



Interpretable machine learning for predicting the derived cetane number of jet fuels using compact TD-NMR

Parker Huggins^{a,b,1}, Jacob S. Martin^{a,c,1}, Austin R.J. Downey^{a,d,*}, Sang Hee Won^a

^a Department of Mechanical Engineering, University of South Carolina, Columbia, 29208, SC, USA

^b Department of Electrical Engineering, University of South Carolina, Columbia, 29208, SC, USA

^c Department of Physics and Astronomy, University of South Carolina, Columbia, 29208, SC, USA

^d Department of Civil and Environmental Engineering, University of South Carolina, Columbia, 29208, SC, USA

ARTICLE INFO

Keywords:

Derived cetane number
Jet fuel
NMR
 T_2 relaxation
Random forest
Explainable AI

ABSTRACT

The derived cetane number (DCN) is a commonly used metric that summarizes fuel ignition characteristics, including ignition propensity and chemical kinetic potential for combustion processes. Traditional methods for determining the DCN of jet fuels are ASTM standards that involve large-scale, laboratory-based experiments. While recent advancements include the estimation of DCN via nuclear magnetic resonance (NMR) and infrared spectroscopy, the search persists for a method capable of real-time and in-situ estimations. This work proposes the use of a compact time-domain NMR (TD-NMR) system for the acquisition of jet fuel T_2 relaxation curves. The system is validated using relaxometric experiments and demonstrates the ability to acquire consistent, structurally viable data on a time-scale of just minutes. Furthermore, an interpretable approach for relaxometric data analysis is presented, allowing for the estimation of a sample's DCN directly from its T_2 relaxation curve. Random forests are trained for DCN prediction on both hydrocarbon and jet fuel samples, and the importance of extracted T_2 curve features are investigated using both the permutation of out-of-bag predictors and partial dependence plots. A model trained on less than 200 total relaxation curves is tested using two novel jet fuel samples, with predictions achieving an RMSE of just 0.96 DCN. Finally, the applicability and limitations of the proposed scheme are discussed.

1. Introduction

Jet fuels play a critical role in the aviation industry, serving as the primary source of energy for aircraft engines. Even with advances in all-electric aircraft, combustion engines that rely on liquid hydrocarbon fuels derived from natural gas, petroleum, coal, and other sustainable sources will remain dominant for decades or perhaps the century [1,2]. One of the key parameters that defines the ignition characteristics and combustion behavior of jet fuels is the derived cetane number (DCN). The DCN provides an indication of a fuel's ignition propensity and chemical kinetic potential for combustion processes, influencing factors such as combustion efficiency, emissions, and overall engine performance [3]. Traditionally, the DCN of jet fuels has been determined using laboratory-based methods, such as ASTM D7768 [4], D7170 [5], and D6890 [6]. These methods utilize combustion chambers and require large sample sizes, time-consuming measurements, and expensive equipment, making them impractical for on-site or real-time monitoring. However, recent advancements in analytical techniques have explored alternative approaches for DCN prediction, including

nuclear magnetic resonance (NMR) spectroscopy [7] and infrared spectroscopy [3,8,9]. Despite their effectiveness, these techniques require bulky instruments and complex data analysis procedures, limiting their applicability in portable and real-time scenarios.

Time-domain NMR (TD-NMR) is a compact and versatile variant of NMR spectroscopy that enables the measurement of the T_2 relaxation curve, a signal that characterizes the decay of nuclear magnetization in a sample [10]. By analyzing T_2 relaxation curves using simple regression techniques and machine learning (ML) algorithms, valuable insights can be gained into the molecular structural characteristics and physical properties of a sample [11,12], including the DCN [13]. The ability to predict the DCN of jet fuels using a custom, compact TD-NMR system offers several significant advantages. Firstly, it provides a rapid and efficient method for evaluating the ignition characteristics of jet fuels, enabling real-time monitoring and quality control during fuel production and use. Secondly, the compact nature of the TD-NMR system allows for portability, making it suitable for on-site or even in-situ measurements within an aircraft. This work demonstrates that

* Correspondence to: 300 Main Street, Columbia SC, 20208, USA

E-mail address: austindowney@sc.edu (A.R.J. Downey).

¹ Equal contribution.




Fig. 1. The compact NMR system used in this work. (a) The desktop setup with primary components and subsystems annotated. (b) The system schematic showing components on custom-printed circuit boards and external components involved in signal processing and amplification.

when compared to ASTM standards [4–6], sample sizes required for a TD-NMR method can be reduced from 100–400 mL to as small as 0.1 mL, and the time for data collection is lowered from 20–30 min to under 1 min.

ML models have been implemented to predict cetane number and DCN for a large array of jet fuels. However, training datasets have consisted of NMR spectroscopic data [14], physical properties of the fuels [15], and functional group data [16]. There are currently no methods that use only the T_2 relaxation curves of jet fuels to predict DCN. While T_2 curves contain structural information, they are difficult to interpret given their low resolution. By applying ML algorithms, however, features can be easily extracted from T_2 relaxation curves, and the relative importance of different features can be estimated using simple analysis techniques [17,18]. Such steps are crucial for prioritizing the most influential variables in a regression model. In this work, random forests are used to regress DCN from T_2 relaxation curve features. Random forests are a class of highly interpretable ML models [19] that exhibit strong performance on small datasets [20,21]. Hence, the decision to utilize random forests over other ML models, specifically artificial neural networks, is motivated by the limited available training data and an attempt to avoid overfitting.

The contributions of this work are threefold: (1) the custom TD-NMR system capable of rapidly acquiring relaxation data from liquid samples containing vital structural information with high repeatability; (2) the feature importance analysis for predicting DCN from T_2 relaxation curves; and (3) the combined T_2 signal truncation and random forest approach for estimating the DCN of jet fuels.

2. Hardware, data, and system validation

This section introduces the custom TD-NMR system developed for this work. NMR relaxometric measurements are discussed, and the datasets generated for analysis are presented.

2.1. Compact time-domain NMR hardware

The compact TD-NMR system and the schematic for the electronics and control subsystems are shown in Fig. 1. The *Van Hove* configuration of this system is open-source [22]. Waveform generation, calibration to the optimal Larmor frequency, and data collection and export are performed by a LabVIEW program. The pictured laptop interfaces with the system's hardware through a Thunderbolt cable, while the remainder of the custom electronics are linked to the PXI chassis via 50 Ω cables. All amplification and signal routing components are mounted on printed circuit boards (PCBs), with the exception of the first-stage low-noise amplifier and the high-power amplifier. The system is designed for analytes to be in standard 5 mm tubes.

The configuration of the permanent magnet utilized can be seen in Fig. 1(a) and consists of two cylindrical dipole magnets enclosed by a steel yolk [23]. The NdFeB permanent magnets are grade N42 and axially magnetized, having a diameter of 76.2 mm (3 in) and a thickness

of 25.4 mm (1 in). When positioned 15 mm apart, the bare magnets generate a measured flux density of 0.5 T. To enhance homogeneity and elevate flux density within the gap, the disks are enclosed by 1018 carbon steel bars, each with a thickness of 19 mm. Further, 7.5 mm thick 1018 steel caps are affixed to the magnet surfaces [22]. While bolstering overall strength and uniformity, the added steel elements markedly improve the safety of the system by restricting field lines from having any significant effect beyond the immediate volume of the magnet. Altogether, a peak flux density of 0.645 T is achieved, a value that corresponds to a Larmor frequency of 27.5 MHz.

