Pre-Aftertreatment Combustion Signal Analysis for Heavy-Duty Diesel Diagnostics

FINAL DRAFT — cleared by counsel; for co-inventor (Steve Forbush) review · 22 June 2026.

Changes applied to this revision: - Internal ‘NOT FOR SUBMISSION’ notes block removed in full (named internal trade-secret and pending-application references) — neutral proprietary statement retained. - Deployment scope reconciled to the validated corpus: 124 engines, 3,058 tests, three continents. - Correspondence routed to starboardresearch.org; entity/byline corrected.

Published by Starboard Research, a segment of Starboard Investment, LLC. © 2026 Starboard Investment, LLC, d/b/a Starboard Research.

SAE Technical Paper — DRAFT V1

Pre-Aftertreatment Combustion Signal Analysis as a Diagnostic Geometry: Detecting Engine Degradation Through SAE J1939 Telemetric Data Independent of Tailpipe Measurement

Jeramiah Forbush, CPA, CGMA — Steve Forbush

Revealiency, an operating name of Emissions-Based Maintenance, LLC | Lehi, Utah, USA

Target Venue: SAE COMVEC (Commercial Vehicle Engineering Congress)

Draft Status: Pre-submission. Internal review only. Not for external distribution.

Abstract

Modern heavy-duty diesel engines compliant with EPA 2007/2010 and later emissions tiers employ aftertreatment systems — diesel oxidation catalysts (DOC), diesel particulate filters (DPF), and selective catalytic reduction (SCR) — that chemically transform the exhaust stream between the combustion event and any tailpipe-located sensor. While effective at reducing regulated emissions, this aftertreatment chemistry masks the underlying combustion signal: tailpipe measurements no longer reflect the cylinder-level combustion state, and conventional engine-control-module (ECM) fault monitoring is calibrated to detect aftertreatment faults rather than combustion-side degradation.

This paper presents a diagnostic geometry that reads combustion signals directly from the pre-aftertreatment SAE J1939 data stream — engine speed, engine load, exhaust gas temperature, intake air parameters, fuel rate, exhaust gas recirculation (EGR) rate, and pre-aftertreatment oxygen concentration — and computes continuous combustion-fingerprint scores against per-vehicle learned tolerance bands. The methodology requires no additional tailpipe sensors and operates entirely on data already available on the J1939 CAN bus.

Anonymized field data from a multi-continent industrial diesel deployment (124 engines and 3,058 ISO 8178 Mode 1 emissions tests across three continents, with OEM field-engineering verification) demonstrates that the methodology detects combustion-degradation patterns associated with specific mechanical conditions — including air intake restriction, turbocharger boost deficiency, injector wear, and timing drift — enabling planned maintenance interventions in advance of in-field component failure events. The paper discusses transferability of the methodology across engine displacement, manufacturer, and EPA emission tier, and identifies the conditions under which the diagnostic geometry remains valid.

Keywords: diesel diagnostics, SAE J1939, combustion fingerprint, aftertreatment, predictive maintenance, ISO 8178, fleet telematics, EPA emission tiers.

1. Introduction

Heavy-duty diesel engines manufactured under EPA 2007, EPA 2010, and EU Stage V emissions regulations incorporate aftertreatment systems engineered to convert engine-out emissions into regulatory-compliant tailpipe emissions. A DOC oxidizes carbon monoxide and unburned hydrocarbons to carbon dioxide and water. A DPF traps particulate matter and is periodically regenerated by injected fuel or active heating. An SCR system injects diesel exhaust fluid (DEF) to convert oxides of nitrogen (NOx) into nitrogen and water. The combined effect is that the gas composition measured at any point downstream of the aftertreatment is chemically different from the gas composition exiting the cylinder.

This chemical transformation is desirable from a regulatory standpoint and reflects the public-health success of modern aftertreatment design. From a diagnostic standpoint, however, it introduces a structural problem: tailpipe measurements increasingly reflect aftertreatment performance rather than combustion performance. A fleet vehicle may operate with progressively degraded combustion — fouled injectors, reduced turbocharger boost, intake-air restriction, timing drift, or accumulated cylinder-pressure imbalance — while tailpipe emissions remain within regulatory bands because aftertreatment compensates for elevated engine-out emissions. The combustion-side degradation manifests as elevated fuel consumption, accelerated component wear, and eventual unplanned failure; it does not, in general, trigger ECM fault codes.

