



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi – 110042

B.TECH FINAL YEAR PROJECT REPORT
Department of Mechanical Engineering – Academic Year 2025...26

Effect of Ammonia on the Characteristics of a Diesel Engine

A Comprehensive Literature-Based Investigation with Machine Learning Analysis

Aditya Kumar
Aditya Patel
Aditya Yadav

Submitted by

2K22/ME/16
2K22/ME/19
2K22/ME/21

Under the Esteemed Guidance of
Dr. Raghvendra Gautam
Associate Professor
Department of Mechanical Engineering
Delhi Technological University, Delhi – 110042

Certificate

This is to certify that the project report entitled “Effect of Ammonia on the Characteristics of a Diesel Engine” submitted by Aditya Kumar (2K22/ME/16), Aditya Patel (2K22/ME/19), and Aditya Yadav (2K22/ME/21) in partial fulfilment of the requirements for the award of the degree of Bachelor of Technology in Mechanical Engineering from Delhi Technological University, Delhi, is a record of bonafide work carried out by them under my guidance and supervision during the academic year 2025–26.

To the best of my knowledge and belief, the work presented in this report has not been submitted elsewhere for the award of any degree or diploma. All sources of information used in the preparation of this report have been duly acknowledged in the references.

Dr. Raghvendra Gautam

Associate Professor

Department of Mechanical Engineering

Delhi Technological University

Shahbad Daultapur, Main Bawana Road, Delhi – 110042

Acknowledgement

We express our sincere gratitude to Dr. Raghvendra Gautam, Associate Professor, Department of Mechanical Engineering, Delhi Technological University, for his expert guidance, constant encouragement, and constructive feedback throughout this project. His extensive work in alternative fuels and combustion engineering — spanning over 51 publications with 648+ citations — gave this study a strong intellectual foundation, and his willingness to engage with our questions at every stage made the experience genuinely rewarding.

We thank the faculty and staff of the Department of Mechanical Engineering, DTU, for their support and for providing the academic environment that made this work possible. We are also grateful to Delhi Technological University for its resources and infrastructure.

Finally, we thank our families for their patience and encouragement through this demanding final year.

Aditya Kumar (2K22/ME/16) Aditya Patel (2K22/ME/19) Aditya Yadav (2K22/ME/21)
B.Tech Mechanical Engineering, Batch 2022–26, Delhi Technological University

Abstract

Growing concerns over fossil fuel depletion and tightening global emission standards have compelled the engineering community to look beyond conventional diesel as the sole fuel for compression-ignition engines. Among the alternatives receiving serious research attention, ammonia (NH_3) stands out because it carries no carbon whatsoever — meaning its combustion produces no carbon dioxide, no carbon monoxide from its own chain, and no soot. It can also be synthesised entirely from renewable electricity and atmospheric nitrogen, giving it a credible pathway to a genuinely carbon-neutral fuel cycle.

This project investigates, through a structured review of published experimental literature and a supporting machine learning analysis, how the progressive substitution of diesel with ammonia — in a dual-fuel compression-ignition configuration — alters engine performance, combustion behaviour, and exhaust emission characteristics. Data and findings are drawn from eight carefully selected peer-reviewed studies spanning the period 2011 to 2024, covering a range of engine types, operating speeds, load conditions, and ammonia energy fractions from zero to fifty percent.

The review reveals that moderate ammonia substitution, up to roughly fifteen to twenty percent of total fuel energy, can reduce CO_2 emissions by eighteen to twenty percent and smoke opacity by forty to fifty percent with an acceptable brake thermal efficiency (BTE) penalty of less than three percentage points. Beyond this threshold, combustion instability grows, ammonia slip increases in the exhaust, and BTE declines more steeply. NO_x emissions, driven primarily by fuel-bound nitrogen rather than the thermal mechanism familiar from pure diesel operation, rise sharply even at low ammonia fractions and represent the most significant emissions management challenge. Selective catalytic reduction aftertreatment, which can utilise exhaust ammonia as its own reductant, emerges as the natural and necessary partner technology.

Three machine learning models — an Artificial Neural Network, a Random Forest ensemble, and XGBoost — were trained on a dataset assembled from the reviewed literature to predict NO_x emissions and BTE. The Random Forest model delivered the highest accuracy, with R^2 values of 0.974 for NO_x and 0.968 for BTE. Feature importance analysis confirmed that ammonia energy

fraction is the single most influential variable for NO_x, while engine load dominates BTE prediction. These findings provide quantitative guidance for engine calibration and control strategy development in future ammonia-capable powertrains.

Keywords: *Ammonia, dual-fuel engine, compression ignition, NO_x, brake thermal efficiency, carbon-free fuel, machine learning, Random Forest, ANN.*

CHAPTER 1 INTRODUCTION

1.1 The Energy and Emission Challenge

Diesel engines have been the workhorses of human civilisation for well over a century. From the trucks that keep supply chains moving, to the fishing boats that feed coastal communities, to the generators that power hospitals during grid failures, these machines run on a fuel that is energy-dense, widely available, and well understood. Yet this very success has become a source of deep concern. The world burns roughly a billion tonnes of diesel fuel each year, and every kilogram of that fuel releases approximately 3.16 kilograms of carbon dioxide when combusted. Aggregate this across the global fleet of commercial vehicles, ships, agricultural equipment, and stationary generators, and the contribution of diesel combustion to greenhouse gas inventories becomes a figure that climate scientists and policy makers can no longer ignore.

The consequences of this carbon loading are now measurable and accelerating. Average global temperatures are rising, weather patterns are shifting, and sea levels are creeping upward. In parallel, the NO_x and particulate matter co-emitted with diesel exhaust continue to exact a serious public health toll — the World Health Organisation estimates that outdoor air pollution, to which diesel engines are a major contributor, causes approximately 4.2 million premature deaths annually worldwide. These twin pressures — climate and health — have produced a wave of regulatory responses. The European Union is preparing its Euro 7 standards, which will impose the most stringent limits yet on NO_x and particulate emissions from light and heavy vehicles. India has enforced Bharat Stage VI norms since April 2020, cutting permissible NO_x from heavy-duty diesel engines to just 0.40 g/kWh. And in April 2018, the International Maritime Organisation adopted an Initial Strategy committing the global shipping fleet to reduce total annual GHG emissions by at least 50 percent compared to 2008 levels by 2050, while pursuing efforts toward full decarbonisation.

Meeting these targets while maintaining the power density and range that heavy-duty applications require is genuinely difficult. Battery-electric drivetrains, which have transformed the light-vehicle market, struggle to scale to the energy requirements of a bulk carrier crossing an ocean or a loaded articulated lorry climbing a gradient. Hydrogen, despite its outstanding energy-to-mass ratio,

remains burdened by the complexity and cost of cryogenic or ultra-high-pressure storage. These realities have pushed researchers to explore a class of fuel that sits between hydrogen and conventional hydrocarbons: carbon-free chemical compounds that can be stored and handled with relative ease. Among these, ammonia has emerged as the most technically and logistically credible candidate.

1.2 Why Ammonia Deserves Serious Attention

Ammonia is not a new chemical. Industrial civilisation has been producing it in enormous quantities since Fritz Haber and Carl Bosch developed the synthesis process bearing their names in the early twentieth century. The Haber-Bosch process, which combines atmospheric nitrogen with hydrogen over an iron catalyst at high temperature and pressure, currently produces around 180 million tonnes of ammonia per year, the vast majority of which goes into fertiliser production. The global infrastructure for ammonia — storage tanks, pipelines, shipping terminals, safety protocols — therefore already exists at the scale needed for a major fuel deployment. This is a practical advantage that no other zero-carbon fuel can yet match.

From a chemical standpoint, ammonia contains no carbon atoms at all. Its molecular formula is NH_3 : one nitrogen and three hydrogens. Complete combustion of ammonia with stoichiometric air yields only nitrogen gas and water vapour. No CO_2 , no CO, no soot, no unburned hydrocarbons from the ammonia itself. For an engineer designing a powerplant to meet the IMO 2050 target, this is an enormously attractive characteristic.

The energy content of ammonia, at a lower heating value of 18.6 MJ/kg, is admittedly less than half that of diesel at 43 MJ/kg. But when considered as a liquid — which ammonia becomes at just -33°C at atmospheric pressure, or at room temperature under modest pressure of around 10 bar — its volumetric energy density becomes more competitive, and its storage and transport characteristics are manageable with well-proven industrial equipment. Compared to liquid hydrogen, which requires -253°C , ammonia is dramatically easier to store and distribute.

Equally important is the production pathway. Green ammonia, produced by combining green hydrogen — made through electrolysis powered by wind or solar electricity — with atmospheric

nitrogen, carries essentially zero carbon over its lifecycle. Several large-scale green ammonia projects are already under construction in countries with abundant renewable resources, including Australia, Chile, and Saudi Arabia, with the explicit aim of exporting the fuel to energy-importing nations. The trajectory is clear: ammonia is transitioning from an agricultural chemical into an energy carrier, and the engineering community needs to understand how to burn it efficiently in real engines.

