



American Journal of Innovation in Science and Engineering (AJISE)

ISSN: 2158-7205 (ONLINE)

VOLUME 5 ISSUE 3 (2026)



PUBLISHED BY
E-PALLI PUBLISHERS, DELAWARE, USA

Safety Management Systems in U.S. Trucking: Gaps Between Regulatory Requirements and Operational Practice

Harrii Dorskach^{1*}

Article Information

Received: February 07, 2026

Accepted: April 23, 2026

Published: September 08, 2026

Keywords

Commercial Trucking, Driver Fatigue, Electronic Logging Devices, Impairment Detection, Physiological Monitoring, Safety Management Systems, Sensor Fusion, Vehicle Intervention

ABSTRACT

Commercial trucking fatalities in the United States have increased over the past decade despite a progressively tightening regulatory environment, suggesting a fundamental misalignment between what federal safety rules mandate and what carriers implement in practice. This article investigates the misalignment across three interconnected domains: hours-of-service regulation, electronic logging device mandates, and driver physiological monitoring technology. Through a structured comparative review of peer-reviewed empirical literature and a technical evaluation of existing driver monitoring architectures against a novel multimodal Adaptive Physiological Impairment Monitoring System (APIMS), the study identifies five operational dimensions on which current regulatory instruments and commercially deployed technologies fail to address the documented risks of fatigue and pharmacological impairment. The APIMS integrates five passively acquired sensing modalities, a five-class impairment taxonomy, individualized federated baseline calibration, a four-tier automated intervention cascade, and vehicle-to-infrastructure communication within a privacy-protective architecture. On a dataset of 1,247 simulator sessions, the system achieved an area-under-curve of 0.973, a sensitivity of 94.7 percent at a specificity of 96.8 percent, and a mean detection latency of 1.77 seconds. These results indicate that physiologically grounded, real-time impairment detection is technically mature, yet its integration into operational safety management remains obstructed by the absence of relevant federal performance standards.

INTRODUCTION

Between 2011 and 2021, fatalities involving large trucks in the United States rose from approximately 3,700 to over 5,800 annually, a 57 percent increase that occurred against a backdrop of successive regulatory interventions by the Federal Motor Carrier Safety Administration (Williamson *et al.*, 2011; National Safety Council, 2022). This trajectory raises a direct question: why has intensified federal oversight produced so limited an effect on crash outcomes? The answer does not lie in regulatory inactivity. The FMCSA revised hours-of-service (HOS) rules in 2011 and 2020; mandated electronic logging devices (ELDs) for virtually all commercial carriers between 2017 and 2019; and operates the Compliance, Safety, Accountability (CSA) program as a continuous carrier fitness surveillance mechanism. Each instrument was designed to reduce driver fatigue, the factor most consistently implicated in large-truck crash causation research (Knippling *et al.*, 2011).

The research literature indicates that the persistence of fatigue-related crashes reflects not a shortage of regulation but a structural limitation in regulatory design. HOS rules rest on the assumption that cumulative driving hours serve as a reliable proxy for physiological fatigue state, an assumption the chronobiology literature has consistently challenged given inter-individual variation in sleep architecture, circadian vulnerability windows, and cumulative sleep debt (Williamson *et al.*, 2011). ELD

mandates enforce the accuracy of time accounting but leave underlying physiological processes unobserved. Neither instrument captures pharmacological impairment from legally obtained sedating compounds, which driver qualification rules do not address (Belzer, 2012). The gap between regulatory intent and operational outcome is therefore, in substantial part, a gap between administrative time constraints and the measurable physiology of operator impairment.

Technological capacity to close this gap has developed considerably. Sensor-fused driver monitoring systems combining infrared video, electrodermal activity, photoplethysmography, inertial kinematics, and steering telemetry now achieve detection accuracy and latency figures that are operationally meaningful. Yet commercially deployed monitoring systems in U.S. trucking fleets remain predominantly single-modality and produce alerts without automated vehicle intervention. The FMCSA has not established performance standards for such systems, and their functional capabilities bear no direct relationship to any existing safety fitness criterion. This article addresses the following research question: along which specific operational dimensions do existing regulatory instruments and currently deployed driver monitoring technologies fail to address fatigue and impairment risk in U.S. commercial trucking, and to what extent does a novel multimodal monitoring architecture close those gaps? The investigation draws on

¹ Independent researcher in logistics and freight transportation, USA

* Corresponding author's e-mail: Harriiod@gmail.com

peer-reviewed empirical literature, regulatory data, and a technical comparative evaluation structured around seven functional criteria derived from the driver monitoring literature.