The diagnostic problem is therefore: how can a fleet operator continuously detect combustion-side degradation when (a) tailpipe sensors are chemically blinded by aftertreatment, and (b) ECM fault monitoring is calibrated to detect aftertreatment system faults rather than combustion-side degradation?

This paper presents one answer: a diagnostic geometry that uses the SAE J1939 CAN-bus data stream directly. Every modern heavy-duty diesel engine broadcasts, on the J1939 bus, a rich set of parameters that together characterize the combustion event before any aftertreatment chemistry intervenes. These parameters are available without additional sensor installation, are broadcast continuously during normal operation, and — when read through an appropriate diagnostic geometry — provide a direct view into combustion-side health that tailpipe measurement cannot offer.

Section 2 establishes the physical basis: which J1939 parameters carry combustion-side information, and why their pre-aftertreatment provenance preserves the diagnostic signal.

Section 3 describes the diagnostic geometry — how the parameters combine into continuous combustion-fingerprint scores, and how per-vehicle tolerance bands are learned. Section 4 presents anonymized field results from a multi-continent industrial diesel deployment. Section 5 discusses the transferability of the methodology across engine displacement, manufacturer, and EPA emission tier. Section 6 addresses limitations and applicable conditions. Section 7 concludes.

2. Physical Basis: J1939 Parameters as Pre-Aftertreatment Combustion Signals

2.1 The aftertreatment compensation problem

Pre-EPA-2007 diesel engines emitted regulated pollutants directly to the atmosphere, with comparatively modest in-engine controls. Diagnostic methods developed in that era — including the ISO 8178 carbon-balance method for fuel-flow calculation and SAE J1003 emissions measurement procedures — were designed around tailpipe measurement of CO, CO2, NOx, total hydrocarbons (THC), and particulate matter. In that regime, tailpipe emissions were a direct expression of combustion chemistry, modified only by exhaust-stream mixing and cooling.

In the post-EPA-2007 regime, three distinct aftertreatment processes modify the exhaust between the turbocharger outlet and the tailpipe:

- The DOC oxidizes engine-out CO and unburned hydrocarbons to CO2 and water. The CO concentration measured downstream of a functioning DOC is therefore not the engine-out CO concentration; it is the residual after catalytic oxidation. Similarly, downstream CO2 is inflated relative to engine-out CO2 by the contribution from oxidized CO and hydrocarbons.
- The DPF accumulates particulate matter and periodically regenerates. During active regeneration, fuel is injected into the exhaust stream and combusted in or upstream of the DPF, producing additional CO2 from a fuel source that did not pass through the cylinder. Tailpipe CO2 during a regeneration event is therefore not representative of cylinder combustion.
- The SCR system injects DEF (urea solution) and converts NOx to N2 and water. Tailpipe NOx is the residual after this conversion and reflects SCR conversion efficiency at least as much as it reflects engine-out NOx.

The aggregate effect of these three processes is that tailpipe measurements, on a Tier 4 or Stage V engine, encode the combined performance of combustion and aftertreatment. Decomposing this combined signal into a pure combustion signal at the tailpipe requires either (a) deactivating the aftertreatment, which is not operationally acceptable, or (b) measuring upstream of the aftertreatment, which requires either custom sensor installation or use of signals already broadcast on the J1939 bus.

2.2 J1939 parameters that carry combustion information

The SAE J1939 protocol defines a broadcast schedule of engine parameters identified by Suspect Parameter Number (SPN) and grouped into Parameter Groups (PGNs). For combustion-side diagnostic purposes, the following parameter groups carry information that is, by construction, pre-aftertreatment:

Parameter Group	Representative SPN(s)	Diagnostic significance
Engine speed and load	SPN 190, SPN 513	Defines the operating point. All other combustion parameters are interpreted relative to load and speed; comparisons are only valid within comparable operating-point windows.
Fuel rate	SPN 183	ECM-reported fuel injected per unit time. Measures fuel dispatched to injectors, not fuel combusted; the discrepancy itself is a diagnostic signal (see Section 3).
Intake manifold parameters	SPN 102, SPN 105, SPN 132	Boost pressure, intake manifold temperature, mass air flow. Together they define the air mass available for combustion at the operating point; deviations indicate turbocharger or intake-restriction faults.
Exhaust gas temperature (pre-aftertreatment)	SPN 173 and related	EGT at the turbocharger outlet or DOC inlet reflects combustion completeness and timing. Elevated EGT at constant load is a primary indicator of late combustion, retarded injection timing, or air-side restriction.
EGR parameters	SPN 2791 and related	EGR mass flow and EGR valve position. EGR modulates intake oxygen and combustion temperature; EGR-rate deviations affect both the air-side and combustion-product side of