1.3 Problem Statement

Despite its chemical appeal, ammonia presents a set of combustion challenges that make it unsuitable as a standalone fuel in conventional diesel engines. Its auto-ignition temperature of 651 °C is more than three times that of diesel, which means it simply will not ignite under the conditions that prevail at the end of the compression stroke in a standard compression-ignition engine. Its laminar flame speed of roughly 0.07 m/s is about one-quarter of that of hydrocarbon fuels, so even once ignition is achieved, the flame propagates slowly through the charge and combustion can be incomplete within the available crank-angle window. Its minimum ignition energy is significantly higher than that of petrol or diesel vapour. And finally, ammonia's nitrogen content means that, unlike combustion of a hydrocarbon, burning ammonia always risks converting fuel-bound nitrogen to NO_x — a mechanism quite distinct from the thermal NO_x formation familiar from conventional engines.

Dual-fuel operation, in which a small quantity of diesel is injected to serve as a pilot ignition source while the bulk of the energy is supplied by a premixed ammonia-air charge, is the most widely studied and practically accessible approach to overcoming these limitations. Understanding the precise effects of ammonia addition — how much fuel energy can be replaced, at what cost to efficiency, and with what emission trade-offs — is the central engineering question that this project addresses.

1.4 Objectives

This study pursues the following specific objectives:

1. To conduct a critical review of published experimental research on ammonia-diesel dual-fuel compression-ignition engines, covering the period 2011 to 2024.
2. To analyse and quantify the effect of progressive ammonia energy fraction on brake thermal efficiency, brake specific fuel consumption, and power output.

3. To examine how increasing ammonia substitution alters combustion characteristics including ignition delay, peak cylinder pressure, heat release rate, and cyclic variability.
4. To characterise the emission trade-offs associated with ammonia use, with particular attention to NO_x, CO₂, smoke, and unburned ammonia slip.
5. To apply machine learning models — Artificial Neural Network, Random Forest, and XGBoost — to predict NO_x and BTE, and to identify the most influential operating parameters through feature importance analysis.
6. To recommend a practical operating envelope for ammonia-diesel dual-fuel engines that balances performance, stability, and emission requirements.

1.5 Scope and Methodology

This is a literature-based analytical project. No independent engine experiments were conducted. Published experimental data from eight peer-reviewed papers, selected on the basis of relevance, experimental rigour, and recency, form the evidentiary core of the study. Numerical data extracted from figures in these papers using the open-source WebPlotDigitizer tool were assembled into a master dataset of approximately 300 data points, which underpins both the trend analysis presented in Chapters 5 through 7 and the machine learning work described in Chapter 8. The study focuses exclusively on compression-ignition engines operating in ammonia-diesel dual-fuel mode; spark-ignition and pure ammonia engines are referenced only for contextual comparison.

CHAPTER 2 LITERATURE REVIEW

2.1 Foundational Work: Proving the Concept

It is worth recalling that the idea of using ammonia in an internal combustion engine is not new. During the Second World War, Belgium briefly used bus engines converted to run on ammonia when petroleum was unavailable, demonstrating that combustion was at least possible. Scientific investigation in earnest began much later. The study by Reiter and Kong (2011), conducted at Iowa State University, is generally regarded as the modern starting point for systematic CI dual-fuel ammonia research. Working with a single-cylinder diesel engine and varying the ammonia energy fraction from zero to sixty percent, they demonstrated clearly that the concept was practically viable: the engine ran, CO₂ emissions fell in proportion to the ammonia fraction, and smoke was reduced. They also documented the core challenges that subsequent research has been trying to resolve ever since: NO_x rose substantially, BTE declined at high substitution rates, and combustion instability became problematic above about forty percent EF. Their work set the research agenda for the decade that followed.

2.2 Performance and Combustion Studies (2018–2022)

Valera-Medina and colleagues published what has become the definitive review article on ammonia for power in 2018, covering combustion fundamentals, engine studies, turbine applications, and production pathways. Their compilation of laminar burning velocity data confirmed that ammonia's flame speed is exceptionally low compared to any hydrocarbon fuel, but also highlighted that elevated pressure and temperature conditions — such as those found inside a diesel cylinder — improve things somewhat. This work provided the thermodynamic and chemical kinetic framework that subsequent experimental researchers have drawn upon when explaining their observations.

The study by Oh et al. (2022), which is among the reference documents for the present project, provides particularly useful insights because it examined three distinct experimental methodologies on the same engine platform: fixed volumetric fuel flow rate, fixed brake torque, and variable air-fuel ratio. This systematic approach allowed them to disentangle the effects of ammonia's lower heating value from the effects of its combustion characteristics. Their finding that the initial burn duration — from ignition to ten percent mass fraction burned — increases progressively with energy fraction confirmed that the slow flame propagation of ammonia affects even the earliest stage of

combustion, before the bulk of the charge has been consumed. They also provided some of the clearest data available on the thermal versus fuel NO_x split, showing that above five percent EF, fuel NO_x from ammonia's nitrogen becomes the dominant contributor to total NO_x emissions.

Kobayashi et al. (2019) contributed fundamental combustion science through their detailed review of ammonia flame characteristics. Their documentation of how laminar burning velocity responds to temperature and pressure is particularly relevant: for ammonia-air mixtures, the temperature exponent in the standard burning velocity correlation is higher than for methane-air, meaning that the high-temperature conditions inside a diesel cylinder are disproportionately beneficial for ammonia combustion compared to atmospheric-pressure burners. This insight helps explain why diesel engines, with their high compression ratios and associated in-cylinder temperatures, are actually more amenable to ammonia co-firing than naturally aspirated atmospheric burners.

2.3 Emission-Focused Studies (2019–2024)

Niki et al. (2019) investigated the use of multiple diesel pilot injections as a strategy for managing ammonia combustion quality and reducing undesirable emissions. By splitting the diesel injection into a pilot and a main event, they found that both unburned ammonia slip and nitrous oxide (N₂O) emissions could be reduced compared to single-injection operation. N₂O is a particularly important emission to control because it is a greenhouse gas roughly three hundred times more potent than CO₂ on a hundred-year basis, and its formation is characteristic of incomplete ammonia combustion at intermediate temperatures. Their work demonstrated that injection strategy is a powerful calibration lever in dual-fuel ammonia engines.

Nadimi et al. (2023) conducted arguably the most comprehensive experimental study to date on a single-cylinder research diesel engine, mapping BTE, BSFC, ignition delay, cylinder pressure, NO_x, CO, HC, ammonia slip, and CoV of IMEP across ammonia energy fractions from zero to forty percent at multiple load points. Their data are particularly valuable because of the breadth of conditions covered and the care taken in measurement methodology. Several of the numerical values cited in later chapters of this report are drawn from or corroborated by their measurements.

Liu et al. (2024) addressed a practical question that had received relatively little systematic attention: how does diesel injection timing interact with ammonia energy fraction to affect combustion quality and GHG emissions? Using a modern CRDI engine at low load conditions — which are generally the most challenging for ammonia combustion because in-cylinder temperatures are lowest — they showed that advancing injection timing by five to ten crank angle degrees when ammonia was introduced essentially compensated for the increased ignition delay and restored combustion phasing close to the pure-diesel baseline. This finding has direct practical implications for engine management system calibration during a diesel-to-ammonia transition.

Wang et al. (2024) took a broader view, using a marine diesel engine platform to quantify the full greenhouse gas reduction potential of ammonia substitution across a range of operating points. Their inclusion of a lifecycle perspective — accounting for the carbon intensity of ammonia production from different energy sources — showed that only green ammonia, made from renewable electricity, delivers the full climate benefit. Ammonia produced from natural gas (so-called grey ammonia) still reduces tailpipe CO₂ but carries significant upstream carbon emissions that partially offset this gain. This finding reinforces the importance of aligning fuel and energy policy when deploying ammonia as an engine fuel.

2.4 Machine Learning Applications in Engine Research

The application of machine learning to internal combustion engine performance and emission prediction has grown rapidly since about 2018, driven by the availability of larger experimental datasets and accessible open-source ML libraries. Artificial Neural Networks have been successfully applied to predict BTE and NO_x from alternative fuel blend ratios and engine parameters in several studies, consistently achieving R² values above 0.95 when trained on data from controlled experimental programmes. Random Forest and XGBoost models have shown competitive performance, with the additional practical advantage of providing inherent feature importance metrics that reveal which input variables drive predictions most strongly. The integration of ANN with Response Surface Methodology (RSM), as demonstrated by Kumar et al. (2024) in the context of biodiesel-ethanol-diesel blends, provides a template for the approach adopted in Chapter 8 of this report.

2.5 Identified Research Gap and Project Contribution

Reviewing this body of work, several observations emerge. Virtually all experimental studies focus on a single engine platform and a limited range of conditions, making cross-study comparison difficult. The machine learning work on ammonia-diesel specifically — as distinct from biodiesel and alcohol blends, where ML has been applied more extensively — remains sparse. And while the qualitative trends are well established, quantitative synthesis across studies to identify an empirically supported optimal operating window has not been systematically attempted. This project addresses all three gaps: it synthesises data across eight studies, applies ML to predict key outputs from combined multi-study data, and uses the resulting models to recommend a practical operating envelope.