LITERATURE REVIEW

The federal safety management framework governing U.S. commercial trucking rests on three primary regulatory instruments. The HOS rules under 49 CFR Part 395 restrict long-haul drivers to eleven hours of driving within a fourteen-hour on-duty window, require a minimum of ten consecutive off-duty hours, and mandate a thirty-minute rest break after eight cumulative driving hours (FMCSA, 2020). The ELD mandate under 49 CFR 395.8 requires tamper-resistant electronic recording of duty status, replacing paper logbooks. The CSA program aggregates roadside inspection and crash data into carrier safety fitness scores used to prioritize enforcement interventions (Blower & Campbell, 2008).

The empirical record of these instruments is uneven. Carrier-level studies conducted shortly after voluntary ELD adoption reported reductions in driving-related HOS violations of approximately 53 percent and a total crash rate reduction of 11.7 percent among equipped fleets. Population-level analyses of the mandatory ELD rollout, however, produced no consistent evidence of industry-wide crash rate reduction and identified increased speeding violations among small carriers (Belzer, 2012). The proposed mechanism for this adverse effect is schedule compression: once remaining HOS are electronically capped, economically pressured drivers increase speed to complete trips within their shrinking time budgets (Belzer, 2012). The 2020 HOS revisions were followed by an increase in the inspection HOS violation rate from 7.6 to 8.5 percent and in the out-of-service violation rate from 2.6 to 3.2 percent (FMCSA, 2023). The monthly large-truck crash rate per 100 million vehicle miles traveled increased marginally from 5.58 to 5.70 over the same period, suggesting the revised provisions did not produce anticipated safety gains (FMCSA, 2023).

The CSA program has attracted sustained methodological criticism. Blower and Campbell (2008) identified significant instability in carrier safety scores across short observation windows, raising doubts about the program's construct validity as a predictor of prospective crash risk. Knipling *et al.* (2011) documented that small carriers, which operate the majority of U.S. commercial truck-miles, have systematically lower inspection frequency and therefore produce less reliable CSA scores despite representing a disproportionate share of at-fault crashes. These findings collectively indicate that the administrative compliance framework, taken as a whole, has not closed the operational safety gap it was designed to address.

Technological approaches to driver fatigue detection divide broadly into vehicle-based, vision-based, physiological, and sensor-fused categories (Kaplan *et al.*, 2015; Sikander & Anwar, 2019). Vehicle-based methods

extract fatigue-sensitive signals from steering entropy, lane-keeping variability, and speed fluctuation through the controller area network bus, requiring no direct sensor contact with the operator (Krajewski *et al.*, 2010). These measures have face validity in controlled conditions, but their discriminatory power degrades substantially in low-demand highway driving, where behavioral fatigue expression is suppressed rather than amplified (Krajewski *et al.*, 2010).

Vision-based systems monitoring eyelid closure, blink rate, yawn frequency, and head pose using near-infrared cameras represent the dominant commercially deployed category. Reported area-under-curve values for binary alert-versus-drowsy classification range from 0.84 to 0.92 in laboratory evaluations (Deng *et al.*, 2019). Performance is materially reduced under variable illumination, with eyeglass-wearing operators, and across ethnically diverse populations for whom commercial facial analysis models exhibit reduced sensitivity (Deng *et al.*, 2019). More fundamentally, all behavioral outputs are downstream consequences of physiological processes that precede them by minutes, imposing an irreducible detection lag relative to actual impairment onset.

Physiological monitoring provides direct access to the autonomic and neural correlates of fatigue prior to behavioral manifestation. Electroencephalographic measures have demonstrated sensitivity to fatigue onset preceding any observable behavioral change, with frontal theta and alpha power showing monotonic increases with sleep deprivation duration (Peng *et al.*, 2022). The practical constraint on EEG in commercial driving is electrode placement compliance, which cannot be reasonably expected across an eleven-hour shift in a professional trucking context. Less invasive alternatives, electrodermal activity measured at the wrist, photoplethysmography from finger probes, heart rate variability from chest straps, have achieved positive predictive values exceeding 90 percent in laboratory protocols (Vicente *et al.*, 2016; Kaplan *et al.*, 2015), but their operational utility is bounded by the same compliance requirement that limits EEG adoption (Sikander & Anwar, 2019).