Parameter Group	Representative SPN(s)	Diagnostic significance
		the carbon balance.
Pre-aftertreatment O2	SPN 3217 and related	Wideband oxygen concentration measured upstream of the DOC. Carries direct stoichiometric information about combustion lambda; together with operating point, sufficient for combustion-side carbon-balance inference.
Barometric pressure and ambient temperature	SPN 108, SPN 171	Establishes ambient reference conditions; required for normalization across deployment sites at different altitudes and ambient temperatures.

Table 1. J1939 parameter groups carrying pre-aftertreatment combustion information.

Three observations about this parameter set bear emphasis. First, every parameter listed is broadcast by the ECM during normal operation; no aftermarket sensor or invasive measurement is required. Second, the pre-aftertreatment oxygen sensor — present on modern engines for closed-loop EGR control and aftertreatment management — provides direct stoichiometric information that, in earlier engine generations, would have required portable exhaust analyzer equipment to obtain. Third, the parameter set is sufficient (as shown in Section 3) to compute combustion-fingerprint scores; no additional measurement is required for the diagnostic to function.

2.3 Why the ECM does not flag combustion degradation

A natural question is why, given that the ECM has access to all the parameters listed in Table 1, the ECM does not itself detect combustion-side degradation and emit a fault code. The answer has three parts.

First, ECM fault monitoring is by design calibrated against the engine’s certified emissions configuration. The ECM emits fault codes when sensor readings fall outside calibration-defined limits or when on-board diagnostic (OBD) routines detect specific named failure modes — failed sensors, plugged DPFs, low DEF level, SCR conversion efficiency below threshold, and so on. Gradual combustion-side degradation — a 5% reduction in turbo boost, a 3% increase in fuel rate at constant load, a 15-degree-Celsius increase in EGT — typically does not violate any calibration-defined limit and therefore does not produce a fault code. The engine continues to operate in compliance, with progressively worse fuel economy and accelerated wear, until a sensor reading actually leaves its calibrated band or a downstream OBD threshold is crossed.

Second, ECM-calibrated thresholds are designed to be conservative — they must not produce false positives across the population of engines in service, including engines that have aged in service but remain operationally sound. The cost of this conservatism is reduced sensitivity to early-stage degradation. The same conservatism is appropriate for a regulator's perspective but inadequate for a fleet maintenance perspective.

Third, the ECM has no concept of per-vehicle learned tolerance. ECM thresholds are set at engine certification and apply identically to every engine of that family. A vehicle that, when new, produced a particular EGT-versus-load characteristic, and now produces an EGT-versus-load characteristic 20 degrees Celsius hotter at every load point, has clearly degraded — but the absolute EGT may still fall within the certification-wide limit. Detecting this kind of degradation requires comparison against the vehicle's own historical operating envelope, which the ECM does not maintain.

These three limitations are not defects of ECM design; they are consequences of the regulatory role the ECM plays. A complementary diagnostic geometry that reads the same J1939 parameters with a different calibration philosophy — per-vehicle, learned, tolerance-band-based — fills the diagnostic gap without conflicting with the ECM's regulatory function.

3. The Diagnostic Geometry

3.1 Operating-point binning

A diesel engine's combustion parameters vary strongly with operating point. Boost pressure, EGT, fuel rate, and pre-aftertreatment oxygen all depend on the combination of engine speed and load. Comparing parameter values across different operating points is not meaningful; comparison must be made within operating-point windows where the parameters are approximately stationary.

The diagnostic geometry therefore begins by partitioning the operating envelope into a set of operating-point bins, defined by ranges of engine speed and engine load. The specific partitioning approach — bin boundaries, count, and adaptation logic — is the subject of an underlying patented methodology. For purposes of this paper it is sufficient to state that the binning is designed to produce, within each bin, a population of J1939 readings that share approximately constant operating conditions and that can be statistically compared across time.

