Algorithmic Detection of Anomalies in Motor Carrier Safety Indicators

Lazizbek Ravshanov
BSc Degree Candidate, IT-Data Technologies, University of Cincinnati;
Associate Degree from Computer Systems Institute, Cincinnati, USA
ORCID ID: 0009-0000-3915-1090;

ABSTRACT: *This article examines approaches to the algorithmic detection of anomalies in motor carrier safety indicators. The analysis is based on a multidimensional feature space combining registration, operational, temporal, and safety characteristics. Expert rules, statistical methods, and machine-learning algorithms, including Isolation Forest, Local Outlier Factor, and One-Class SVM, are considered. Particular attention is paid to threshold selection, false-positive reduction, and the use of an integrated anomaly score to prioritize additional review. Entity-level anomalies and carrier re-identification signals are also examined. The findings indicate that reliable anomaly detection requires the combined use of indicator normalization, peer-group comparison, temporal analysis, and confirmation of deviations through multiple independent features.*

KEYWORDS: *Anomaly Detection, Anomaly Score, Carrier Re-Identification, False-Positive Alerts, Machine Learning, Motor Carriers, Transportation Safety.*

I. INTRODUCTION

The digitalization of the transport sector has generated large volumes of carrier data, including registration records, vehicle characteristics, inspection results, violations, crashes, and other safety indicators. However, risk assessment still often relies on individual metrics and predefined thresholds. This may be insufficient because elevated risk can arise from an unusual combination of several characteristics even when each indicator remains within an acceptable range.

Algorithmic anomaly detection offers a way to construct multidimensional carrier profiles using registration, operational, temporal, and safety-related features. The main challenges involve selecting appropriate detection methods, defining thresholds, reducing false-positive alerts, and interpreting complex behavioral and identification anomalies.

The aim of this study is to systematize and substantiate approaches to the algorithmic detection of anomalous motor carrier profiles based on combined registration, operational, and safety indicators, and to define key principles for threshold selection and false-positive reduction.

II. FEATURE ENGINEERING FOR ALGORITHMIC ANOMALY DETECTION IN MOTOR CARRIER PROFILES

Algorithmic detection of anomalous motor carrier behavior requires transforming heterogeneous registration and safety data into a structured feature space. The underlying data may describe registration status, fleet size, operational intensity, inspection outcomes, violations, and crash history. However, the analytical value of an individual indicator depends on context: the same number of violations may have different

implications for small and large carriers. Therefore, absolute values should be supplemented with relative, temporal, and comparative features that account for differences in operating scale and observed activity.

Existing research supports the use of multidimensional feature sets for carrier safety analysis. Gyawali et al. used explainable machine-learning models to identify factors associated with commercial motor vehicle crash severity, demonstrating that the contribution of individual features varies across operating environments [1]. Wang et al. analyzed data from 596 commercial trucks and constructed 22 operational and behavioral features. Their best-performing model achieved an accuracy of 96.1% and an AUC of 0.866 [2]. These findings support the use of combined characteristics rather than a single safety indicator for identifying atypical carrier profiles.

For the purposes of the present study, the feature system is structured into registration, operational, safety, temporal, comparative, and combined characteristics. This classification reflects both the source variables and the derived indicators used to account for operating scale, temporal dynamics, peer-group deviation, and interactions between individual features, as summarized in **Table 1**.

Table 1. Main Feature Groups for Algorithmic Assessment of Motor Carriers

Feature group	Source indicators	Derived features	Analytical purpose
Registration	Registration date, operating status, type of activity, location.	Company age, frequency of registration changes.	Detection of atypical profile dynamics.
Operational	Number of vehicles and drivers.	Drivers/power units ratio, fleet growth rate.	Adjustment for operating scale.
Safety	Inspections, violations, out-of-service events, crashes.	Violations per inspection, OOS rate, crash rate.	Comparison of safety outcomes independently of the absolute number of events.
Temporal	Dates of inspections, crashes, and profile changes.	Rate of change, intervals between events.	Detection of persistent adverse trends.
Comparative	Indicators of similar carriers.	Percentile, deviation from peer-group median.	Identification of atypicality relative to comparable carriers.
Combined	Features from multiple groups.	Ratios and feature combinations.	Detection of multivariate anomalies.

