: every signal is scored against the athlete's own baseline, and return-to-competition criteria are percentages of the athlete's own pre-injury values (see *In Practice*).

Judgment noise can also be measured: give the same case to several staff independently and record how far apart their calls are (a "noise audit"). That is a direction for a future prototype, not a current feature, but it would sit alongside Decision Latency as a second system-level measure: one for how fast the system decides, one for how consistently. The evidence for noise comes from insurance, law, and medicine rather than sport, so it is cited here as a framework that transfers, not as sport-specific evidence.

Choosing What to Build

Decision triage routes decisions about athletes. A performance department also decides which projects, tools, and vendors to take on, and the usual problem is not a shortage of ideas but a list of twenty that nobody can execute. Adapted from Ajzenstadt (2026a), four questions cut that list down:

- **Would the athlete or coach notice?** Start where the decision they live with gets better (translation is the measure of the output, Section 5), ahead of internal efficiency for its own sake.
- **Will someone put their name on one number?** Every project gets a named owner and one measure it is meant to move. If nobody will own it, it comes off the list.
- **Can it run as a one-block pilot?** Roughly a month, or one training block, in a sandbox first. A project that needs a new platform before it can be tested comes off for now, with one exception below.
- **What happens when it's wrong?** Start where a person can catch the error before it reaches the athlete. This is the reversibility test from decision triage, applied to the tool rather than the call.

The exception: automating the complicated (Section 2) sometimes means the plumbing is the innovation. Infrastructure work stays on the list when it shortens decision latency for a named owner, and the owner and the number still apply. Check first whether connecting tools you already own solves the problem; most delay comes from tools that don't talk to each other, not from missing ones. Building it yourself is a separate question. Own the intelligence layer: the decision rules, triage, baselines, models, and translation, because that is the work that compounds. Buy commodity infrastructure (storage, matching athlete IDs across vendors, security, uptime) where a good option exists, but keep the definitions and quality checks in house even when the database is bought, since they decide what the system is allowed to say. Build plumbing only where nothing suitable exists or where it directly cuts decision latency (Siscoe, 2026, a practitioner view from a vendor). Staff at MLS clubs describe the same hybrid from the buyer's side: a commercial athlete management system with in-house planning and analysis tools connected to it, and security as a reason to keep sensitive systems in house (Blin et al., 2026). The list usually ends at two or three. Pick one, ship it, and review it with the two-question review before starting the next: did it change the decision, and was adopting it a sound call given what was known at the time? Then one more test, a block later: is the named owner still using it without the builder, and could they spot and correct a bad output on their own? Building is the easy half; getting a system adopted by the people whose work it changes is the other half (Bates and Hawkinson, in comments on Cheng, 2026). The first pilot also leaves the scaffolding the others will need. This is a practitioner heuristic from outside sport, cited as such.

Order matters across pilots too. Clancy and Peric (2026) describe four steps, each taken on the evidence of the one before: connect the data (one athlete ID, one clock, one vocabulary); link every decision to its outcome; hand routine work to agents, with a person releasing it; and only then predict, on the club's own history. Prediction comes last because nothing can grade it until the first two steps exist. On complex problems, a club's own history informs prediction but does not hand over a recipe, because conditions change (Snowden & Boone, 2007); the written expectations are there to calibrate judgment, not to replace it. The quick wins are unglamorous: a return-to-play pack assembled overnight, or a post-game pack ready before the plane lands.

5. Decision Architecture and Distributed Performance

At the highest level of sport, opportunity is perishable. **Decision Latency**, the space between knowing and acting, is the primary KPI of our entire system. Speed is not a product of "urgency," which breeds panic; it is a product of **Architecture** (Wagle).

Decision Latency is measured by triage route, because a single number would reward rushing complex problems:

- **Discipline-led and consultation:** time from flag to action. Lower is better.
- **Integrated:** time from flag to convening the right people. The decision itself is allowed to take longer.

One guardrail keeps the measure honest: "urgency changes the cadence, not the minimum expertise or authority required" (Rebelo et al., 2026). A time-critical decision can use a brief huddle or a direct conversation instead of a meeting, but it never gets faster by skipping a required expert or owner.

Latency measures time; it does not show the path. So once a block, run a route audit: trace one finding from the analytics group to the athlete or coach who had to change something because of it, and count the genuine conversations it passed through on the way (Thomas, 2026). A fast response with no conversation may mean nobody translated the finding; a long chain of hand-offs means the routing is broken. At the club with baseball's best 2026 record, "a finding becomes a conversation with a person before it becomes a change to a pitch" (Thomas, 2026, on the Milwaukee Brewers).