This bin-wise comparison structure is the central methodological move. Rather than tracking a parameter's absolute value over time — which is dominated by operating-point variation — the diagnostic tracks the parameter's distribution within each operating-point bin. Degradation manifests as a shift in the bin-wise distribution that cannot be explained by operating-condition variation, because operating-condition variation has been controlled for by the binning.

3.2 Per-vehicle learned tolerance bands

For each operating-point bin, the diagnostic geometry establishes a tolerance band: the range of values that each J1939 parameter is expected to take, for this vehicle, when the vehicle is operating in healthy condition. The tolerance band is learned, not pre-specified — it is derived from the vehicle's own historical operating data across a learning period sufficient to capture the vehicle's healthy operating envelope.

Per-vehicle learning is essential for three reasons. First, manufacturing variation across nominally identical engines is sufficient that a fleet-wide tolerance band would either be too wide to catch early degradation or too narrow to avoid false positives. Second, application-specific factors — fuel quality, ambient conditions, duty cycle, accessory loads — produce vehicle-specific operating-envelope shifts that are not degradation and should not trigger alerts. Third, the diagnostic geometry's value lies in detecting changes from a vehicle's own healthy baseline; absolute thresholds, even if accurately specified, would miss the early-stage degradation that is the primary diagnostic objective.

The specific mathematics by which tolerance bands are computed, updated, and progressed through learning phases is the subject of separate patent applications and is not disclosed here. What is disclosed is the structural role of the learned tolerance band: it is the per-vehicle, per-operating-point reference against which incoming J1939 readings are scored.

3.3 The combustion-fingerprint score

Each incoming reading is scored against the applicable per-vehicle, per-operating-point tolerance band. The scoring produces, for each tracked J1939 parameter, a continuous deviation measure expressing how far the reading falls from the center of the learned tolerance band. Deviation measures across parameters are then combined into a single combustion-fingerprint score for the reading.

The combustion-fingerprint score is continuous, not binary. Healthy operation produces fingerprint scores clustered near a baseline value; degradation produces fingerprint scores progressively further from baseline. The score's continuous character is critical: it enables trending, prediction, and intervention scheduling, rather than only post-hoc fault detection. A vehicle whose fingerprint score has been climbing slowly over six months — even if no individual reading has crossed a fault threshold — is a vehicle with developing combustion-side degradation that warrants inspection before a catastrophic failure event.

The pattern of deviation across parameter groups, rather than the magnitude of any single deviation, is the diagnostic signature. Different mechanical conditions produce different signature patterns. Air-side restriction (clogged air filter, intake leak, turbocharger boost deficiency) produces one pattern; injector-side degradation (fouling, wear, calibration drift) produces another; timing-side degradation (injection timing retard, EGR-rate drift) produces a third. The diagnostic geometry distinguishes these patterns by the joint deviation profile across the J1939 parameter set, not by any single parameter's deviation in isolation.

3.4 Relationship to underlying carbon-balance methodology

The combustion-fingerprint approach is consistent with — and informed by — the carbon-balance method codified in ISO 8178, 40 CFR Part 1065, and SAE J1003 [1][2][3]. The carbon-balance method establishes that, at constant operating conditions, exhaust CO2 concentration is proportional to fuel mass flow rate, with the proportionality constant determined by fuel chemistry. The J1939 parameter set described in Section 2 — specifically, the combination of pre-aftertreatment oxygen, fuel rate, EGR rate, intake mass air flow, and exhaust gas temperature — is sufficient to perform a carbon-balance inference without requiring direct measurement of exhaust CO2.

This inferential carbon-balance path is significant because it eliminates the need for an in-exhaust CO2 sensor while preserving the stoichiometric rigor of the carbon-balance method. The fuel rate inferred from the J1939 parameter set can be compared against the ECM-reported fuel rate (SPN 183); systematic disagreement between the two — when one is derived from combustion stoichiometry and the other from injector control — is itself a diagnostic signal. The author's earlier work [4] establishes the conditions under which the carbon-balance inference is valid and the magnitude of the structural error in ECM-reported fuel rate.

