



MACHINES THAT LISTEN TO ENGINES: AI-DRIVEN OPTIMIZATION OF COMBUSTION USING ALTERNATIVE FUELS

The internal combustion engine (ICE) remains the dominant mover across ground transportation, maritime shipping, aviation, and industrials machinery, and will continue to be for decades. The combination of increasingly strict global emissions regulations and a growing arsenal of alternative fuels is forcing engines to operate in ways their original design was never meant to accommodate. Traditional engine control systems, built around fixed lookup tables with closed loop feedback calibrated for conventional fuels, cannot accommodate the wide compositional variability of alternative fuels such as hydrogen, ammonia, and high blend biofuels in real time. Artificial intelligence is currently filling that gap. This article covers the rapidly advancing field of AI-driven combustion optimization. It will break down all the alternative fuels reshaping engine chemistry, the machine learning methods currently being used to predict and control the combustion reaction, the emergence of reinforcement learning as an autonomous control mechanism, and the development of digital twins as real time engine management tools. The barriers that still lie ahead will be analyzed including data scarcity, interpretability constraints, deployment challenges, and regulatory uncertainty. The central argument of this article is that AI is not a supplementary tool for engine development but a necessary one. Without adaptive, data driven optimization, the simultaneous demands of fuel flexibility, emissions compliance, and real time performance control cannot be met by a traditional calibration approach. The technical trajectories most likely to define engine optimization over the coming decade are examined in that context.

Section 1: introduction- The Accelerating Demand for Intelligent Engine Control

Whenever a freight ship crosses the Pacific, a tractor cuts through a midwestern field, or a truck climbs over a mountain pass, an internal combustion engine is doing the work. Even as electric vehicles have commanded significant industry and media focus in recent years, the ICE continues to power approximately 95% of the world's transportation energy [1,2]. The vehicle fleet globally numbers over 1.4 billion units, and the majority of those will be running on the combustion process for decades. The transition to electrification is real, but it is not right around the corner as many believe it is. It is not fast enough to meet the emissions targets that governments around the world have already written into law.

The regulatory environment surrounding ICEs has tightened substantially in recent years. Euro 7 takes effect in November 2026, the European Union has mandated a 90% CO₂ reduction from new vehicles by 2035, and the International Maritime Organization has committed to net zero shipping emissions by 2050 [1, 3, 4, 5]. These targets are not projections. They are strict legal requirements that engine manufacturers must meet using fuels and technologies available to them. That constraint is the immediate context for what this paper covers.

The problem is two sided, and both sides are ramping up their intensity. On one hand, the fuels entering the combustion chambers are changing. Hydrogen, ammonia, biofuels, and synthetic e-fuels are being developed for specific hard to electrify sectors including heavy transport, maritime, and aviation rather than as universal gasoline replacements. Biofuels and ethanol blends are more broadly applicable but still introduce important variation in combustion properties. Modern engines already compensate for this variability, with spark ignition engines adjusting ignition timing through

knock sensors and maintain stoichiometric air fuel ratios through lambda sensor feedback. The challenges are these existing mechanisms were designed for modest variation within an already known fuel group, not for wider ranged combustion properties that hydrogen, ammonia, and multifuel blend strategies introduce [2, 6].

On the other hand, the emissions compliance window is narrowing. Modern engines already operate away from their peak thermal efficiency points specifically to keep NO_x and CO within acceptable limits. This is a deliberate tradeoff that regulations are compressing further. As NO_x limits tighten and CO₂ targets become more aggressive, the margin available for that tradeoff shrinks, making the optimization harder to solve with fixed systems. Modern engine control systems are not completely static. Closed loop feedback from lambda sensors continuously adjusts the air fuel ratio around the stoichiometric target, and knock sensors allow the ECU to adjust ignition timing in real time when combustion instability is detected. These feedback systems provide meaningful adaptability within a known fuel envelope. What they cannot do is handle much wider compositional variability of hydrogen, ammonia, or variable biofuel blends. They also cannot optimize across the full set of conflicting emissions, efficiency, and durability objectives that future regulations demand [2, 6].

Artificial intelligence is emerging as a bridge between where engines currently are and where they need to be in the future to keep up with the rapidly evolving fuel landscape. Machine learning models trained on engine sensor data can generalize across a far wider range of operating conditions and fuel compositions than the calibration maps conventional ECUs rely on, capturing nonlinear combustion behavior that map based interpolation cannot resolve. Reinforcement learning operating within validated simulation environments, can explore injection and timing strategies across operating conditions that manual calibration cannot. Digital twins can maintain a real-time virtual copy of a running engine and adjust its control parameters as engine conditions rapidly change. All these tools are beginning to transform the ICE from a static mechanical system into an adaptive self-optimizing machine [2, 6, 7].

This article examines that transformation. It begins with the new fuels entering service and the combustion challenges they create. It then surveys the AI and machine learning toolkit being used to optimize engines, from artificial neural networks and long short-term memory networks to surrogate models and physics informed methods. It covers reinforcement learning as an autonomous control architecture, digital twins as the platform that integrates all these tools together. It also covers the substantial barriers that still separate current research from future deployment. The goal is to give a technically grounded but accessible account of an active and consequential area of engine research.

Section 2: A Landscape of Alternatives- The Fuels of the Future

Not all fuels are created equal, and that has never mattered more than it does today. Each alternative fuel under active development for ICE applications introduces combustion properties that fall outside the range for conventional ECU feedback. The control challenge is different in each case, but the common thread is that traditional engine development strategies are insufficient for any of them [8, 9, 10].

Understanding this nuance is the foundation for why AI is necessary. A single fuel with stable, predictable properties can be calculated with a fixed map. A diverse and variable fuel portfolio is something that cannot.

The combustion challenges are substantial. Hydrogen's laminar flame speed is 170 to 200 centimeters per second in stoichiometric conditions. The extreme flame speed creates a tendency toward two distinct but connected ignition problems in spark ignition engines. Backfire occurs when hydrogen ignites in the intake port from hot surfaces or leftover combustion gases before the intake valve closes. Pre-ignition occurs inside the cylinder when hot spots including spark plug electrodes, exhaust valves, or carbon deposits ignite the charge before the intended spark event.

Both represent control challenges that go further than what a conventional ECU calibration can manage and that require active cycle level monitoring and response [8, 10]. Hydrogen has a wide flammability range from 4% to 75% by volume in air which makes mixture formation management even more complicated.

Figure 1 compares the energy densities of the major alternative fuels against traditional diesel and gasoline.




Figure 1: Energy density comparison of alternative and conventional fuels on a gravimetric and volumetric basis [10, 13, 15].

Ammonia is gaining consideration as a zero-carbon fuel candidate for heavy transport and maritime applications. The combustion challenges center on two main properties of ammonia. Ammonia's laminar flame speed is lower than that of hydrogen or diesel, at a rate of seven to nine centimeters per second at stoichiometric conditions, making ignition difficult and combustion unstable at high loads without a co-fuel or ignition promoter [8, 12]. Ammonia has some other products including nitrogen oxides (NOx), nitrous oxide (N₂O), and if combustion is incomplete unburned ammonia in the exhaust. N₂O is a potent greenhouse gas with a warming potential about 265 times greater than CO₂ over a 100-year period. The most practical approach is dual fuel operation. This involves pairing ammonia with a small quantity of diesel or hydrogen to stabilize the ignition stage. Precisely managing the blend ratio and injection timing across varying load and temperature conditions is a multivariable optimization problem well-suited to AI-driven control [8, 13].

Figure 2 shows a radar chart comparing key combustion properties across the major alternative fuels and traditional diesel and gasoline. It makes clear why each fuel requires a specific optimization approach.




Figure 2: Combustion property comparison across hydrogen, ammonia, bioethanol, synthetic e-diesel, and conventional fuels. It shows laminar flame speed, minimum ignition energy, auto ignition temperature, research octane number, and stoichiometric air fuel ratio [8, 9, 11].

