



Data-driven Detection and Early Prediction of Thermoacoustic Instability in a Multi-nozzle Combustor

Chandrachur Bhattacharya^{a,b}, Jacqueline O'Connor^a, and Asok Ray^{a,c}

^aDepartment of Mechanical Engineering, The Pennsylvania State University, University Park, PA, USA;

^bDepartment of Electrical Engineering, The Pennsylvania State University, University Park, PA, USA;

^cDepartment of Mathematics, The Pennsylvania State University, University Park, PA, USA

ABSTRACT

Thermoacoustic instability (TAI) is a critical issue in modern lean-burn gas-turbine combustors, which is induced by a strong coupling between the resonant combustor acoustics and fluctuations in the heat release rate. This instability may lead to high-amplitude pressure waves that generate undesirable noise levels as well as fatigue stresses in mechanical structures of the combustor. The intense pressure fluctuations due to TAI may also cause large flow perturbations and possibly flow reversal that may lead to flame oscillations, flame liftoff, and even flame blow-out. Hence, there is a strong need for exercising control actions in a timely fashion to mitigate the TAI phenomena. Anomaly detection is an essential prerequisite to the design of a good controller and such a detector must be able to reliably predict a forthcoming TAI. To detect and predict the onset of a TAI from an ensemble of pressure time series, this paper investigates three data-driven methods: Fast Fourier transform (FFT), symbolic time series analysis (STSA), and hidden Markov modeling (HMM). The main focus of the paper is to make a comparative evaluation of these three anomaly detection methods for classification of the current regime of operation into stable and unstable categories as well as for real-time identification of precursors to impending instabilities with short-length time series of measured variables (e.g., pressure oscillations). The results, generated on experimental data from a multi-nozzle combustor apparatus, have been compared to evaluate the performance of FFT, STSA, and HMM methods for TAI analysis.

ARTICLE HISTORY

Received 24 February 2020

Revised 3 September 2020



Accepted 3 September 2020

KEYWORDS

Thermoacoustic instability; symbolic time-series analysis; hidden Markov modeling; data-driven anomaly detection; instability onset prediction; multi-nozzle combustor

Introduction

Combustion systems in modern-day applications are designed to operate fuel-lean so as to allow complete combustion and reduce NO_x emissions. These lean-burning combustors, however, are often prone to thermoacoustic instability (TAI) (Rayleigh 1845), which is a consequence of the coupling between the unsteady heat release rate from the flame and the natural acoustics of a (typically confined) combustor. The TAI phenomena lead to large and periodic pressure fluctuations, which can cause thermo-mechanical stresses on the structure as well as excessive vibration (Lieuwen and Yang 2005). The adverse pressure waves also disrupt the bulk flow, and may lead to flow choking and flow reversal in the combustion system, which may in turn cause further flame instability and even flame extinction or blow-

CONTACT Asok Ray  axr2@psu.edu  Department of Mechanical Engineering, The Pennsylvania State University, PA 16802, USA

The article has republished with minor changes. These do not impact the academic content of the article.

© 2020 Taylor & Francis Group, LLC

off. Along this line, Candel (2002) showed that TAI is further aided by the low damping typically seen in many combustors (e.g., in gas-turbine engines).

Much research effort has been expended to understand the physics of TAI phenomena (Lieuwen 2005) with the objective of designing combustion systems (e.g., Chattopadhyay et al. (2017)) that would be less prone to TAI. Along this line, significant research on combustor system models has identified the pertinent mechanisms by which combustion instability occurs (Lee and Santavicca 2003; O'Connor, Acharya, Lieuwen 2015; Richecoeur et al. 2008; Sen et al. 2016), for a range of combustor geometries, flow conditions, and other parameters related to TAI phenomena. Despite better understanding of instability mechanisms, many combustion systems use passive techniques to suppress instability, which include resonators (Noiray and Schuermans 2012) and fuel staging (Samarasinghe et al. 2017). A viable alternative is to design active controllers that can enhance mitigation of TAI. Much of the combustion control research has been conducted on a device known as the Rijke tube (Rijke 1859), which is a laboratory-scale apparatus that is capable of emulating pertinent acoustic properties of real-life combustors. In this regard, several researchers have reported their research work on open-loop and closed-loop control of TAI in a Rijke tube (e.g., Kabiraj and Sujith (2012); Zhao (2012); Zhao and Chow (2013)).