Normalization of safety indicators according to carrier operating scale is essential for avoiding biased comparisons. Relative measures based on inspections, fleet size, vehicle miles traveled, or other exposure indicators provide a more meaningful basis for comparison than absolute event counts.

The reliability of such indicators also depends on the number of observations. Lack et al. reported a pronounced class imbalance in FMCSA vehicle-level data, with crashes accounting for approximately 2% of observations [3]. Therefore, the same relative rate should not be interpreted identically when it is derived from a few observations and when it is based on a substantially larger sample.

Temporal and comparative features provide additional context by indicating whether an observed deviation persists over time and differs from the pattern observed among comparable carriers. The sequence used to construct the resulting feature space is shown in **Fig. 1**.

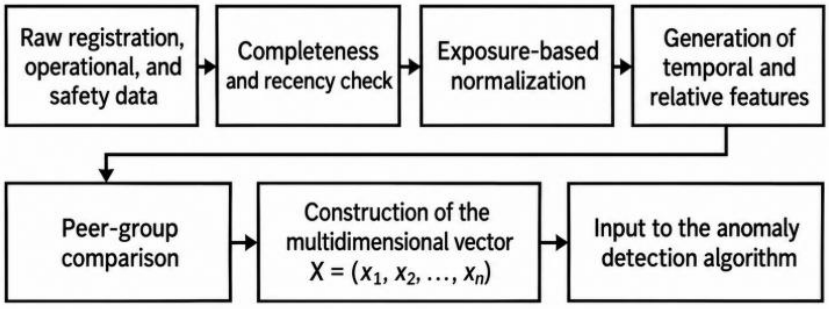


Fig. 1. Feature Space Construction for Motor Carrier Anomaly Detection

The growing availability of digital data creates opportunities to incorporate additional operational features into carrier risk analysis. According to the 2026 Fleet Technology Trends Report, GPS fleet tracking was used by 80% of surveyed fleet professionals, while 46% reported using video telematics; among video telematics users, 74% reported improvements in driver safety [4]. Although these data do not replace federal safety indicators, they illustrate the expanding information base available for subsequent algorithmic analysis.

Thus, the feature space for anomaly detection should include not only source indicators but also their relative, temporal, comparative, and combined transformations. This makes it possible to identify cases in which individual values remain within typical ranges while their combination forms an atypical carrier profile.

III. Algorithms for Detecting Atypical Combinations Of Indicators

Once the feature space has been constructed, anomaly detection focuses on determining how strongly an individual carrier profile deviates from patterns observed among comparable carriers. Depending on data availability and interpretability requirements, this task can be addressed using rule-based, statistical, or machine-learning methods.

When labeled examples of high-risk or atypical carriers are limited, unsupervised and one-class approaches become particularly relevant. Methods such as Isolation Forest, Local Outlier Factor, and One-Class SVM differ in the way they define abnormality and in their sensitivity to local density, feature scaling, and model parameters [5, 6]. Their main characteristics in the context of motor carrier analysis are summarized in **Table 2**.

Table 2. Comparison of Methods for Detecting Anomalous Motor Carrier Profiles

Method	Detection principle	Main advantage	Main limitation
Rule-based	Predefined conditions and threshold combinations.	High interpretability and direct control over detection logic.	Limited ability to identify previously unknown feature combinations.
Statistical methods	Deviation from an expected distribution or reference range.	Transparent interpretation and relatively low computational complexity.	Reduced effectiveness for high-dimensional and nonlinear patterns.
Isolation Forest	Isolation of rare observations through recursive random partitioning.	Suitable for multidimensional data and does not require labeled anomalies.	Results depend on model parameters and contamination assumptions.
Local Outlier Factor	Comparison of local observation density with neighboring data points.	Effective for detecting local anomalies within heterogeneous groups.	Sensitive to neighborhood size and local data density.

