

FROM COMPLIANCE TO PREVENTION

A Predictive Analytics Framework for Construction Safety Management

By Tabrez Zahoor

Construction remains among the most hazardous industries in the U.S. The U.S. Bureau of Labor Statistics (BLS, 2024) reported a construction industry total recordable incident rate (TRIR) of 2.3 per 100 full-time workers in 2023.

Yet high-performing contractors in the Construction Industry Institute cohort averaged just 0.27 during the same period (CII, 2025). This gap, a ten-fold performance differential that cannot be attributed to project type or workforce composition alone, is explained in large part by the analytics infrastructure behind the number.

The dominant paradigm in construction safety management remains reactive. Organizations collect operational safety data through inspection platforms, audit systems and incident reporting tools, then aggregate that data into lagging metrics that describe what has already happened. TRIR, days away, restricted or transferred (DART) rates, and OSHA 300 logs are essential for regulatory compliance, insurance benchmarking and prequalification, but they share a structural limitation that no methodological refinement can overcome: they only describe the past. They predict nothing.

The published leading indicator literature has increasingly recognized this gap. Hallowell et al. (2013) demonstrated that integrating proactive safety controls alongside lagging measures measurably reduces incident rates. Hinze et al. (2013) established that treating all inspection findings as equivalent data points suppresses the analytical signal of high-severity observations. The National Safety Council's Campbell Institute (2018, 2022) formalized serious injury and fatality potential (SIF-P) tracking as a predictive discipline. Despite this research foundation, most construction organizations still lack a practical, deployed system that integrates these concepts into a unified portfolio-level decision-support tool.

This article presents the predictive safety analytics framework (PSAF), an original methodology developed through applied data science practice across an active construction portfolio, and documents its system architecture, current deployment status, retrospective evaluation results and implications for construction safety management practice.

Structural Limitation of Lagging Indicators

A zero-incident month is not evidence of safety; it is evidence of the absence of recorded incidents. The distinction is operationally critical. Under-inspection, inconsistent severity classification and near-miss underreporting can all produce artificially favorable lagging metrics while actual site risk accumulates undetected. As Xu et al. (2021) established through a systematic review of the construction safety leading indicator literature, active leading indicators correlate strongly with near-term project safety performance when properly structured, but most organizations lack the analytical infrastructure to exploit this relationship.

The gap between available data and actionable intelligence is not a data collection problem. Every active construction project already generates daily inspection findings, labor records, near-miss reports and audit observations through platforms deployed for compliance purposes. The problem is that standard safety systems collect this data and report it in ways that make it operationally useless for prevention. PSAF is designed to close that gap by transforming compliance data into forward-looking portfolio intelligence.

Predictive Safety Analytics Framework

PSAF integrates three analytically distinct leading indicator components, each addressing a specific limitation of conventional safety reporting, into a single composite score. The framework was developed through direct data science practice across an active mid-sized general contractor's construction portfolio.

Component 1: Labor-Normalized Findings Rate

Raw inspection counts are analytically misleading without exposure-hour normalization. A project recording 50 findings during a low-activity month with 5,000 labor hours presents an identical count to a project recording 50 findings during a high-intensity month with 100,000 labor hours, yet the safety implications

fundamentally differ. The first project is under-inspected relative to its exposure; the second is actively monitored.

PSAF normalizes inspection findings against labor-hour exposure, applying the same principle that underpins OSHA's incident rate methodology so that projects of any size, phase or workforce intensity can be compared on a risk-adjusted basis.

Equation 1:

$$LNFR = \frac{\text{Total inspection findings}}{TLH \div 1,000}$$

where:

LNFR = Labor normalized findings rate

TLH = Total labor hours

The rate expresses inspection findings per 1,000 work hours, the scaling used in the deployed implementation. The same quantity may equivalently be expressed per 200,000 exposure hours—the convention underlying OSHA incident rate methodology—by multiplying by 200. Either scaling eliminates the most common source of misleading portfolio safety comparisons and enables meaningful trend analysis across projects at different phases of construction. The deployed SIF-P rate described in Equation 3 uses the 200,000-hour convention.