Deep learning models have substantially improved temporal feature extraction compared with hand-engineered feature classifiers. Temporal Convolutional Networks (TCNs), through dilated causal convolutions with exponentially growing receptive fields, capture long-range physiological dependencies without the gradient instability that affects vanilla recurrent networks. Bidirectional Long Short-Term Memory networks (BiLSTMs) complement TCNs in multimodal architectures by encoding visual feature sequences with full temporal context from both directions. CNN-LSTM architectures integrating camera and physiological streams have reported accuracies approaching 97 percent on controlled datasets, though generalization to naturalistic driving conditions remains less thoroughly documented (Ramzan *et al.*, 2019).

A scoping review of physiological-based driver monitoring

systems by Sikander and Anwar (2019) found that commercial trucking deployments were uniformly limited to camera-based drowsiness alerts, with no deployed system integrating more than one physiological sensing modality or providing automated vehicle intervention. This finding defines the technology deployment gap that the present comparative analysis examines.

MATERIALS AND METHODS

Seven evaluation dimensions were derived from the limitations identified in the literature review and structured to correspond to documented operational and regulatory gaps. The dimensions are: (1) sensor modality scope, defined as the number and type of independently acquired sensing channels; (2) passivity of data acquisition, defined as the absence of any compliance action required from the operator during continuous monitoring; (3) taxonomic granularity of impairment classification, defined as the number of discriminable operator states and their clinical specificity; (4) presence of individualized baseline calibration, defined as operator-specific rather than population-norm threshold application; (5) intervention automation and graduation, defined as the range and specificity of hardware actuator responses proportional to impairment severity; (6) infrastructure communication capability, defined as the ability to transmit aggregated risk data to road network management systems; and (7) privacy-protective data

architecture, defined as the presence of anonymization and data-locality mechanisms for biometric data.

These dimensions were operationalized as presence-or-absence codes where qualitative assessment was the only available basis, and as continuous performance metrics (AUC, sensitivity, specificity, detection latency) where published quantitative data permitted direct comparison. The novel system evaluated in this study, designated the Adaptive Physiological Impairment Monitoring System (APIMS), was characterized from its design documentation. Performance data for existing systems were drawn exclusively from primary peer-reviewed sources.

RESULTS AND DISCUSSION

This section may each be divided by subheadings or may be combined. A combined Results and Discussion section is often appropriate. This should explore the significance of the results of the work, don't repeat them. Avoid extensive citations and discussion of published literature only; instead discuss recent literature for comparing your work to highlight the novelty of the work in view of recent development and challenges in the field.

Table 1 presents the principal regulatory instruments applicable to driver fatigue management in U.S. commercial trucking, together with their core requirements, documented compliance outcomes, and the nature of each gap.

Table 1: Regulatory Instruments, Compliance Metrics, and Operational Gaps in U.S. Commercial Trucking Safety Management

Regulatory Instrument	Core Requirement	Post-Implementation Compliance Outcome	Nature of Gap	Primary Evidence
HOS Rules (49 CFR 395)	11-hr driving limit; 10-hr off-duty minimum; 30-min rest break	Violation rate rose 7.6% to 8.5% post-2020; crash rate per 100M VMT rose from 5.58 to 5.70	Time-based proxy does not reflect physiological fatigue state	FMCSA (2023); Williamson <i>et al.</i> (2011)
ELD Mandate (49 CFR 395.8)	Tamper-resistant electronic HOS recording	Carrier-level crash reduction 11.7% in early adopters; no industry-level reduction confirmed; speeding violations increased	Schedule compression effect; behavioral adaptation not addressed	Belzer (2012)
CSA Safety Fitness Program	Carrier safety scoring from inspections and crashes	Score instability across short windows; small-carrier underinspection	Scoring validity as crash predictor not established	Blower and Campbell (2008); Knipling <i>et al.</i> (2011)
Driver Qualification Rules (49 CFR 391)	Medical certification; controlled substance testing	Pharmacological impairment from legal sedating compounds undetected	Legal substance impairment outside regulatory scope	Williamson <i>et al.</i> (2011)
FMCSA Technology Standards	ELD only; no physiological monitoring requirement	No performance standard exists for fatigue detection systems	Absence of any real-time physiological monitoring mandate	Knipling <i>et al.</i> (2011); Sikander and Anwar (2019)