A prerequisite to TAI control is the capability to detect the very occurrence of instability. Several methods have been reported in literature for detecting TAI from time series data generated from combustion systems. More recently, data-driven methods have emerged for detecting and predicting the onset of TAI. It is noted that data-driven techniques do not require the knowledge of the physics of the system, and therefore are more versatile and need relatively less human intervention in both training and testing phases of the algorithm as well as for the application of the algorithm in a real-time setting. For example, Nair et al. (2013) have shown that, under a TAI operation, the pressure signal tends to become sinusoidal in nature, as compared to a more disordered and chaotic signal in the stable regime. Murugesan and Sujith (2015) have applied this knowledge to distinguish between stable and unstable operations in a complex network (Lacasa et al. 2008) setting. A similar approach has also been shown by Mondal et al. (2017) who introduced a more robust Fast Fourier Transform (FFT)-based method; this method relies on the fact that, during TAI, as the signal is sinusoidal, a major portion of the system's energy goes into a single frequency (and its harmonics), which manifests as a high-magnitude spike in the FFT. This method is further expanded upon by Bhattacharya et al. (2020a) where the method was shown to be effective at discriminating between lean blow-out, thermoacoustic instability, and stable operation by using a very simple scalar metric based approach. Recurrence analysis (Sen et al. 2018) has also been used as a means to distinguish stable signals from unstable ones. However, methods like recurrence analysis and complex networks are too slow to be used in an online detection or analysis setting; therefore, other faster and more accurate data-driven techniques have been explored. For example, high-speed flame image analysis along with a neural network have been used by several investigators (e.g., Sarkar et al. (2015); Hauser et al. (2016)); however, such a technique requires the mounting of cameras in the combustor, which is plausible for laboratory scale combustors, but may not be so for real-life applications (e.g., gas turbine engines).

Recently a fast and efficient technique has emerged, namely, symbolic time series analysis (STSA). This method is built upon a finite-memory Markov model, called the D -Markov machine (Mukherjee and Ray 2014; Ray 2004), which is constructed in the setting of

probabilistic finite state automata (PFSA). The concept of *D*-Markov machines has been used by several researchers (e.g., (Sarkar et al. 2016; Unni, Mukhopadhyay, Sujith 2015)) to distinguish between stable and unstable operations; since the execution time of this method is fast, it is a good choice for online operations. More recently, (Ghalyan et al. (2019) and Mondal et al. (2019) have reported usage of hidden Markov models to distinguish combustion regimes using short-length data of pressure time series from a Rijke tube (Bhattacharya et al. 2020b; Matveev 2003; Rijke 1859). Another popular method for time-series classification is neural networks (NN) in their various configurations, such as deep neural networks (DNN) (Wang, Yan, Oates 2017), recurrent neural networks (RNN) (Husken and Stagge 2003), and long short-term memory (LSTM) networks (Hochreiter and Schmidhuber 1997). However, in this paper, the authors have chosen not to broach the NN-based methods for the following reasons: (i) many of these NN techniques may not be suitable for online detection in dynamical systems that are often restricted to use short-length time-series, (ii) NN-based methods need significantly larger amount of data to accurately train, (iii) although shallow nets are decently fast to train and test, RNNs or LSTMs have extremely high training times (much longer than either PFSA or HMMs), and (iv) a major contribution of this paper is identification of “optimal” parameters of the underlying algorithms, which for NNs, are very different and elaborate and have a significantly different design approach.

Two major issues arise in tackling the problem of TAI detection in transient data. The first issue is accurate detection of changes from stable to unstable regimes and vice versa in the offline setting. In this detection mode, the aim is to reliably classify the transient time-series to match the experimental ground truth as much as possible. The second issue is early detection of the onset of TAI for the purpose of triggering preventive control actions in an online closed-loop (e.g., real-time monitoring and active control) setting. In this mode, the aim is to reliably capture precursors to changes from the stable to the unstable regime as early as possible, even before the ground truth shows such changes. Furthermore, a good algorithm is one that can predict an impending change into instability by using a short length of data from a sensor (i.e., to provide high-accuracy classification/detection using as few data points as possible). This will cause the algorithm to be faster (because it has to process less data to reach a decision), as well as help make an earlier decision (because fewer data points mean reduced delays in the data recording time). Both of the above characteristics are highly desirable, especially in the TAI context as the growth from stable to full-blown instability happens within a very short time interval (e.g., in the order of ~ 1 second).