Component 2: Severity-Weighted Audit Score

Standard audit systems treat all inspection findings as equivalent data points. Research confirms that this equivalence assumption suppresses the analytical signal of high-severity observations, burying the most critical risk information beneath routine administrative findings (Hinze et al., 2013; Winge et al., 2019). PSAF addresses this through an escalating severity weighting schema (Table 1, p. 13).

The weighted audit score integrates these weights into a ratio metric.

Equation 2:

$$WAS = \frac{PPF}{PPF + \sum(SW \cdot NFC)}$$

where:

WAS = Weighted audit score

PPF = Positive or proactive findings

SW = Severity weight

NFC = Negative finding count

Including positive findings in both numerator and denominator produces a self-normalizing ratio bounded between 0 and 1. A score approaching 1.0 indicates that inspection activity is dominated by proactive, positive observations; a score approaching 0 indicates that severity-weighted negative findings are overwhelming proactive activity. The formulation eliminates division-by-zero naturally (as long as any inspection activity occurs, the denominator is non-zero) and requires no arbitrary constant. The bounded output is directly comparable across projects of different sizes and phases. A declining WAS across consecutive reporting periods indicates that the ratio of proactive activity to severity-weighted risk observations is deteriorating, regardless of the absolute volume of findings.

Component 3: SIF-P Rate

Campbell Institute (2018) research establishes that serious injuries and fatalities (SIFs) are preceded by identifiable precursor conditions (e.g., unprotected fall exposures, unsupported excavation walls, uncontrolled energy sources, suspended loads over occupied areas). PSAF operationalizes this research by tracking SIF-P observations as a normalized rate metric.

Equation 3:

$$SIF-P \text{ rate} = \frac{OLSE \cdot 200,000}{\text{Total labor hours}}$$

where:

OLSE = Observations with life-safety exposure

SIF-P observations are identified through the safety platform's severity classification system and supplemented by field supervisor flagging protocols. A rising SIF-P rate occurring alongside a declining WAS is the co-movement PSAF is designed to surface: two independent components deteriorating together, which neither lagging metrics nor raw finding counts would show.

Composite Safety Analytics Score

All three components are designed to integrate into a single composite safety analytics score (CSAS), updated at a common reporting cadence across the project portfolio.

Equation 4:

$$CSAS = (0.40 \cdot BNWAS) + (0.35 \cdot BNLNFR) + (0.25 \cdot IBNSIF-P \text{ rate})$$

where:

BNWAS = Baseline-normalized WAS

BNLNFR = Baseline-normalized LNFR

IBNSIF-P = Inverted baseline-normalized SIF-P

Each component is normalized against a project-specific rolling baseline before weighting, so that the composite reflects directional trend relative to that project's own established performance rather than absolute values. This makes scores comparable across projects of different scales, types and construction phases. The SIF-P rate component is inverted; a rising SIF-P rate reduces the CSAS score, reflecting increasing risk. The composite output maps to a three-tier executive classification (Table 2). Computing the composite continuously requires that all three components be produced at the same cadence; where audit scoring is less frequent than findings and SIF-P reporting, the composite cannot be calculated as a continuous series and the components must be tracked individually.

Weight Derivation

The composite weighting schema assigns the highest weight to the WAS (0.40) on the basis that severity-weighted audit trends are the earliest-changing of the three signals, followed by the LNFR (0.35), which provides the exposure-normalization foundation for portfolio comparison, and the SIF-P rate (0.25). SIF-P receives the lowest composite weight not because it is least important in absolute terms—an immediately dangerous to life and health (IDLH) observation is the highest-priority single data point in the system—but because SIF-P events are infrequent, making the normalized rate noisier at short reporting intervals. These weights reflect design judgment informed by the published leading indicator literature. They have not been derived empirically from outcome data, and organizations implementing PSAF should calibrate them against their own history.