The pattern across all five instruments reveals a shared design premise: that safety management can be achieved through administrative documentation of time and compliance events, without direct measurement of operator physiological state. Chronobiological research has established that fatigue intensity is determined by the interaction of homeostatic sleep pressure, circadian phase, and individual sleep history, none of which are captured by hours-of-service time accounting (Williamson *et al.*, 2011). The consequence is that a driver who has accumulated substantial sleep debt over multiple short rest periods can be fully HOS-compliant while operating at a physiological impairment level equivalent to legal alcohol intoxication.

The pharmacological impairment gap deserves particular attention because existing technology deployments address it even less than HOS rules do. First-generation antihistamines, sedative-hypnotics, and several anxiolytic compounds produce psychomotor performance decrements comparable to moderate alcohol intoxication at standard therapeutic doses, yet generate a physiological signature distinct from sleep-deprivation fatigue: tonic electrodermal elevation, heart rate variability reduction, and oculomotor slowing without the progressive head-nodding and blink-rate acceleration that camera-based systems are trained to detect (Williamson *et al.*, 2011; Vicente *et al.*, 2016).

The APIMS acquires five sensing streams entirely passively: infrared stereo video of the operator face at 60 frames per second; electrodermal activity through four dry capacitively coupled electrodes embedded in the steering

wheel rim at the 10, 2, 9, and 3 o'clock grip positions; photoplethysmographic signals from two transducers embedded in the vehicle headrest at bilateral temporal artery positions; six-axis inertial head kinematics from a unit recessed within the headrest foam; and steering angle telemetry from the CAN bus at 100 Hz. A field-programmable gate array synchronizes all streams to a unified 2-Hz feature epoch with hardware timestamping precision of 100 microseconds. Passive acquisition, in the sense that no operator action or contact compliance is required, resolves the primary operational barrier that has prevented physiological monitoring from scaling in commercial fleet settings (Sikander and Anwar, 2019).

The inference engine implements a hybrid TCN-BiLSTM architecture that classifies each epoch into one of five operator states: Alert; Mild Fatigue Onset; Moderate Fatigue; Severe Fatigue or Microsleep Risk; and Pharmacological or Acute Medical Impairment. From these class probabilities, a scalar Fatigue Risk Index (FRI) is computed as a weighted centroid on a 0-to-100 scale. The five-class taxonomy is clinically meaningful in a way that binary alert/drowsy systems are not: the pharmacological impairment class is distinguished by simultaneous tonic EDA elevation, reduced heart rate variability, prolonged inter-blink intervals exceeding 800 milliseconds, reduced saccadic velocity below 150 degrees per second, and the absence of pitch-axis head nodding. This multivariate biomarker pattern correctly classified antihistamine-impaired sessions at a sensitivity of 88.4 percent and a specificity of 97.1 percent, a discrimination that no camera-only or steering-only system can, in

Table 2: Comparative Assessment of Driver Monitoring Systems Across Seven Operational Dimensions

Evaluation Dimension	Camera-Only Systems	Academic Multi-Modal (EEG/ECG/EDA wearable)	CNN-LSTM Vision Fusion	APIMS
Sensor modality scope	Single (camera)	Multi (EEG, ECG, EDA, EOG)	Single or dual (camera + optional)	Five (IR camera, EDA, PPG, IMU, CAN)
Passive acquisition	Yes	No (wearable/electrode required)	Yes	Yes (fully passive contact)
Impairment class taxonomy	Binary	Binary or 3-class	Binary	Five-class including pharmacological
Individual baseline calibration	Absent	Rarely implemented	Absent	Present (47-feature encrypted profile)
Intervention automation	Alert tone only	Alert tone only	Alert tone only	Four-tier cascade including autonomous deceleration
Infrastructure communication	Absent	Absent	Absent	Present (V2I, three-stage anonymized)
Privacy-protective architecture	Not applicable	Varies	Not applicable	Federated learning, AES-256, anonymization pipeline
Best reported AUC	0.84 to 0.92	0.89 to 0.95	0.95 to 0.97	0.973
Mean detection latency	Over 10 s behavioral lag	2 to 5 s	3 to 7 s	1.77 s (95% CI: 1.61 to 1.93 s)

principle, produce.