The core electronics essential for NMR signal excitation and detection are depicted in the schematic of Fig. 1(b). Powering the entire system is a single 24 V power supply that is channeled into a network of linear regulators. The regulators step down the 24 V potential to three levels: 12 V for the initial stage low-noise amplifier, 5 V for the second stage low-noise amplifier, and 1.8 V for the switch. The excitation process initiates with a sinusoidal waveform set at -5 dBm from the waveform generator that is matched to the Larmor frequency. This signal is divided into two branches using a 2-way, 0° power divider. One branch is routed to the local oscillator (LO) port of the frequency mixer with a conversion loss of 4.6 dB, while the other is directed to the absorptive switch equipped with shunt legs terminated at 50 Ω . The switch's state is under the control of a pulse generator, carefully synchronized with the Carr–Purcell–Meiboom–Gill (CPMG) pulse sequence [24].

The train of pulsed RF signals is passed through a high-power amplifier that boosts the power to 35 dBm. The RF pulses are then directed into the duplexer circuit comprising a pi filter and crossed diodes. The duplexer, created with lumped elements calibrated to resonate at the Larmor frequency, serves to (1) shield the sensitive low-noise amplifiers from potential damage due to high-power pulses, and (2) direct power towards the probe to reduce losses. The probe configuration includes a solenoidal coil along with two adjustable tuning and matching capacitors. The solenoid is composed of 8 turns of hand-wound copper wire insulated with Kapton and has an internal diameter of 5 mm. Ceramic trimmer capacitors are placed both in series with and across the probe to match its impedance to 50 Ω . The probe's quality factor (Q value), which is a common metric for determining efficient power delivery, was obtained using a network analyzer and the formula

$$Q = \frac{f_c}{\Delta f}, \quad (1)$$

where f_c is the resonant frequency of the circuit and Δf is the resonance width at half-power. The Q value for the probe was determined to be 70, a number that balances efficient power transmission without complicating frequency optimization during scan preparation. Following sample excitation, the microvolt-level NMR response traverses the duplexer and receives a 40 dB amplification at the first stage low-noise amplifier. The signal is then subjected to frequency mixing with the original sinusoid, resulting in a decaying waveform positioned within the audio frequency range. This waveform undergoes an additional




Fig. 2. Spin echos from a sample of DI water following the application of a CPMG pulse sequence. The circled peaks are used to construct the sample's T_2 relaxation curve.

40 dB amplification before being passed through a 5–15 kHz active band-pass filter.

Within the NI PXI chassis, three components constitute the signal generation, control, and data acquisition subsystems. These components include a pulse train generator (PXIe-5413), an arbitrary waveform generator (PXI-5421), and a 16-bit digitizer (PXI-5922). The timing and coordination of waveforms are under the control of a custom LabVIEW program. The initiation of a scan depends on the rising edge of the pulse train generator, an event that triggers the switch and initiates the sinusoidal signal that is directed into the power splitter. For this system, a pulse duration of 7 μ s corresponds to a 90° flip of the sample's magnetization. The time gap separating 90° and 180° pulses, often referred to as the τ value, is 0.625 ms. Each scan encompasses a total of 3955 pulses, collectively establishing a relaxation window spanning 5 seconds. For consistency, the decay curve of a sample is averaged from four consecutive scans, and a 10-second delay is observed between scans to ensure complete relaxation of the sample. The acquisition of four averages of a sample's decay curve is thus accomplished in less than one minute. If it is desired that total acquisition time be further reduced, this can be achieved by shortening the delay between successive scans and/or acquiring fewer than four averages of a sample's decay curve.

The LabVIEW program is configured to identify and graph the peak voltage value associated with each spin echo. These voltages decay exponentially over time and form the T_2 relaxation curve of the sample. Fig. 2 displays three spin echos captured following the application of one 90° and two 180° RF pulses to a sample of DI water. A signal-to-noise (SNR) ratio of 25 dB was measured for a single scan. While one might anticipate diminishing amplitudes for successive echos in Fig. 2, initial echos exhibit amplitude fluctuations due to phase discrepancies between the 90° and 180° RF pulses. However, due to phase normalization, these fluctuations diminish considerably after several successive 180° pulses. Furthermore, the process is refined by averaging multiple scans, ultimately culminating in the synthesis of the final decay curve. To streamline user interaction, a graphical user interface (GUI) has been developed. The GUI serves as the system's front end, enabling users to adjust parameters, view acquired decay curves, and export data for further analysis.

It is important to note that the system's magnetic field strength is influenced by a temperature shift gradient of -800 ppm/K. Therefore, an essential task before each scan is to identify the optimal

Table 1

The 12 pure hydrocarbon and 17 jet fuel samples used for dataset generation and analysis. Reported DCN values are subject to small uncertainties (approximately ± 1).

Hydrocarbons		Jet fuels		
Name	DCN	Name	POSF	DCN
toluene ^a	6.0	Gevo-ATJ	10 151	15.5
1,3,5-trimethylbenzene	8.0	JP-8/Gevo-ATJ ^a	10 153	30.5
iso-cetane ^a	14.2	Sasol IPK ^a	7629	31.3
iso-octane ^a	18.9	Shell CPK	13 690	37.2
n-propylbenzene	19.5	JP-8/IPK	7718	40.0
methylcyclohexane	22.5	JP-5	10 289	40.9
n-butocyclohexane ^a	47.8	CN40 ^b	N/A	42.0
n-heptane ^a	56.0	CN50 ^b	N/A	53.4
n-octane	64.4	Jet-A	4658	47.1
n-decane	66.4	JP-8 ^a	6169	47.3
n-hexadecane	73.5	JP-8	10 264	49.6
n-dodecane	100.3	Jet-A	10 325	50.0
-	-	JP-8/HRJ Tallow	7719	53.3
-	-	HRJ Tallow ^a	6308	58.1
-	-	Shell SPK	5729	58.4
-	-	S-8 ^a	4734	58.7
-	-	HRJ Camelina	7720	58.9

^a Reprobed for model validation.

^b Reserved for model testing.

operating frequency. This is accomplished through a LabVIEW program which performs both coarse and fine estimations of the Larmor frequency. For coarse estimation, ambient temperature is measured with a thermocouple, and the Larmor frequency is estimated using a temperature-to-frequency profile developed through extensive testing. To refine this estimate, a frequency sweep is conducted around the estimated Larmor frequency, and the optimal operating frequency is selected as the one maximizing SNR. Beyond optimal frequency identification, no additional shimming is required as the system acquires data in the time domain. In contrast to the high-resolution data obtained via NMR spectroscopy, the low-resolution of the presented open-source, compact TD-NMR system [22] is a key aspect of this work.

2.2. Considered datasets

Table 1 summarizes the two sets of samples considered for analysis. The first set consists of 12 pure hydrocarbons, including aromatics, cycloalkanes, and alkanes, with DCNs ranging from approximately 6 to 100. The second set is a collection of 17 jet fuels with DCNs in the range of approximately 15 to 60. A training dataset was generated by acquiring the T_2 relaxation curve of each sample three times. However, the T_2 curves of CN40 and CN50 were not acquired for training, as these samples were reserved for testing the ability of models to generalize to new data. To augment the training dataset, the three T_2 curves of each sample were averaged pair-wise, synthesizing additional relaxation curves that were assigned the same DCN as their constituents. Altogether, 72 hydrocarbon and 90 jet fuel T_2 curves were produced for training. A validation dataset was generated by probing ten randomly selected hydrocarbon and jet fuel samples from Table 1 five times. While the samples used for validation are present in the training dataset, the T_2 curves used for validation were acquired separately and thus are distinct from those used for training. All of the data used in this work is provided as a supplement to this article and is publicly available [25] under the "0.645_Tesla_magnet" sub-dataset.