Biofuels are currently the most commercially mature alternative fuels. They hold a significant advantage of drop-in compatibility with existing engines and fueling infrastructure. Existing engine calibrations can sometimes accommodate normal blend levels; however, optimization for higher blend ratios and variable feedstock compositions requires AI based control which can adapt to the specific blend ratios and composition of the fuel [8, 16].

Figure 3 ranks the major alternative fuels by commercial maturity and technology readiness. It ranges from fully deployed first-generation biofuels on one end to early-stage microalgae derived fuels and green ammonia on the other.




Figure 3: Technology readiness and commercial maturity ranking of major alternative fuels for ICE application [10, 14, 15].

E-fuels, also called power-to-liquid fuels or electrofuels, are synthetic hydrocarbons. Their most practical advantage is drop-in capability. E-fuels can be used in existing engines and existing fuel infrastructure without modification. They each have minimal differences in their combustion processes requiring tailored optimization [9, 10].

The most immediately practical pathway involves dual-fuel or blended systems that combine conventional fuel with an alternative one in a controlled ratio. Ammonia-diesel dual-fuel systems are under active development for maritime and road applications [8, 12, 13]. Hydrogen biodiesel blends for compression ignition engines improve flame speed and thermal efficiency while lowering net carbon emissions [17].

Blended configurations introduce a second dimension of variability into the control problem. Not only does each fuel have unique combustion properties, but the properties change based on the ratio of the blend. A conventional lookup table has no ability to account for this variability. It is precisely this challenge that makes machine learning the defining technology for the next generation of engine control [2, 6].

Section 3: The Limits of Traditional Engine Control

To appreciate the innovation AI brings to engine optimization, it is crucial to understand what was there before. The conventional approach to managing an ICE relies on what is called a calibration map, also known as a lookup table. During development, engineers run an engine through thousands of operating points on a test bench, systematically varying parameters such as fuel injection timing, injection pressure, exhaust gas recirculation rate, and air fuel ratio. At every point they measure the performance and emissions of the engine and record the optimal settings. The measurements are stored in an engine control unit (ECU) after being compiled into a multidimensional table.

In operation the ECU reads inputs from a range of sensors and applies the corresponding calibrated parameters. Modern production engines are more instrumented than this basic description suggests. Lambda sensors measure exhaust oxygen content continuously, allowing the ECU to adjust fuel injection in real time to maintain the target air fuel ratio. Knock sensors that are mounted on the engine block detect the high frequency vibrations characteristic of abnormal combustion and trigger immediate ignition timing to protect the engine. Mass airflow sensors, manifold pressure sensors, coolant and intake air temperature sensors, cam position sensors, and throttle position sensors all feed into the ECU at the same time. This enables a degree of closed loop adaptability that goes well beyond a simple lookup table [2, 6]. This closed loop feedback is what allows modern engines to compensate for changes in conditions, fuel quality variations, within a known fuel family and gradual wear over the engine's service life.

Figure 4 illustrates the three generations of engine control. It goes from the static lookup table to the modern closed loop ECU with lambda and knock sensors to the AI adaptive system. It shows precisely where conventional closed-loop correction ends and where AI optimization becomes necessary: the point at which fuel chemistry variability exceeds the correction range that lambda and knock sensor feedback was calibrated to handle.




Figure 4: Three generations of engine control architecture showing the basic lookup table, the modern closed loop ECU with lambda and knock sensors containing within a pre-mapped correction range, and the AI adaptive system capable of handling wider compositional variability of alternative fuels [2, 6].

This approach is effective for a single fuel with a stable and predictable combustion process. It was designed around the assumption that fuel properties stay within a narrow and well characterized range. Gasoline and diesel do vary geographically and seasonally in properties including octane rating, cetane number, energy content, and volatility. Modern closed loop feedback through lambda and knock sensors allow the ECU to compensate for most of this variation. The limits of this approach become apparent when the fuel properties shift outside the range the closed loop system was designed to manage. An AI model faces an analogous boundary because a model trained on conventional fuel data cannot reliably extrapolate to hydrogen or ammonia combustion without retraining on representative data. The distinction is that machine learning models can be retrained, expanded, and updated as new fuel data becomes available. Alternative fuels such as hydrogen, ammonia, and variable biofuels introduce combustion property variation that are much larger than pump to pump variation in conventional fuels. Lambda and knock sensors can correct for modest deviation but they cannot reoptimize the whole combustion strategy. They cannot do this when flame speed, ignition energy, and stoichiometry of the fuel change greatly [2, 18].

Alternative fuels expose all the limitations of the static design of lookup tables. The combustion properties of hydrogen enriched natural gas blend shift as the hydrogen fraction changes from 5% to 20% to 50%. A batch of biodiesel derived from soybean oil combusts differently than one derived from palm oil or waste cooking oil, because the fatty acid composition of the individual oils differs [8, 16]. Ammonia dual-fuel combustion is sensitive to the diesel pilot quantity, the injection timing, the engine load, and the ambient temperature in ways that cannot be pre-programmed in a lookup table.

Modern engines with closed-loop lambda and knock sensor feedback handle most of this variation automatically. The ECU continuously trims fuel injection and ignition timing in response to real-time sensor data. This is why a modern gasoline engine runs acceptably at high altitude or in extreme heat without requiring a new calibration. The problem is not that conventional engines are completely blind to operating conditions. It is that the closed-loop correction mechanisms are designed to operate within a defined window around a known fuel and a known combustion regime. When ammonia replaces a significant fraction of diesel, or when a hydrogen fraction shifts from 5% to 50% in a blended fuel, the combustion reaction changes in ways that fall outside the correction window. The lambda sensor can detect that the mixture is wrong, but the ECU has no calibrated response for a combustion process it was never mapped for [2, 18]. Alternative fuel integration involves challenges beyond engine control which include fuel system materials compatibility, injector design, and in some scenarios fundamental engine geometry modifications.

The result is a trade-off that engine engineers have dealt with for decades. Calibrations are necessarily conservative, programmed to avoid the worst outcomes rather than trying to optimize the engine for the best possible performance. This causes efficiency to be left on the table. Conservative calibrations are shaped by emissions requirements. Ignition timing and mixture strength are set to keep NOx, CO, and particulate emissions within regulatory limits. The thermodynamic optimum and the emissions compliant operating point are not always the same and regulations take priority. The result is that the engine operates within regulatory emissions limits but not at peak efficiency. The gap between where the engine operates and where it could operate shows performance that is recoverable which static calibration has no mechanism to fix. For conventional fuels in stable operating environments the cost of conservative practices is manageable; however, for alternative fuels which are under tighter regulations it is not feasible to leave that efficiency on the table [6, 18].

Figure 5 illustrates the multi-objective nature of the engine calibration problem. The fundamental trade-offs between power, emissions, and efficiency are set by combustion chemistry and thermodynamics. They cannot be eliminated by any control system. What AI changes is the ability to generalize across a parameters space that no calibration table could cover. A machine learning model trained on a representative sample of data that learns the underlying nonlinear relationships between inputs and outputs and can predict accurately at operating conditions it was never shown.

The scores assigned to each strategy on each dimension are relative rankings on a scale of one to ten. They are derived from the published performance data and calibration trade-off relationships documented in the engine optimization literature, where ten represents the best achievable outcome for that specific dimension and scores reflect the well-established engineering reality that advancing one objective such as NOx reduction or peak power requires sacrificing another such as fuel efficiency or durability [2, 6, 18].

The scores assigned to each strategy on each dimension are relative rankings on a scale of one to ten. They are derived from the published performance data and calibration trade-off relationships documented in the engine optimization literature, where ten represents the best achievable outcome for that specific dimension and scores reflect the well-established engineering reality that advancing one objective such as NOx reduction or peak power requires sacrificing another such as fuel efficiency or durability [2, 6, 18].