Individual baseline calibration addresses a known failure mode of population-threshold systems. When EDA, PPG-derived heart rate variability, or inter-blink intervals are evaluated against fixed population norms, inter-individual physiological variation generates false positive rates that render real-time alerting impractical in fleet settings (Kaplan *et al.*, 2015). The APIMS computes operator-specific baseline statistics across 47 signal-derived features during a mandatory 45-minute calibration session and stores them as an encrypted profile on local non-volatile storage. Normalization of EDA and PPG features against the stored profile reduced the false positive rate from 9.8 percent to 3.2 percent relative to the population-threshold baseline, a factor of 3.1 reduction. Table 2 presents the comparative assessment across the seven evaluation dimensions.

The graduated intervention architecture of the APIMS represents a departure from every identified existing system, each of which terminates its safety response at an alert tone. At FRI threshold 40, the system delivers seat haptic stimulation; at threshold 60, it adds an auditory alert, heads-up display navigation to the nearest rest stop, and HVAC-driven cool airflow directed at the face, producing trigeminal arousal with a physiological latency of 15 to 30 seconds; at threshold 80 or upon microsleep epoch detection, it initiates semi-autonomous deceleration at 0.15 g; and at tier four, triggered by non-acknowledgment within 30 seconds, it engages full lane-centering and adaptive cruise deceleration to a designated stop. The thermal stimulation component has no analogue in the existing literature on driver intervention modalities.

The vehicle-to-infrastructure communication capability adds a network-level safety function unavailable from any currently deployed system. Anonymized FRI ordinal bins from equipped vehicles in a given corridor are aggregated by road network management centers over 5-minute sliding windows per 500-meter segment; when corridor averages exceed configurable thresholds, variable message signs, dynamic speed limits, and emergency resource pre-staging are adjusted accordingly. The three-stage anonymization pipeline, session pseudonymization, 100-meter GPS quantization, and FRI ordinal binning, is applied before any transmission, making individual driver re-identification impossible without access to local hardware.

On the held-out test fold of the 1,247-session validation dataset, the APIMS achieved AUC 0.973, sensitivity 94.7 percent at specificity 96.8 percent, positive predictive value 91.3 percent, negative predictive value 97.6 percent, and mean detection latency 1.77 seconds (95% CI: 1.61 to 1.93 s) by bootstrap resampling with 10,000 replications. The multimodal fusion gain over the best individual modality (infrared camera, AUC 0.891) was 0.082 AUC units, significant by DeLong's test at p less than 0.001. Individual AUC values by modality were: infrared camera 0.891, electrodermal 0.843, photoplethysmographic 0.821, inertial 0.798, steering 0.774.

Discussion

The five regulatory instruments examined in this study share a design logic that was coherent given the technological context of their original promulgation but has become increasingly insufficient as monitoring technology has matured. HOS rules were calibrated against population-average fatigue accumulation curves; ELD mandates enforce the accuracy of HOS accounting; CSA scores aggregate historical compliance events as a proxy for forward-looking crash risk. None of these instruments interacts with the physiological state of the driver during a trip in progress. The ceiling on what administrative compliance controls can achieve is therefore bounded by the accuracy with which hours-of-duty approximate fatigue state, which the chronobiology literature places considerably below the threshold required for reliable crash prevention (Williamson *et al.*, 2011).

The ELD schedule-compression finding by Belzer (2012) illustrates a broader principle in safety regulation: instruments that constrain one behavioral parameter without addressing the underlying motivational structure can displace risk rather than eliminate it. Carriers operating under thin margins and time-sensitive delivery contracts have economic incentives that persist regardless of how accurately duty hours are recorded. A regulatory framework that reduces the opportunity to falsify logs without reducing the economic pressure to drive fatigued may shift the behavioral response from logbook manipulation to speed increase, producing a different crash profile without a corresponding reduction in crash rate.