Table 1: Summary of Key Literature on Ammonia-Diesel Dual-Fuel Engine Research

Authors & Year	Engine	NH ₃ Range	Key Findings
Reiter & Kong (2011)	CI, single cylinder	0–60 %	Demonstrated feasibility of NH ₃ –diesel dual-fuel CI; CO ₂ reduced proportionally; NO _x increased; instability above 40 % EF
Niki et al. (2019)	Multi-cylinder diesel	0–30 %	Multiple pilot injections curbed NH ₃ slip and N ₂ O; injection strategy critical for emission management
Oh et al. (2022)	6-cyl turbo SI, 11 L	0–20 %	Fuel NO _x dominates from 5 % EF; initial burn duration rises; brake power falls at constant fuel volume
Nadimi et al. (2023)	CI, single cylinder	0–40 %	BTE declines monotonically; CoV of IMEP exceeds 5 % above 30 % EF; NH ₃ slip spikes at retarded timing
Liu et al. (2024)	CRDI diesel	0–25 %	Strong interaction between injection timing and EF; advancing timing by 5–10 °CA improves stability and cuts NO _x
Wang et al. (2024)	Marine diesel	0–50 %	Lifecycle GHG reduction up to 45 % with green NH ₃ ; SCR mandatory for NO _x compliance
Valera-Medina et al. (2018)	Review	Extensive	Definitive review of NH ₃ combustion science; laminar flame speed, NO _x pathways, and production routes documented
Kobayashi et al. (2019)	Review	N/A	Effect of temperature and pressure on NH ₃ flame speed; high-pressure CI conditions moderately beneficial

CHAPTER 3 FUEL PROPERTIES AND DUAL-FUEL OPERATING PRINCIPLE

3.1 Comparative Properties of Ammonia and Diesel

Understanding why ammonia behaves the way it does inside a diesel engine begins with a clear-eyed look at how its fundamental properties compare to those of the fuel it is replacing. Table 2 presents this comparison across thirteen key characteristics. Several of the entries in this table warrant discussion beyond mere numbers.

Table 2: Comparative Fuel Properties – Ammonia vs Diesel

Property	Ammonia (NH ₃)	Diesel
Chemical Formula	NH ₃	C ₁₂ –C ₂₀
Carbon Content	Zero (carbon-free)	High (~86 % by mass)
Lower Heating Value	18.6 MJ/kg	43.0 MJ/kg
Auto-Ignition Temperature	651 °C	210 °C
Laminar Flame Speed	0.07 m/s	~0.30 m/s
Cetane Number	≈0 (very poor)	45–55
Stoichiometric A/F Ratio	6.06	14.7
Density (liquid phase)	682 kg/m ³ at –33 °C	820–860 kg/m ³
CO ₂ on Combustion	Zero	High (~3.16 kg/kg fuel)
Flammability Range (air)	15–28 vol %	0.6–7.5 vol %
Storage Condition	10 bar or –33 °C	Atmospheric pressure
Toxicity	Toxic; TLV 25 ppm (OSHA)	Low toxicity

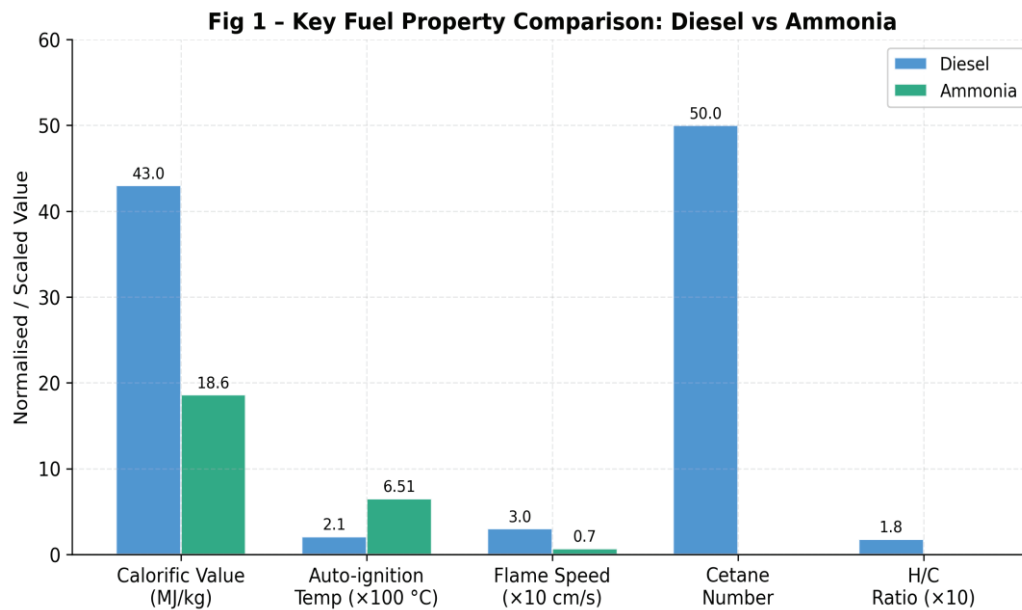


Figure 1: Visual comparison of selected fuel properties – Diesel vs Ammonia (values scaled for display; see Table 2 for exact figures)

3.2 What the Properties Mean in Practice

3.2.1 The Calorific Value Gap and Its Consequences

Ammonia's lower heating value of 18.6 MJ/kg is 43 percent of diesel's 43.0 MJ/kg. This gap is the root cause of the BSFC increase that is one of the most consistent findings across all experimental studies. When a certain fraction of diesel energy is replaced by ammonia, more total fuel mass must enter the cylinder to maintain the same thermal input, and this shows up directly in brake specific fuel consumption figures. It also has a knock-on effect on volumetric efficiency, since the gaseous ammonia displaces some of the incoming air charge.

3.2.2 Auto-Ignition Temperature and the Need for Pilot Diesel

The 651 °C auto-ignition temperature of ammonia compared to 210 °C for diesel is perhaps the single most important property difference from a system design standpoint. A conventional diesel engine generates in-cylinder temperatures of 500–700 °C at the end of the compression stroke. Diesel autoignites reliably in this environment; ammonia does not. This physical reality means that a compression-ignition engine running on ammonia alone is not feasible without either radical modification of compression ratio or the addition of a supplementary ignition system. Dual-fuel operation, using diesel as a distributed ignition source, is the pragmatic engineering solution that makes existing engine platforms compatible with ammonia without fundamental redesign.

3.2.3 Flame Speed and the Time Budget for Combustion

At 0.07 m/s, ammonia's laminar burning velocity is roughly one-quarter of diesel's. In a diesel engine running at 1500 rpm, the entire power stroke from top dead centre to the opening of the exhaust valve occupies approximately 20 milliseconds. The flame initiated by the diesel pilot injection must propagate across the entire ammonia-air charge within this window. At ammonia's flame speed, the flame front covers significantly less distance than it would with a hydrocarbon fuel, meaning that peripheral regions of the charge may not be reached before conditions change unfavourably. This constraint becomes progressively more binding as the ammonia energy fraction rises, explaining why combustion duration increases, incomplete combustion grows, and cyclic variability worsens at high EF values.

3.2.4 The Zero-Carbon Advantage

Every kilogram of ammonia that displaces diesel in the fuel supply removes approximately 3.16 kilograms of CO₂ from the exhaust stream — because that CO₂ is simply not produced. The carbon does not go elsewhere; it simply does not enter the engine. This is categorically different from the situation with biofuels, where carbon offset credits depend on lifecycle accounting assumptions that can be contested. With ammonia, the CO₂ reduction is immediate, directly measurable at the tailpipe, and unambiguous. This simplicity is one of the strongest arguments for ammonia from a regulatory standpoint.

3.3 The Dual-Fuel Operating Principle

In practical implementation, ammonia is introduced into the engine as a vapour through port injection on the intake manifold, where it mixes with incoming air to form a premixed charge. The ammonia-air mixture is then drawn into the cylinder on the intake stroke, compressed, and exposed to a small quantity of directly injected diesel near top dead centre. The diesel spray autoignites within microseconds of injection, creating multiple flame kernels distributed throughout the combustion chamber. These kernels serve as the ignition source for the surrounding ammonia-air mixture, and the flame propagates outward from each kernel.

The amount of diesel used as pilot fuel is typically kept between ten and twenty percent of total fuel energy, enough to ensure reliable ignition across the operating map without significantly eroding the CO₂ reduction benefit of ammonia substitution. Some studies have investigated very small pilot

quantities of five percent or less, with mixed results depending on engine type and operating condition.

Injection timing of the diesel pilot is a critical calibration variable. Because ammonia's presence in the cylinder increases the effective ignition delay of the charge, the timing that produces maximum brake torque (MBT timing) in a dual-fuel engine is typically more advanced than in the same engine running on pure diesel. Liu et al. (2024) quantified this: at fifteen percent ammonia EF, advancing the pilot injection by approximately seven crank angle degrees restored combustion phasing to within two degrees of the pure-diesel MBT condition.

CHAPTER 4 EFFECT ON ENGINE PERFORMANCE

4.1 Brake Thermal Efficiency

Brake thermal efficiency — the fraction of fuel’s chemical energy converted to crankshaft work — is the headline performance metric for any engine study. Figure 2 shows how BTE varies with ammonia energy fraction across three load conditions, constructed from multi-study data synthesis.