Figure 5. Conceptual multi-objective engine optimization radar chart comparing a maximum power strategy, a minimum NOx strategy, a maximum efficiency strategy, and a balanced AI-optimized strategy across seven performance dimensions. NOx and CO control strategies differ significantly between spark ignition and compression ignition engines; the trade-offs shown here are specific to spark ignition operation and are not directly transferable to diesel combustion. The chart does not suggest that AI eliminates the underlying trade-offs between power, emissions, and efficiency. The chemistry and thermodynamics of combustion impose real constraints that no control system can override. What the chart illustrates is that the operating point a conventional single-objective calibration selects leaves significant performance available on other dimensions, and that a machine learning system searching for a higher-dimensional parameter space can find operating points that are better balanced across all objectives simultaneously than any single-objective strategy would locate. The gain is not magic. It is the result of searching for a larger and more complex optimization space than manual calibration can practically explore [2, 6, 18].

Optimization of the engine is a structurally difficult challenge to overcome. Engine calibration involves simultaneous management of multiple output variables that often conflict with each other. Advancing injection timing improves thermal efficiency and power but it also increases the peak temperature of the cylinders and NOx formation. Increasing exhaust gas recirculation reduces NOx but raises particulate matter and hydrocarbon emissions in addition to reducing efficiency. Leaning the mixture toward excess air improves fuel economy but can cause combustion issues such as misfiring and instability. Enriching the mixture stabilizes combustion but raises CO and hydrocarbon emissions. Every setting that helps one metric hurts another. Finding the operating settings that balance all these parameters simultaneously across the full range of engine speeds, loads, temperatures, and fuel compositions is a huge optimization problem of scale that manual calibration cannot solve on its own [6, 18].

The gap between the benefits that static calibration can deliver and what modern regulations and alternative fuels require is a space that machine learning now occupies. Machine learning is what will make the optimization of blended alternative fuels a possibility in the future.

Section 4: The AI Toolkit- How Machine Learning Reads an Engine

The combustion process inside a running engine generates continuous streams of information. In research and motorsport applications, in-cylinder pressure sensors can track the pressure wave of every combustion event. The cost and durability requirements of production vehicles have prevented their widespread adoption in commercial engines. Lambda sensors measure exhaust oxygen content continuously. They provide the closed loop air fuel ratio feedback that allows the ECU to maintain stoichiometric combustion in real time. Knock sensors mounted on the engine block detect the characteristic high frequency vibrations of abnormal combustion and feed that signal directly to the ignition timing control loop. Temperature sensors track combustion heat at multiple locations in the exhaust stream. NOx sensors measure nitrogen oxide production. Engine speed, injection timing, fuel rail pressure, air temperature, and numerous additional parameters are logged at rates of thousands of samples per second. A modern-day engine produces more data in single drive cycle than a human engineer could manually look through in days [2, 6].

Machine learning exists to find structure in exactly this kind of data. A trained neural network can identify the nonlinear relationships between engine inputs and combustion outcomes that no analytical equation can fully capture. Critically, different engine problems call for different AI tools. Artificial neural networks reduce the number of bench experiments required by predicting steady state performance at operating conditions that were never directly tested. This allows engineers to focus physical testing on validation rather than exhaustive parameter sweeps. Long short-term memory networks manage the time-dependent nature of transient combustion, where each cycle depends on the one before it. Surrogate models stand in for computationally expensive CFD simulations, compressing days of computation into milliseconds. Reinforcement learning agents discover optimal control strategies through direct interaction with the engine rather than from pre-labeled data. Physics-informed models ensure that predictions remain physically valid even when training data is scarce, which is the defining challenge for alternative fuel development. Together these methods collectively replace or augment every stage of the conventional calibration process, from bench testing through ECU programming to real-time control [2, 6, 21].

ANNs are the most widely used machine learning tool in engine optimization research. An ANN consists of layers of interconnected nodes, each performing a weighted sum of its inputs followed by a nonlinear activation function. When trained on a sufficient dataset of engine operating conditions and measured data, an ANN learns to approximate the mapping between inputs including injection timing, fuel pressure, and blend ratio. It also can map NOx emissions, brake specific fuel consumption, and indicate mean effective pressure [2, 16].

Figure 6 shows the architecture of a feedforward ANN application. It traces the path from sensor inputs through hidden layers to combustion output predictions.




Figure 6: Architecture of a feedforward ANN. Displaying sensor inputs, hidden layers, and combustion output predictions [2, 16].

The value of ANNs in engine development is not that they eliminate experimental testing but that they change how the testing is used. A trained ANN can interpolate across a parameter space from a more strategically selected set of data points which reduces the total number of bench runs needed to characterize a given operating region rather than eliminating them. The reduction in required test points depends heavily on the complexity of the combustion system and the accuracy demands of the application, but the ability to fill in the gaps between the measured points with a trained model rather than running all the combinations shows a meaningful gain in efficiency [2, 16]. Research published in 2025 applying ANNs to hydrogen enriched ICEs demonstrated strong predictive accuracy for both performance and emissions metrics across a wide range of variable input parameters including hydrogen fraction and load conditions [16]. Case studies from Bosch, Toyota, and Mahindra have demonstrated AI-powered engine management systems that reduce tailpipe emissions by up to 18% and improve fuel efficiency by approximately 10% compared to conventional lookup table calibration, with results measured across standard drive cycle testing

on gasoline and mild hybrid powertrains [2]. It is worth noting that hybrid powertrains represent one of the strongest use cases for AI based engine control because the interaction between the combustion engine and the electric motor creates an optimization space with more degrees of freedom than a conventional drivetrain. Machine learning is well suited to managing that complexity in real time.

One of the limitations of ANNs is that they treat each prediction independently of one another. This is a massive limitation because engine combustion is not independent. Each combustion cycle depends on the thermodynamic state left by the previous one. The behavior of the engine is based on the history of the past several seconds of operation and changes based on the load and environment changes. Long short-term memory (LSTM) networks are a type of recurrent neural network. They are specifically designed to learn from the previous results and data. They learn from sequential time dependent data by maintaining an internal memory state that carries information across each step [22, 23].

LSTMs have proven valuable for real time NOx estimation. Physical models of NOx formation are accurate but require solving complex chemical kinetics equations. These equations cannot be computed fast enough to run on an ECU during normal engine operation [23]. Lookup tables are fast but inaccurate, especially under rapidly evolving environments where the engine is changing load and speed. LSTM networks trained on experimental NOx data can estimate engine-out NOx in real time with accuracy that far outweighs lookup tables. It additionally can do this at a low enough cost for ECU deployment. A 2025 study applying a hybrid LSTM and multi-head attention model to selective catalytic reduction (SCR) after treatment systems demonstrated simultaneous prediction of NOx outlet concentration and ammonia slip under short driving conditions [24].

Figure 7 quantifies the performance advantage of LSTMs compared with static lookup tables. It compares NOx prediction error under both steady state and changing engine operating conditions across four control strategies.




Figure 7: NOx prediction error comparison across fixed lookup tables, standard ANNs, LSTM networks, and LSTM plus attention hybrid models [22, 23, 24].

Deep neural networks (DNNs) extend the ANN architecture by adding more layers. They can learn hierarchical representations of complex combustion properties because of the added layers. They have been applied to multi-output prediction problems where an individual model must simultaneously estimate exhaust gas temperature, NOx concentration, and after treatment conversion efficiency from a set of engine operating input parameters [22]. Deep learning has also been applied to predicting pollutant emissions from ICEs under transient conditions in a 2024 study that showed the ability to predict CO₂, NOx, and various other outputs from real world driving data with high accuracy [2].

Convolutional neural networks (CNNs) were originally developed for image recognition, are being adapted for combustion analysis because of their ability to detect spatial and temporal patterns within structured data. The accuracy these networks achieved is sufficient to guide real time operational adjustments, one of the main reasons for AI being implemented to combustion optimization [21]. The architectural principles of CNNs transfer nicely to ICE combustion optimization.

Surrogate Models for Engine Design Optimization

Before the engine reaches the road, machine learning plays a distinct role in the design and calibration process itself, separate from the real-time control methods discussed elsewhere in this section. The surrogate model is one of the most impactful applications of machine learning for engine development. Surrogate models are data efficient approximations of high-fidelity simulations that can be evaluated in milliseconds rather than hours. The standard workflow for these models involves running a carefully selected set of CFD simulations and the results are used to train a machine learning model. Then that model is used to search the optimization space are a scale that a full CFD model would never allow [18, 19, 20].

