

Toward an Intelligence Layer for Trucking Operations: A Practice-Informed Conceptual Framework For AI-Enabled Fleet Decision Support

Abdulkhamid Abdullaev

Master's Student, University of South Florida, USA

Received: 12 February 2026; **Accepted:** 08 March 2026; **Published:** 31 March 2026

Abstract: The trucking industry generates large volumes of operational data through fuel transactions, GPS and telematics systems, electronic logging devices, maintenance records, dispatch systems, safety platforms, and financial tools. Yet these data are frequently fragmented across functional departments, limiting managers' ability to identify cross-system anomalies and prioritize emerging operational risks. This paper develops a practice-informed conceptual framework for an artificial-intelligence-enabled 'intelligence layer' in trucking operations. The framework is grounded in operational observations documented during the early development of Dieselly.ai and is analytically connected to recent literature on machine learning in transportation, fuel-consumption modeling, predictive maintenance, anomaly detection, explainable artificial intelligence, and human-AI decision support.

A structured conceptual analysis was used to organize recurring operational problems into four design requirements: cross-functional data integration, anomaly-based attention prioritization, human-in-the-loop decision support, and context-aware explanation. The analysis suggests that the primary value of AI in fleet operations may lie less in replacing existing transportation-management systems than in connecting their outputs, detecting inconsistencies, and directing managerial attention toward the vehicles, drivers, transactions, or maintenance conditions most likely to require investigation. Fuel efficiency, suspicious fuel transactions, maintenance risk, and safety monitoring are discussed as illustrative use cases.

The paper proposes a layered architecture that separates data integration, contextual inference, anomaly detection, explanation, and human decision authority. Because the study is conceptual and does not report a controlled field trial, no causal claims about operational performance are made. The framework is intended to guide future product development and empirical validation of AI-supported decision systems in trucking.

Keywords: Artificial intelligence; trucking; fleet management; telematics; anomaly detection; predictive maintenance; fuel efficiency; explainable AI; decision support; human-in-the-loop.

Introduction: Digital transformation in road freight has produced an expanding ecosystem of specialized software. Fleet operators may use separate systems for vehicle location, fuel-card transactions, hours-of-service compliance, maintenance, dispatch, accounting, safety, and recruiting. Each system can provide useful information within its functional boundary, yet the operational reality of trucking is inherently cross-functional. A maintenance problem

can affect fuel consumption; dispatch choices can influence mileage and fuel use; driver behavior can affect safety and asset wear; and safety performance can affect insurance, freight access, and profitability. When these relationships remain distributed across disconnected systems, managers must manually reconstruct the operational picture before they can decide what requires attention.

The original operational problem that motivated

Dieselly.ai emerged from this fragmentation. From the operator's perspective, the issue was not an absence of software or data, but an absence of a shared, timely interpretation of information already produced by existing systems. The founder's early observations emphasized that dispatch, fuel, safety, fleet maintenance, and accounting often hold different pieces of the same operational event, while no single layer continuously compares those pieces for inconsistency or emerging risk. This observation led to a central design question: can software connect heterogeneous operational signals and identify what deserves attention before a manager discovers the problem manually?

Recent research supports the technical plausibility of this direction. Machine-learning applications in transportation have demonstrated value in predictive maintenance, energy management, fault diagnosis, and system optimization, while continuing to face challenges related to data quality, real-time implementation, model complexity, and integration (Zhang, Zhang, & Dinavahi, 2025). In heavy-duty trucking specifically, machine-learning and deep-learning models have been used to estimate fuel consumption from real-world operating data (Katreddi & Thiruvengadam, 2021; Fan et al., 2024), identify influential operational features (Zhu et al., 2024), and detect anomalous high-fuel-consumption behavior (Turan et al., 2024). Predictive-maintenance research likewise shows that vehicle sensor data can support earlier diagnostics, maintenance prioritization, and reliability management (Mahale, Kolhar, & More, 2025).