This paper has investigated three data-driven detection algorithms, namely, a fast Fourier transform (FFT) based method, symbolic time series analysis (STSA), and hidden Markov modeling (HMM). Subsequently, these three methods have been applied on transient data of combustor pressure time series, generated from an experimental apparatus. The apparatus is a four-around-one multi-nozzle can-combustor that mimics the complexities seen in can-annular power-generation gas turbines by capturing flame and flow interaction phenomena. These methods are compared for their efficacy of early detection of TAI, along with a detailed insight into the effect of the model hyperparameters and the required data length on the combustor performance. The involved costs of CPU execution time are also investigated. Thus, depending on the intended purpose and the criteria of demanded performance, the algorithm must be appropriately selected by

the user. In essence, this paper acts as a guideline as to the choice of the design methodology and selection of the most suitable algorithm for the purpose at hand.

Contributions: Summarized below are major contributions of the paper toward comparison of three data-driven anomaly detection methods, namely, fast Fourier transform (FFT), symbolic time series analysis (STSA), and hidden Markov modeling (HMM):

- (a) Comparison of the methods for accuracy in online detection of anomalies using short-length time-series data, including the ability to predict onset of instability and the computational time needed for online execution.
- (b) Identification of key algorithm parameters, including window length and down-sampling rate of time series, for enhancement of both accuracy and execution time in both training and testing phases.
- (c) Experimental validation of the results of analysis on an ensemble of time series data collected from a laboratory-scale multi-nozzle combustion apparatus at different operating conditions.

Organization of the paper: The paper is organized in five sections, including the present one. Section II describes experimental apparatus and the procedure for generating the data sets to validate the underlying algorithms. Section III briefly describes the mathematical principles of the three data-driven methods: FFT, STSA, and HMM, for anomaly detection. Section IV makes comparative evaluation of the above three methods by validation with the same experimental data. Section V summarizes and concludes the paper with recommendations for future research.

Description of experimentation

An ensemble of experimental data for this analysis was obtained over a three-year period on a laboratory-scale apparatus, which is built upon a four-around-one multi-nozzle can combustor, for a range of operating conditions. The combustor, shown in [Figure 1](#), consists of five industrial fuel-injector nozzles each with a swirl number of 0.7. The flames attach to the centerbodies of the nozzles, which are recessed from the dump plane. The optically-accessible quartz combustor has a diameter of 260 mm and a length of 300 mm, and is open to the atmosphere at the exit, which enforces a pressure release boundary condition at the end of the combustor and ensures that combustor operation takes place at atmospheric pressure at all times. More details on the experimental setup are discussed by Samarasinghe et al. (2017).

A fully premixed mixture of preheated air at 473 K and natural gas are fed into the combustor system and split evenly among the five nozzles, resulting in a nozzle exit velocity of ~ 26 m/s. Of all 487 test cases available for analysis, most of them were transient tests where the amount of fuel was varied in order to change the stability of the combustor, while the others were either completely stable or completely unstable. At this preheat temperature and flow rate, the combustor experiences thermoacoustic instability for equivalence ratios at and above $\phi = 0.7$ whereas the system is stable for lower equivalence ratios. It was found that thermoacoustic instability could be suppressed through the use of fuel staging (Samarasinghe et al. 2017), where the fuel was unevenly distributed among nozzles while