The ActivO framework, developed at Argonne National Laboratory and recognized with an R&D100 Award, demonstrated that this approach can reduce the number of function evaluations needed to reach a global optimum by 80% compared to genetic algorithms and other conventional methods. Not only could it do that but it also maintained and improved solution quality. The AutoML-GA approach further advanced the process by automating the selection of machine learning hyperparameters through Bayesian optimization, removing the need for expert tuning of the surrogate model. It also made the workflows accessible to engineers without deep machine learning expertise [19]. A 2025 study at AIAA developed a bi-fidelity ANN surrogate that combined massive quantities of low-fidelity CFD data with a small number of high-fidelity experimental measurements. It achieved an accurate combustion pressure prediction while drastically reducing experimental testing requirements [25].

By contrast, the AI methods used for real-time, in-operation control, including LSTMs for NOx estimation and reinforcement learning for injection timing, operate under entirely different computational constraints and serve a different purpose [18, 19].

Figure 8 compares the time needed for engine combustion system design optimization from a range of five methods. It does from full CFD as the baseline through progressively more efficient machine learning surrogate models.




Figure 8: Reduction in engine design optimization time using machine learning surrogate methods relative to full CFD baseline [19, 25, 26].

For alternative fuel applications, experimental data is scarce and CFD models of hydrogen or ammonia combustion are computationally expensive, the surrogate model approach provides a method to achieve optimization from limited amounts of data.

AI in Real-Time Engine Control

Gaussian process regression (GPR) is a probabilistic machine learning method that produces not only a prediction but also an estimate of the confidence for every single prediction. GPR's ability to quantify uncertainty is well suited for experimental optimization problems where every single data point is expensive to obtain, and the designer needs to know which data points the most accurate and which ones are the least accurate [27].

In a 2024 study published in the journal Data-Centric Engineering, GPR surrogate models were applied to a gas turbine combustor optimization problem. While gas turbines and internal combustion engines differ substantially in their combustion architecture and operating regimes, the multi-objective optimization challenge they share, simultaneously managing emissions, stability, and efficiency across a continuous parameter space, makes Bayesian surrogate modeling approaches applicable to ICE calibration for alternative fuels. In this study, the surrogate model simultaneously addressed NOx reduction, thermoacoustic stability, and lean extinction limits. The Bayesian approach allowed the optimizer to find the most informative experimental data along each step, converging on a Pareto-optimal combustor design. This approach only needed a small amount of trial runs compared with the likes of a traditional grid search approach [27].

An ongoing limitation of data driven by ML models is their tendency to produce physically unrealistic predictions in areas where the training data is poorly represented. For traditional fuels with large experimental datasets, this is manageable. For hydrogen and ammonia combustion, where experimental data sets are limited, it can turn into a large problem. Models trained on limited ammonia engine data may extrapolate to conditions it has never seen and produce predictions that violate thermodynamic or chemical kinetics constraints [21].

Physics informed machine learning (PIML) also known as physics informed neural networks (PINNs) when implemented as neural networks, addresses this by storing physical laws directly into the model's training objective. The model is punished during training not only when its predictions do not reflect the data but also when they violate known physical relationships. These relationships include basic laws of nature such as energy conservation or reaction equilibrium constraints. The result is a model that respects the underlying physics even when only given small datasets [21]. For alternative fuel combustion, where new chemical reactions are still being characterized, physics informed models provide more reliable extrapolation than solely data driven models. They are emerging as the preferred method for applications where safety and accuracy are critical to the underlying systems.

Figure 9 shows the relative strengths and limitations of each major AI method covered within this section. It compares all six approaches across five key metrics in a radar chart to guide understanding where each method is applied to the best of its capability.




Figure 9: AI method comparison radar chart scoring ANNs, LSTM, DNN/CNN hybrids, surrogate models, GPR, and PIML [2, 6, 16, 21, 24, 27].

Section 5: Reinforcement learning- Teaching Engines to Optimize Themselves

The machine learning methods described in Section 4 are all supervised. They learn from labeled datasets of inputs and outputs produced by experiments. Reinforcement learning (RL) is completely different. It learns by interacting with the environment, taking actions, observing the consequences of those actions, and adjusting its behavior to maximize a cumulative reward signal. It requires no prior knowledge of the system it controls, no labeled training data, and no assumption about the mathematical form of the relationship between its actions and their outcomes [28, 29, 30]. In practice, however, RL agents cannot freely experiment on production engines. Unrestricted exploration during training could produce dangerous combustion events, hardware damage, or emissions violations. For this reason, RL agents for engine control are almost always trained first in high-fidelity simulation environments and then transferred to real hardware under carefully constrained conditions, a process known as sim-to-real transfer. Safety boundaries are enforced throughout training to prevent the agent from exploring control actions outside the acceptable operating envelope [29, 30].

A supervised model can predict what will happen if an engineer sets a specific setting such as injection timing. An RL agent can discover what injection timing strategy should be used across all scenarios. It does this with the goal of optimally balancing NOx reduction against fuel efficiency, without being told what the answer is beforehand [29, 30].

Figure 10 shows the RL agent training loop, from the agent's action selection through the environmental response and reward signal to the updating of the model which drives learning.




Figure 10: RL agent training loop diagram showing the cyclical interaction between the agent action selection, engine, or simulation [28, 29].

The foundational RL algorithm applied to engine control is Q-learning, a model free method that learns the value of each action in every state and uses each of those values to select the most optimal action. In a 2024 study published in the International Journal of Engine Research, Q-learning was applied to the dilute combustion limit problem in spark ignition engines, where the challenge is to run the engine as lean as possible for maximum efficiency. The challenge was to prevent cycle-to-cycle combustion variability (CCV) that results in rough running or misfire [28].

The Q-learning agent was trained to adjust the fuel injection quantity on a cycle-by-cycle basis, using the previous cycle's combustion result as part of its state representation. The agent learned to maintain stoichiometric combustion at the lean limit while minimizing CCV. It achieved a net increase in fuel conversion efficiency of 1.33% compared to conventional control. The result demonstrates that RL based cycle-by-cycle control can find efficiency improvements that static calibration cannot [28].

Deep reinforcement learning (DRL) combines RL with DNNs as the function approximator for the agent's policy and value estimates. The combination of these two allows the agent to manage continuous, high-dimensional state and action space that Q-learning cannot manage. DRL enables optimization over injection strategies involving multiple sequential fuel pulses per cycle. Each has its own timing and quantity which creates a control problem that traditional methods cannot effectively manage [29].

A landmark study from Argonne National Laboratory and the National Renewable Energy Laboratory applied DRL to the control of a compression-ignition engine using a multi-pulse injection strategy. The DRL agent was required to simultaneously maximize engine work while minimizing the NOx emissions. The agent discovered a control policy that lowered NOx emissions threefold compared to the baseline injection strategy while engine work decreased by only 2%. The DRL approach achieved this result without needing any prior knowledge of the engine system without assuming differentiability of the objective function [29].

Figure 11 directly compares the performance of a DRL agent against conventional engine calibration across six key metrics, showing the scale of improvement DRL delivers without requiring any prior knowledge of the engine system.




Figure 11: Head-to-head comparison of deep reinforcement learning against conventional engine calibration across six performance dimensions [28, 29, 30].

The DRL framework used transfer learning across multiple physical models to rapidly accelerate training. It started with a simplified combustion model and progressively refined the policy on high-fidelity simulations. This approach significantly lowered the number of real engine interactions needed to achieve a well-trained model. This addresses one of the fundamental barriers to realistic RL deployment in engine development programs [29].

The maritime sector has emerged as one of the most active testbeds for DRL based engine optimization. This is driven by the International Maritime Organization's strict Carbon Intensity Indicator (CII) and Energy Efficiency Existing Ship Index (EEXI) regulations. A 2024 study applied DRL to optimize energy efficiency across five vessel types. These included cruise ships, oil tankers, bulk carriers, container ships, and car carriers. They operated with bio-LNG and hydrogen as alternative fuels [31].