4. Field Results

4.1 Deployment overview

The diagnostic geometry has been deployed across a multi-continent industrial diesel fleet comprising 124 engines and 3,058 ISO 8178 Mode 1 emissions tests across three continents. The deployment population includes:

- Surface mining haul trucks, displacement range 50 to 175 liters, V12 and V16 configurations, EPA Tier 1 through Tier 2.
- Commuter rail and transit vehicles, mixed engine populations.
- Power-generation engines operated at fixed speed and load.

Two major engine manufacturers' field-engineering teams supervised independent verification activity in the course of the deployment. Site identification, customer identification, and engine-OEM identification are anonymized in this paper consistent with the data-handling protocols established at the deployment outset.

Three measurement methods were operated in parallel during the verification activity: (a) the J1939 combustion-fingerprint geometry described in Section 3, (b) fill-to-fill volumetric fuel measurement supervised by the OEM field-engineering team, and (c) ECM-reported fuel rate via SPN 183. The three-way comparison provides a structural validation of the J1939-derived diagnostic signals against physical ground-truth and against the conventional ECM-reported baseline.

4.2 Degradation patterns detected

Four mechanical-condition categories were associated with distinct combustion-fingerprint signatures in the deployment data:

Mechanical condition	J1939 signature pattern	Typical intervention
Air-side restriction (clogged filter, intake leak, low boost)	Elevated EGT at constant load; reduced intake MAF; reduced boost pressure relative to load; shifted pre-aftertreatment O2 toward leaner-than-expected for the bin.	Air filter replacement; intake-system inspection; turbocharger inspection.
Injector-side degradation (fouling, wear, calibration drift)	ECM-reported fuel rate (SPN 183) elevated relative to inferred fuel rate; bin-wise fuel-rate distribution widened; EGT pattern shifted without corresponding boost-pressure change.	Injector calibration; injector replacement; ECM recalibration.
Timing drift (injection timing, EGR-rate)	EGT deviation pattern with relatively preserved boost pressure and intake MAF; pre-aftertreatment O2 shifted in characteristic direction depending on whether timing or EGR is driving the drift.	Timing solenoid inspection or replacement; EGR valve and cooler inspection; ECM calibration verification.
Aggregate combustion-quality drift (multiple co-occurring contributors)	Combined deviation across all parameter groups; fingerprint score climbs progressively without a single parameter crossing an ECM fault threshold.	Staged inspection following the prescriptive maintenance workflow; sequential intervention with re-scoring after each stage.

Table 2. Mechanical-condition categories and associated J1939 combustion-fingerprint signatures observed in the deployment data.

4.3 Detection lead time relative to ECM fault codes

Across the deployment population, combustion-fingerprint scores trended upward in advance of ECM fault codes in cases where ECM fault codes ultimately fired, and identified degraded operating envelopes in cases where ECM fault codes did not fire during the observation period. In several specifically documented cases, vehicles with elevated fingerprint scores were inspected at the diagnostic system’s recommendation and were found to have the predicted mechanical conditions, despite the ECM reporting no active fault codes for those vehicles at the time of inspection.

A representative example, anonymized: a 60-liter V16 EPA Tier 1 haul truck completed a four-stage maintenance intervention (ECM calibration, timing solenoid replacement, injector adjustment, and verification) directed by the diagnostic geometry. Following the intervention, the fingerprint score returned to baseline. The vehicle was returned to service. Subsequent monitoring detected a renewed upward fingerprint trend over a period of weeks, indicating fresh degradation in the air-intake-side parameter group; the ECM reported no fault codes during this period. Inspection at the diagnostic system's recommendation identified air-intake-side issues consistent with the predicted signature. The case illustrates the geometry's role as a continuous monitoring tool rather than a single-event diagnostic.

4.4 Three-way fuel-rate comparison

In a controlled verification activity, fuel rates from the three measurement methods — J1939-derived (per the inferential carbon balance described in Section 3.4), fill-to-fill volumetric (supervised), and ECM-reported (SPN 183) — were compared across a five-vehicle subset of the deployment population. The J1939-derived fuel rates tracked the supervised fill-to-fill measurements across the population, while ECM-reported fuel rates diverged upward from the fill-to-fill measurements by margins reaching the high-teens-percent range. The structural reason for the divergence — that SPN 183 reports fuel injected rather than fuel combusted, and is therefore insensitive to return-flow fuel volumes, injector calibration drift, and combustion-efficiency variation — is examined in the author's earlier work [4].