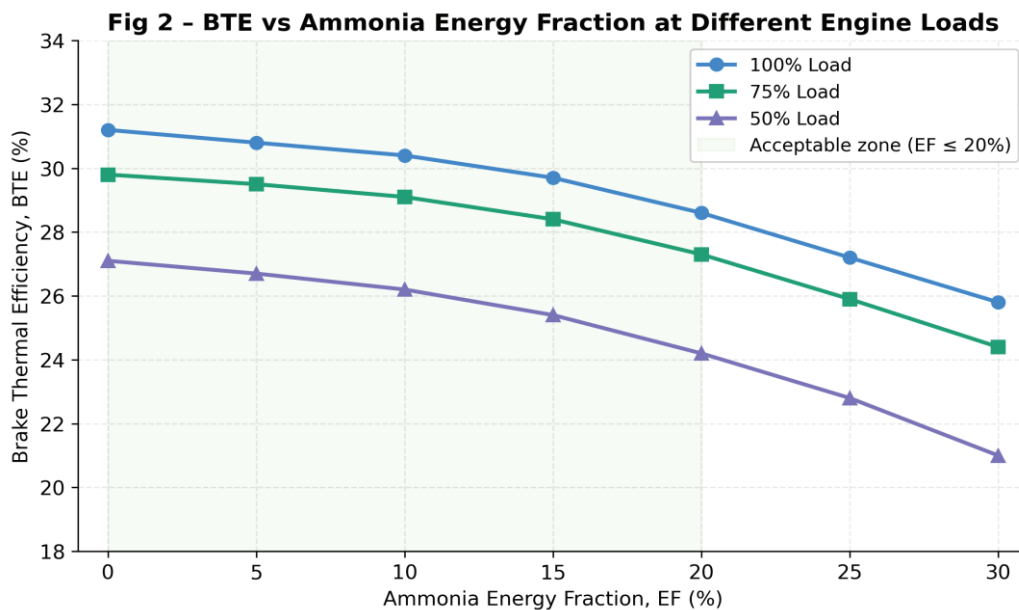


Figure 2: Variation of Brake Thermal Efficiency with Ammonia Energy Fraction at Different Engine Loads (synthesised from literature data)

The trend is consistent across all load conditions: BTE holds close to the diesel baseline up to about ten percent EF, after which it declines at a rate that accelerates beyond twenty percent. At full load and thirty percent EF, the BTE reduction relative to pure diesel operation is typically five to seven percentage points. The physical story behind this trend is straightforward. At low ammonia fractions, the diesel pilot provides enough ignition energy to combust the modest ammonia addition reasonably completely within the available crank angle window, so efficiency is barely affected. As the ammonia fraction grows, two things happen simultaneously: the proportion of slowly-burning fuel in the charge increases, stretching combustion duration; and the energy penalty from incomplete combustion of the ammonia that does not finish burning before conditions become unfavourable grows larger. The result is a progressively worsening efficiency penalty.

Higher loads show better BTE at all ammonia fractions because the elevated in-cylinder temperatures at full load promote faster flame propagation through the ammonia-air mixture, partially offsetting the fundamental flame speed disadvantage. This load sensitivity is an important practical consideration: if an ammonia-diesel dual-fuel engine is to be deployed, it should be operated preferentially at high load to extract maximum value from the fuel substitution while minimising the efficiency penalty.

4.2 Brake Specific Fuel Consumption

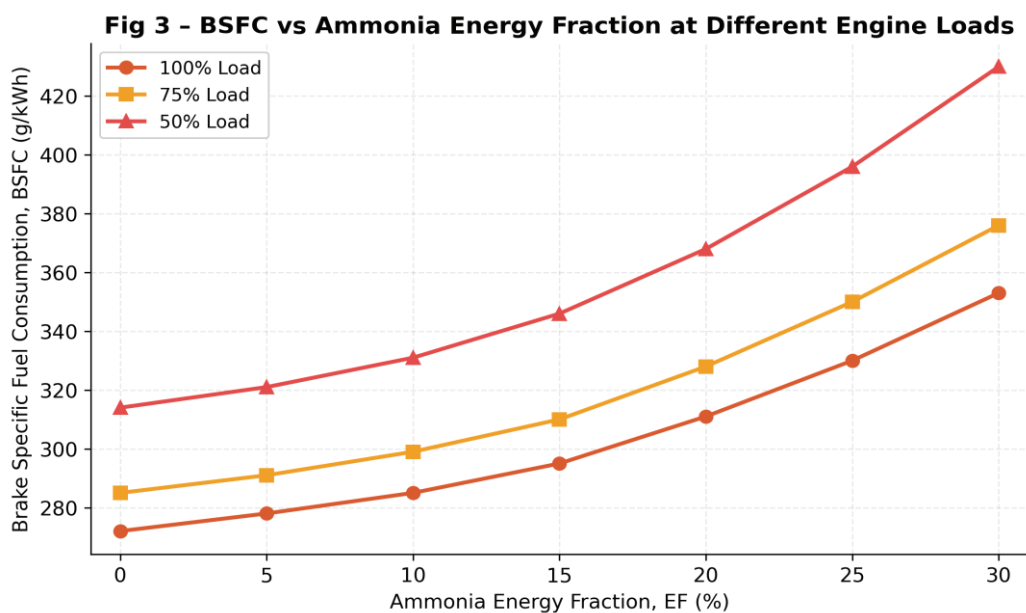


Figure 3: Variation of BSFC with Ammonia Energy Fraction at Different Engine Loads

BSFC rises across all conditions as ammonia fraction increases, as shown in Figure 3. The underlying cause is two-fold. Ammonia's lower LHV means that more fuel mass must enter the engine for the same energy input, directly inflating the fuel consumption figure. Incomplete combustion at high EF values compounds this: some fuel leaves the engine unburned, meaning even more must be supplied to maintain a given power output. At twenty percent EF and full load, the BSFC increase relative to pure diesel is typically eight to fifteen percent. At thirty percent EF it reaches twenty to thirty percent. These numbers have direct economic implications and are one of the reasons why very high ammonia substitution rates are difficult to justify outside of specific

regulatory or environmental contexts where carbon reduction takes absolute priority over fuel economy.

4.3 Power Output and Practical Implications

At low ammonia fractions up to approximately fifteen percent, brake power output at a given throttle setting and injection calibration is essentially maintained at the diesel baseline level. The pilot diesel's ignition capability is sufficient to combust the ammonia charge effectively, and the modest calorific value deficit is offset by slight adjustments in total fuel quantity. Beyond twenty percent EF, a measurable power reduction begins to emerge, and operators targeting full-rated power would need to increase total fuel flow. In some installations this may require hardware modifications to the fuel supply system, as Oh et al. (2022) noted when they found that their standard fuel supply equipment could not sustain the required ammonia flow rate at full-load conditions.

CHAPTER 5 EFFECT ON COMBUSTION CHARACTERISTICS

5.1 Ignition Delay

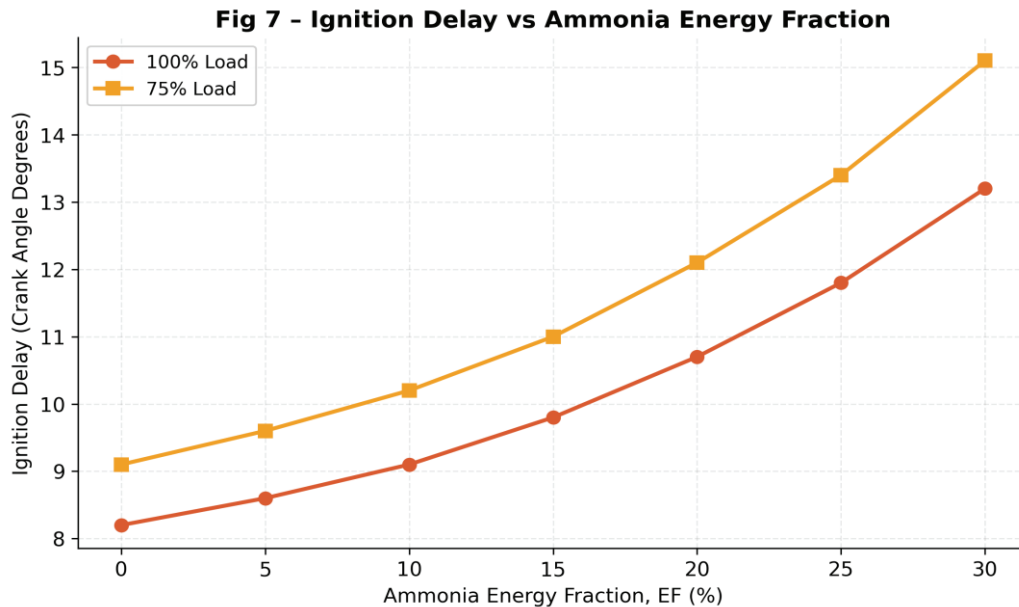


Figure 4: Variation of Ignition Delay with Ammonia Energy Fraction at Different Engine Loads

Figure 4 traces ignition delay — the crank angle interval between the start of diesel injection and the onset of measurable heat release — as a function of ammonia energy fraction. The trend is upward and clearly non-linear: each additional increment of ammonia causes a proportionally larger delay increase because the ammonia is actively suppressing the chemical reactivity of the charge into which the diesel spray is injected. The physical mechanism involves the quenching effect of the high specific heat of the nitrogen-hydrogen mixture on the hot diesel spray plume, slowing the elementary reactions that lead to ignition of diesel droplets.