The DRL algorithm was trained to dynamically adjust engine parameters across six operational situations, each representing unique combination of cargo load and weather conditions. Fuel efficiency improved by up to 10%, EEXI values decreased by 8% to 15%, and CII ratings improved by 10% to 30% depending on the specific scenario. Under large cargo loads, the DRL optimized vessel achieved a fuel efficiency of 7.2 nautical miles per ton compared to 6.5 nautical miles per ton with conventional methods [31].

Figure 12 breaks down the DRL performance improvements by vessel type across fuel efficiency, EEXI reduction, and CII rating improvement, displaying how magnitude of benefit varies with vessel size and specifics about the operation.




Figure 12: DRL based engine optimization performance improvements across five vessel types operating on bio-LNG and hydrogen [31].

Research at Ain Shams University applied deep Q-network (DQN) reinforcement learning to the problem of idle speed control in ICE. They compared RL performance against conventional proportional-integral-derivative (PID) control. The RL approach achieved an improved average fuel consumption reduction of around two grams per minute at idle. They did this by optimizing throttle position and ignition timing in real time which is not possible with conventional methods. The study validated this result on both simulation and physical engine hardware, proving that the RL controller's superior performance in simulation translated directly into real world circumstances [30].

Deploying RL agents on production ECU hardware remains a large challenge. Modern ECUs have limited memory and computational bandwidth compared to the servers RL agents are trained on. Research in the areas of model compression, neural network quantization, and pruning are actively working to reduce the computational storage of trained RL agents to make them deployable onto modern ECUs. The combination of these techniques with ever evolving ECU design gives the belief that RL based engine control will be commercially feasible within the next decade [2, 30, 32].

Section 6: Digital Twins- A Living Model of Every Engine

A simulation of an engine is built once, validated against test data, and then used to answer questions about how the engine behaves under conditions that have not yet been evaluated. A digital twin is something completely different. A digital twin is a computational model that maintains continuous, real-time alignment with a specific physical engine through a live data connection. As the physical engine operates and changes, its sensor outputs data directly into the digital twin, which

is then able to update its state to reflect the changes of the real engine in every moment [33, 34].

The distinction matters because engines do not behave the same way across their time in operation. Wear accumulates in the injectors and valves, changing the fuel spray characteristics from the values assumed in the original calibration. Carbon deposits on the injector tips alter the spray pattern. Engine oil viscosity changes with temperature and degradation. Fuel composition varies between the different batches. Any of these changes can shift the combustion properties away from the conditions in which the calibration was based on. This reduces efficiency and increases emissions without causing a fault code that would notify the driver. A digital twin tracks these shifts and minute adjustments in real time. It can compensate for them by adjusting control parameters before the performance of the engine significantly decreases [33, 34, 35].

A combustion digital twin normally integrates three components. A physics-based model provides the thermodynamic framework governing the combustion process. It usually reduced order mathematical correlations and empirical approximations rather than detailed chemical kinetics. A data driven machine learning layer learns the deviations between the physics model and the engine's real behavior. It compensates for modeling errors and physical changes that the first-principles model does not account for. A real time data assimilation layer continuously ingests sensor data from the actual engine and updates both model components to keep up with the ever-changing state of the engine [34, 36].

Figure 13 shows the architecture of a combustion digital twin, tracing the real time data path from the actual engine through the sensor network to the physics model, machine learning correction layer, and control output feedback loop.




Figure 13: Digital twin architecture showing real time data from an actual engine and sensor network through a physics-based model and machine learning layer [34, 36].

The most complicated current implementations use physics-constrained neural networks. These embed the governing equation of fluid mechanics and thermodynamics right into the neural network architecture. They make sure that the data-driven layer never produces predictions that violate the physical laws of fluid mechanics or thermodynamics [36]. This is valuable for alternative fuel engines where the chemical kinetics of hydrogen and ammonia combustion are still being learned about and studied. The physics model may have regions of uncertainty that purely data-driven models could exploit in not physically manners.

The most mature commercial deployment of combustion digital twin technology is in the aerospace sector. Rolls-Royce has integrated digital twin technology with its flight data monitoring and health management system for the Trent XWB turbofan engine. The digital twin constantly analyzes real time engine data, simulates operational behavior, and identifies deviations from the expected performance which may indicate faults within the system. In 2024, the Trent XWB digital twin successfully predicted around 100 potential faults in advance of their occurrence. This enabled preventive maintenance on the systems and avoided additional costs [37]. While the Trent XWB is a gas turbine rather than an ICE, it demonstrates the operational maturity of the real-time sensor fusion, physics-based modeling, and machine learning anomaly detection architecture that ICE digital twin research is now actively replicating for alternative fuel combustion management [33, 35].

The principles of this aerospace application transfer to ICE optimization. An ICE running a variable blend of ammonia and diesel is a system whose behavior is constantly changing in blend ratio, temperature, and wear state. A digital twin that tracks all these slight changes and movements within the engine can optimally adjust injection timing, pilot fuel quantity, and EGR rate. This leads to near-optimal combustion across varying conditions that static calibration has no ability to account for [37].

A 2024 paper published in iScience by researchers at the Universite libre de Bruxelles and Cambridge University outlined the critical requirements for real world deployment of digital twins in industrial combustion systems. They paid particular attention to the challenges caused by renewable synthetic fuels [35]. The authors identified five key enablers. They included high-quality training data, robust feature extraction from high dimensional sensor streams, adaptive simulation frameworks that update their sub grid models based on more recent data, reduced order models that approximate high-fidelity simulations at ECU compatible computational costs, and lifelong learning capabilities that enable the twin to improve continuously from experience operating rather than the twin needing to be retrained often [34].

Figure 14 shows the digital twin lifecycle from the design phase through commissioning to field deployment. It shows how it continuously learns and the twin's role as it evolves across each stage of the engine's life.




Figure 14: Digital twin lifecycle diagram showing progression from design phase through field deployment [34, 35].

A similar line of work at ECCOMAS 2025 presented a digital twin specifically designed for combustion system design. It combined automated CFD data generation with machine learning based surrogate optimization. The system used open-source CFD tools to create training data across a range of injector designs and combustion chamber designs. It then trained machine learning surrogates to guide the design optimization process and achieved high predictive accuracy while reducing the number of required CFD simulations [35]. The designed design-phase digital twin is complementary with the operational digital twin that runs on deployed hardware. Together they cover the full engine lifecycle from development to operation in the field.

The development of AI based injection pressure optimization for hydrogen-spirohyra biodiesel dual-fuel engines, published in Nature Scientific Reports in 2025, represents one of the clearest examples of real hardware integration of digital twins [17]. The study showed that expanding machine learning based optimization using larger datasets, real time adaptive control, and digital twin capabilities support intelligent combustion management across a range of operating circumstances that static calibration has no ability to compensate for. It has direct use for broader industrial integration.

Section 7: What is still in the way

The case for AI driven combustion optimization is strong and the results in the preceding sections are real. The distance between a promising research result and a widely deployed production system is always longer than it appears from within the laboratory. Engine optimization involves safety critical hardware operating in demanding environments, regulated by governments that require demonstrations of repeated reliability, and integrated into vehicles and vessels that must function reliably for years without expert maintenance. All these requirements place intense constraints on AI deployment within engines that do not exist in a research environment.

Data Scarcity

The performance of a machine learning model is directly tied to the quality and quantity of its training data. For traditional diesel and gasoline engines decades of experimental testing have produced large, well-defined datasets that provide a solid foundation for machine learning model training. For hydrogen ICEs, ammonia dual-fuel engines, and many complicated biofuel blends, experimental databases are rare [21]. Hydrogen combustion at high engine loads, ammonia ignition under cold start conditions, and biofuel performance across the entire range of feedstock composition are all areas where training data required for dependable machine learning models does not yet exist at the scale necessary for production deployment.