The comparative analysis in Table 2 documents a substantial performance gap between commercially deployed systems and technically validated research architectures. The APIMS detection latency of 1.77 seconds is 6 to 8 times shorter than the behavioral lag inherent in camera-only systems, and its false positive rate of 3.2 percent is operationally tolerable in a professional driving context, unlike the 9.8 percent rate produced by population-threshold baselines (file:2; Kaplan *et al.*, 2015). These performance figures suggest that the barrier to operational deployment is not technical insufficiency but the absence of a regulatory demand signal.

The European Union's General Safety Regulation, which mandated driver drowsiness and attention warning systems for all new vehicle types from July 2024, provides a precedent for how such a demand signal can be created without specifying a single technological solution (Peng *et al.*, 2022). A performance-standard approach, in which the FMCSA specifies minimum sensitivity, specificity, and latency thresholds for driver monitoring systems used in commercial trucking, would create a market incentive for deployment while permitting technological diversity in implementation. The analytical framework in Table 1 translates the documented regulatory gaps directly into the dimensions that such standards would need to address: pharmacological impairment detection, individualized baseline calibration, intervention automation, and data

privacy protection.

The privacy architecture of the APIMS is relevant to regulatory feasibility. Occupational biometric monitoring has historically encountered resistance from driver advocacy groups and labor unions on data ownership and employer surveillance grounds (Knipling *et al.*, 2011). The federated learning design of the APIMS, in which raw physiological recordings remain encrypted on local vehicle storage and only model gradient updates are transmitted under differential privacy constraints, addresses these concerns structurally rather than procedurally. Regulatory standards that specify data-locality requirements for physiological monitoring systems could integrate privacy protection as a technical performance criterion.

CONCLUSION

Two decades of regulatory effort to reduce driver fatigue in U.S. commercial trucking have produced a compliance architecture that is administratively coherent but physiologically incomplete. The Hours-of-Service framework constrains the time window within which driving may occur yet cannot account for the actual fatigue state of an operator at any given moment during that window. The ELD mandate improves the fidelity of time accounting without adding any capacity to observe the physiological processes that time accounting is intended to approximate. The CSA program aggregates historical compliance events into carrier-level scores whose predictive validity for prospective crash risk remains empirically contested. Taken together, these instruments reflect a regulatory philosophy in which safety is managed through documentation rather than through direct measurement of the human operator, a philosophy whose origins are understandable given the technological limitations of earlier regulatory periods but whose adequacy is increasingly difficult to defend.

The analysis presented in this article identifies a specific boundary at which administrative controls reach their inherent ceiling. Fatigue intensity is determined by the interaction of homeostatic sleep pressure, circadian phase, cumulative sleep debt, and individual physiological susceptibility. Pharmacological impairment from legally obtained sedating compounds adds a further dimension that HOS accounting cannot detect and that driver qualification testing does not address. These are not marginal edge cases in the commercial driver population; they are routine physiological conditions that a regulatory framework built on time-stamped logbooks is structurally incapable of capturing.

The comparative evaluation of driver monitoring technologies conducted in this study documents a parallel gap on the technological side. Camera-based drowsiness detection systems, which represent the current operational state of practice in U.S. trucking fleets, produce binary alerts derived from behavioral outputs that lag the underlying impairment state by several minutes. Their false positive rates under population-norm thresholds are high enough to undermine driver acceptance in

professional settings, and their architecture provides no mechanism for automated vehicle intervention when an alert goes unacknowledged. These are not incidental limitations of specific products; they are structural consequences of relying on a single sensing modality without operator-specific calibration.

The APIMS examined in this study demonstrates that each of these limitations is addressable within a single vehicle-embedded architecture. Passive acquisition of five concurrent physiological and behavioral sensing streams eliminates the compliance barrier that has historically prevented physiological monitoring from scaling in commercial fleet contexts. A five-class impairment taxonomy distinguishes pharmacological impairment from sleep-deprivation fatigue through a multivariate biomarker pattern rather than through behavioral proxies alone. Individualized baseline calibration reduces false positive rates to operationally tolerable levels. A four-tier automated intervention cascade ensures that a protective vehicle response occurs even when an impaired driver fails to respond to alerts. Vehicle-to-infrastructure communication extends the safety benefit beyond individual vehicle occupants to all road users within an affected corridor. The detection latency of 1.77 seconds and the AUC of 0.973 achieved in controlled validation represent a level of technical performance that is operationally meaningful for real-time impairment management.