This result has two implications for the present diagnostic geometry. First, it confirms that the J1939-derived fuel-rate inference is a structurally accurate representation of actual fuel combustion, not merely a re-derivation of the ECM-reported value. Second, it identifies the disagreement between J1939-derived fuel rate and ECM-reported fuel rate as itself a diagnostic signal: a vehicle whose ECM reports a fuel rate systematically higher than the J1939-derived rate is a vehicle with one of the structural sources of SPN-183 error, which is itself diagnostically informative.

5. Transferability Across Engine Displacement, Manufacturer, and Emission Tier

5.1 Theoretical transferability

The diagnostic geometry presented in this paper rests on three foundations: (a) the J1939 protocol's parameter definitions, which are standardized across heavy-duty diesel manufacturers; (b) the stoichiometric carbon-balance relationships, which are properties of the fuel chemistry and are independent of engine design; and (c) the per-vehicle learned tolerance-band structure, which is by construction adapted to each individual vehicle. None of these foundations is specific to a particular engine displacement, manufacturer, or emission tier.

On theoretical grounds, the methodology should therefore transfer across engine families, OEMs, and emission tiers without re-derivation of underlying constants. Verification of this theoretical transferability is, however, an empirical question.

5.2 Empirical evidence within the deployment

The deployment population spans engine displacements from 50 to 175 liters, V12 and V16 configurations, EPA Tier 1 and Tier 2 certifications, and multiple major engine manufacturers. Across this range, the combustion-fingerprint geometry has produced diagnostically actionable signals in each engine family and tier. The transferability claim within this empirical range is supported by the deployment data.

Two areas of transferability warrant explicit further validation:

- Extension to EPA Tier 4 and Stage V engines with full aftertreatment (DOC, DPF, SCR). The geometry's pre-aftertreatment positioning is well-suited to these engines in principle — the aftertreatment chemistry has no effect on the upstream J1939 parameters that the geometry consumes — but empirical validation on Tier 4 / Stage V deployments is part of planned future work.
- Extension to on-highway Class 8 fleets. The current deployment is predominantly off-highway (mining, transit, power generation). On-highway operation involves a different duty cycle and a different mix of operating-point bins; the geometry should adapt naturally through per-vehicle learning, but on-highway deployment data has not been included in this paper.

5.3 Engine-fingerprint stability

A separate transferability question is whether the per-vehicle learned tolerance bands themselves remain stable over the vehicle's service life, modulo genuine degradation. The deployment data is consistent with stable tolerance bands across the learning-phase progression for vehicles that have not undergone major component replacement; vehicles that have undergone overhauls or major-component replacement require re-initialization of the learning phase, which is handled by the underlying methodology's phase-progression structure.

6. Limitations and Applicable Conditions

The diagnostic geometry rests on a set of conditions that should be made explicit:

- J1939 broadcast completeness. The geometry assumes that the parameter set described in Section 2 is broadcast by the ECM at adequate cadence. Engines or ECM configurations that do not broadcast pre-aftertreatment oxygen, exhaust gas temperature, or EGR rate cannot be fully diagnosed by the present geometry; the geometry degrades gracefully — diagnostic specificity decreases — as parameter availability decreases, but a complete diagnostic requires complete parameter availability.

- Operating-point coverage. The per-vehicle learning phase must observe the vehicle across a sufficient range of operating-point bins to establish tolerance bands in each bin that the vehicle subsequently operates in. Vehicles operated in narrow duty cycles establish narrow tolerance-band coverage; degradation that manifests only in unobserved operating-point regions is not detected until the vehicle enters those regions.

- Fuel composition consistency. The carbon-balance inference described in Section 3.4 assumes consistent fuel composition. Fuel-composition changes — biodiesel blending, fuel-quality variation, contamination — can shift the inferred fuel rate without corresponding combustion-side degradation, and would require fuel-composition input to interpret correctly.

- Ambient-condition normalization. Operation at significantly different altitude or ambient-temperature conditions than the vehicle's learning-phase envelope requires ambient-condition normalization to maintain diagnostic specificity. The specific normalization approach is the subject of separate patent applications and is not described here.

- Non-degradation explanations for fingerprint changes. A rising fingerprint score is a signal that the vehicle's combustion behavior has shifted from its learned tolerance band. The shift may be caused by degradation, by fuel-composition change, by accessory-load change, by environmental change, or by a measurement artifact. Diagnostic interpretation requires correlation with operating context, not application of the fingerprint score in isolation.

