

Be-Healthy.AI

Live EPL Monitoring Shows Meaningful Separation of Injury Exposure

Case Study: Application of the Be-Healthy.AI Early-Warning System in the 2025/26 English Premier League Season

Publication Date: 11 March 2026.

Players are anonymised for publication.

Bottom line: In live, out-of-sample EPL monitoring, the Be-Healthy.AI system is most useful as an attention-ranking layer. It surfaces player-level escalation patterns ahead of some injury events and separates league-level injury exposure most clearly at the top end of the risk distribution.

Executive Summary

This case study reviews the live operation of the **Be-Healthy.AI Early-Warning System** during the 2025/26 Premier League season. The model was trained only on observations available up to 30 June 2025 and is therefore running out-of-sample on the live season. This specific deployment uses public data only.

Two results matter. First, several player trajectories show the same operational pattern: risk builds ahead of certain injury events and falls during absence or recovery windows. Second, the league-level segmentation is not cosmetic. Players entering the Critical level record injuries **5.6 times more frequently** than the league average in this sample, while the Low level records no injuries.

That combination makes the system useful as a ranking tool rather than a binary oracle. The practical value is not in declaring with certainty that a given player will be injured. It is in **helping performance and medical staff decide** who deserves earlier review, closer monitoring, or targeted preventive action.

1. Scope and Evidence Frame

This document is designed to answer a narrow question: under live conditions, does the system surface usable early-warning patterns at player level and separate injury exposure at league level?

Designed to show	Not designed to claim
Whether selected live player trajectories show recurring pre-injury escalation and recovery resets.	Deterministic injury prediction for every individual player.
Whether current risk classes correspond to materially different observed injury frequencies across the league.	A full technical validation pack with every denominator, cohort rule, and labelling decision reproduced in-line.
How the model should be used operationally: as a triage and attention-ranking layer inside existing workflows.	Replacement of medical, performance, or coaching judgement.

2. Live Monitoring Setup

The Case Study setup is given in the table below.

Competition	English Premier League (EPL)
Live season monitored	2025/26
Training data cut-off	30 June 2025
Deployment mode	Live monitoring on season data not seen during model development
Data boundary	Publicly available data only
Scoring cadence	Daily player-level injury risk estimates
Views used in this document	Prediction evolution (time series) and relative risk index
Publication note	Selected player examples are anonymised for publication

The two chart types serve different purposes. The time-series view shows whether risk is building, flattening, or resetting. The relative risk index expresses current risk against the player's own recent baseline, which makes it easier to spot when the profile has moved outside its normal range.

This publication reports selected trajectories and relative injury frequencies by risk class. It should be read as a case study rather than a full methodological appendix.

3. Flagship Case: Recurrent Muscular Injury Pattern

Player 1 is the clearest single example in this document because the pattern repeats. The player sustained three muscular injuries across the season. In each cycle the risk curve rises materially before the injury, drops during the absence window, and then rebuilds after return.

Takeaway: The value here is not certainty. It is the repeatability of the escalation pattern across multiple injury cycles.




Figure 1. Player 1 showing repeated pre-injury escalation before three muscular injury episodes

That is the kind of behaviour a monitoring system should surface. A single spike can be noise. A recurring build-up, followed by a recovery reset and then another build-up, is operationally different. It gives staff a reason to concentrate attention earlier on a player whose exposure profile is drifting in the wrong direction.

The latest portion of the trajectory matters as well. After the most recent return, the curve is already moving upward again. The point is not to declare the next event inevitable. The point is that the player has re-entered a state that previously preceded injury.

4. Supporting Cases

Player 2 – Accumulation Between Injuries

Player 2 shows a slower, steadier pattern. Between an unidentified injury on 3 October 2025 and a muscular injury on 15 December 2025, the risk curve trends in one direction: upward. The rise is not explosive, but it is persistent, and it reaches the upper bands before the second event.