Practitioners name the cost on both sides: discussion that slows everything ("the boat can go a bit slow because everybody's having a chat about who's going to be doing what when," King et al., 2024b, p. 158), and waiting for data that will never make the answer certain. The stance here: prepare the car (the athlete, the plan, the information) as well as you can, then drive the trail, bumps and all. No amount of preparation produces a smooth road, so the system is built to support adjustment on the way rather than certainty before starting; as one practitioner put it, "It's only the decision until it's not, and then we change" (King et al., 2024b, p. 164).

Every flag type also has a named decision owner. Rebelo et al. (2026) separate four roles that are often blurred: **providing information**, **recommending**, **holding final authority**, and **implementing**. One person can hold several, but they are named before options are chosen, and "a recommendation remains advisory unless its author also holds the relevant authority." Where the team does not agree, material dissent is recorded and travels with the options to the owner. Speed without clear ownership produces the question King et al.'s participants asked, "Who's actually gonna get the finger wagged at them if it goes wrong?" (2024b, p. 163).

The athlete has a seat in the decision. Athletes are actors in the system (Section 1), and some decisions are theirs to share. Preference-sensitive choices about treatment, rehabilitation, return to sport, or health and performance trade-offs should normally use shared decision-

making: options, benefits, harms, and uncertainty explained in plain terms, with the athlete's goals and risk tolerance shaping the deliberation, starting before a preferred option has effectively been settled (Rebelo et al., 2026). Consent, where required, is voluntary, never inferred from silence or performance pressure. Autonomy does not entitle an athlete to demand clearance or selection, and professional authority does not justify overriding a valid refusal.

By establishing a **Distributed Situation Awareness (DSA)**, we ensure the right actors have the right information at the right time. No single specialist holds the "whole" truth; instead, the "Hive Mind" connects partial perspectives to create total system awareness.

The ultimate goal of this decentralized performance model, as proposed by Zach Higginbotham, is **"making the sports scientist unnecessary."** By distributing expertise to the coaches and athletes closest to the action, we move the point of decision-making as far "downstream" as possible. When the coach is equipped with the mental models to interpret GPS and force plate data in real time, the need for a scientific intermediary vanishes, and decision latency reaches its theoretical minimum.

Translation is the measure of the output. Craig Duncan (2026) states the test plainly: "The value of a sports scientist is not measured by the complexity of a spreadsheet. It is measured by the simplicity of the decision it enables." He contrasts the *Technician*, who hoards forty GPS and force-plate metrics and expects the coaching staff to adapt to their terminology, with the *Translator*, who filters the noise and delivers one plain-English recommendation in fifteen seconds, summarised as "Performance = Capacity - Noise." This is the principle behind the WAIMS coach view (four tabs of traffic lights rather than ten tabs of detail) and behind moving decisions downstream. Performance Systems Science adds two qualifications:

- **Build the translation into the system, not into one person.** A single translator filtering "behind closed doors" is one person's judgment, carrying the judgment noise of Section 4 and a single point of failure when that person is unavailable. The filtering role moves into the architecture (Section 3): the coach view is designed once to show the few signals that matter, and the full detail stays available to the staff who need it.
- **Simple must still be honest.** A fifteen-second answer has to carry its uncertainty. "There isn't enough data yet" is a legitimate plain-English answer, which is why WAIMS reports N/A rather than a guess. Simplifying the message must not mean overstating the confidence.

6. Tactical Execution: Dominating the Edge

Elite execution is the final proxy for performance. Our architecture manifests in the "details of the edge," where we maintain **"Undeniable Hospitality."** This is the human interface, the trust and relationship-building, that ensures athletes buy into the rigorous, data-led interventions required for success.

Return-to-Competition (RTC) Framework

We move away from time-based templates toward **criterion-based progressions**. Athletes must clear objective exposures, each measured against their own baseline, before advancing. The measures follow the demands of the sport:

- **Sport-specific output:** for example, throwing velocity in baseball, or jump and change-of-direction capacity in basketball.
- **Running exposure:** sprint distance and speed against GPS or tracking thresholds, to ensure tissue readiness.
- **Competition volume:** progressive exposure to the density of the schedule, such as minutes, back-to-backs, and travel.