2.3. Hydrogen-to-carbon ratio

The hydrogen content in fuel and hydrocarbon analysis holds a large significance as it directly influences combustion characteristics. For example, the hydrogen-to-carbon (H/C) ratio serves as a key indicator of specific energy density for jet fuels [26]. In the case of

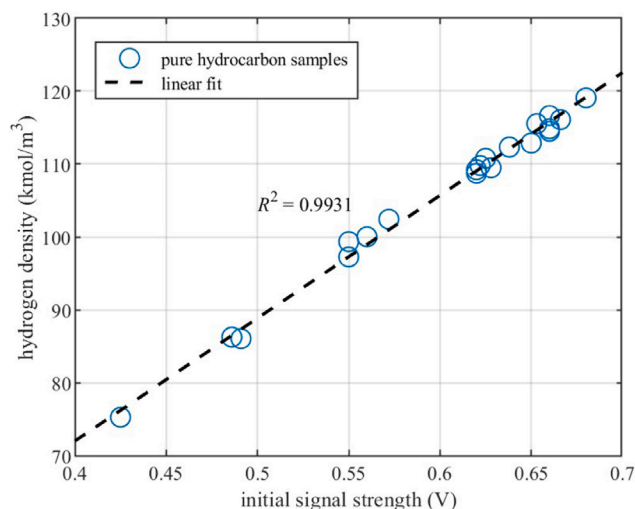


Fig. 3. Hydrogen density vs initial signal amplitude for 12 pure hydrocarbons and eight mixtures of hydrocarbons with 1:1 mass ratios.

petroleum-derived jet fuels, higher hydrogen content (H/C ratio) corresponds to less aromatic content, a balance indicative of improved ignition propensity and reduced particulate emissions. Furthermore, the blending of sustainable aviation fuels with petroleum-derived jet fuels yields increased hydrogen content, thus reducing particulate emission considerably [27]. A thorough understanding of hydrogen content is therefore essential for optimizing fuel formulations, ensuring efficient combustion, and minimizing environmental impacts.

For a general T_2 relaxation signal, the signal-to-noise ratio is described by

$$\text{SNR} \propto N A T^{-1} B_0^{3/2} \gamma_{\text{exc}}^{3/2} \gamma_{\text{obs}}^{3/2} T_2^* n_s^{1/2}, \quad (2)$$

where N is the number of spins, A is the abundance of active spins, T is temperature, B_0 is the static magnetic field strength, γ_{exc} and γ_{obs} are the gyromagnetic ratios of the excited and observed spins, T_2^* is the effective transverse relaxation time, and n_s is the number of scans for a given sample [10,28]. Note that the only sample-dependent parameter is the number of spins, or hydrogen density, so the initial amplitude of a T_2 relaxation signal is directly proportional to the sample's hydrogen content. To validate the developed compact NMR system, a plot of hydrogen density against initial signal strength was generated using the 12 hydrocarbons in Table 1 and eight mixtures with 1:1 mass ratios. The theoretical hydrogen density of each sample was calculated according to

$$\rho_{\text{H}} = \frac{\rho_s N_{\text{H}}}{M_w}, \quad (3)$$

where ρ_s is mass density, N_{H} is hydrogen number, and M_w is molecular weight. Five scans (averages) were performed with the NMR system for consistency. The resulting plot is shown in Fig. 3, where the clear linear relationship is consistent with the proportionality expressed in Eq. (2).

2.4. T_2 relaxation

T_2 relaxation data plays a key role in unraveling the structural intricacies and chemical bond groupings within hydrocarbon samples. Because hydrogen nuclei within a hydrocarbon molecule experience distinct magnetic environments, T_2 relaxation time can provide insights into molecular mobility and even interactions within the molecular structure itself. Shorter T_2 relaxation times are indicative of restricted molecular motion, often associated with denser or more complex structures. Conversely, longer T_2 relaxation times suggest greater molecular mobility, a characteristic of less constrained or simpler structures.




Fig. 4. T_2 relaxation curves for various pure hydrocarbons constructed from a CPMG sequence comprised of 3955 total spin echos.

Hence, by analyzing T_2 relaxation data, one can detect subtle variations in a hydrocarbon's molecular composition, identify specific chemical moieties, and correlate relaxation times with particular structural features or bond groupings. This information is invaluable for characterizing hydrocarbons, understanding their behavior in diverse environments, and optimizing their applications in various industrial processes, including the formulation of advanced fuels and petrochemical products.

A T_2 relaxation curve can be modeled as

$$M_{xy}(t) = M_0 \exp(-t/T_2), \quad (4)$$

where $M_{xy}(t)$ is the net magnetization of spins in the transverse plane, M_0 is the thermal equilibrium magnetization, and T_2 is the primary relaxation time. For hydrocarbons, it is well established that T_2 time correlates with hydrogen number [29]. In particular, increased molecular size corresponds to reduced spacing between spins, an effect that decreases relaxation time. To further validate the proposed TD-NMR system, the T_2 time of each hydrocarbon in Table 1 was plotted against its hydrogen number. Primary T_2 relaxation times were identified by fitting first-order exponentials to the acquired decay curves. Fig. 4 pictures the relaxation curves of four representative hydrocarbon samples, while Fig. 5 illustrates the dependence of T_2 time on hydrogen number. It can be seen that the measured relaxation times decrease linearly with hydrogen number, a result that is consistent with established research [10].

3. Methodology

This section introduces the interpretable machine learning techniques employed for relaxometric data analysis. A procedure for estimating feature importance is summarized, and random forests are trained for DCN prediction on both pure hydrocarbon and jet fuel samples.

3.1. Explainable machine learning approach

Fig. 6(a) outlines the utilized process for estimating DCN from raw T_2 relaxation data. In this pipeline, the random forest is the interpretable machine learning methodology that takes as input T_2 relaxation curve features and outputs an inferred DCN [19]. The scheme requires minimal data preprocessing and is overall computationally straightforward. Fig. 6(b) illustrates the associated training procedure

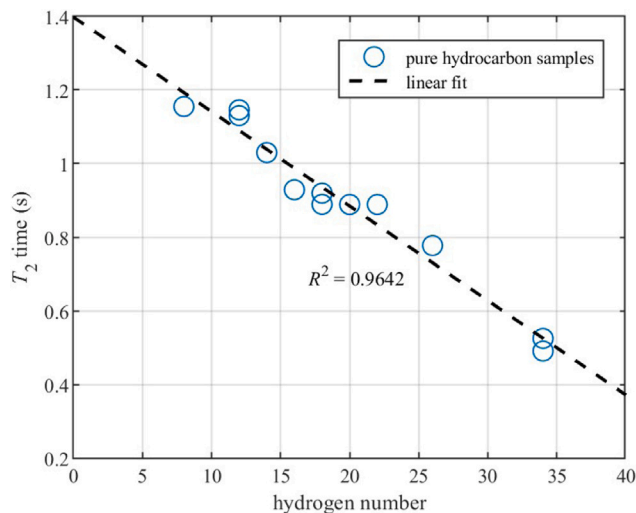


Fig. 5. Relationship between T_2 time and hydrogen number for 12 pure hydrocarbon samples.

for fitting random forests to T_2 relaxation data. The specifics of hyper-parameter tuning and feature importance estimations are detailed in later sections. For comparison, models are trained for DCN prediction using various subsets of the training data feature space, and performance is analyzed on the basis of root mean square error (RMSE). Given N observations, the RMSE between true DCN and model output is computed using

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}, \quad (5)$$

where y_i and \hat{y}_i are true DCN and model-projected DCN for the i th observation, respectively.