However, technical prediction accuracy alone does not define operational usefulness. Fleet managers work in high-frequency environments characterized by limited attention, competing priorities, uncertain information, and accountability for final decisions. Recent research on explainable AI and decision support emphasizes the need to align AI outputs with human judgment, organizational processes, and transparency requirements rather than treating prediction as an end in itself (Kuznietsov et al., 2024; Abdulrashid et al., 2026). This consideration is especially important in trucking, where an anomalous fuel transaction, maintenance pattern, or safety signal may justify investigation without necessarily proving misconduct,

failure, or imminent danger.

Accordingly, this paper does not present Dieselly.ai as a validated autonomous management system. Instead, it converts an early-stage founder perspective into a research-oriented conceptual framework. The objective is to articulate how an AI-enabled intelligence layer could integrate existing operational signals, prioritize anomalies, preserve human decision authority, and support more timely investigation across fleet functions. The paper addresses three research questions: (1) Which recurring information-management problems in trucking are most suitable for cross-functional AI decision support? (2) What architectural and governance principles should guide an intelligence layer that assists rather than replaces fleet managers? and (3) Which use cases provide the clearest pathway for empirical validation?

METHODS

This study uses a practice-informed conceptual case-study design. The purpose is theory building and design clarification rather than hypothesis testing. The core source material is the founder-authored account of the first four months of Dieselly.ai development (August-December 2025), which documents operational problems observed from within the trucking industry, proposed AI use cases, and an emerging product vision. The source material is treated as practitioner evidence: it provides direct insight into operational pain points but does not constitute an independent empirical evaluation of the product.

To strengthen analytical rigor, the operational observations were interpreted alongside peer-reviewed research on intelligent transportation, heavy-duty vehicle fuel modeling, anomaly detection, predictive maintenance, explainable AI, and decision-centered human-AI interaction. The resulting framework therefore combines practice-based problem formulation with literature-supported technical and governance principles.

The analysis proceeded in three stages. First, recurring operational statements in the founder account were grouped by functional domain: fuel, maintenance, safety, dispatch, finance, and workforce-related operations. Second, cross-cutting mechanisms were identified, including data fragmentation, abnormal-pattern detection, contextual comparison, attention

prioritization, and human investigation. Third, these mechanisms were mapped to the recent literature to distinguish product-specific intuition from broader research-supported design principles.

The conceptual synthesis focused on relationships among signals rather than on any single predictive model. For example, a fuel transaction can be interpreted more meaningfully when compared with GPS location, mileage, historical fuel efficiency, vehicle condition, and route context. Similarly, tire-pressure values become more useful when analyzed as trends and combined with maintenance history, mileage, operating conditions, and diagnostic information. This relational perspective was used to derive the proposed intelligence-layer architecture.

No proprietary fleet dataset, randomized experiment, customer trial, or validated production benchmark is reported in this paper. The study therefore does not estimate cost savings, fraud-reduction rates, safety effects, maintenance improvements, or fuel-efficiency gains attributable to Dieselly.ai. Numerical findings cited in the discussion are taken from external peer-reviewed studies and are used only to establish technical plausibility and research context.

The analysis also adopts a human-in-the-loop assumption: AI outputs are treated as decision-support signals rather than final determinations. An anomaly should indicate that an event deserves investigation; it should not automatically accuse a driver of fraud, declare a vehicle unsafe, or authorize a consequential action without appropriate human review.

RESULTS

The conceptual analysis identified four interdependent design requirements for an AI-enabled intelligence layer in trucking operations. These requirements recur across fuel, safety, maintenance, dispatch, and financial use cases and form the core result of the study.

The first finding is that trucking's central information problem is relational rather than purely informational. Individual departments often possess useful data, but operational significance emerges when those data are connected. Vehicle location alone may not explain an unusual fuel purchase; fuel-card data alone may not reveal whether the transaction is plausible; maintenance records alone may not explain a decline

in fuel efficiency. The intelligence layer should therefore integrate signals across existing systems rather than attempt to replace all operational software with a single monolithic platform.

This finding implies a modular architecture in which existing systems remain systems of record while the intelligence layer functions as a system of interpretation. Data connectors ingest relevant signals, a contextual layer aligns them by vehicle, driver, trip, time, and location, and analytical services evaluate whether cross-system relationships are normal, unusual, or worthy of review.