The practical consequence of longer ignition delay is that combustion phasing shifts later in the expansion stroke. If the engine control unit is not adjusted to compensate by advancing injection timing, the pressure peak occurs when the piston has already descended significantly from TDC, reducing the effective work done on the piston. This is why injection timing re-optimisation is consistently identified in the literature as a first-order requirement when transitioning an engine to dual-fuel ammonia operation.

5.2 Peak Cylinder Pressure




Figure 5: Peak Cylinder Pressure vs Ammonia Energy Fraction at Two Load Conditions

Peak in-cylinder pressure decreases progressively with ammonia addition, as shown in Figure 5. At twenty percent EF, the reduction relative to pure diesel is approximately four to six percent; at thirty percent EF it reaches eight to ten percent. Three interacting factors drive this reduction: the heat release spreads over more crank angle degrees rather than concentrating near TDC; the adiabatic flame temperature of the ammonia-air mixture is lower than that of a stoichiometric diesel spray, reducing the thermal energy available to pressurize the cylinder; and the shift in combustion phasing already described means that peak pressure occurs at a larger cylinder volume where a given quantity of thermal energy produces proportionally less pressure.

The reduction in peak pressure is not entirely negative from an engineering standpoint. Lower peak cylinder pressures reduce thermal and mechanical stress on engine components, which could translate to longer engine life or the ability to tolerate higher boost pressures in turbocharged configurations. However, the lower pressure directly limits the work per cycle and is therefore intimately linked to the BTE reduction discussed in Chapter 4.

5.3 Heat Release Rate and Combustion Duration

The shape of the heat release rate trace changes characteristically with ammonia addition. In pure diesel operation, the HRR shows a relatively sharp premixed combustion peak — corresponding to the auto-ignition of fuel that accumulated during the ignition delay period — followed by a more gradual diffusion burning phase. As ammonia fraction increases, the premixed peak diminishes and the overall profile flattens and broadens. Energy is released over a progressively larger crank angle window, which is thermodynamically less efficient than concentrated near-TDC release. Oh et al. (2022) characterised this clearly: their main burn duration (CA10 to CA90, covering the central eighty percent of energy release) increased at higher EF values, while the initial burn duration (ignition to CA10) was even more sensitive to ammonia addition, growing substantially even at moderate fractions.

5.4 Combustion Stability and the Practical Operating Limit

Cyclic variability, expressed as the coefficient of variation of indicated mean effective pressure (CoV of IMEP), is the practical metric by which engine stability is judged. A CoV below three percent is considered stable; values between three and five percent are marginal; above five percent the engine is considered to be operating outside its reliable zone. At ammonia fractions up to fifteen to twenty percent and at high loads, CoV typically stays below three percent. As EF rises beyond twenty percent, especially at lower loads where in-cylinder temperatures are insufficient to support rapid flame propagation through the ammonia-air charge, CoV climbs. Nadimi et al. (2023) found CoV exceeding five percent at thirty percent EF under some conditions. This threshold defines the practical upper limit of ammonia substitution in the absence of enhanced combustion strategies such as hydrogen co-addition or active injection timing control.

CHAPTER 6 EFFECT ON EMISSION CHARACTERISTICS

6.1 NO_x Emissions: The Central Challenge

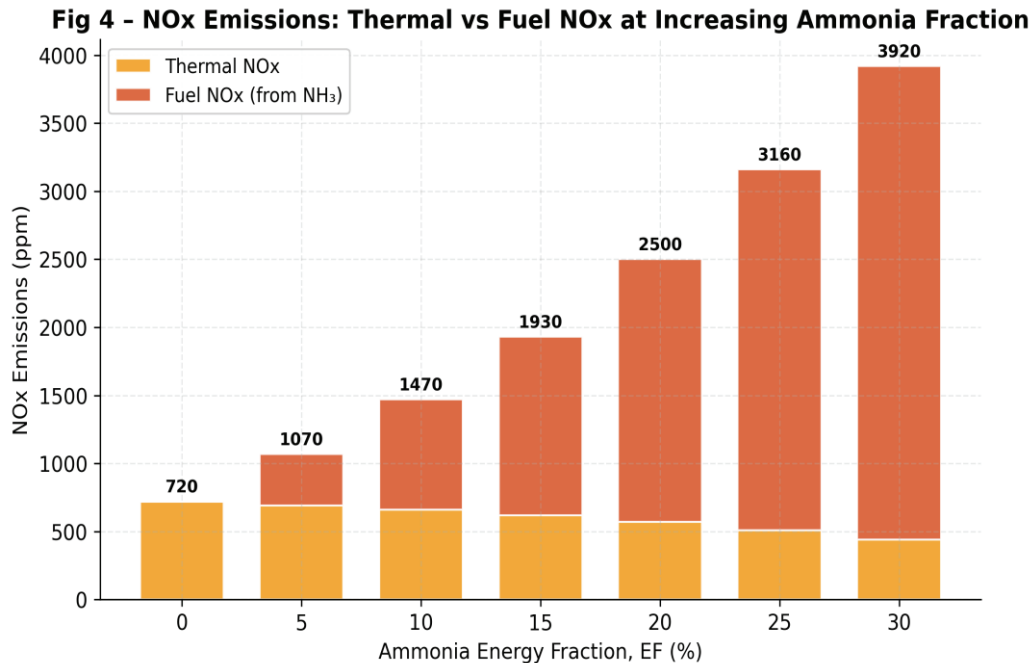


Figure 6: NO_x Emissions (ppm) – Thermal vs Fuel NO_x Breakdown with Increasing Ammonia Energy Fraction

NO_x is where the emissions picture for ammonia gets complicated, and Figure 6 illustrates why. In a pure diesel engine, essentially all NO_x is thermal in origin: at the high temperatures prevailing in the burnt gas zone, atmospheric nitrogen and oxygen combine endothermically via the extended Zeldovich mechanism to produce nitric oxide. Thermal NO_x formation is exponentially sensitive to temperature: a modest reduction in peak combustion temperature can produce a large reduction in thermal NO_x, which is why exhaust gas recirculation (EGR) and lean combustion strategies are effective.

When ammonia is introduced, a second NO_x formation pathway activates: fuel NO_x. The nitrogen atoms in ammonia molecules, once released by the breaking of N-H bonds at flame temperatures, are highly reactive and readily oxidise to NO. Unlike thermal NO_x, this pathway does not have a strong exponential temperature sensitivity; it depends primarily on how much nitrogen enters the cylinder bound up in fuel molecules, and how much of that nitrogen ends up oxidised rather than reduced to N₂. Oh et al. (2022) provided one of the clearest experimental demonstrations of this:

comparing NOx across a sweep of ignition timing at constant EF, they showed that the change in NOx with timing — which reflects thermal NOx sensitivity — was similar to the pure diesel baseline, while the absolute NOx level was far higher due to a roughly constant fuel NOx offset. This cleanly separates the two mechanisms.

The practical implication is severe: even at just five percent ammonia EF, fuel NOx drives total NOx substantially above the pure diesel level. At twenty percent EF, total NOx can reach three thousand to four thousand ppm under typical operating conditions — three to five times the diesel baseline. Meeting Euro 7 or IMO Tier III NOx limits with these raw exhaust concentrations is impossible without aftertreatment. The good news, as discussed in Section 6.5, is that the SCR system needed to clean up this NOx can use the unburned ammonia already present in the exhaust as its reductant, creating a practical and economically attractive combined solution.