Table 2 reports the ten standard, physically interpretable time-domain signal metrics investigated in this work. The features are computed using first-order exponentials fit to the relaxation curves and include exponential amplitude and decay rate, statistical averages, higher-order moments, and peak-related metrics. Note that not the entire duration of a signal is considered for feature computation. This choice was motivated by Fig. 7, where the natural logarithm of select hydrocarbon and jet fuel T_2 curves are shown to flatten around -5 on a log-scale (<10 mV); it is at this point that the NMR signal becomes indistinguishable from noise. To exclude this latter portion of signals from influencing feature values and consequently model predictions, T_2 relaxation curves are truncated after decaying to a fraction of their initial amplitude. For a relaxation curve with $M_{xy}(0) = M_0$, only the first t_c -seconds of the signal are utilized for feature computation, where $M_{xy}(t_c) = \alpha M_0$ for some predetermined α satisfying $0 < \alpha < 1$. During training, the truncation fraction α is treated as an additional parameter to be optimized.

Feature-feature plots for all T_2 curves within the training dataset were generated for initial feature exploration and are shown in the appendix (Fig. A.1). Plots of a feature against itself have been replaced by a histogram of the feature's distribution. In general, the histograms of pure component feature observations appear more disjoint than their jet fuel counterparts. Hence, it is expected that the regression of DCN for hydrocarbons will be more straightforward than for jet fuels, a projection that follows from the comparatively complex structure of hydrocarbon-derived fuels. See the appendix (Appendix A) for additional information.

3.2. Data-driven feature analysis

To estimate the relative importance of the ten signal features considered for DCN regression, separate ensembles of bootstrap aggregated

Table 2

Considered T_2 relaxation curve features with associated interpretations and formulas. Here x represents the data points constituting a sample's T_2 relaxation curve.

Feature	Interpretation	Formula
amplitude	initial signal strength	M_0
decay rate	T_2 relaxation rate	$1/T_2$
mean	average value	$\frac{1}{N} \sum_{i=1}^N x_i$
standard deviation	spread around the mean	$\sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2}$
root mean square	average power	$\sqrt{\frac{1}{N} \sum_{i=1}^N x_i^2}$
shape factor	signal shape	x_{rms}/\bar{x}
kurtosis	tail length	$\frac{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^4}{[\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2]^2}$
skewness	signal asymmetry	$\frac{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^3}{[\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2]^{3/2}}$
impulse factor	ratio of amplitude to mean	M_0/\bar{x}
crest factor	ratio of amplitude to RMS	M_0/x_{rms}

(bagged) regression trees were trained for DCN prediction using the hydrocarbon and jet fuel training data. Bagging is a method in which a random forest is trained by fitting weak learners (individual trees) to randomly sampled subsets of the training data. The feature observations not included in these subsets are called out-of-bag observations and can be used to evaluate model performance during training. Additionally, permuting out-of-bag observations is a common technique for feature importance analysis [30]. The process is straightforward: permute the out-of-bag observations for each feature and observe the resulting changes in model response. The greater the impact of feature permutation on model error, the more important that feature is in the regression model. This procedure is applied across all weak learners and is summarized in Algorithm 1.

Training loops were implemented to identify the optimal truncation amplitude for feature computation, and Bayesian optimization algorithms from MATLAB's *Statistics and Machine Learning Toolbox* [31] were used to automatically tune the following hyperparameters: forest size, learning rate, minimum leaf size, maximum number of splits, and number of predictors to sample. Following the procedure shown in Fig. 6(b), models were trained twice on both the hydrocarbon and jet fuel training data. Between the first and second rounds of training, the only parameter that was varied was the node-splitting technique, which was switched from a curvature test to an interaction test to discern the potential for interactions among the investigated features. The curvature algorithm conducts a chi-square test of independence between each feature and the response, splitting nodes according to

Algorithm 1 Estimate out-of-bag predictor importance.

```

1: procedure ESTIMATEPREDICTORIMPORTANCE( $R, D$ )
2:    $R$ : random forest
3:    $D$ : training data
4:    $I_i$ : importance of  $i$ th predictor
5:   for tree  $t \in \{1, \dots, T\}$  do
6:     obtain subset of predictors used to fit  $t$ ,  $p \subseteq \{1, \dots, P\}$ 
7:      $\epsilon_t \leftarrow$  out-of-bag error
8:     for  $p_i \in p$  do
9:        $r_i \leftarrow$  randperm( $p_i$ )
10:       $\epsilon_{ti} \leftarrow$  model error
11:       $\delta_{ti} \leftarrow \epsilon_{ti} - \epsilon_t$ 
12:    end for
13:     $\bar{\delta}_i \leftarrow$  mean( $\delta_i$ )
14:     $\sigma_i \leftarrow$  stdev( $\delta_i$ )
15:     $I_i \leftarrow \bar{\delta}_i/\sigma_i$ 
16:  end for
17: end procedure

```

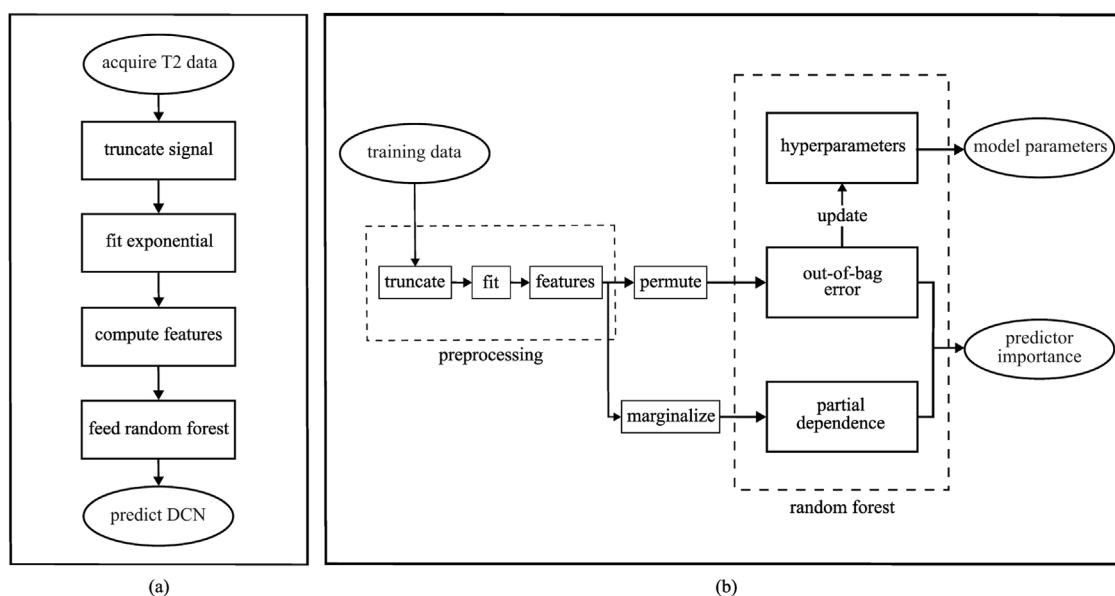


Fig. 6. Visual summary of the interpretable machine learning approach. (a) Flowchart depicting steps for inferring DCN from raw T_2 relaxation data. (b) Random forest training workflow for hyperparameter tuning and feature importance estimations.

Table 3

Averaged out-of-bag feature importance estimates obtained via fitting loops employing both curvature and interaction node-splitting techniques, where a higher number infers greater importance. The features are ordered in alignment with Table 2.