System Architecture & Deployment

PSAF is designed around a four-layer architecture: a cloud-based construction safety platform as the primary data source, an API-based automated data pipeline, an Excel computation engine, and a Power BI dashboard. Deployment status differs by component. At the study organization, the LNFR, the SIF-P rate and the severity-weighted audit score are each produced and published to the dashboard; the integrated composite is not computed operationally, for the cadence reason previously described. The architecture was designed to maximize analytical value from existing operational infrastructure without requiring

new software procurement or custom application development, a design principle central to PSAF's practical deployability across mid-sized construction organizations.

Data Source Layer

The primary data source is a cloud-based construction safety management platform used across all active projects for field inspection management, incident reporting and audit documentation. The platform captures inspection findings with standardized severity classifications, labor hour entries at the project and trade level, near-miss reports and SIF-P flagged observations. All data elements required by PSAF's three analytical components are captured through the organization's existing inspection workflows, requiring no additional field data collection burden.

Data Pipeline Layer

PSAF uses the platform's REST API to automate data extraction across the full project portfolio. The pipeline is implemented through Microsoft Power Automate, which executes scheduled API calls daily at 11 p.m. so that current data is available for each reporting cycle. The pipeline accesses the four primary endpoints of inspections, incidents, observations and labor hours, and it parses responses into structured files deposited to a SharePoint folder. Automated validation logic flags records with missing project identifiers, null severity classifications or anomalous labor hour entries before the score is finalized.

Computation Layer

The Excel computation engine consists of interconnected worksheets implementing discrete analytical modules: labor normalization, severity weighting and SIF-P rate calculation. Power Query connects the computation engine to the SharePoint outputs, enabling automatic data refresh when the pipeline completes. Each module applies the relevant PSAF formula and publishes its output to the dashboard. A composite module, applying the weighting schema and classifying each project against the red/yellow/green thresholds, is specified in the framework but is not in operation at the study organization for the cadence reason previously described.

Visualization & Reporting Layer

The Power BI dashboard connects to the computation engine through a

TABLE 1
PSAF SEVERITY WEIGHTING SCHEMA

Severity classification	Examples	Weight
Low	Housekeeping, minor PPE, administrative	1
Medium	Minor injury potential, immediately correctable	3
High	Serious injury potential, requires prompt attention	5
IDLH (immediately life-threatening)	Unprotected fall exposure, uncontrolled energy, excavation failure	10

TABLE 2
CSAS THRESHOLD CLASSIFICATION & OPERATIONAL RESPONSE

CSAS score	Status	Operational response
> 85%	Strong	Inspection discipline is strong; maintain current protocols
70% to 85%	Watch	Adverse trend detected; increase inspection frequency
< 70%	Act now	Immediate executive review and targeted intervention required

SharePoint file connection and refreshes on completion of the pipeline run. For each active project, the deployed dashboard presents dedicated views for the LNFR, the SIF-P rate, the severity-weighted audit score, and incident and DART rates. Each view provides a period-by-period table of the underlying counts and the derived metric, a trend chart across the project's reporting history with threshold boundaries overlaid, and a score-to-date summary. Project selection and date-range filtering allow review of any project across any period.

Data Quality & Inspector Consistency

PSAF's analytical outputs depend on the consistency of severity classifications assigned by field inspectors. Three mechanisms address classification consistency: inspector calibration training using standardized hazard scenario sets with defined correct classifications, supervisor review of High and IDLH classifications before they are finalized in the system, and periodic analyst review flagging projects with unusual classification distributions relative to their own baselines. EHS professionals implementing PSAF should note that inter-rater reliability has not been formally quantified through a controlled agreement study in this deployment; this is a limitation that future implementations should address through

structured Cohen's kappa assessment across inspector pairs.

Retrospective Evaluation Project & Data Description

PSAF was evaluated retrospectively against data from a single active project—a municipal public safety headquarters—across 26 consecutive months from June 2024 through July 2026. Over that period the project accumulated 128,497 labor hours, 3,178 recorded inspection findings and 21 SIF-P observations, with monthly labor hours rising from under 500 at mobilization to a peak of 11,120. Two of the three PSAF components are available at monthly resolution across the full period: the LNFR and the SIF-P rate. The severity-weighted audit score is available at five points only—April 2025, September 2025, December 2025, March 2026 and June 2026—reflecting the organization's audit cadence. Because the composite requires all three components at a common cadence, CSAS could not be computed as a continuous series for this project; the components are therefore reported individually and their co-movement assessed.