The second finding is that fleet managers face an attention-allocation problem. Managers may be responsible for dozens or hundreds of trucks while simultaneously responding to calls, repairs, dispatch changes, compliance issues, fuel transactions, invoices, driver concerns, and unexpected events. Under these conditions, the objective of AI should not necessarily be to automate every decision. A higher-value objective is to rank the vehicles, drivers, transactions, or conditions that deserve attention first.

This shifts the design target from dashboard proliferation to anomaly-oriented prioritization. Rather than requiring a manager to inspect every asset equally each morning, the system should surface a manageable queue of cases with supporting context. Such prioritization is consistent with anomaly-detection approaches in heavy-duty vehicle fuel monitoring, where unsupervised and ensemble methods can identify records that deviate from expected operating patterns (Turan et al., 2024).

The third finding is that an abnormal signal is not equivalent to a verified problem. Fuel consumption above expectation may result from route grade, vehicle weight, traffic conditions, driving behavior, weather, maintenance state, or data error. Research on heavy-duty fuel modeling confirms that fuel consumption is influenced by multiple interacting variables and that richer operational context improves predictive performance (Katreddi & Thiruvengadam, 2021; Fan et al., 2024; Zhu et al., 2024).

The same principle applies to suspicious transactions. A fuel-card purchase that does not correspond with available GPS information should be treated as a discrepancy requiring review, not as automatic

evidence of fraud. A robust intelligence layer should therefore combine anomaly scores with contextual evidence and communicate uncertainty explicitly.

The fourth finding is that explainability and human authority are operational requirements, not optional interface features. Fleet decisions can affect employment, safety, maintenance expenditure, service continuity, and customer relationships. For these reasons, the system should explain why a case was flagged, which signals contributed to the flag, how

unusual the pattern is relative to relevant baselines, and what evidence remains uncertain.

This approach aligns with research emphasizing explainable AI for trustworthy transportation systems and decision-centered AI in logistics (Kuznietsov et al., 2024; Abdulrashid et al., 2026). The manager should be able to review the evidence, request additional information, dismiss a false positive, escalate an investigation, or override an automated recommendation.

Table 1. Proposed layered architecture for AI-enabled fleet decision support.

Layer	Primary function	Illustrative inputs / outputs
1. Data integration	Connect existing systems of record and normalize identifiers and timestamps.	GPS/telematics, fuel cards, ELD/HOS, maintenance, dispatch, safety, finance.
2. Contextualization	Link signals by vehicle, driver, trip, route, time, location, and operating state.	Vehicle-trip histories, maintenance state, mileage, route and load context.
3. Analytics and anomaly detection	Estimate expected behavior and detect deviations or inconsistent cross-system patterns.	Fuel-efficiency deviation, location mismatch, unusual purchase behavior, deteriorating equipment patterns.
4. Explanation and prioritization	Rank cases by urgency and present interpretable reasons for attention.	Priority queue, contributing signals, uncertainty indicators, comparison with baseline.
5. Human decision and feedback	Enable managers to investigate, confirm, dismiss, escalate, and provide feedback.	Human disposition, notes, corrective action, validated labels for future model improvement.

Fuel is a major operational cost in trucking, making it a practical starting point for intelligence-layer validation. The system can compare recent fuel consumption with vehicle-specific historical baselines while incorporating mileage, route, operating conditions, vehicle weight, and maintenance context. Prior studies show that machine-learning models can estimate heavy-duty vehicle fuel consumption from operational signals and that vehicle weight, speed, throttle behavior, and other contextual variables materially affect prediction (Fan et al., 2024; Zhu et al., 2024). The goal in the proposed framework is not merely to predict fuel use, but to identify when a deviation is sufficiently unusual to justify managerial investigation.

A second use case is cross-validating fuel transactions against vehicle location, timing, expected tank

capacity, recent mileage, and historical purchase behavior. The appropriate output is an inconsistency alert rather than a fraud verdict. This distinction reduces the risk of over-automation and supports due process in operational investigations.

Maintenance signals can be interpreted as patterns rather than isolated readings. For example, slowly declining tire pressure may be more informative than a single pressure value, especially when combined with mileage, maintenance history, diagnostics, ambient conditions, and operating context. Recent review evidence indicates that AI-driven predictive maintenance can support vehicle health monitoring, maintenance scheduling, and fault diagnosis, while also highlighting the importance of model explainability and data quality (Mahale et al., 2025).