Feature	Importance			
	Hydrocarbons		Jet fuels	
	Curvature	Interaction	Curvature	Interaction
amplitude	0.887	0.801	0.576	0.921
decay rate	2.039	2.242	1.256	2.307
mean	0.503	0.418	0.556	0.756
std. dev.	0.336	0.299	0.424	0.543
RMS	0.160	0.132	0.417	0.497
SF	0.070	0.062	0.111	0.064
kurtosis	0.512	0.602	0.359	0.429
skewness	0.236	0.238	0.241	0.281
IF	0.079	0.057	0.078	0.063
CF	0.077	0.045	0.063	0.040

the predictor that minimizes p -value. The interaction algorithm operates similarly but also performs a chi-square test of independence between each feature and all other features, splitting nodes based on the predictor that minimizes p -value for both tests.

Upon completing a full training loop, feature importance estimates were averaged across all models to capture general trends in the data, ensuring that the results are unbiased by the truncation fraction used for feature computation. The averaged feature importance estimates are presented in Table 3, where larger values indicate greater importance. For both hydrocarbons and jet fuels, T_2 decay rate appears as the most important predictor of DCN. This result is expected as the relaxation rate of a sample is closely tied to molecular structure. Initial signal strength (amplitude) and sample mean follow as the second and third most influential variables, respectively. As anticipated, highly correlated features such as shape, impulse, and crest factor receive similarly low importance estimates. For hydrocarbons, there is a tendency for predictors to be assigned higher importance when using the curvature test for node splitting. However, this is not the case for jet fuels, where features are systematically assigned greater importance when using the interaction test to split nodes. This suggests that complex interactions occur upon the conglomeration of pure components for fuel synthesis, a result reflected in the interaction among extracted features.

To further assess feature importance, partial dependence plots were generated for all models trained on the hydrocarbon and jet fuel T_2

relaxation data. A partial dependence plot visually depicts how the output of a regression model varies with each of its input features [32]. Letting X denote the set of feature inputs to a regression model $f(\cdot)$, the partial dependence of $f(\cdot)$ on a subset of its input features, $X^S \subset X$, is computed by holding the predictors in X^S constant and marginalizing over all other predictors in X^C (i.e., the complementary set of X^S). For N total feature observations, assuming that each observation is equally likely, partial dependence is estimated as

$$f^S(X^S) = \frac{1}{N} \sum_{i=1}^N f(X^S, X_i^C), \quad (6)$$

where $f(X^S, X_i^C)$ is the output of the regression model for the inputs X^S and the i th observation within X^C [31,32]. Fig. 8 shows normalized partial dependence plots generated from averaged model outputs on the hydrocarbon and jet fuel T_2 relaxation data. The horizontal axes display feature values, normalized so that differing variances among the features are not exacerbated. The vertical axes depict regressed DCN, also normalized to capture trends in model outputs. The plots elucidate findings in Table 3 and Fig. A.1. Namely, T_2 decay rate not only varies the most among samples, but it also performs the best at driving changes in estimated DCN. Moreover, there exists a clear disparity in feature variance between the hydrocarbon and jet fuel relaxation data.

4. Results and discussion

This section reports results and provides a discussion of their interpretation.

4.1. Model DCN predictions

Fig. 9 displays model predictions on the validation dataset, where the sample mean of predictions have been plotted along with standard deviation (error) bars. Notably, the most inaccurate DCN predictions are for those fuels with low DCN. This inaccuracy is likely due to the distribution of jet fuel DCNs in the training dataset, where only four of the 15 fuel samples have a DCN below 40. Thus, via one or more of the following approaches, future work will focus on improving model predictions for jet fuels with low DCN: (1) extending the training dataset to include additional fuel samples with small DCN; (2) generating synthetic T_2 relaxation data having features statistically similar


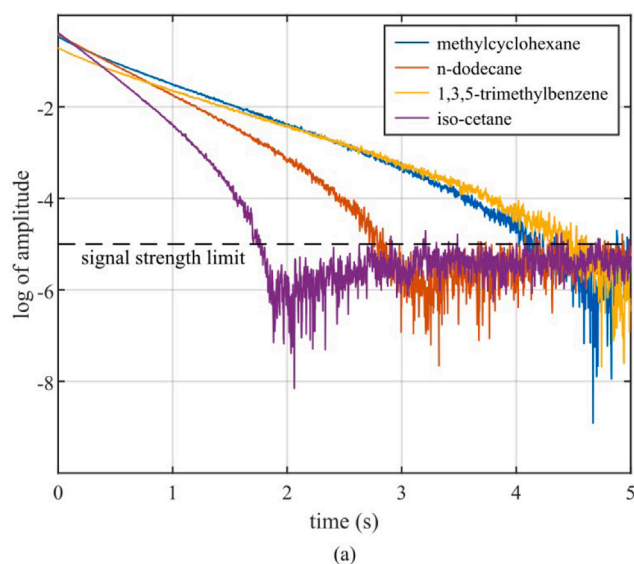


Fig. 7. Natural logarithm of select T_2 curves. (a) Four pure hydrocarbon samples. (b) Four jet fuel samples.

to those fuels with low DCN; or (3) retraining models using a stratified sampling approach, where additional weight is given to models trained on low-DCN subsets of the training data.

The proposed interpretable ML approach proves highly effective at inferring DCN from visually similar T_2 curves. For example, consider Fig. 10 which shows the relaxation curves acquired from CN40 and CN50, i.e., the two test fuel samples not used during training. The decay curves of the two samples not only appear similar, but the features computed from their relaxation curves vary by (on average) less than one part-per-hundredth. Yet, using the compact TD-NMR system in conjunction with a random forest trained on only 162 T_2 curves, model predictions for both samples are within ± 2 of true DCN, achieving a total RMSE of just 0.96 DCN. These results are detailed in Table 4 and indicate the ability of the model to infer the DCN of novel jet fuel samples.

4.2. Discussion on DCN regression

Compared to predicting on pure components, considering a smaller fraction of T_2 curves improves performance for jet fuel DCN regression.


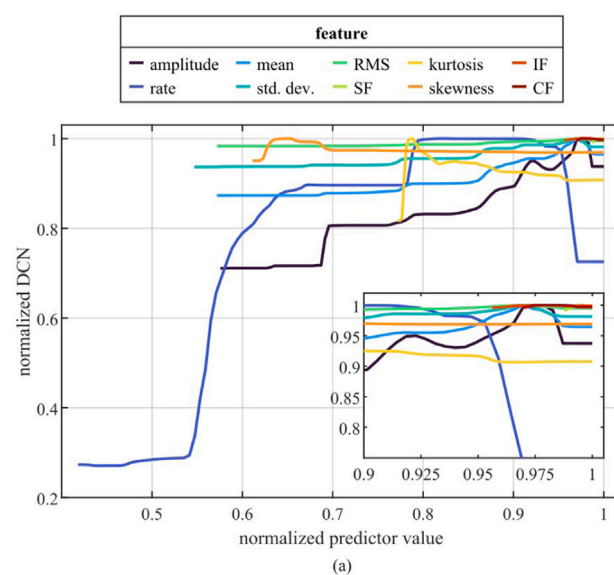


Fig. 8. Normalized partial dependence plots generated from averaged model outputs. (a) Random forests trained on hydrocarbon T_2 data. (b) Random forests trained on jet fuel T_2 data.

Table 4

Summary of model predictions on the validation and test jet fuels.