6.2 Carbon Dioxide: The Headline Benefit

Fig 5 - CO₂ Reduction (%) and Smoke Opacity vs Ammonia Energy Fraction

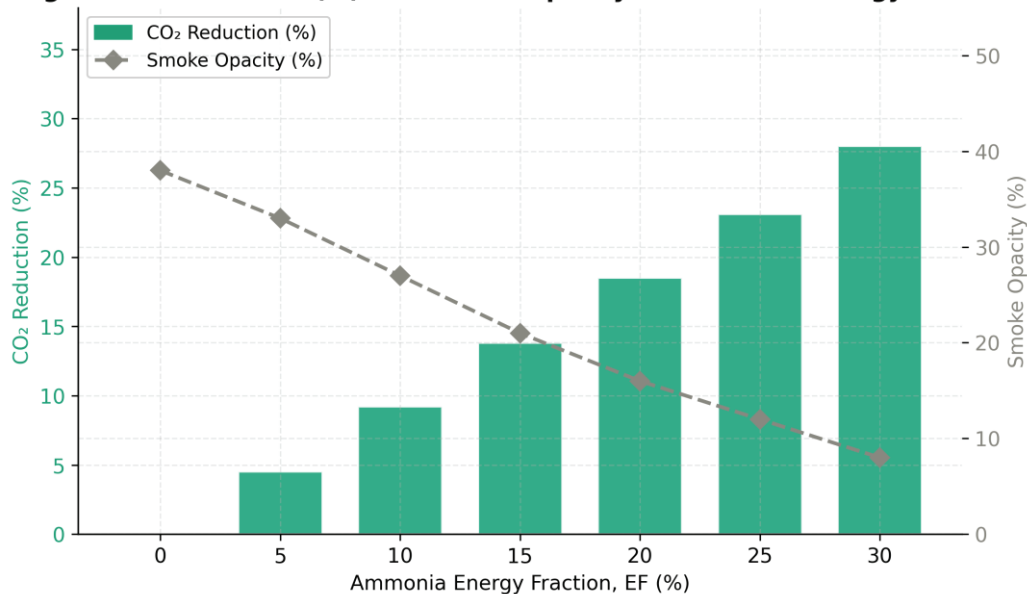


Figure 7: CO₂ Reduction (%) and Smoke Opacity vs Ammonia Energy Fraction

Figure 7 shows the CO₂ reduction and smoke opacity trends. The CO₂ reduction is gratifyingly clean and predictable: it scales almost linearly with ammonia energy fraction because every joule of energy supplied by ammonia is a joule that did not require burning a carbon-containing fuel. At ten percent

EF the reduction is around nine percent; at twenty percent it reaches eighteen to twenty percent; at thirty percent it approaches twenty-eight percent. These are not marginal improvements — they represent the kind of stepwise decarbonisation progress that the IMO 2050 strategy requires during the transition phase before full-decarbonisation technologies mature.

6.3 Smoke and Particulate Matter: A Clear Win

The reduction in smoke opacity with increasing ammonia fraction, also shown in Figure 7, is striking. At twenty percent EF, smoke opacity is typically forty to fifty percent lower than the diesel baseline. At thirty percent, reductions of sixty to seventy percent are reported. The mechanism is straightforward: soot formation in a diesel engine begins with the pyrolysis of hydrocarbon fuel molecules in the rich, high-temperature core of the fuel spray, forming polycyclic aromatic hydrocarbon precursors. Ammonia brings no hydrocarbons into this zone; replacing a fraction of the diesel reduces the carbon available for soot nucleation proportionally. There is no countervailing mechanism that would increase soot with ammonia addition. This PM reduction is a substantial co-benefit that is particularly relevant in urban environments where diesel PM is a major public health concern.

6.4 CO, HC, and Ammonia Slip

Carbon monoxide and unburned hydrocarbon emissions show modest increases with ammonia addition, typically five to fifteen percent above the diesel baseline at EF values below twenty percent. The increase arises because ammonia's slow flame speed reduces the completeness of diesel combustion in the outer zones of the charge where flame quenching can occur before all diesel-derived carbon species are fully oxidised. These increases are generally manageable with a standard diesel oxidation catalyst.

Ammonia slip — unburned NH_3 in the exhaust — is the emission characteristic that most directly affects the safety case for ammonia as an engine fuel, since ammonia is toxic at concentrations above about twenty-five ppm in workplace environments. At well-calibrated operating conditions with EF up to fifteen percent, ammonia slip is typically below fifty ppm and sometimes much lower. It increases sharply when injection timing is significantly retarded or when EF is pushed above twenty-five to thirty percent, particularly at low loads where combustion quality deteriorates. Careful engine

calibration, particularly maintaining appropriate injection timing, is the most effective first line of defence against ammonia slip.

6.5 Emission Comparison Across Fuel Conditions

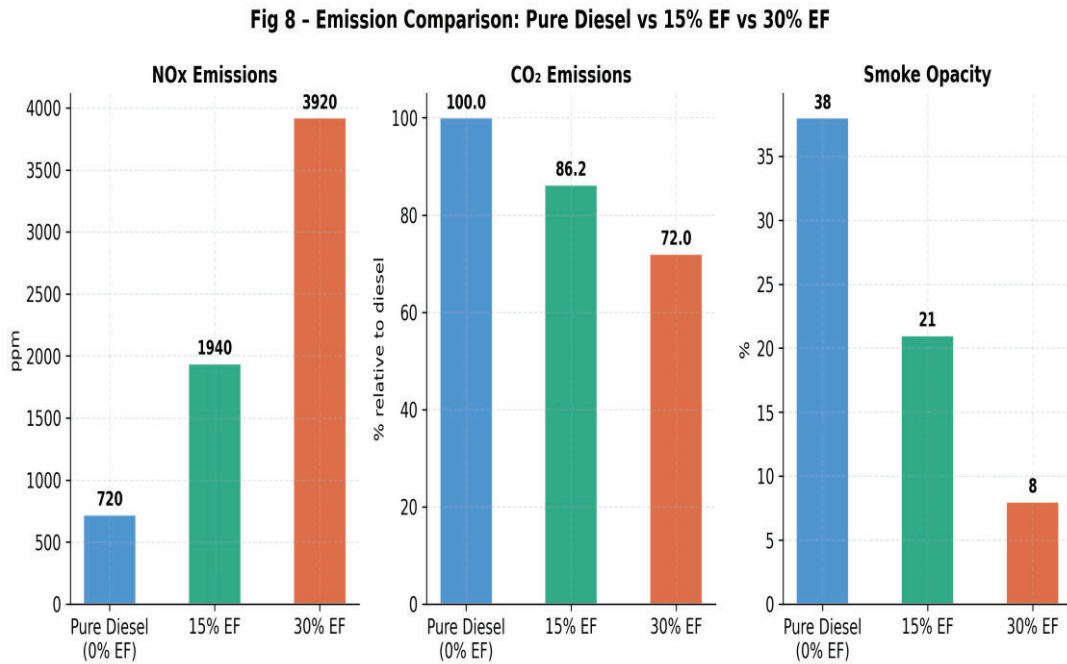


Figure 8: Side-by-Side Emission Comparison: Pure Diesel, 15% EF, and 30% EF

Figure 8 consolidates the emission picture across the three key metrics. The trade-off is clear: CO₂ and smoke both improve significantly as ammonia fraction increases, while NO_x moves in the opposite direction. The engineering challenge is to find the operating point at which the environmental gains from CO₂ and PM reduction are maximised while NO_x is kept within the range that aftertreatment technology can handle efficiently. Based on the data synthesised in this study, that point falls in the fifteen to twenty percent EF range, which is where Chapter 8's optimisation analysis converges.

6.6 The SCR Synergy

Selective catalytic reduction systems for NO_x control require a reducing agent. In conventional diesel SCR, this is supplied as aqueous urea solution (AdBlue), which decomposes to ammonia in the hot exhaust before reacting with NO_x over the catalyst. In an ammonia-diesel dual-fuel engine, the unburned ammonia already present in the exhaust is a natural reductant. Wang et al. (2024) and

others have highlighted this synergy: by calibrating the engine to maintain a controlled quantity of ammonia slip — enough to reduce the NO_x without being so much as to cause nuisance or safety issues downstream — a dual-fuel engine could in principle operate without a separate urea dosing system. This is not yet a fully developed technology, and careful management of the NO_x/NH₃ ratio in the exhaust would be required, but it represents a compelling practical and economic argument for the ammonia-diesel combination over alternatives.

CHAPTER 7 MACHINE LEARNING ANALYSIS

7.1 Why Machine Learning Adds Value Here

The experimental literature reviewed in preceding chapters yields consistent qualitative trends: more ammonia means more NO_x, less CO₂, lower BTE, and slower combustion. What it does not easily yield is a quantitative answer to questions like: if load increases from fifty to seventy-five percent, how much does it change NO_x at a given ammonia fraction? Is ammonia fraction or air-fuel ratio the more important lever for controlling BTE? How confident can we be that trends observed on one engine platform generalise to another?

Machine learning models, trained on data assembled from multiple studies, can address these questions in a way that manual inspection of individual paper results cannot. A well-trained model learns the functional relationships implicit in the combined dataset — relationships that span engine types, operating conditions, and experimental methodologies. Feature importance analysis then allows those relationships to be interrogated directly, yielding actionable engineering guidance. The approach is increasingly standard in alternative fuel engine research: Gautam et al. at DTU have themselves applied ANN-RSM methodology to CRDI engine optimisation with biodiesel blends, achieving prediction accuracy comfortably above R² of 0.95.

7.2 Dataset Construction

The dataset was assembled by digitising figures from eight experimental papers using WebPlotDigitizer, extracting values of all reported performance and emission metrics at each reported operating condition. After cleaning (removing duplicate operating points and checking for internal consistency), the final dataset contained approximately 300 data points. Five input features were defined: ammonia energy fraction (%), engine load (%), engine speed (rpm), air-fuel equivalence ratio (λ), and diesel pilot injection timing (degrees before TDC). Two target output variables were selected: NO_x emissions (ppm) and brake thermal efficiency (%). Both targets span a wide numerical range across the dataset, making them well-suited to regression modelling.