The same intelligence-layer logic can be extended to safety-related signals from ELDs, telematics, inspections, maintenance systems, and driver-behavior data. The system should not promise accident prevention; instead, it can identify combinations of conditions associated with elevated operational risk and prioritize them for review. In safety-sensitive settings, explanations and human oversight are especially important because opaque or overly confident recommendations can undermine trust and accountability (Kuznietsov et al., 2024).

DISCUSSION

The proposed framework reframes the value proposition of AI in trucking. The problem is not simply that managers need more dashboards or more predictions; rather, they need better interpretation of relationships among existing signals. This distinction is strategically important. A new dashboard can increase cognitive load if it adds another interface that must be monitored. An intelligence layer should instead reduce the amount of manual comparison required across systems and produce a prioritized, explainable set of cases for human review.

This perspective is consistent with broader transportation research showing that AI creates value through prediction, diagnosis, optimization, and decision support, but that real-world adoption depends on data integration and operational usability (Zhang et al., 2025). It also aligns with decision-centered research arguing that AI systems should be evaluated in terms of how they interact with organizational decision processes, not only in terms of model performance (Abdulrashid et al., 2026).

For an early-stage platform such as Dieselly.ai, the framework suggests a phased development strategy. The first phase should emphasize reliable data integration and a narrow set of high-frequency use cases where ground truth can be obtained relatively quickly. Fuel anomalies are attractive because fuel-card transactions, GPS traces, mileage, and historical consumption can often be matched at the vehicle level. Maintenance alerts provide another pathway where subsequent inspection outcomes can be used to validate whether alerts were useful.

The second phase should focus on explanation quality and workflow integration. A technically accurate alert

can still fail if managers do not understand why it was generated or if responding to the alert requires excessive navigation across systems. Each alert should therefore include the relevant evidence, the baseline against which the event was compared, and a clear distinction between observed facts and model-based inference.

The third phase should build feedback loops. Manager dispositions - confirmed issue, false positive, acceptable exception, maintenance action, policy violation, or unresolved case - can become labels for model refinement. This can turn fleet operations into a learning system while preserving human authority over consequential decisions.

- Proposition 1: Cross-functional models that combine fuel, location, vehicle, route, and maintenance context will produce more operationally useful anomaly alerts than single-source thresholds.
- Proposition 2: Attention-prioritization interfaces will reduce managerial review burden more effectively than additional standalone dashboards, provided that alert precision remains acceptable.
- Proposition 3: Explanations that identify contributing signals and uncertainty will increase manager trust and appropriate reliance compared with opaque anomaly scores.
- Proposition 4: Human feedback on alert outcomes will improve model calibration and reduce false positives over time, particularly when feedback is captured in structured categories.
- Proposition 5: The business value of an intelligence layer will depend more on workflow integration and decision latency than on marginal improvements in predictive accuracy alone.

The principal limitation is that this study is conceptual and practice-informed. The founder account reflects operational experience and product-development reasoning but is not an independent observational dataset. The framework therefore should be interpreted as a set of design propositions requiring field validation rather than as evidence of product effectiveness.

A second limitation concerns generalizability. Trucking companies differ in fleet size, operating region, equipment, telematics vendors, fuel-card providers,

dispatch practices, driver policies, and data quality. An intelligence-layer architecture must be adaptable to these differences, and model performance observed in one fleet may not transfer directly to another.

Third, AI-supported monitoring creates governance questions related to privacy, cybersecurity, employee monitoring, data ownership, bias, false positives, and the consequences of automated suspicion. Future deployments should define role-based access, audit trails, retention policies, appeal mechanisms, and clear boundaries on automated action.

Future research should evaluate the framework using prospective fleet data. A practical pilot design would begin with one or two use cases - for example fuel-efficiency anomalies and fuel-location mismatches - and compare AI-prioritized cases with existing manual review procedures. Evaluation metrics should include precision, recall, false-positive rate, time-to-detection, manager review time, resolution time, and the percentage of alerts that lead to verified operational action.