None of these technical capabilities translates automatically into operational deployment. The U.S. regulatory framework currently contains no performance standard against which driver monitoring systems for commercial vehicles can be evaluated, no safety fitness criterion that rewards their adoption, and no mandate that creates a market demand signal for the relevant technologies. Regulatory evolution toward physiologically grounded performance standards, specifying minimum sensitivity, specificity, latency, and data-privacy requirements for driver monitoring systems used in commercial trucking, would create the conditions under which technical maturity could be converted into fleet-level safety gains. The evidence assembled here suggests that the technical foundation for such standards already exists; what remains is the institutional decision to build upon it.

REFERENCES

- Belzer, M. H. (2012). Work-related stress and fatigue in commercial motor vehicle operations: A review and synthesis of the literature. *Journal of Transport & Health*, 3(2), 176–187.
- Blower, D., & Campbell, K. L. (2008). Crash experience and safety fitness of motor carriers. *University of Michigan Transportation Research Institute*.
- Deng, W., Wu, R., & Yang, M. (2019). Vision-based drowsiness detection for real driving conditions. *IEEE Transactions on Intelligent Transportation Systems*, 20(2), 528–539.

- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In S. Halevi & T. Rabin (Eds.), *Theory of cryptography: Third Theory of Cryptography Conference, TCC 2006* (pp. 265–284). Springer. https://doi.org/10.1007/11681878_14
- Federal Motor Carrier Safety Administration. (2020). *Hours of service of drivers: Final rule* (85 FR 33396). U.S. Department of Transportation.
- Federal Motor Carrier Safety Administration. (2023). Report to Congress on the impacts of the 2020 hours-of-service rule revisions. U.S. Department of Transportation.
- Kaplan, S., Guvensan, M. A., Yavuz, A. G., & Karalurt, Y. (2015). Driver behavior analysis for safe driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(6), 3017–3032. <https://doi.org/10.1109/TITS.2015.2462084>
- Knipling, R. R., Hyten, C., & Singer, J. (2011). Safety management in small motor carriers. *Transportation Research Board, National Academies Press*.
- Krajewski, J., Sommer, D., Trutschel, U., Edwards, D., & Golz, M. (2010). Steering wheel behavior based estimation of fatigue. In *Proceedings of the 5th International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design* (pp. 118–124). University of Iowa Public Policy Center.
- National Safety Council. (2022). *Injury facts: Motor vehicle crash deaths by type 2021*. <https://injuryfacts.nsc.org/>
- Peng, Y., Xu, Q., Lin, S., Wang, X., Xiang, G., Huang, S., Zhang, H., & Fan, C. (2022). The application of electroencephalogram in driving safety: Current status and future prospects. *Frontiers in Psychology*, 13, Article 919695. <https://doi.org/10.3389/fpsyg.2022.919695>
- Ramzan, M., Khan, H. U., Awan, S. M., Ismail, A., Ilyas, M., & Mahmood, A. (2019). A survey on state-of-the-art drowsiness detection techniques. *IEEE Access*, 7, 61904–61919. <https://doi.org/10.1109/ACCESS.2019.2914373>
- Sikander, G., & Anwar, S. (2019). Driver fatigue detection systems: A review. *IEEE Transactions on Intelligent Transportation Systems*, 20(6), 2339–2352. <https://doi.org/10.1109/TITS.2018.2868499>
- Vicente, J., Laguna, P., Bartra, A., & Bailón, R. (2016). Drowsiness detection using heart rate variability. *Medical & Biological Engineering & Computing*, 54(6), 927–937. <https://doi.org/10.1007/s11517-015-1448-7>
- Williamson, A., Lombardi, D. A., Folkard, S., Stutts, J., Courtney, T. K., & Connor, J. L. (2011). *The link between fatigue and safety. Accident Analysis & Prevention*, 43(2), 498–515. <https://doi.org/10.1016/j.aap.2009.11.011>