None of these limitations is unique to the geometry presented here; any combustion-side diagnostic that operates on field data is subject to comparable constraints. The geometry's value lies not in being unconstrained but in making its constraints explicit and operating reliably within them.

7. Conclusions

Modern aftertreatment systems chemically transform exhaust between the combustion event and any tailpipe sensor, masking combustion-side degradation from tailpipe measurement and from ECM fault monitoring calibrated to the certification-emissions configuration. Detecting combustion-side degradation in this regulatory regime requires a measurement geometry that reads the combustion signal upstream of the aftertreatment chemistry.

The SAE J1939 protocol broadcasts a parameter set — engine speed and load, fuel rate, intake parameters, exhaust gas temperature, EGR parameters, and pre-aftertreatment oxygen — that together carry combustion-side information sufficient for diagnostic inference. The parameter set is broadcast during normal operation, requires no additional sensor installation, and is by construction pre-aftertreatment.

The diagnostic geometry presented here partitions the operating envelope into operating-point bins, learns per-vehicle tolerance bands within each bin from the vehicle's own healthy operating data, and computes continuous combustion-fingerprint scores against those bands. The geometry detects combustion-side degradation patterns associated with air-side restriction, injector-side degradation, timing-side drift, and aggregate combustion-quality drift, with detection lead times in advance of ECM fault-code activation.

Field data from an industrial-scale deployment across three continents, encompassing 124 engines and 3,058 ISO 8178 Mode 1 tests, supports the methodology across the engine displacement, manufacturer, and EPA-tier range represented. Extension to EPA Tier 4 and Stage V aftertreatment-equipped engines, and to on-highway Class 8 fleets, is part of planned future work. The methodology is consistent with the carbon-balance principles codified in ISO 8178, 40 CFR Part 1065, and SAE J1003, and complements rather than replaces existing emissions-certification measurement frameworks.

For fleet operators, the practical implication is that combustion-side degradation in modern aftertreatment-equipped engines is detectable without aftermarket sensor installation, using data already broadcast on the J1939 bus, with sufficient lead time for planned maintenance intervention. For OEMs and regulators, the implication is that the J1939 parameter set carries diagnostic information beyond what current ECM monitoring extracts from it, and that complementary diagnostic geometries can extract additional value from existing data infrastructure without modification to the engine or the regulatory framework.

References

[1] International Organization for Standardization. ISO 8178-1:2020, *Reciprocating internal combustion engines — Exhaust emission measurement — Part 1: Test-bed measurement systems of gaseous and particulate emissions*.

[2] U.S. Environmental Protection Agency. 40 CFR §1065.655, *Carbon-based chemical balances of fuel, DEF, intake air, and exhaust*.

[3] SAE International. SAE J1003, *Diesel Engine Emission Measurement Procedure*.

[4] Forbush, J. and Forbush, S. *Exhaust CO_2-Based Fuel Consumption Derivation for Diesel Fleet Maintenance Diagnostics*. Revealiency, an operating name of Emissions-Based Maintenance, LLC technical white paper, April 2026.

[5] SAE International. SAE J1939, *Serial Control and Communications Heavy Duty Vehicle Network*.

[6] U.S. Environmental Protection Agency. 40 CFR Part 86 and 40 CFR Part 1036, *Heavy-Duty Engine Certification and Compliance*.

[7] Heywood, J.B. *Internal Combustion Engine Fundamentals*, 2nd ed. McGraw-Hill Education, 2018.

[8] Forbush, J. and Forbush, S. *Emissions-Based Maintenance System and Method*. U.S. Patent No. 10,718,284 B2. Issued July 21, 2020.

Disclosure

The methodology described in this paper is implemented in commercial diagnostic platforms operated by the authors' organization and is the subject of issued and pending patent protection. Specific numerical values, coefficient sets, scoring-function parameters, tolerance-band-learning parameters, and operating-point-bin boundaries are proprietary and are not disclosed in this paper. The paper describes the diagnostic geometry at a level sufficient for technical evaluation and reproduction-in-principle without disclosing trade-secret implementation details.

The authors are the co-founders of the implementing organization. Correspondence: jforbush@starboardresearch.org.