Fuel	POSF	DCN		
		True	Predicted	Error
HRJ Tallow	6308	58.1	57.4	0.7
JP-8/Gevo-ATJ	10 153	30.5	37.5	7.0
JP-8	6169	47.3	47.0	0.3
S-8	4734	58.7	58.4	0.3
Sasol IPK	7629	31.3	40.6	9.3
CN40	N/A	42.0	42.6	0.6
CN50	N/A	53.4	54.6	1.2

For example, truncating signals after decaying to 13%–15% yielded the lowest RMSE for hydrocarbon predictions, while a range of 23%–25% was optimal for jet fuels. This difference can be ascribed to the disparity in variance between hydrocarbon and jet fuel T_2 decay rates. To allow for more resolution in the estimate of a sample's relaxation time, a longer duration of the signal should be considered. Consequently, the optimal truncation fraction for pure hydrocarbons is less than that for jet fuels.

Lastly, DCN regression from T_2 relaxation data exhibits sensitivity to same-sample feature differences arising from small irregularities in setup conditions. In fact, instances were observed during training and

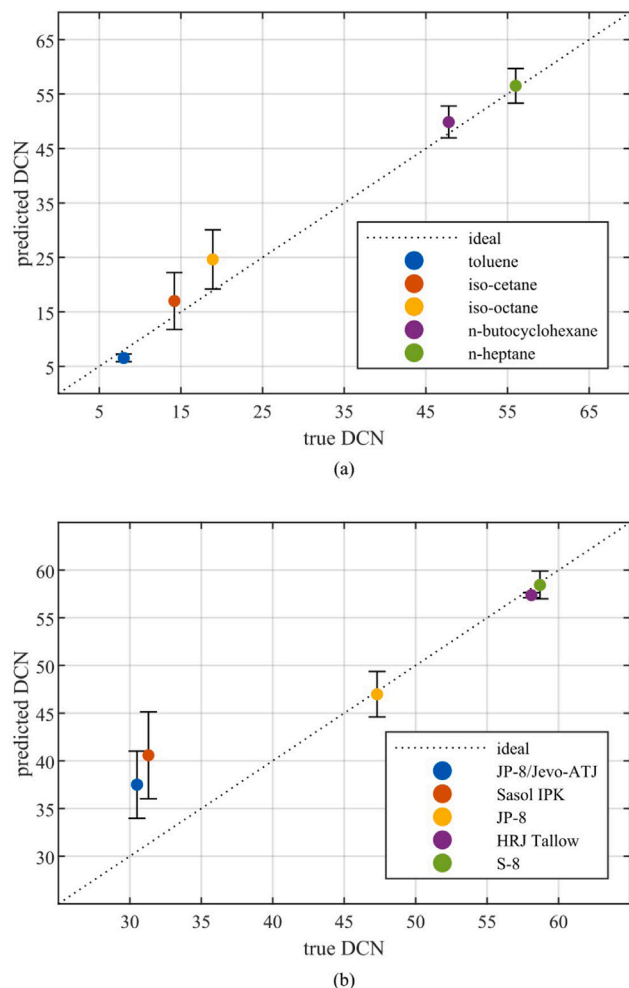


Fig. 9. True and predicted DCNs for the validation samples, where each data point corresponds to the average of five individual predictions. (a) Hydrocarbon RMSE: 3.08. (b) Jet fuel RMSE: 5.21.




Fig. 10. T_2 relaxation curves probed from two jet fuel samples not exposed to models during training.

validation where a model would output different estimates of DCN for the same sample. The lack of model robustness against such feature differences is likely due to the limited available training data. There are two natural steps towards achieving more consistent predictions: (1) enhancing system regularity through improved temperature control, hardware standardization, and noise reduction; and (2) extending the existing training and validation datasets to include the T_2 relaxation curves of additional jet fuel samples.

5. Conclusions

This work presents a compact TD-NMR system for acquiring jet fuel T_2 relaxation curves. The system integrates a permanent magnet with custom electronics and control hardware, including a GUI for easy Larmor frequency tuning and signal acquisition. Validation experiments using pure hydrocarbon samples demonstrated the system's ability to rapidly acquire T_2 relaxation data containing valuable structural information. Datasets were generated for DCN regression from T_2 relaxation data, and the metric of out-of-bag error was used to assess the relative importance of investigated T_2 curve features. A model trained on fewer than 200 observations was then tested using two novel jet fuel samples, yielding an RMSE of just 0.96 DCN. While the proposed approach for DCN prediction shows significant promise, it is sensitive to small feature differences and thus prone to inconsistency. To that end, future work will focus on training models using a larger array of jet fuels to improve prediction consistency, especially for fuels with comparatively low DCN. Additionally, system enhancements will be explored, including the incorporation of a flow-through module for accelerated, real-time probing.

CRediT authorship contribution statement

Parker Huggins: Writing – review & editing, Writing – original draft, Software, Methodology, Formal analysis, Data curation. **Jacob S. Martin:** Writing – original draft, Methodology, Data curation. **Austin R.J. Downey:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Sang Hee Won:** Writing – review & editing, Validation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This material was sponsored by the ARO, United States under Grant No. W911NF 21-1-0306 and the National Science Foundation, United States under Grant Nos. 2152896 and 2344357. The views and conclusions contained within this document are those of the authors and should not be interpreted as representing the official policies of the ARO, NSF, or U.S. Government.

Appendix A

Fig. A.1 shows feature-feature plots for all hydrocarbon and jet fuel feature observations within the training dataset, where plots of a feature against itself have been replaced by a histogram of the feature's distribution. For reference, a measure of the correlation between each pair of features is reported via the Pearson correlation coefficient defined as