Physics informed machine learning provides partial mitigation by constraining model predictions to be physically consistent even with sparse datasets. Synthetic data generation using lower-fidelity simulation, validated against the limited available data is another approach being investigated. Transfer learning, where a model trained on traditional fuel data is fine-tuned on a smaller alternative fuel dataset has shown promise in a few recent studies [21]. However, none of these approaches fully solve the problem for the kind of comprehensive, well-instrumented experimental campaign that produced high quality training data. Building datasets like that for alternative fuels is an infrastructure investment that the research community and industry are still in the initial stages of making.

Interpretability and Certification

A DNN with millions of parameters is a black box. It receives inputs and produces outputs through a chain of mathematical operations that human engineers cannot directly understand. The opacity of this process is acceptable in many applications, but engine control is not one of them. When an AI controller makes a decision that results in unusual combustion, a regulator needs to understand why that decision was made. When a manufacturer seeks type approval for an AI based ECU a certification authority needs to verify that the controller will perform safely in all circumstances. When a maintenance engineer discovers a problem, they need insights about the problem that a black box model cannot provide [38].

Explainable AI (XAI) is an active research field specifically addressing this problem. Techniques such as SHAP (Shapley Additive exPlanations) values, attention mechanism that highlight which inputs most influenced a given output, and gradient based saliency maps that visualize which traits drive a model's predictions are being applied to engine machine learning models [38]. Physics informed models provide natural explanations of their behavior better than data driven models because they are constrained by physical laws of nature. The fundamental interpretability tradeoff is the most accurate models tend to be the least interpretable and the certification process for AI based safety critical automotive control does not allow for this inability to interpret why the model made specific decisions [21, 38].

Figure 15 ranks the six primary deployment barriers by severity, technicality needed, regulatory, safety, and standardization categories. It identifies data scarcity and combustion safety validation as the most critical constraints currently.

What Is Blocking AI-Driven Engine Optimization From Reaching Production




Figure 15: Severity ranking of key barriers to widespread deployment of AI based combustion optimization [2, 5, 21, 38].

Real-Time Deployment of Embedded Hardware

The ECU is not a server. Production automotive ECUs have memory measured in megabytes, processor clock speeds measured in the hundreds of megahertz, and real time operating requirements measured in microseconds per combustion cycle. The DRL agents and DNN surrogate models that achieve the best results in the research settings are trained and evaluated on GPU clusters with hundreds of gigabytes of memory. Deploying a model trained in that environment onto an ECU requires techniques for model compression, quantization, and pruning that reduce model size and computational cost without unacceptable degradation in accuracy [2, 30].

Edge AI research is specifically targeting this gap by developing neural network architectures designed for embedded hardware from the start rather than retrofitted after training in unrealistic deployment environments. Lightweight ANN architectures trained with knowledge distillation have shown that substantial compression is achievable while preserving most of the performance benefit [2, 32]. In these models a large high-performance model teaches a smaller deployable model in the form of mega-knock, a severe low speed pre-ignition event driven by the interaction of ethanol's high charge cooling effect, elevated cylinder pressures, and oil dilution. This can cause catastrophic engine damage and traditional knock sensor feedback alone is not sufficient to prevent effectively. The low-probability nature of these failure modes compounds the challenge. A machine learning model trained mainly on normal operating data will not encounter these events frequently enough to learn a reliable detection and response strategy, leaving a gap that neither conventional control nor current AI approaches fully address [10, 11].

Combustion Safety and Alternative Fuel Risk

Alternative fuels introduce combustion failure modes that traditional engine control was not designed to manage and that AI system trained mostly on traditional fuel data may not be able to properly interpret. Hydrogen ICEs are prone to pre-ignition and backfire, especially at high loads because hydrogen can ignite from hot surfaces in the intake port before the intended spark event occurs. Ammonia presents toxicity risks if leaked in an unburned form making incomplete combustion a safety concern rather than just an emissions issue. The wide flammability range of hydrogen creates the chance of combustion events outside the intended operating system if the mixture of the control system fails. High concentration ethanol blends introduce a specific failure mode in the form of mega-knock, a severe low speed pre-ignition event driven by the interaction of ethanol's high charge cooling effect, elevated cylinder pressures, and oil dilution. This can cause catastrophic engine damage and traditional knock sensor feedback alone is not sufficient to prevent effectively. The low-probability nature of these failure modes compounds the challenge. A machine learning model trained mainly on normal operating data will not encounter these events frequently enough to learn a reliable detection and response strategy, leaving a gap that neither conventional control nor current AI approaches fully address [10, 11].

An AI controller that has learned to optimize NOx emissions and efficiency under standard operating conditions may not have been exposed to enough training examples of more obscure and rare operating conditions that do not respond appropriately to being exposed to those conditions. Formal safety validation methods, including worst case scenario analysis, hardware in the loop testing across the full operating envelope, and conservative fallback control laws that activate when the AI system detects that it has left its training distribution are all active areas of research and development [2, 21]. Regulatory agencies in Europe and the United States are starting to develop frameworks for certifying and approving AI based safety critical automotive systems. However, those frameworks are not completed yet and the validation burden they will inflict on manufacturers remains unclear.

Regulatory Uncertainty

The regulatory environment for both alternative fuels and AI based engine control is actively changing. The EU's revision of its 2035 ICE targets, announced in December 2025, changed the compliance framework in ways that affect which fuels manufacturers have an incentive to develop, and which AI optimization capabilities are needed [4, 39]. The EPA's Tier 5 rulemaking for non-road engines is progressing through a political environment that has introduced considerable uncertainty about timing and stringency [5]. The IMO's 2050 net zero

This uncertainty trickles down into AI development decisions. A manufacturer considering a major investment in ammonia dual-fuel optimization needs regulatory confidence that ammonia will be a viable compliance pathway before committing themselves to developing the fuel. The absence of universal benchmarks for evaluating AI based engine control systems compounds the problem. Separate manufacturers, research groups, and regulatory bodies are using completely different metrics to evaluate performance. This makes it difficult to compare results across the field and establish industry wide standard which corporations can base their investment off [2, 5]. Harmonization of evaluating the engine optimization frameworks across the world is necessary for widespread AI driven combustion optimization deployment and unfortunately it is currently missing.

Section 8: The Road Ahead- Intelligence Meets Alternative Fuels

The trajectory of AI driven engine optimization is not simply more than what already exists. The research frontier is moving toward architecture and methodologies that are qualitatively more capable than the supervised learning and standalone RL systems that currently dominate. A couple of deployments within the next decade are likely to define the field.

Physics informed neural networks embedded in real time digital twins represent what is becoming the dominant computational architecture for the next generation of engine optimization. Rather than choosing between physics-based models that are interpretable and data-driven models which are more accurate, PIML digital twins combine both ideas within a single framework that learns from data while abiding by thermodynamic and chemical kinetic constraints [21, 34, 36]. For alternative fuel engines where training data is limited this hybrid approach provides both the reliability and adaptability that purely data driven models cannot deliver on.

Multi-objective optimization is evolving from a research capability to normal engineering practice. Current production ECU calibrations typically optimize a small number of objectives, normally just optimizing NOx and fuel consumption, with other emissions species managed by dedicated after treatment systems. Emerging AI optimization frameworks simultaneously manage power output, NOx, CO, unburned hydrocarbons, particulate matter, combustion stability, and long-term hardware durability within a unified reward structure [2, 40]. As these frameworks develop and deployment hardware becomes more capable, the gap between research grade multi-objective optimization and production grade engine control will continue to shrink [6, 40].

Federated learning offers a path toward solving the data scarcity problem without needing centralized data sharing. In a federated learning framework, each vehicle or vessel trains a local model update on its own data and shares only the model parameters, not the raw sensor data, with a central server that aggregates the updates into and improves global model [44]. Applied to alternative fuel combustion, this strategy could allow a fleet of hydrogen ICE vehicles to collectively build training datasets for rare operating circumstances. Some of these circumstances include a cold start at high altitude with a high hydrogen fraction which no individual vehicle would encounter frequently enough to learn on its own from the data. The privacy and security properties of federated learning make it compatible with the data governance requirements of both automotive manufacturers and maritime operators.