One-Class SVM	Estimation of a boundary enclosing typical observations.	Applicable when examples of normal behavior dominate the dataset.	Requires feature scaling and careful parameter tuning.

A key distinction for this study is between univariate and multivariate anomalies. A univariate anomaly is associated with an extreme value of a single indicator, whereas a multivariate anomaly may arise when individual variables remain within typical ranges but their joint configuration is unusual. The latter case is especially important for carrier analysis because fixed thresholds may fail to detect atypical combinations of registration, operational, temporal, and safety characteristics [5].

The model output can therefore be represented as an anomaly score reflecting the degree of deviation of a carrier profile from typical observations. This score should not be interpreted as direct evidence of unsafe or non-compliant behavior. Its primary function is to rank profiles and identify cases requiring additional review. Consequently, the practical usefulness of anomaly detection depends not only on the selected algorithm but also on threshold calibration and control of false-positive alerts.

IV. THRESHOLD CALIBRATION AND FALSE-POSITIVE CONTROL IN ANOMALY DETECTION

Anomaly detection requires a decision threshold that separates routine variation from profiles requiring additional review. A universal cutoff is rarely appropriate because carrier-level indicators depend on operating scale, inspection frequency, activity type, and the amount of accumulated data. Thresholds should therefore be calibrated against the anomaly-score distribution within a relevant comparison group rather than applied uniformly across all carriers.

Threshold performance should be assessed using precision, recall, false positive rate, and F1-score. Their relative importance depends on the analytical objective: precision becomes more important when only a limited number of profiles can be reviewed manually, whereas recall is critical when missing potentially anomalous cases is costly. Under strong class imbalance, accuracy alone may provide a misleading estimate of model performance [7]. In addition to metric-based evaluation, threshold calibration should account for data and profile characteristics that can generate false-positive alerts. The main sources of such alerts and corresponding mitigation approaches are summarized in **Table 3**.

Table 3. Main Sources of False-Positive Alerts and Mitigation Approaches

Source of false positives	Potential effect	Mitigation approach
Small number of inspections or events	Unstable relative rates.	Minimum observation requirements before scoring.
Short-term fluctuation	Temporarily inflated anomaly score.	Time-window aggregation and smoothing.
Differences in carrier operating scale	Biased comparison of absolute values.	Exposure normalization and peer-group comparison.
Isolated registration change	Misclassification of a routine update as an anomaly.	Confirmation through repeated or persistent changes.
Correlated features	Artificial amplification of the anomaly signal.	Feature-dependence analysis and dimensionality control.
Incomplete or asynchronous data	Spurious combinations of abnormal values.	Data-quality and temporal-consistency checks.

A practical approach is to use several priority levels rather than a single binary decision. The transition between levels can depend on the anomaly score, persistence of the deviation, and the number of independent confirming features. Adaptive thresholding is particularly useful when the score distribution changes over time, because static cutoffs may generate excessive false alarms or suppress relevant deviations. Recent anomaly-detection research therefore increasingly considers dynamic calibration, sliding-window statistics, and target false-alarm rates when selecting operational thresholds [8, 9].

In operational use, the anomaly score should therefore serve as the starting point of a sequential assessment rather than as a standalone decision criterion. Threshold evaluation can be followed by confirmation from independent features and assignment of a review priority, as shown in **Fig. 2**.