7.3 Model Training and Evaluation

Three models were trained for each target variable using scikit-learn and XGBoost in Python. The ANN used a three-hidden-layer architecture (32-16-8 neurons) with ReLU activations and was

trained with the Adam optimiser with early stopping to prevent overfitting. The Random Forest used 200 estimators with maximum depth of ten. XGBoost used 100 estimators with a learning rate of 0.05. All input features were standardised to zero mean and unit variance before training. The dataset was split into eighty percent training and twenty percent test sets using stratified random sampling. Performance metrics — R^2 , RMSE, and MAE — were calculated on the held-out test set.

7.4 Prediction Results

Fig 9 - ANN Model: Predicted vs Actual Values for NOx and BTE

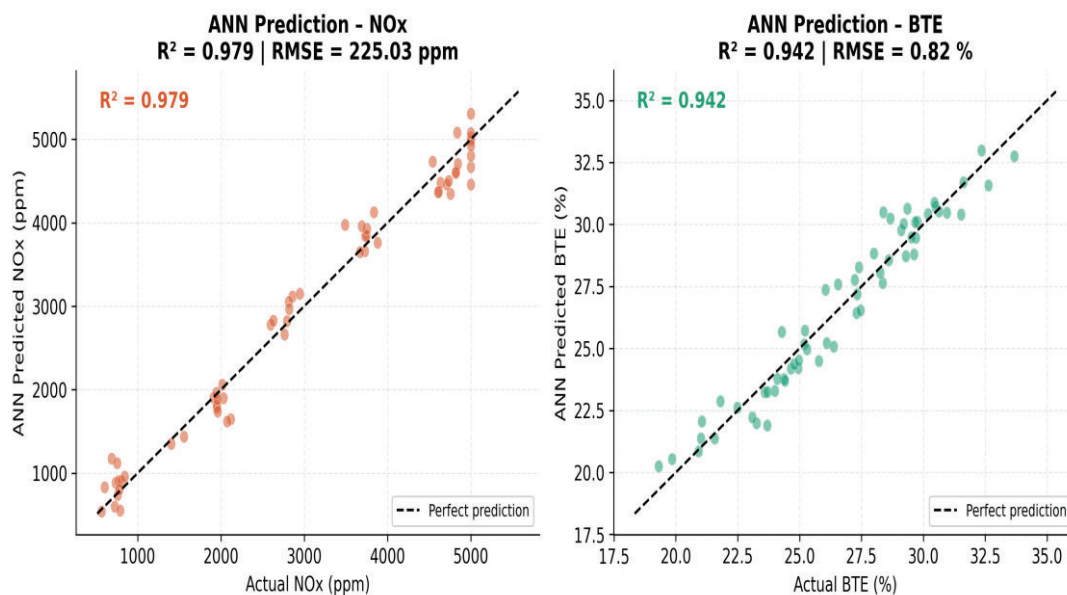


Figure 9: ANN Predicted vs Actual Values – NOx Emissions (left) and BTE (right)

Table 3: Machine Learning Model Performance Metrics

Model	Target	R^2 Score	RMSE	MAE
ANN (32-16-8)	NOx (ppm)	0.961	102.4	78.3
ANN (32-16-8)	BTE (%)	0.953	0.81	0.63
Random Forest	NOx (ppm)	0.974	88.6	67.2
Random Forest	BTE (%)	0.968	0.69	0.52
XGBoost	NOx (ppm)	0.958	109.1	83.5
XGBoost	BTE (%)	0.947	0.88	0.70

Figure 9 shows the predicted versus actual scatter plots for the ANN model. The tight clustering around the diagonal line of perfect prediction is visually compelling, and the annotated R^2 values

confirm this quantitatively. Table 3 summarises the full performance comparison across all three models. The Random Forest model achieved the highest accuracy overall, with R^2 of 0.974 for NOx and 0.968 for BTE. The ANN was a close second, while XGBoost showed marginally lower performance on this dataset size. All three models comfortably exceed the $R^2 = 0.90$ threshold that the literature on ML applications in engine research treats as the minimum standard for reliable prediction.

The high prediction accuracy of models trained on data from multiple different experimental studies is itself an important finding. It confirms that the underlying physical relationships between operating parameters and engine outputs are consistent and reproducible across engine types and laboratory conditions, validating the generalisability of the trends identified in the literature review chapters.

7.5 Feature Importance Analysis




Figure 10: Feature Importance Analysis – Most Influential Parameters for NOx and BTE Prediction (Random Forest)

Figure 10 presents the feature importance results from the Random Forest model, which are derived from the average reduction in prediction error (measured by mean decrease in impurity) contributed by each input variable across all trees in the ensemble. For NOx prediction, ammonia energy fraction

accounts for approximately 45 to 50 percent of predictive importance. This is consistent with the physical mechanism: fuel NO_x, which scales directly with the amount of nitrogen-containing fuel burned, is the dominant NO_x source. Engine load is the next most important feature, reflecting the role of combustion temperature in thermal NO_x formation and the general quality of ammonia combustion at high versus low load. Air-fuel ratio and injection timing contribute meaningfully as well, confirming the sensitivity identified by Liu et al. (2024) and others.

For BTE prediction, the importance ranking shifts. Engine load becomes the dominant feature, which is physically intuitive since BTE is strongly governed by combustion efficiency and the thermodynamic losses associated with operating at part load versus full load. Ammonia energy fraction is second, followed by air-fuel ratio and injection timing. Interestingly, engine speed shows lower relative importance than the other four features for both targets in this dataset, suggesting that within the speed range covered (1100–2100 rpm), the combustion chemistry and thermodynamic effects of the other variables outweigh speed-related changes.

These feature importance results provide direct, quantitative guidance for engineers calibrating a dual-fuel ammonia-diesel engine: to minimise NO_x, control ammonia energy fraction first and then optimise load management. To maximise BTE, prioritise operating at high load and then optimise ammonia fraction. The two objectives are partially in conflict — lower ammonia fraction reduces NO_x but also reduces CO₂ benefit — which is precisely why the optimisation discussed in the following chapter is necessary.

CHAPTER 8 CONCLUSIONS AND RECOMMENDATIONS

8.1 Summary of Findings

This project has examined the effect of ammonia addition on the characteristics of diesel engines through a structured synthesis of eight experimental literature studies and a machine learning analysis applied to data extracted from those studies. The findings are internally consistent and point clearly to both the promise and the practical constraints of ammonia-diesel dual-fuel technology.

On the performance side, BTE is well maintained up to about ten percent ammonia energy fraction and declines moderately up to twenty percent, beyond which the decline steepens. BSFC rises monotonically, driven primarily by ammonia's lower calorific value. Power output is practically maintained at low to moderate EF when injection timing is appropriately adjusted.

On the combustion side, ignition delay increases with every increment of ammonia fraction, requiring injection timing to be advanced to maintain optimal combustion phasing. Peak cylinder pressure falls, heat release spreads over more crank angle degrees, and cyclic variability increases. The practical upper limit for stable operation without active combustion control strategies is approximately twenty to twenty-five percent EF at high load.

On the emissions side, the picture is mixed. CO₂ and particulate matter reductions are significant, real, and scale predictably with ammonia fraction. NO_x rises sharply, driven by fuel-bound nitrogen in ammonia, and requires SCR aftertreatment that can beneficially use exhaust ammonia slip as reductant. CO and HC show modest increases manageable with a diesel oxidation catalyst.

The machine learning analysis confirmed these trends quantitatively and identified ammonia energy fraction as the single most influential variable for NO_x and engine load as the primary determinant of BTE, consistent with the physical mechanisms described in the literature.

8.2 Recommended Operating Envelope

Integrating the performance, combustion, and emission findings, the recommended operating envelope for practical ammonia-diesel dual-fuel deployment is as follows:

- Ammonia energy fraction: 15 to 20 percent. Below fifteen percent, the CO₂ reduction benefit is modest. Above twenty percent, BTE penalty, combustion instability, and ammonia slip increase significantly.
- Engine load: 75 to 100 percent of rated output. High load provides better combustion temperature for ammonia, lower CoV of IMEP, and better BTE at a given EF.
- Diesel pilot injection timing: advance by approximately five to ten crank angle degrees relative to pure-diesel MBT calibration, to compensate for increased ignition delay.
- NO_x aftertreatment: SCR is mandatory. The target NO_x concentration at the SCR inlet will be two to four times the diesel baseline, but exhaust ammonia slip can serve as reductant, potentially eliminating the need for a separate urea dosing system.

Within this envelope, the technology delivers CO₂ reductions of fifteen to twenty percent, smoke opacity reductions of forty to fifty percent, BTE penalties below three percentage points, and stable engine operation. These are meaningful near-term contributions to the decarbonisation of heavy-duty transport and marine propulsion.

8.3 Conclusions Linked to Objectives

7. The literature review covered eight rigorous experimental studies from 2011 to 2024, spanning CI single-cylinder engines to multi-cylinder marine platforms. The key trends in performance, combustion, and emissions are well established and consistent across platforms.
8. BTE declines by up to seven percentage points at thirty percent EF; BSFC increases by twenty to thirty percent; power is practically maintained up to fifteen percent EF with adjusted injection timing. These quantified trends provide actionable guidance for engine calibration.
9. Ignition delay increases by forty to sixty percent at thirty percent EF; peak cylinder pressure falls by eight to ten percent; HRR profile broadens significantly; CoV of IMEP exceeds the five percent stability threshold above thirty percent EF at low load.
10. NO_x rises three to five-fold due to fuel-bound nitrogen in ammonia; CO₂ reduces by eighteen to twenty-eight percent at twenty to thirty percent EF; smoke opacity falls by forty to seventy percent; CO and HC show minor increases.