Generative AI is starting to appear within fuel design applications. In 2023, a research team published framework where deep learning models trained on molecular property data was used to design fuel blend formulation with specified combustion targets. These included octane rating, sooting tendency, and energy density [7]. The capability of using AI to propose new fuel formulas and then proceed by using machine learning surrogate models to rapidly predict their engine performance without experimental testing could potentially compress development timelines for the future of alternative fuel formulations. This creates a feedback loop between fuel chemistry and engine optimization that did not exist previously.

Edge AI hardware is rapidly advancing to the degree that the gap between what can be done in a research setting and what can be performed on a production ECU is closing. Automotive grade neural processing units with matrix multiplication hardware and low power consumption are entering production for driver assistance. The same silicon is also being evaluated for engine management applications. Within the current development cycle of six years per vehicle generation, it is possible that DNN based combustion control will be production ready for commercial applications [2, 32].

The broader implication of these developments is that the combustion engine is not coming anywhere near its technological limit. For over a century, engine optimization has been constrained by the speed at which engineers could characterize, model, and calibrate combustion behavior. AI removes that limitation. A digital twin running on physics informed machine learning models can evaluate combustion conditions in a single second than a calibration engineer could in over a year. A DRL agent can discover injection methods that no engineer would ever think to try. The result is not a replacement for engineering judgement but an amplifier and evolution of it. An evolution that will allow engines to burn new fuels in new environments to meet emissions targets that would be unfathomable by conventional practices.

Section 9: Conclusion

The ICE faces a more demanding environment today than it ever has. Regulatory targets that were ambitious when they were originally drafted are now getting signed into law. Alternative fuels with combustion properties that original engine design was never meant for are entering the industry. And the expectation that engines must optimize themselves in real time across all blends of fuel, load, temperature, and operating condition has made the traditional lookup table obsolete and an artifact of a simpler era.

What makes this moment different from previous inflection points in engine development is that the optimization challenge has outgrown what any human-designed calibration system can solve. The number of fuel types, blend ratios, operating conditions, and emissions constraints now intersecting simultaneously exceeds the capacity of lookup tables, fixed maps, and rule-based control logic. AI does not merely improve on conventional calibration. It addresses a class of problems that conventional calibration was structurally incapable of solving.

The barriers to full commercial deployment are real and cannot be ignored. Training data for alternative fuel combustion is minimal. Deep learning models are tough to interpret in the way safety certifications require. Real-time deployment of ECU hardware remains a technically demanding challenge. Regulatory frameworks for AI based safety critical engine control are not universal and are still being developed. The policy regarding the environment for alternative fuels is constantly changing and makes long-term commitment difficult for corporations.

The convergence of regulatory pressure, fuel diversity, and AI capability is not coincidental. Each force is accelerating the others. Regulations tighten the performance envelope until only adaptive systems can operate within it. Alternative fuels expand the combustion parameter space until no static calibration can cover it. And AI provides the only optimization framework that scales with both. The direction of travel is not in question. What remains to be determined is how quickly the data infrastructure, certification frameworks, and edge deployment hardware can catch up to what the research already demonstrates is possible.

The combustion engine has been refined by generations of engineers working at the constraints of manual calibration and physical testing. AI does not replace that engineering tradition, but it extends its reach into operating spaces and fuel chemistries that no lookup table can ever cover. The engine that burns hydrogen, ammonia, or any alternative fuel reliably and efficiently a decade from now will owe that performance to both the engineers who designed it and the machine learning systems that learned, in real time, how to make it run.

Citations:

1. Interregs. (2024). European emission standards. Wikipedia. https://en.wikipedia.org/wiki/European_emission_standards

2. Sen, B. (2025). AI-powered emission control and fuel optimization in IC engines. International Journal of Science, Engineering and Technology, 13(4). https://www.ijset.in/wp-content/uploads/_IJSET_V13_issue4_154.pdf

3. RAC. (2025). Euro 7 emissions standard: What is it and when does it come into effect? <https://www.rac.co.uk/drive/advice/emissions/what-is-euro-7-and-when-does-it-start/>

4. Euronews. (2025, December 16). EU carmakers to comply with 90% emissions reduction by 2035 as full combustion engine ban scrapped. <https://www.euronews.com/my-europe/2025/12/16/eu-carmakers-to-comply-with-90-emissions-reduction-by-2035-as-full-combustion-engine-ban-s>

5. Association of Equipment Manufacturers. (2025). Navigating the global emissions regulations for non-road engines and powertrains. <https://www.aem.org/news/navigating-the-global-emissions-regulations-for-nonroad-engines-and-powertrains>

6. Aliramezani, M., Koch, C. R., & Shahbakhti, M. (2022). Modeling, diagnostics, optimization, and control of internal combustion engines via modern machine learning techniques: A review and future directions. Progress in Energy and Combustion Science, 88, 100967. <https://www.sciencedirect.com/science/article/abs/pii/S0360128521000654>

7. Schweidtmann, A. M., et al. (2023). Artificial intelligence-driven design of fuel mixtures. Cell Reports Physical Science, 4(1). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9814251/>

8. Khujamberdiev, R., & Cho, H. M. (2025). Hybrid fuels for CI engines with biofuel hydrogen ammonia and synthetic fuel blends. Energies, 18(11), 2758. <https://www.mdpi.com/1996-1073/18/11/2758>

9. Uchida, N., et al. (2025). E-fuels in IC engines: A key solution for a future decarbonized transport. International Journal of Engine Research. <https://journals.sagepub.com/doi/10.1177/14680874251325296>

10. Ismael, M. A., et al. (2025). Recent fuel-based advancements of internal combustion engines: Status and perspectives. Energy & Fuels. <https://pubs.acs.org/doi/10.1021/acs.energyfuels.5c00057>

11. Choi, B., et al. (2025). Sustainable carbon-free fuels: The potential of hydrogen/ammonia and future industrialization strategies. International Journal of Automotive Technology. <https://link.springer.com/article/10.1007/s12239-025-00399-2>

12. Zhang, C., et al. (2025). A comprehensive comparison of green ammonia and green methanol from a full chain. Carbon Neutral Systems. <https://link.springer.com/article/10.1007/s44438-025-00009-9>

13. Boztas, E., et al. (2025). Techno-economic and life-cycle assessment of hydrogen-ammonia fuel blends in tugboat engines for sustainable port operations. Sustainability, 17(22), 10285. <https://www.mdpi.com/2071-1050/17/22/10285>

14. IDTechEx. (2024). Sustainable biofuels & e-fuels market 2025-2035: Technologies, players, forecasts. <https://www.idtechex.com/en/research-report/sustainable-biofuels-and-e-fuels-market/1039>

15. Sustainability Global. (2025). Top ten sustainable fuels 2025. <https://sustainabilityglobal.org/article/top-ten-sustainable-fuels-2025/>

16. Zbikowski, M., & Teodorczyk, A. (2025). Machine learning for internal combustion engine optimization with hydrogen-blended fuels: A literature review. Energies, 18(6), 1391. <https://www.mdpi.com/1996-1073/18/6/1391>

17. Ramalingam, K., et al. (2025). AI based optimization of injection pressure for hydrogen and spirogyra biodiesel dual fuel engine to enhance combustion performance and emission characteristics. Scientific Reports, 15. <https://www.nature.com/articles/s41598-025-34179-w>

18. Perez-Gordillo, D. S., et al. (2025). Evaluation of machine learning models for optimization of internal combustion engines: A comparative study. International Journal of Engine Research. <https://journals.sagepub.com/doi/abs/10.1177/14680874251377140>

19. Owoyele, O., & Pal, P. (2022). A novel machine learning-based optimization algorithm (ActivO) for accelerating simulation-driven engine design. Applied Energy, 285, 116455. https://www.researchgate.net/publication/357861324_Artificial_Intelligence_and_Data_Driven_Optimization_of_Internal_Combustion_Engines