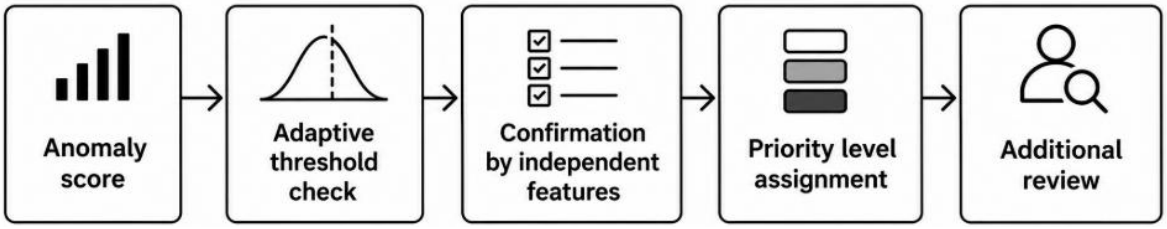


Fig. 2. Decision Logic for Anomaly Prioritization and Review

Thus, threshold calibration should be treated as a separate analytical stage rather than as a secondary model parameter. Combining peer-group calibration, temporal persistence, data-quality checks, and multi-feature confirmation makes it possible to use anomaly detection as a prioritization mechanism while limiting false-positive alerts.

V. DETECTION OF COMPLEX BEHAVIORAL ANOMALIES AND CARRIER RE-IDENTIFICATION SIGNALS

A more complex form of anomaly detection arises when atypical behavior is expressed not within a single carrier profile but through relationships between multiple entities and changes in their characteristics over time. Relevant patterns may include abrupt registration changes, the appearance of a new carrier shortly after another ceases operations, recurring contact or operational attributes, and continuity in safety-related characteristics. Such cases require entity-level analysis in which the unit of assessment includes both individual records and the relationships connecting them.

One practical application is the identification of chameleon, or reincarnated, carriers. FMCSA uses the Utility for Risk-Based Screening and Assessment (URSA) to compare new applicant data with existing carrier records and estimate the likelihood of attempted reincarnation or other improper acquisition of operating authority. The Vetting Management Tool documentation describes an Affinity Score, which reflects the degree of similarity between an applicant and existing carriers, and a Motivation Score, which reflects the potential safety risk and incentive to reincarnate. FMCSA reports that this mechanism is used to screen all new applicants seeking operating authority [10].

For algorithmic analysis, a single matching attribute should not be treated as sufficient evidence of re-identification. Shared addresses, telephone numbers, vehicles, or associated persons may occur for legitimate reasons. Greater analytical value is provided by combinations of independent matches, their temporal sequence, and associated safety characteristics. The principal signal groups are summarized in **Table 4**.

Table 4. Feature Groups for Detecting Complex Carrier Identification Anomalies

Feature group	Example of observed match or change	Analytical interpretation
Registration	Similar addresses, contact details, or associated persons.	Possible relationship between carrier profiles.

Operational	Shared vehicles or similar fleet composition.	Possible continuity of operational structure.
Temporal	Short interval between termination of one profile and emergence of another.	Possible sequence of carrier reincarnation.
Safety	Similar patterns of violations, inspections, or other adverse indicators.	Potential continuity of the risk profile.
Combined	Simultaneous presence of several independent matches.	Higher priority for additional review.

An example of an analytical environment supporting this type of assessment is FleetSight, which combines carrier, inspection, vehicle, and safety information and includes mechanisms for identifying cross-carrier relationships and potential chameleon signals [11]. Its public description reports coverage of approximately 4.4 million registered carriers, more than 8 million inspection records, and over 13 million VIN observations. In the present study, FleetSight is considered only as an implementation example of the analytical stage rather than as an object of architectural analysis.

The output of entity-level matching should nevertheless be interpreted as a basis for additional review rather than as automatic evidence of unsafe or non-compliant behavior. A comparable principle is applied in FMCSA's Safety Measurement System, where performance data are used to prioritize carriers for monitoring and interventions but do not themselves constitute a federal safety rating [12]. Consequently, the most appropriate approach is to use combinations of registration, operational, temporal, and safety signals to establish analytical priority, followed by additional verification before any substantive conclusion is made.