Evaluation Results

The severity-weighted audit score declined across four consecutive audits and then recovered. Table 3 presents a summary of the retrospective evaluation, including the following data:

TABLE 3
EVALUATION SUMMARY

Metric	Result
Project evaluated	1 active project (municipal public safety headquarters)
Evaluation period	26 months (June 2024 through July 2026)
Total labor hours	128,497
Total inspection findings	3,178
Total SIF-P observations	21
Severity-weighted audit observations	5 (April 2025, Sept. 2025, Dec. 2025, March 2026, June 2026)
Audit score trend	0.872 to 0.754 across four audits (decline of 13.5%), then 0.911 (recovery)
SIF-P rate peak	106.5 per 200,000 hours (Oct. 2025)
Recordable-classified events	6 total; 4 clustered Feb. 2026 through May 2026
Component deterioration versus incident cluster	3 to 5 months earlier

•WAS: 0.872 (April 2025), 0.845 (Sept. 2025), 0.803 (Dec. 2025), 0.754 (March 2026), 0.911 (June 2026); a decline of 13.5% across 11 months, followed by recovery

•LNFR: 31.4 findings per 1,000 work hours in Sept. 2025, falling to 19.4 in Dec. 2025 and 17.1 in March 2026, while monthly labor hours rose from 7,516 to 11,120 across the same interval; inspection activity was declining relative to a growing exposure base

•SIF-P rate: 48.2 observations per 200,000 work hours in July 2025, 79.8 in Sept. 2025, and a peak of 106.5 in Oct. 2025; rose through the second half of 2025 before falling back below 41 for the remainder of the period

Recordable-classified incidents on this project clustered in Feb. 2026 through May 2026: four events across four consecutive months, including the project's only DART case in April 2026. The preceding 12 months recorded a single event in Aug. 2025, and the project's only other event occurred in July 2024. All three components therefore deteriorated during the second half of 2025, 3 to 5 months before the incident cluster began. The Sept. 2025 and Dec. 2025 audits both recorded declines, the findings rate fell as exposure grew, and the SIF-P rate peaked in Oct. 2025, at a point when the project's recordable incident count over the preceding 12 months stood at one.

This is a single project and a single episode, evaluated retrospectively. Five audit observations cannot establish a detection interval, no intervention was recorded in response to these signals, and no claim is made that PSAF prevented or would have prevented any event. What the evaluation demonstrates is that the framework's three components moved together and moved early on real portfolio data, in a direction consistent with the leading indicator hypothesis. Establishing whether that co-movement supports reliable early intervention requires prospective testing across multiple organizations.

Implications for EHS Practice

PSAF is built on the premise that the data that construction organizations already collect through inspection platforms deployed for compliance purposes can support forward-looking risk indicators when properly normalized, weighted and integrated. Four operational principles characterize this shift from compliance reporting to prevention intelligence:

- Compare by labor intensity, not raw counts.** A project recording 50 findings during a 5,000-hour month and a project recording 50 findings during a 100,000-hour month present identical raw numbers but fundamentally different risk profiles. Labor-hour normalization is the prerequisite for meaningful portfolio comparison.

- Weight hazards by severity, not frequency.** Treating a housekeeping finding and an immediately life-threatening fall exposure as equivalent data points buries the most critical safety signals beneath

routine administrative observations. Severity weighting redirects analytical attention to what could cause serious harm.

- Track SIF-P precursors before incidents occur.** SIFs are preceded by identifiable precursor conditions that are visible in inspection data before any injury occurs. SIF-P tracking converts this research finding into an operational early warning signal.

- Produce every component at the same cadence.** A composite score can only be computed when all its inputs are available at a common reporting interval. Where audit scoring occurs quarterly while findings and SIF-P data are produced monthly, the composite cannot be calculated at all; the components must either be aligned or tracked separately. Establishing that cadence is the practical prerequisite for portfolio-level composite scoring, and the most common obstacle to it.