Proactive Workload Management

We define **Injury Prevention as Performance Enhancement**. By monitoring "drifts" in recovery markers and fatigue trends before they lead to failure (Dekker), we maximize player availability. This is achieved through **Unique Performance Profiles** rather than templates. "Unique" does not mean ignoring the population: a decision starts from how similar athletes usually do and is then adjusted for this athlete (Section 4, *Two Kinds of Noise*). Profiles utilize:

- **Movement Assessments:** Quantifying functional limiters.
- **Strength Asymmetries:** Targeting force-production gaps.
- **Power Metrics:** Utilizing force plates to track explosive capacity and nervous system state.

In Practice

These principles are not only on paper. WAIMS, a readiness and injury-risk decision-support prototype I built, puts several of them into working software: athletes are judged against their own baselines rather than a population average; a flag leads to a logged action and a recorded outcome, closing the loop; decision latency, the time from a flag to a staff response, is measured; coaches see a few traffic lights while staff keep the full detail; and return to competition is criterion-based, each stage cleared against the athlete's own pre-injury values, with a staff member deciding when to advance. WAIMS is a prototype, not a validated or production system, and it reports "not enough data" rather than guessing. Before WAIMS, I built individual performance plans in R for every athlete in a Power 4 wrestling program, combining blood biomarkers, genetic profiles, cognitive testing and athlete context, and went over them with the full staff.

Frequently Asked Questions

How is this different from an athlete management system or a dashboard? An athlete management system stores data. This is about the decision: who sees what and when, who

owns the call, and whether the action was logged and reassessed. The dashboard is the smallest part.

How would you use AI? Rule first, model second, language model last. AI surfaces, summarizes, and translates; a named person decides. Every model is graded in the open, and no AI agent runs under a person's login (see *Where AI Fits*).

How accurate is it? I don't quote an accuracy number that hasn't been checked. The system grades its own flags against what actually happened and says "not enough data" when it doesn't have enough. The fair test is not perfection but whether staff plus the model make fewer mistakes than staff alone, and I would hold any vendor's model to the same standard.

How would you structure a department? Analytics reduces noise, performance science targets and implements solutions, and research and development reassesses them. The structure is flat, objectives are shared, and information is routed by design.

How do you know whether a support team is working? Decision latency, measured by triage route, plus whether logged interventions worked and whether the calls were sound given what was known at the time.

Where would you start? With one pilot: something the athlete or coach would notice, with one named owner, one number, one training block, and a person who can catch a mistake before it reaches the athlete. Then the next one, in order: connect the data, link each decision to its outcome, hand routine work to AI with a person releasing it, and predict last.

We do not manage athletes; we architect the ecosystems that allow elite performance to emerge as an inevitability.

7. Integrated Bibliography and References

- Ajzenstadt, M. (2026a). Four questions for cutting a list of innovation use cases down to one pilot [X post].
- Ajzenstadt, M. (2026b). Deterministic code, not language-model calls, for decisions that can already be written as a rule [X post].
- Ajzenstadt, M. (2026c). Never give an AI agent your login [X post].
- Anand, A. (Decision Science). Systems engineering and automation in high-performance data pipelines.
- Bittencourt, N. F. N., et al. (2016). *Complex systems approach for sports injuries: Moving from risk factors to pattern recognition.* British Journal of Sports Medicine.
- Blin, J. (2026). If teams are paying for AI injury prediction, how accurate should it be? [LinkedIn post, with practitioner comments including D. Kerin and P. Cox].
- Blin, J., Norman, D., Moseley, C., Messer, M., & Proessl, F. (2026). Inside MLS: how clubs are rethinking performance, AI, athlete management systems and technology for 2027 [Upside Group podcast, summarized on LinkedIn].