$$\rho(x, y) = \frac{\text{cov}(x, y)}{\sigma_x \sigma_y}, \quad (\text{A.1})$$

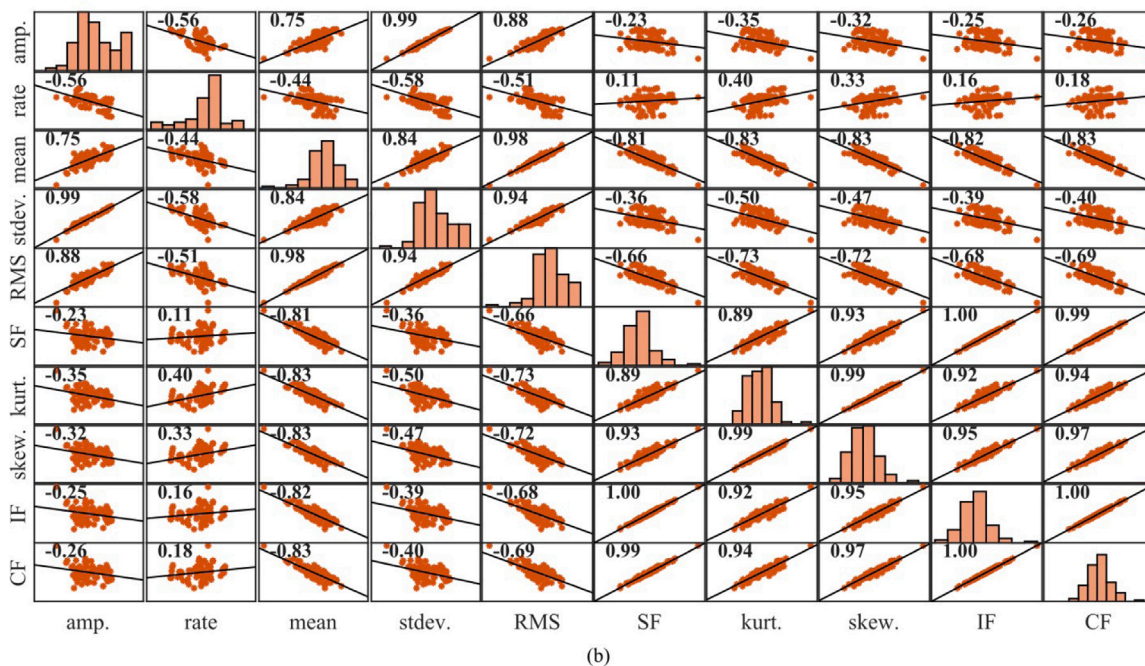
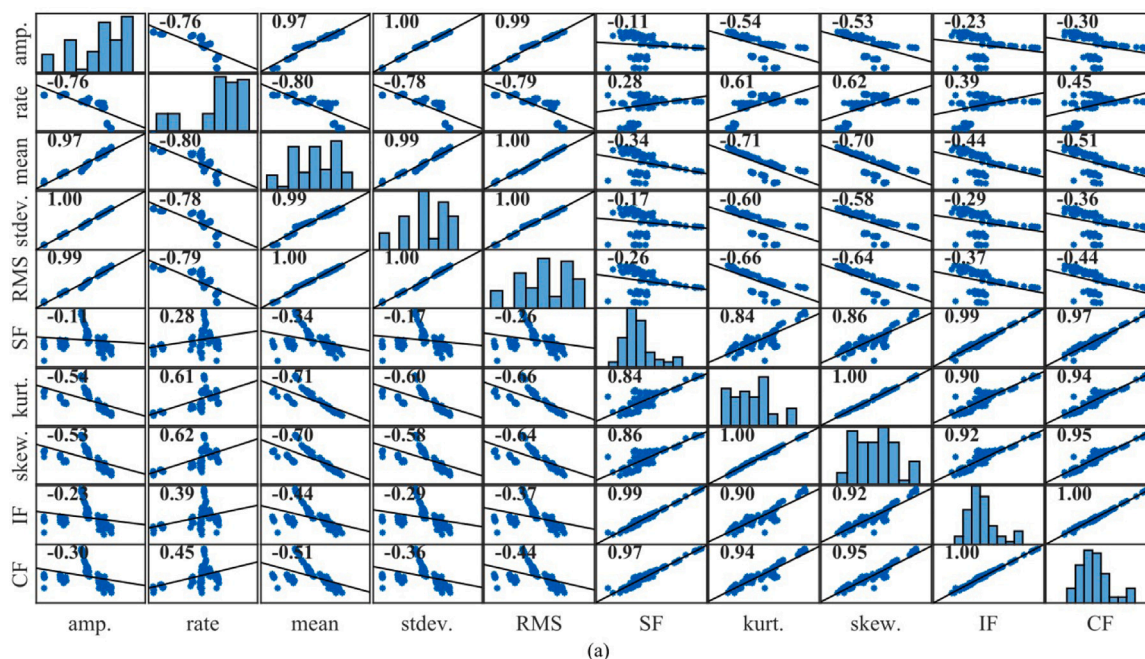


Fig. A.1. Visualization of T_2 curve feature distributions and correlations. (a) Hydrocarbon training data. (b) Jet fuel training data.

where x and y are feature observations, $\text{cov}(x, y)$ is the covariance of x and y , and σ_x and σ_y are, respectively, the standard deviation of x and y . While many of the considered features are highly correlated and consequently redundant predictors, all features are considered for analysis as to (1) discern the potential for interaction between pairs of predictors and (2) validate the results of feature importance estimates. Attention is drawn to the clear differences in feature distributions between the hydrocarbon and jet fuel T_2 data.

Appendix B. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.snb.2024.137018>.

Data availability

The data is uploaded with the paper and shared on GitHub.

References

- [1] D. Carpenter, S. Nates, F.L. Dryer, S.H. Won, Evaluating ignition propensity of high cycloparaffinic content alternative jet fuel by a chemical functional group approach, *Combust. Flame* 223 (2021) 243–253, <http://dx.doi.org/10.1016/j.combustflame.2020.09.024>.
- [2] F.L. Dryer, Chemical kinetic and combustion characteristics of transportation fuels, *Proc. Combust. Inst.* 35 (1) (2015) 117–144, <http://dx.doi.org/10.1016/j.proci.2014.09.008>.