Implementing PSAF does not require new software, additional field staff or changes to existing inspection workflows. Every metric is calculable from data that mid-to-large contractors already collect through their existing platforms. The technology stack (a commercial safety platform API, Excel and Power BI) is available within most construction organizations' existing software licenses.

EHS professionals considering PSAF implementation should note that the composite weighting schema and threshold boundaries presented in this article reflect design judgment rather than empirical derivation from outcome data. Organizations implementing PSAF should treat these parameters as

starting points requiring calibration against their own historical data rather than universal constants.

Comparison With Existing Leading Indicator Approaches

The published leading indicator literature has established the theoretical value of each PSAF component individually. Hallowell et al. (2013) demonstrated the statistical relationship between proactive leading indicator monitoring and incident rate reduction. Hinze et al. (2013) established the signal-suppression problem created by equivalent-weight audit scoring. The Campbell Institute (2018, 2022) formalized SIF-P tracking as an operational discipline. OSHA's (n.d.) leading indicators guidance encourages the integration of proactive safety measures alongside lagging metrics.

PSAF's contribution relative to this literature is not the novelty of any individual component. It is their integration into a unified, portfolio-level framework with a documented evaluation against real construction data. Prior frameworks have established statistical or qualitative relationships between leading indicators and safety outcomes. This article reports a retrospective single-project evaluation in which all three components deteriorated concurrently in the months preceding a cluster of recordable incidents and identifies the cadence prerequisite that continuous composite scoring imposes. Table 4 summarizes how this article's contribution relates to published leading indicator approaches.

EHS professionals should note four limitations. First, this evaluation covers one project and one episode; five audit observations cannot establish a detection interval or a false-positive rate. Second, the composite was not computed operationally and no intervention was recorded, so nothing can be concluded about intervention effectiveness. Third, inter-rater reliability of severity classification has not been formally quantified. Fourth, the weighting schema and threshold boundaries reflect design judgment rather than empirical derivation from outcome data. Prospective, multi-organization replication is the priority next step in PSAF's empirical validation.

Conclusion

Construction organizations are not facing a data shortage. Every active project

SIX ACTIONS FOR PREDICTIVE SAFETY

- Normalize by labor exposure.** Compare inspection findings relative to labor hours, not raw counts, to make meaningful risk comparisons across projects of different sizes, phases and workforce intensity.

- Weight findings by severity.** Give greater analytical weight to high-severity and life-threatening hazards so critical risks are not buried beneath routine observations.

- Track SIF-P precursors.** Monitor conditions with serious injury and fatality potential, such as fall exposures and uncontrolled energy, to identify emerging danger before an injury occurs.

- Review risk signals weekly.** Use composite-score reviews rather than relying on monthly lagging metrics so teams can detect deteriorating conditions early enough to intervene.

- Build on existing systems.** Turn inspection, labor, near-miss and audit data already captured through safety platforms into predictive intelligence instead of adding new field-data requirements.

- Calibrate to the specific organization.** Treat published score weights and thresholds as starting points, then adjust them using the organization's own historical data, portfolio characteristics and safety baseline.

TABLE 4
COMPARISON OF PSAF WITH
PUBLISHED LEADING INDICATOR APPROACHES

Source	Contribution established	Study design	Reported relationship	Deployed system
Hinze et al. (2013)	Leading indicators of construction safety performance; equal-weight audit scoring suppresses high-severity signal	Multi-project field study	Statistical association between proactive activity and incident rates	No
Hallowell et al. (2013)	Proactive safety control: monitoring leading indicators alongside lagging measures	Multisite study	Lower recordable rates where proactive indicators are monitored	No
Campbell Institute (2018, 2022)	SIF-P precursor concept; cumulative risk and safety networks	Practitioner consensus research	Qualitative precursor identification preceding serious events	No
Xu et al. (2021)	Systematic review of construction safety leading indicators	Review of prior empirical studies	Active leading indicators correlate with near-term safety performance when properly structured	No
Winge et al. (2019)	Causal factors and connections in construction incidents	Retrospective incident analysis	Causal pathways in recorded incidents	No
PSAF	Integration of exposure normalization, severity weighting and SIF-P precursor tracking into one portfolio-level composite	Single-organization retrospective evaluation; one project, 26 months, 128,497 labor hours	All three components deteriorated concurrently 3 to 5 months before a cluster of recordable incidents	Components in production; composite not computed operationally (audit cadence)