20. Badra, J., et al. (2021). Engine combustion system optimization using computational fluid dynamics and machine learning: A methodological approach. Journal of Energy Resources Technology, 143(2). <https://asmedigitalcollection.asme.org/energyresources/article/143/2/022306/1086007>

21. Ihme, M., & Chung, W. T. (2024). Artificial intelligence as a catalyst for combustion science. Proceedings of the Combustion Institute, 40, 105730. <https://waitong94.github.io/assets/ihme2024ai.pdf>

22. Quach, N. Y., et al. (2023). NOx emissions prediction in diesel engines: A deep neural network approach. Environmental Science and Pollution Research. <https://link.springer.com/article/10.1007/s11356-023-30937-3>

23. MathWorks. (n.d.). Using deep learning networks to estimate NOx emissions. <https://www.mathworks.com/company/technical-articles/using-deep-learning-networks-to-estimate-nox-emissions.html>

24. Shi, Y., et al. (2025). Neural network-based control optimization for NH3 leakage and NOx emissions in SCR systems. Processes, 13(7), 2029. <https://www.mdpi.com/2227-9717/13/7/2029>

25. Pavalavanni, P. K., & Yang, S. (2026). A feature-focused neural network based bi-fidelity surrogate model for combustion prediction in aerial internal combustion engines. AIAA SciTech Forum 2026. <https://arc.aiaa.org/doi/10.2514/6.2026-2035>

26. Argonne National Laboratory. (n.d.). HPC modeling and artificial intelligence to advance hydrogen-fueled gas turbines for power generation. <https://www.anl.gov/partnerships/hpc-modeling-and-artificial-intelligence-to-advance-hydrogenfueled-gas-turbines-for-power-generation>

27. Reumschuessel, J. M., et al. (2024). Data-driven optimization of a gas turbine combustor: A Bayesian approach addressing NOx emissions, lean extinction limits and thermoacoustic stability. Data-Centric Engineering. [https://www.cambridge.org/core/services/aop-cambridge-core/content/](https://www.cambridge.org/core/services/aop-cambridge-core/content/view/0415A99D66D539ACBB9173AE2A1CD7A6/S2632673624000297a.pdf)

[view/0415A99D66D539ACBB9173AE2A1CD7A6/S2632673624000297a.pdf](https://www.cambridge.org/core/services/aop-cambridge-core/content/view/0415A99D66D539ACBB9173AE2A1CD7A6/S2632673624000297a.pdf)

28. Maldonado, B. P., Kaul, B. C., Schuman, C. D., & Young, S. R. (2024). Reinforcement learning applied to dilute combustion control for increased fuel efficiency. International Journal of Engine Research. <https://journals.sagepub.com/doi/10.1177/14680874241226580>

29. Henry de Frahan, M. T., et al. (2022). Deep reinforcement learning for dynamic control of fuel injection timing in multi-pulse compression ignition engines. International Journal of Engine Research, 23(5). <https://journals.sagepub.com/doi/full/10.1177/14680874211019345>

30. Omran, I., et al. (2024). Deep reinforcement learning implementation on IC engine idle speed control. Ain Shams Engineering Journal, 15(5), 102670. <https://www.sciencedirect.com/science/article/pii/S2090447924000455>

31. Karagiannidis, N., et al. (2024). Optimizing maritime energy efficiency: A machine learning approach using deep reinforcement learning for EEXI and CII compliance. Sustainability, 16(23), 10534. <https://www.mdpi.com/2071-1050/16/23/10534>

32. Liu, Z., et al. (2025). Artificial-intelligence-based energy management strategies for hybrid electric vehicles: A comprehensive review. Energies, 18(14), 3600. <https://www.mdpi.com/1996-1073/18/14/3600>

33. RTInsights. (2026, January 3). Digital twins transition to intelligent, AI-driven systems in 2026. <https://www.rtinsights.com/digital-twins-in-2026-from-digital-replicas-to-intelligent-ai-driven-systems/>

34. Parente, A., & Swaminathan, N. (2024). Data-driven models and digital twins for sustainable combustion technologies. iScience, 27(4), 109349. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10946323/>

35. Centini, M., Cozza, I. F., & McAlpine, R. (2025). Development of a digital twin for combustion systems using automated machine learning. ECCOMAS 2025. https://congressarchive.cimne.com/dte_aicomas_2025/abstracts/fbaa86c18fb411efb344000c29ddfc0c.pdf

36. Lim, J. T., et al. (2025). Physics-informed machine learning for intelligent gas turbine digital twins: A review. Preprints.org. <https://www.preprints.org/manuscript/202509.1360>

37. Xiao, L., et al. (2025). Simulation modeling method for aero engines based on digital twins. In Advances in Intelligent Systems and Computing. Springer. https://link.springer.com/chapter/10.1007/978-981-95-0942-3_2

38. Balasubramaniam, S., et al. (2024). Explainable AI: Applications, challenges, current solutions and future research directions. Proceedings of the 12th International Conference on Information Technology (ICIT). <https://dl.acm.org/doi/10.1145/3718391.3718408>

39. National Law Review. (2025). EU launches call for evidence: Revision of combustion engine and vehicle emissions regulations. <https://natlawreview.com/article/eu-launches-call-evidence-revision-combustion-engine-and-vehicle-emissions>

40. Jamal, T., et al. (2025). An artificial intelligence-based strategy for multi-objective optimization of internal combustion engine performance and emissions. Expert Systems with Applications. <https://www.sciencedirect.com/science/article/abs/pii/S0957417425000946>

41. Interregs. (2025). First two Euro 7 emissions implementing regulations published. <https://www.interregs.com/articles/spotlight/321/first-two-euro-7-emissions-implementing-regulations-published>

42. Abdelhafez, A., et al. (2025). Development of intelligent learning-based surrogate models for predicting diesel ignition at different thermodynamic conditions. Applied Thermal Engineering. <https://www.sciencedirect.com/science/article/abs/pii/S0196890425011628>

43. Sarkar, A. (2024, January 15). Understanding feedforward neural networks. LearnOpenCV. <https://learnopencv.com/understanding-feedforward-neural-networks/>

44. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 54, 1273–1282.

Biographies

Dr. Raj Shah is a Director at Koehler Instrument Company in Holtsville, New York, where he has served for over three decades, contributing to the advancement of petroleum, fuels, lubricants, and analytical instrumentation technologies worldwide. Over the course of his distinguished career in the energy and chemical engineering industries, he has become widely recognized for both his technical leadership and sustained service to global professional societies. Dr. Shah is an elected Fellow by his peers at ASTM International, the Institute of Chemical Engineers (IChemE), the Chartered Management Institute (CMI), the Society of Tribologists and Lubrication Engineers (STLE), the American Institute of Chemists (AIC), the National Lubricating Grease Institute (NLGI), the Institute of Measurement and Control (InstMC), the American Oil Chemists’ Society (AOCS), the Institute of Physics (IOP), The Energy Institute (EI), and The Royal Society of Chemistry (RSC). These fellowships reflect his multidisciplinary impact across chemical engineering, tribology, measurement science, energy technology, and applied chemistry. He is also the recipient of the prestigious ASTM Eagle Award and ASTM’s highest honor, the Award of Merit (Fellow), recognizing more than 30 years of leadership and contribution to Committee D02 on Petroleum Products, Liquid Fuels, and Lubricants. He recently co-edited the bestseller, Fuels and Lubricants Handbook: Technology, Performance, Properties, and Testing, a major reference work for the industry. Dr. Shah has now authored and co-authored over 750 technical publications, conference papers, and industry articles, and continues to be an active contributor to the scientific and engineering literature. Further information regarding his work and recognitions can be found at <https://shorturl.at/I7000>.



Nate Polishook

Mr. Nate Polishook is a intern at Koehler Instrument Company in Holtsville, New York where he researches fuel and petroleum as part of a thriving internship program

Author Contact Details

Dr. Raj Shah, Koehler Instrument Company
• Holtsville, NY11742 USA
• Email: rsah@koehlerinstrument.com
• Web: www.koehlerinstrument.com