- [3] Y. Wang, Y. Cao, W. Wei, D.F. Davidson, R.K. Hanson, A new method of estimating derived cetane number for hydrocarbon fuels, *Fuel* 241 (2019) 319–326, <http://dx.doi.org/10.1016/j.fuel.2018.12.027>.
- [4] ASTM D7668-17 Standard Test Method for Determination of Derived Cetane Number (DCN) of Diesel Fuel Oils—Ignition Delay and Combustion Delay Using a Constant Volume Combustion Chamber Method, American Society for Testing Materials, West Conshohocken, PA, 2017.
- [5] ASTM D7170-16 Standard Test Method for Determination of Derived Cetane Number (DCN) of Diesel Fuel Oils—Fixed Range Injection Period, Constant Volume Combustion Chamber Method (Withdrawn 2019), American Society for Testing Materials, West Conshohocken, PA, 2016.
- [6] ASTM D6890-22 Standard Test Method for Determination of Ignition Delay and Derived Cetane Number (DCN) of Diesel Fuel Oils by Combustion in a Constant Volume Chamber, American Society for Testing Materials, West Conshohocken, PA, 2022.
- [7] A.G. Abdul Jameel, N. Naser, A.-H. Emwas, S. Dooley, S.M. Sarathy, Predicting fuel ignition quality using ¹H NMR spectroscopy and multiple linear regression, *Energy Fuels* 30 (11) (2016) 9819–9835, <http://dx.doi.org/10.1021/acs.energyfuels.6b01690>.
- [8] E. Al Ibrahim, A. Farooq, Prediction of the derived cetane number and carbon/hydrogen ratio from infrared spectroscopic data, *Energy Fuels* 35 (9) (2021) 8141–8152, <http://dx.doi.org/10.1021/acs.energyfuels.0c03899>.
- [9] A. Dalmiya, M. Sheyyab, J.M. Mehta, K. Brezinsky, P. Lynch, Derived cetane number prediction of jet fuels and their functional group surrogates using liquid phase infrared absorption, *Proc. Combust. Inst.* 39 (1) (2023) 1495–1504, <http://dx.doi.org/10.1016/j.proci.2022.08.104>.
- [10] J. Martin, A. Downey, S.H. Won, Compact time domain NMR design for the determination of hydrogen content in gas turbine fuels, in: *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 86205, American Society of Mechanical Engineers, 2022, <http://dx.doi.org/10.1115/DETC2022-90023>, p. V001T01A017.
- [11] P. Huggins, W. Janvrin, J. Martin, A. Womer, A.R. Downey, J.L. Ferry, M. Baalousha, J. Yan, Assessing magnetic particle content in algae using compact time domain nuclear magnetic resonance, in: *Ocean Sensing and Monitoring XVI*, vol. 13061, SPIE, 2024, 1306102, <http://dx.doi.org/10.1117/12.3013987>.
- [12] Y. Li, Z. Shi, L. Shang, Q. Tao, Q. Tang, H.-J. Krause, S. Yang, G. Ding, H. Dong, Graphene quantum dots-based magnetic relaxation switch involving magnetic separation for enhanced performances of endoglin detection using ultra-low-field nuclear magnetic resonance relaxometry, *Sensors Actuators B* 380 (2023) 133389, <http://dx.doi.org/10.1016/j.snb.2023.133389>.
- [13] P.M. Santos, R.S. Amais, L.A. Colnago, A. Rinnan, M.R. Monteiro, Time domain-NMR combined with chemometrics analysis: An alternative tool for monitoring diesel fuel quality, *Energy Fuels* 29 (4) (2015) 2299–2303, <http://dx.doi.org/10.1021/acs.energyfuels.5b00017>.
- [14] A.G. Abdul Jameel, V. Van Oudenhoven, A.-H. Emwas, S.M. Sarathy, Predicting octane number using nuclear magnetic resonance spectroscopy and artificial neural networks, *Energy Fuels* 32 (5) (2018) 6309–6329, <http://dx.doi.org/10.1021/acs.energyfuels.8b00556>.
- [15] R. Li, J.M. Herreros, A. Tsolakis, W. Yang, Machine learning regression based group contribution method for cetane and octane numbers prediction of pure fuel compounds and mixtures, *Fuel* 280 (2020) 118589, <http://dx.doi.org/10.1016/j.fuel.2020.118589>.
- [16] A.G. Abdul Jameel, A comprehensive model for cetane number prediction using machine learning, 2021, <http://dx.doi.org/10.1115/GT2021-60097>, p. V03BT04A030.
- [17] Y. Date, F. Wei, Y. Tsuboi, K. Ito, K. Sakata, J. Kikuchi, Relaxometric learning: a pattern recognition method for T2 relaxation curves based on machine learning supported by an analytical framework, *BMC Chem.* 15 (2021) 1–8, <http://dx.doi.org/10.1186/s13065-020-00731-0>.
- [18] J.S. Cooper, H. Kiiveri, E. Chow, L.J. Hubble, M.S. Webster, K.-H. Müller, B. Raguse, L. Wiczorek, Quantifying mixtures of hydrocarbons dissolved in water with a partially selective sensor array using random forests analysis, *Sensors Actuators B* 202 (2014) 279–285, <http://dx.doi.org/10.1016/j.snb.2014.05.094>.
- [19] M. Haddouchi, A. Berrado, A survey of methods and tools used for interpreting random forest, in: *2019 1st International Conference on Smart Systems and Data Science, ICSSD, IEEE*, 2019, <http://dx.doi.org/10.1109/icssd47982.2019.9002770>.
- [20] R.L. Marchese Robinson, A. Palczewska, J. Palczewski, N. Kidley, Comparison of the predictive performance and interpretability of random forest and linear models on benchmark data sets, *J. Chem. Inf. Model.* 57 (8) (2017) 1773–1792, <http://dx.doi.org/10.1021/acs.jcim.6b00753>.
- [21] M. Liu, M. Wang, J. Wang, D. Li, Comparison of random forest, support vector machine and back propagation neural network for electronic tongue data classification: Application to the recognition of orange beverage and Chinese vinegar, *Sensors Actuators B* 177 (2013) 970–980, <http://dx.doi.org/10.1016/j.snb.2012.11.071>.
- [22] ARTS-Lab, Compact-NMR, 2024, GitHub, URL <https://github.com/ARTS-Laboratory/Compact-NMR>.
- [23] R.S. Sahebjavaher, K. Walus, B. Stoeber, Permanent magnet desktop magnetic resonance imaging system with microfabricated multiturn gradient coils for microflow imaging in capillary tubes, *Rev. Sci. Instrum.* 81 (2) (2010) 023706, <http://dx.doi.org/10.1063/1.3280171>.
- [24] S. Meiboom, D. Gill, Modified spin-echo method for measuring nuclear relaxation times, *Rev. Sci. Instrum.* 29 (8) (1958) 688–691, <http://dx.doi.org/10.1063/1.1716296>.
- [25] P. Huggins, J. Martin, A. Downey, S.H. Won, Dataset-hydrocarbon-and-fuel-processing, 2024, GitHub, URL <https://github.com/ARTS-Laboratory/Dataset-hydrocarbon-and-fuel-processing>.
- [26] S.H. Won, P.S. Veloo, S. Dooley, J. Santner, F.M. Haas, Y. Ju, F.L. Dryer, Predicting the global combustion behaviors of petroleum-derived and alternative jet fuels by simple fuel property measurements, *Fuel* 168 (2016) 34–46, <http://dx.doi.org/10.1016/j.fuel.2015.11.026>.
- [27] R.H. Moore, K.L. Thornhill, B. Weinzierl, D. Sauer, E. D'Ascoli, J. Kim, M. Lichtenstern, B. Michel, P. Beaton, A.J. Beyersdorf, et al., Biofuel blending reduces particle emissions from aircraft engines at cruise conditions, *Nature* 543 (7645) (2017) 411–415.
- [28] T.D. Claridge, *High-resolution NMR Techniques in Organic Chemistry*, vol. 27, Elsevier, 2016.
- [29] S.M. Nagel, C. Strangfeld, S. Kruschwitz, Application of ¹H proton NMR relaxometry to building materials—a review, *J. Magn. Reson. Open* 6 (2021) 100012, <http://dx.doi.org/10.1016/j.jmro.2021.100012>.
- [30] B. Gregorutti, B. Michel, P. Saint-Pierre, Correlation and variable importance in random forests, *Stat. Comput.* 27 (2017) 659–678, <http://dx.doi.org/10.1007/s11222-016-9646-1>.
- [31] *Statistics and Machine Learning Toolbox: MATLAB (R2023a)*, The MathWorks Inc., Natick, Massachusetts, United States, 2023, URL <https://www.mathworks.com/products/statistics.html>.
- [32] J. Moosbauer, J. Herbinger, G. Casalicchio, M. Lindauer, B. Bischl, Explaining hyperparameter optimization via partial dependence plots, *Adv. Neural Inf. Process. Syst.* 34 (2021) 2280–2291, URL <https://proceedings.neurips.cc/paper/2021/hash/12ced2db6f0193dda91ba86224ea1cd8-Abstract.html>.

Parker Huggins is currently an undergraduate research assistant in the Adaptive Real-Time Systems Laboratory (ARTS) Lab at the University of South Carolina, where he is pursuing an electrical engineering degree. His research is focused on signal processing and he has published two conference papers as first author. In this work, Parker Huggins contributed to experimental work and data collection, machine learning model development, data analysis and interpretation, writing and drafting the manuscript.

Jacob S. Martin Graduated from the University of South Carolina in December 2023 with an MS in Physics. He obtained a B.S. in electrical engineering from the University of South Carolina in 2021. His graduate research has been focused on the development of a tabletop nuclear magnetic resonance (NMR) system for fuel and environmental analysis, which has been the subject of one Journal paper and multiple conference papers. In this work, Jacob Martin contributed to hardware system design, experimental work and data collection, writing and drafting the manuscript.

Austin R.J. Downey is an Associate Professor of Mechanical Engineering for the College of Engineering and Computing at the University of South Carolina. He received a B.S. in Civil Engineering and a Ph.D. in Engineering Mechanics; and Wind Energy Science, Engineering, and Policy (dual majors) from Iowa State University, USA. His expertise and research interests include sensor development, in situ sensing, and online machine learning. In this work, Dr. Downey contributed to conceptualization, hardware system design, machine learning model development, writing and drafting the manuscript, data analysis and interpretation, project supervision and administration.

Sang Hee Won is a Professor of Mechanical Engineering and Interim Director of Aerospace Engineering at the University of South Carolina. He earned his B.S., M.S., and Ph.D. in Mechanical and Aerospace Engineering from Seoul National University, South Korea. Dr. Won's research focuses on the combustion chemistry of transportation fuels, the development of experimental platforms for studying combustion dynamics, plasma-assisted combustion techniques, and novel combustion diagnostic methods. In this work, Dr. Won contributed to writing and drafting the manuscript, data analysis and interpretation, project supervision and administration.