- Carbone, L. (2026). The statistical and epistemological limits of injury-prediction modeling [LinkedIn post].
- Cheng, L. (2026). Asking every candidate what they have built with AI [LinkedIn post, with practitioner comments including A. Bates and C. Hawkinson].
- Clancy, D. Concepts on Leadership Density and cognitive capacity in volatile environments.
- Clancy, D. (2026). The feeling of an elite environment: vigor, rigor, and signals to notice [LinkedIn post].
- Clancy, D., & Peric, I. (2026). Most sports teams don't have an AI problem: mapping performance and intelligence departments task by task [LinkedIn post].
- Dekker, S. (2011). *Drift into Failure: From Hunting Causes to Understanding Determinants of System Safety.* Ashgate Publishing.
- Diamond, A. (2026). Decision hygiene for return to play, after Kahneman et al.'s *Noise* [LinkedIn post]. Senior Physical Therapist, Boston Red Sox.
- Duncan, C. (2026). The Technician and the Translator: the value of a sports scientist is measured by the simplicity of the decision it enables [LinkedIn post]. Human Performance Strategist.
- Ekstrand, J., Lundqvist, D., Davison, M., D'Hooghe, M., & Pensgaard, A. M. (2019). *Communication quality between the medical team and the head coach/manager is associated with injury burden and player availability in elite football clubs.* British Journal of Sports Medicine, 53(5), 304-308. PMID 30104210.
- Epstein, D. (2019). *Range: Why Generalists Triumph in a Specialized World.* Riverhead Books.
- Graham, W. The Closed-Loop Framework: Integrating reductionist and systems thinking in athletic performance.
- Greenberg, W., & Clubb, J. (2021). *Why 'best practice' is not always best in sport.* British Journal of Sports Medicine, 55(22), 1249-1250.
- Heydari et al. (2026). *The anatomy of a personal health agent.* Nature, 659, 197-207. https://doi.org/10.1038/s41586-026-11161-8
- Higginbotham, Z. Decentralized Performance: Models for distributing expertise to the point of action.
- Kadish, I. HP Alignment and the hospitality of elite-level execution in Major League Baseball.
- Kahneman, D., & Klein, G. (2009). *Conditions for intuitive expertise: A failure to disagree.* American Psychologist, 64(6), 515-526.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment.* Little, Brown Spark.
- Kerth, S. Strategic Architecture: Building the technical foundations for integrated performance teams.
- King, R. (2026). *Blended Intelligence: A framework for how professionally diverse multi-disciplinary teams solve problems in elite sport.* Professional Doctorate thesis, University of Lancashire. Framework summary: https://blended-intelligence.com/research-and-ideas
- King, R., McHugh, D., Alexander, J., Kiely, J., Yiannaki, C., & Rhodes, D. (2024a). *Multidisciplinary team practitioners working in high performance sport: Skilled intuitive*

'doers' or novel problem-solving innovators. European Journal of Sport Sciences, 3(2), 15-26. https://doi.org/10.24018/sport.2024.3.2.143

- King, R., Yiannaki, C., Kiely, J., Rhodes, D., & Alexander, J. (2024b). *Multi-disciplinary teams in high performance sport, the what and the how: A utopian view or a darker reality.* Journal of Expertise, 7(4). https://www.journalofexpertise.org/articles/volume7_issue4/JoE_7_4_King_etal.pdf
- McLean, S. Systems Thinking, Distributed Situation Awareness (DSA), and the "Bicycle" Mechanistic Fallacy in elite sport. Also the source of the Challenger O-ring example in Section 2 (podcast discussion).
- Pryor, R. (2026). Don't replace the database: put AI on top of it [LinkedIn post].
- Rebelo, A., King, R., Valente-dos-Santos, J., & Thorpe, R. T. (2026). *Integrated decision-making under uncertainty in high-performance sport support systems.* Sports Medicine - Open, 12, 175. https://doi.org/10.1186/s40798-026-01149-z
- Rogers Commission (1986). *Report of the Presidential Commission on the Space Shuttle Challenger Accident.* Washington, DC.
- Siscoe, N. (2026). Own the intelligence layer, buy the infrastructure [X post].
- Snowden, D. J., & Boone, M. E. (2007). *A leader's framework for decision making.* Harvard Business Review, 85(11), 68-76.
- Stewart, P., Fletcher, D., Arnold, R., & McEwan, D. (2024). *Exploring perceptions of performance support team effectiveness in elite sport.* Sport Management Review, 27(2), 300-321. https://doi.org/10.1080/14413523.2023.2284987
- Thomas, S. (2026). *The people behind the wins: what the Rays, Brewers, Dodgers and Guardians do better than most of MLB.* ST Performance written piece, October 2026.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA.* University of Chicago Press.
- Wagle, J. P. The architecture of decision latency and idea meritocracy in professional sport.
- Zaccaro, S. J., Marks, M. A., & DeChurch, L. (Eds.) (2012). *Multiteam systems: An organization form for dynamic and complex environments.* Routledge.