VI. CONCLUSION

Algorithmic detection of anomalies in motor carrier safety indicators should be based on a multidimensional feature space combining registration, operational, temporal, and safety characteristics. Reliance solely on individual absolute indicators or fixed thresholds limits analytical accuracy, since atypical behavior may emerge from a combination of several moderate deviations. A more robust approach combines normalization, peer-group comparison, temporal analysis, expert rules, and machine-learning methods such as Isolation Forest, Local Outlier Factor, and One-Class SVM.

The practical effectiveness of this approach depends not only on detection sensitivity but also on the ability to control false-positive alerts. This requires accounting for observation volume, persistence of deviations, source-data quality, and confirmation by multiple independent signals. A promising direction is the development of entity-level analysis that evaluates not only individual carrier characteristics but also relationships among organizations, vehicles, and registration profiles over time. Such methods can support early detection of complex behavioral anomalies, including cases that warrant additional review for potential chameleon carrier patterns.

REFERENCES

[1] N. Gyawali, S. Khanal, D. Caragea, H.M.A. Aziz, and E.J. Fitzsimmons, Predicting commercial motor vehicle crash severity in Kansas using explainable machine learning, *Journal of Transportation Safety & Security*, 17(12), 2025, 1580–1611. https://doi.org/10.1080/19439962.2025.2540388.

[2] X. Wang, X. Tang, T. Fan, Y. Zhou, and X. Yang, Commercial truck risk assessment and factor analysis based on vehicle trajectory and in-vehicle monitoring data, *Transportation Research Record: Journal of the Transportation Research Board*, 2678(12), 2024, 1428–1443. https://doi.org/10.1177/03611981241252148.

[3] C.D. Lack, K.S. Berkow, and Y. Gao, Insights into motor carrier crashes: A preliminary

investigation of FMCSA inspection violations, *Accident Analysis & Prevention*, 155, 2021, Article 106105. https://doi.org/10.1016/j.aap.2021.106105.

[4] Verizon Connect, 2026 Fleet Technology Trends Report: AI-powered technology improves safety, productivity and data-driven decision-making, https://www.verizon.com/about/news/2026-verizon-connect-fleet-technology-trends-report (accessed 24 September 2026).

[5] D. Fährmann, L. Martín, L. Sánchez, and N. Damer, Anomaly detection in smart environments: A comprehensive survey, *IEEE Access*, 12, 2024, 64006–64049. https://doi.org/10.1109/ACCESS.2024.3395051.

[6] T. Hayashi, D. Cimr, H. Fujita, and R. Cimler, Critical review for one-class classification: Recent advances and reality behind them, *WIREs Data Mining and Knowledge Discovery*, 16(1), 2026, Article e70058. https://doi.org/10.1002/widm.70058.

[7] M. Altalhan, A. Algarni, and M. Turki-Hadj Alouane, Imbalanced data problem in machine learning: A review, *IEEE Access*, 13, 2025, 13686–13699. https://doi.org/10.1109/ACCESS.2025.3531662.

[8] U.A. Usmani, I. Abdul Aziz, J. Jaafar, and J. Watada, Deep learning for anomaly detection in time-series data: An analysis of techniques, review of applications, and guidelines for future research, *IEEE Access*, 12, 2024, 174564–174590. https://doi.org/10.1109/ACCESS.2024.3495819.

[9] L. Correia, J.-C. Goos, P. Klein, T. Bäck, and A.V. Kononova, Online model-based anomaly detection in multivariate time series: Taxonomy, survey, research challenges and future directions, *Engineering Applications of Artificial Intelligence*, 138(Part A), 2024, Article 109323. https://doi.org/10.1016/j.engappai.2024.109323.

[10] Federal Motor Carrier Safety Administration, Vetting Management Tool: User Manual, https://csa.fmcsa.dot.gov/Documents/VMT_User_Manual-508.pdf (accessed 29 September 2026).

[11] FleetSight, FMCSA Carrier Intelligence Search, https://fleetsight.vercel.app/ (accessed 29 September 2026).

[12] Federal Motor Carrier Safety Administration, Safety Measurement System, https://ai.fmcsa.dot.gov/SMS/ (accessed 29 September 2026).