Takeaway: Not every useful signal is a sharp spike. Some players deteriorate through a slower accumulation that still warrants earlier review.




Figure 2. Player 2 showing a gradual accumulation between two injury events

Player 2 – Relative Risk View of Player 2

The relative risk index tells the same story from a different perspective. Before both injuries, the player's risk sits above the baseline range derived from recent observations. During the injury periods the index falls sharply, reflecting the interruption in exposure, then rises again after return.

Takeaway: The index is useful when the question is not simply 'how high is the score?' but 'how abnormal is this player's current state relative to their own norm?'.




Figure 3. Player 2 showing their relative risk index against the player's own recent baseline

Player 3 – Stable Low-Risk Counterexample

A credible case study needs a counterexample. Player 3 remains in the Low to Medium-Low range across the full season with no sustained build-up. The curve moves, but it never develops the kind of loading pattern seen in the higher-risk cases.

Takeaway: The model is not locked into a state of perpetual alarm. It also identifies players whose observed profile remains comparatively quiet.




Figure 4. Player 3 showing stable low-risk profile without sustained escalation

Player 4 – Persistently Elevated Profile

Player 4 illustrates a different risk state again. This player has recorded five injuries across different injury types and spends much of the season in the upper bands. Here the signal is not a short-lived spike. Elevated exposure appears closer to the player's normal operating state.

Takeaway: Some players do not cycle cleanly between calm and alarm. Persistent elevation is itself a meaningful state that can justify closer monitoring.




Figure 5. Player 4 showing persistently high-risk profile across multiple injury events

5. League-Level Evidence: Risk Class Separation

This is the strongest evidence in the document because it moves beyond anecdotes. The question is simple: once players are grouped by current risk class, do those groups show meaningfully different observed injury frequencies?




Figure 6. Observed injury frequency relative to the league average, by risk class

Yes. Separation is clearest at the ends of the distribution. The Critical band records injuries 5.6 times more frequently than the league average in this sample. The High band sits at 1.8 times the league average. At the other end, the Low band records no injuries.

For reference, the full ordering in this sample is: Low 0.0x, Medium-Low 0.6x, Medium-High 0.9x, High 1.8x, and Critical 5.6x relative to the league average.

The middle bands behave as they should in a ranking system rather than a binary classifier. Medium-Low sits below the global rate, Medium-High sits close to it, and the upper tail carries the clearest concentration of injury exposure. Movement across risk bands is not just a presentational choice. In this sample it corresponds to materially different observed injury exposure.

6. Operational Meaning

For a club, the practical use is straightforward. The model should improve the order in which staff review players, not dictate treatment decisions.

1. Start with the upper bands. High and Critical are where the observed exposure is most concentrated, so those players deserve earlier discussion.
2. Read the signal with local context. Match load, training history, diagnostics, player-reported symptoms, and tactical plans still matter.
3. Use lower bands mainly for de-prioritisation, not for complacency. A lower current score can justify less immediate attention, but it does not eliminate risk.

Used this way, the system can fit into an existing medical or performance workflow without demanding that the club outsource judgement to the model.

7. Limitations

- This deployment uses public data only. Club-internal load, wellness, screening, and treatment data are outside the scope of this case study.
- The player examples are selected and anonymised. They illustrate live model behaviour, but they are not intended to replace a full cohort-level appendix.
- The model is not a causal engine. High risk does not guarantee injury, and the Low band's zero-injury result applies only to the observed sample reported here.
- This publication reports relative segment performance and selected trajectories. Full denominators, labelling rules, and statistical detail are beyond its scope.

8. Conclusion

In live EPL monitoring, the Be-Healthy.AI system is most credible as an attention-ranking layer. It does not need to be a binary predictor to be useful.

The player trajectories show recurring pre-injury accumulation and recovery resets. The league-level segmentation shows that the upper bands carry materially higher observed exposure, especially in the Critical class.

That is enough to justify practical use. Review earlier. Prioritise better. Apply the strongest local context around the subset of players who most warrant it.