generates daily streams of inspection findings, labor records, near-miss reports and audit observations—the raw material of predictive safety intelligence. The challenge is not collection; it is conversion. PSAF provides the analytical architecture to convert compliance data into prevention intelligence at portfolio scale, using tools and data sources that most mid-to-large contractors already have.

The framework’s core contribution to EHS practice is operational: a composite score designed to tell safety leaders and project executives not what went wrong last month, but where risk is building now—early enough to intervene before any recordable event occurs. The concurrent component deterioration observed in this retrospective evaluation, in the months preceding an incident cluster, suggests that information relevant to prevention may already be present in the data that organizations collect

today. Realizing it in practice requires producing all three components at a common reporting cadence and testing the composite prospectively across multiple organizations.

Future research should extend PSAF’s validation to multi-organizational portfolios, formally quantify inspector classification consistency through inter-rater reliability studies, and explore the application of the same analytical architecture (normalization, severity weighting, composite integration, threshold classification) to construction schedule performance, quality management and financial risk domains. **PSJ**

Disclaimer

During the preparation of this article, the author used ChatGPT to refine language, improve clarity and assist with formatting. The author has reviewed and

edited and takes full responsibility for the content, accuracy and originality of the final publication. For updated ASSP AI policy information, see <https://assp.us/AI>.

References

Bureau of Labor Statistics (BLS). (2024). Injuries, illnesses and fatalities. Table 1: Incidence rates of nonfatal occupational injuries and illnesses by industry and case types, 2023. U.S. Department of Labor. <https://bit.ly/4AovgOK>

Campbell Institute. (2018). Serious injury and fatality prevention: Perspectives and practices. <https://bit.ly/4xhV2kH>

Campbell Institute. (2022). Serious injury and fatality prevention: Leading indicators, cumulative risk and safety networks. <https://bit.ly/4yrVGNA>

Construction Industry Institute (CII). (2025). 2024 safety summary report. <https://bit.ly/4xDYDKo>

Hallowell, M.R., Hinze, J.W., Baud, K.C. & Wehle, A. (2013). Proactive construction safety control: Measuring, monitoring and responding to safety leading indicators. *Journal of Construction Engineering and Management*, 139(10), 04013010. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000730](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000730)

Hinze, J., Thurman, S. & Wehle, A. (2013). Leading indicators of construction safety performance. *Safety Science*, 51(1), 23-28. <https://doi.org/10.1016/j.ssci.2012.05.016>

OSHA. (n.d.). Using leading indicators to improve safety and health outcomes. www.osha.gov/leading-indicators

Winge, S., Albrechtsen, E. & Mostue, B.A. (2019). Causal factors and connections in construction accidents. *Safety Science*, 112, 130-141. <https://doi.org/10.1016/j.ssci.2018.10.015>

Xu, J., Cheung, C., Manu, P. & Ejohwomu, O. (2021). Safety leading indicators in construction: A systematic review. *Safety Science*, 139, 105250. <https://doi.org/10.1016/j.ssci.2021.105250>

Tabrez Zahoor, M.S., is a senior data scientist at W.T. Rich Company, a construction management firm based in Massachusetts, where he specializes in construction analytics and portfolio-level performance measurement. He developed the predictive safety analytics framework, an original methodology for converting construction inspection and labor data into predictive weekly safety intelligence. His work focuses on building practical analytics frameworks that help construction organizations convert operational data into forward-looking decisions across safety, schedule, quality and financial performance domains.

Cite this article

Zahoor, T. (2026, Oct.). From compliance to prevention: A predictive analytics framework for construction safety management. *Professional Safety*, 71(10), 11-15.