For maintenance and safety use cases, longitudinal validation is required. Models should be tested against subsequent inspections, repairs, roadside events, or other verifiable outcomes. Importantly, studies should report not only model accuracy but also whether alerts are interpretable, whether managers act on them appropriately, and whether the system reduces cognitive burden without encouraging automation bias.

CONCLUSION

The central argument of this paper is that trucking does not primarily suffer from a shortage of operational data; it suffers from fragmentation and limited capacity to interpret relationships among data in real time. An AI-enabled intelligence layer can address this problem by connecting existing systems, contextualizing signals, detecting anomalies, prioritizing managerial attention, and explaining why a case deserves review. The most appropriate role for AI in this setting is decision support rather than unqualified automation. Fuel efficiency, transaction consistency, predictive maintenance, and safety monitoring provide practical domains in which this approach can be tested. The next research step is empirical: the framework should be evaluated in live fleet environments using transparent metrics for

detection quality, decision latency, human workload, and operational outcomes. If validated, the intelligence-layer model could offer a pathway toward more efficient, transparent, and safer fleet management without requiring trucking companies to abandon the systems on which they already depend.

REFERENCES

1. Abdulrashid, I., Mohsin, M. T., Khalafalla, M., & Delen, D. (2026). A decision-centered assessment of explainable AI in transport logistics. *Journal of Computer Information Systems*, 1-28. <https://doi.org/10.1080/08874417.2026.2648967>
2. Fan, P., Song, G., Zhai, Z., Wu, Y., & Yu, L. (2024). Fuel consumption estimation in heavy-duty trucks: Integrating vehicle weight into deep-learning frameworks. *Transportation Research Part D: Transport and Environment*, 130, 104157. <https://doi.org/10.1016/j.trd.2024.104157>
3. Gong, J., Shang, J., Li, L., Zhang, C., He, J., & Ma, J. (2021). A comparative study on fuel consumption prediction methods of heavy-duty diesel trucks considering 21 influencing factors. *Energies*, 14(23), 8106. <https://doi.org/10.3390/en14238106>
4. Katreddi, S., & Thiruvengadam, A. (2021). Trip based modeling of fuel consumption in modern heavy-duty vehicles using artificial intelligence. *Energies*, 14(24), 8592. <https://doi.org/10.3390/en14248592>
5. Kuznietsov, A., Gyevar, B., Wang, C., Peters, S., & Albrecht, S. V. (2024). Explainable AI for safe and trustworthy autonomous driving: A systematic review. *IEEE Transactions on Intelligent Transportation Systems*, 25(12), 19342-19364. <https://doi.org/10.1109/TITS.2024.3474469>
6. Mahale, Y., Kolhar, S., & More, A. S. (2025). A comprehensive review on artificial intelligence driven predictive maintenance in vehicles: Technologies, challenges and future research directions. *Discover Applied Sciences*, 7, 243. <https://doi.org/10.1007/s42452-025-06681-3>
7. Turan, B. B., Genc, E., Akcig, I. N., Goztepe, N., Mumcuoglu, M. E., & Unel, M. (2024). Detecting high fuel consumption in HDVs with ensemble of anomaly detection models. In 2024 IEEE 22nd

International Conference on Industrial Informatics
(INDIN) (pp. 1-6). IEEE.
<https://doi.org/10.1109/INDIN58382.2024.10774415>

8. Wiggerthale, J., & Reich, C. (2024). Explainable machine learning in critical decision systems: Ensuring safe application and correctness. *AI*, 5(4), 2864-2896. <https://doi.org/10.3390/ai5040138>
9. Zhang, Y., Zhang, S., & Dinavahi, V. (2025). A survey of machine learning applications in advanced transportation systems: Trends, techniques, and future directions. *eTransportation*, 24, 100417. <https://doi.org/10.1016/j.etrans.2025.100417>
10. Zhu, X., Shen, X., Chen, K., & Zhang, Z. (2024). Research on the prediction and influencing factors of heavy duty truck fuel consumption based on LightGBM. *Energy*, 296, 131221. <https://doi.org/10.1016/j.energy.2024.131221>