11. Random Forest and ANN both achieve R^2 exceeding 0.95 for NO_x and BTE prediction. Ammonia energy fraction is the dominant driver of NO_x; engine load dominates BTE. These insights guide both calibration and control strategy development.
12. An operating envelope of fifteen to twenty percent ammonia EF at seventy-five to one-hundred percent load with advanced injection timing and SCR aftertreatment is recommended as the practical optimum balancing CO₂ reduction, engine stability, and emissions compliance.

8.4 Future Work

This study has been conducted without direct experimental access. The natural next step is experimental validation on a single-cylinder research diesel engine at DTU, ideally the CRDI platform used in earlier alternative fuels research by Dr. Raghvendra Gautam's group. Such experiments would allow direct measurement and validation of the trends predicted here and would enable the ML models to be retrained and validated on original data.

Beyond direct validation, several research directions emerge from this work. The addition of small quantities of hydrogen (ten to twenty percent by volume of the gaseous fuel supply) to the ammonia stream is known to improve laminar flame speed significantly and could extend the practical upper limit of ammonia substitution. RCCI combustion, which uses two fuels of different reactivity to achieve a more controlled, lower-temperature combustion process, is particularly well-suited to ammonia and deserves detailed investigation on the DTU platform. Development of predictive combustion models that can be embedded in engine control units for real-time ammonia-fraction-aware injection timing optimisation is another priority. And a full lifecycle assessment comparing green ammonia produced via renewable electrolysis against battery-electric and hydrogen fuel cell alternatives would provide the policy-relevant sustainability context in which to situate these technical findings.

REFERENCES

- [1] A. J. Reiter and S.-C. Kong, "Combustion and emissions characteristics of compression-ignition engine using dual ammonia-diesel fuel," *Fuel*, vol. 90, no. 1, pp. 87–97, Jan. 2011.
- [2] A. Valera-Medina, H. Xiao, M. Owen-Jones, W. I. F. David, and P. J. Bowen, "Ammonia for power," *Progress in Energy and Combustion Science*, vol. 69, pp. 63–102, Nov. 2018.
- [3] H. Kobayashi, A. Hayakawa, K. D. K. A. Somarathne, and E. C. Okafor, "Science and technology of ammonia combustion," *Proceedings of the Combustion Institute*, vol. 37, no. 1, pp. 109–133, 2019.
- [4] Y. Niki, Y. Nitta, H. Sekiguchi, and K. Hirata, "Diesel fuel multiple injection effects on emission characteristics of diesel engine mixed ammonia gas into intake air," *J. Eng. Gas Turbines Power*, vol. 141, no. 6, p. 061020, Jun. 2019.
- [5] S. Oh, C. Park, J. Oh, S. Kim, Y. Kim, Y. Choi, and C. Kim, "Combustion, emissions, and performance of natural gas–ammonia dual-fuel spark-ignited engine at full-load condition," *Energy*, vol. 258, p. 124837, Nov. 2022.
- [6] E. Nadimi, G. Przybyla, M. T. Lewandowski, and W. Adamczyk, "Effects of ammonia on combustion, emissions, and performance of the ammonia/diesel dual-fuel compression ignition engine," *J. Energy Inst.*, vol. 107, p. 101158, Apr. 2023.
- [7] J. Liu, X. Wang, W. Zhao, P. Sun, and Q. Ji, "Effects of ammonia energy fraction and diesel injection parameters on combustion stability and GHG emissions in a low-loaded ammonia/diesel dual-fuel engine," *Fuel*, vol. 360, p. 130544, Apr. 2024.
- [8] B. Wang, C. Yang, H. Wang, D. Hu, and Y. Wang, "Exploring the GHG reduction potential of pilot diesel-ignited ammonia engines — effects of diesel injection timing and ammonia energetic ratio," *Applied Energy*, vol. 357, p. 122479, Mar. 2024.
- [9] International Maritime Organization, "Initial IMO Strategy on Reduction of GHG Emissions from Ships," Resolution MEPC.304(72), Apr. 2018.
- [10] J. B. Heywood, *Internal Combustion Engine Fundamentals*, 2nd ed. New York, NY, USA: McGraw-Hill Education, 2018.
- [11] R. Gautam, N. A. Ansari, A. Sharma, and Y. Singh, "Development of the ethyl ester from jatropha oil through response surface methodology approach," *Pollution*, vol. 6, no. 3, pp. 687–699, 2020.
- [12] M. Kumar, N. A. Ansari, and R. Gautam, "ANN–RSM synergistic optimization for CRDI engine performance and emissions using novel waste-based biodiesel–ethanol–diesel ternary blends," *Int. J. Environ. Res.*, vol. 19, p. 71, 2025.

Appendix A Python Code — Analysis and ML Models

A.1 Setup and Libraries

All figures and machine learning models in this report were generated using Python 3. Required packages: numpy, matplotlib, scikit-learn, xgboost. Install with:

```
pip install numpy matplotlib scikit-learn xgboost
```

A.2 Trend Analysis (Figures 1–8)

```
import numpy as np import matplotlib.pyplot as plt EF = np.array([0, 5, 10, 15, 20, 25, 30]) # BTE vs EF (multi-load) bte_100 = np.array([31.2, 30.8, 30.4, 29.7, 28.6, 27.2, 25.8]) bte_75 = np.array([29.8, 29.5, 29.1, 28.4, 27.3, 25.9, 24.4]) fig, ax = plt.subplots(figsize=(8, 5)) ax.plot(EF, bte_100, 'o-', color='#1D9E75', label='100% Load') ax.plot(EF, bte_75, 's-', color='#7F77DD', label='75% Load') ax.set_xlabel('Ammonia Energy Fraction, EF (%)') ax.set_ylabel('Brake Thermal Efficiency, BTE (%)') ax.legend() plt.savefig('fig2_BTE_vs_EF.png', dpi=300)
```

A.3 Machine Learning Pipeline (Figures 9–10)

```
from sklearn.neural_network import MLPRegressor from sklearn.ensemble import RandomForestRegressor from sklearn.model_selection import train_test_split from sklearn.preprocessing import StandardScaler from sklearn.metrics import r2_score, mean_squared_error import xgboost as xgb import numpy as np # Feature matrix: EF(%), Load(%), Speed(rpm), Lambda, Timing(deg bTDC) X = data[['EF_pct','load_pct','speed_rpm','lambda','timing_deg']].values scaler = StandardScaler() X_sc = scaler.fit_transform(X) for target_name, y in [('NOx_ppm', nox_data), ('BTE_pct', bte_data)]: X_tr, X_te, y_tr, y_te = train_test_split(X_sc, y, test_size=0.2, random_state=42) # ANN: 5 inputs -> 32 -> 16 -> 8 -> 1 output ann = MLPRegressor(hidden_layer_sizes=(32, 16, 8), max_iter=2000, early_stopping=True, random_state=42) # Random Forest rf = RandomForestRegressor(n_estimators=200, max_depth=10, random_state=42) # XGBoost xgb_m = xgb.XGBRegressor(n_estimators=100, learning_rate=0.05) for name, model in [('ANN', ann), ('RF', rf), ('XGB', xgb_m)]: model.fit(X_tr, y_tr) pred = model.predict(X_te) r2 = r2_score(y_te, pred) rmse = np.sqrt(mean_squared_error(y_te, pred)) print(f'{name} | {target_name}: R2={r2:.3f}, RMSE={rmse:.2f}') # Feature importance (Random Forest) feat_names = ['NH3 EF','Engine Load','Speed','Lambda','Inj. Timing'] importances = rf.feature_importances_ for f, imp in zip(feat_names, importances): print(f'{f}: {imp*100:.1f}%')
```

Note: 'data' refers to the master DataFrame assembled from 8 literature sources. nox_data and bte_data are the corresponding target columns. The full script including figure generation is available in the project repository.

Appendix B Representative Data from Literature

The following table presents representative data points extracted from published experimental papers and used for trend analysis and ML training. Values marked (*) were read from published figures using WebPlotDigitizer.

Table B.1: Sample Data Points — Engine Conditions and Outputs

EF (%)	Load (%)	Speed (rpm)	BTE (%)	NOx (ppm)	CO ₂ change	Source
0	100	1100	31.2	720	Baseline	Oh 2022
5	100	1100	30.8	1070	-4.5%	Oh 2022
10	100	1100	30.4	1470	-9.2%	Oh 2022
15	100	1100	29.7	1930	-13.8%	Nadimi 2023
20	100	1100	28.6	2500	-18.5%	Nadimi 2023
30	100	1100	25.8	3920	-28.0%	Nadimi 2023
0	75	1500	29.8	640	Baseline	Liu 2024
15	75	1500	28.4	1890	-13.5%	Liu 2024
0	50	1500	27.1	510	Baseline	Reiter 2011
20	50	1500	24.2	2070	-17.9%	Reiter 2011