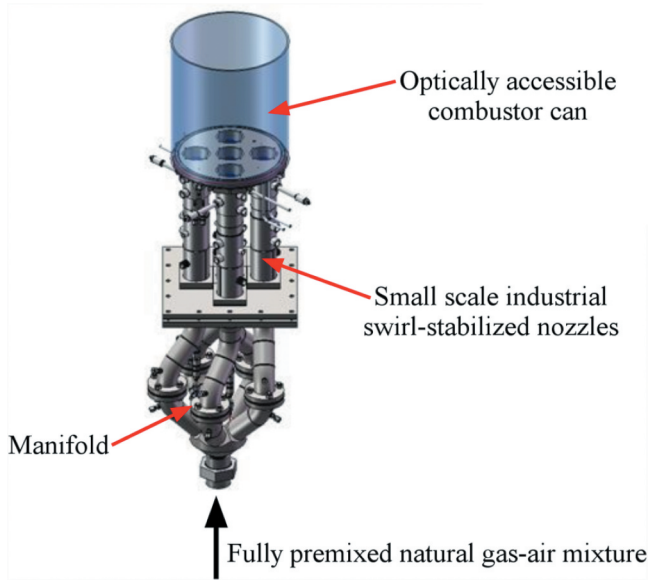


Figure 1. Multi-nozzle combustor apparatus

the global equivalence ratio was maintained. For example, by increasing the equivalence ratio of the center nozzle to $\phi = 0.85$ and decreasing the equivalence ratio of the outer nozzles to $\phi = 0.68$, the instability could be completely and repeatably suppressed (Culler et al. 2018b).

The staging fuel was added in a transient manner through a partially-premixed fuel circuit located in the injector nozzles, as shown in Figure 2. It is noted that the staging fuel can be added to any of the five injectors; previous work has shown that each injector has a slightly different suppression efficacy, which has been quantified using what is termed as

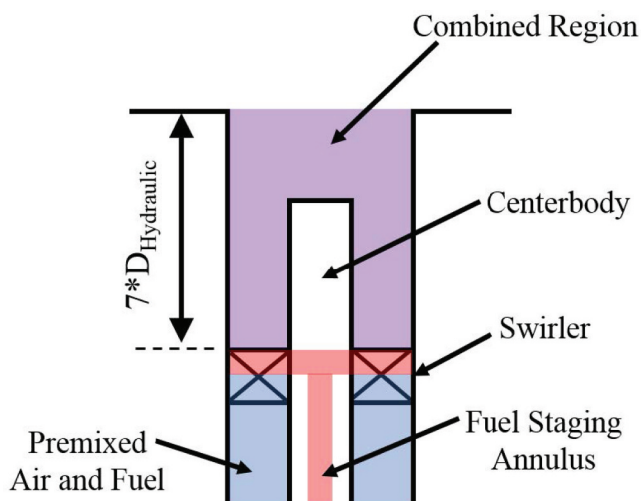


Figure 2. Schematic diagram of fuel-injector nozzle with staging-fuel flow paths.

the “bifurcation equivalence ratio,” or the equivalence ratio at which the instability is suppressed (Culler et al. 2018a). The staging fuel is drawn from the same fuel manifold as the fully-premixed fuel and is added to the staged nozzle so that the fuel flow rate to the fully premixed mixture decreases. Staged fuel flow rate is controlled by using a proportional control valve located just upstream of the staging-fuel injection point. The objective is to control the amount of fuel added or subtracted and the timescale over which that fuel flow rate is changed. This “transient timescale” determines the timescale of the instability onset or decay process, depending on whether the staging fuel is subtracted or added (Culler et al. 2018b). Cross-spectral analysis of the pressure traces in the combustor and in the fuel system show that oscillations in the fuel system as a result of the opening and closing of the transient valve do not impact combustion instability behavior (Howie et al. 2020).

Several diagnostic devices have been used to monitor the behavior of the combustor. For example, K-type thermocouples are located at the tip of the centerbodies and on the dump plane; these measurements of metal temperatures are used to monitor the thermal state of the combustion system and to ensure repeatability of transient tests. High-speed piezo-electric pressure transducers (PCB Model 112A22) are coupled to a signal conditioner (PCB model 482A16) to measure pressure fluctuations in the nozzles and in the combustor. The combustor pressure fluctuations, analyzed in this work, have been generated from the transducer located in the dump plane. The transducer is not directly in contact with the combustion gas but located in a cooling jacket that has a stand-off length of ~ 0.125 inches (~ 3 mm) from the combustion gas. The resonant frequency of that cavity volume is almost an order of magnitude higher than the upper response frequency of the transducer and so does not impact the measurements. All data have been acquired at the sampling frequency of 16,384 Hz for a duration of 8 or 16 s, depending on the duration of the transient test. No analog or anti-aliasing filters are deemed necessary while obtaining this data as the sampling frequency was significantly higher (on the order of ~ 20 times) than any significant frequency content in the signal (significant frequencies of interest were in the range of 400–700 Hz).

Experimental results for algorithm validation

The ensemble of data described above, has been used in this analysis. Table 1 shows the range of operating conditions at which the data were collected. In these tests, the inlet flow velocity (~ 26 m/s), inlet temperature (473 K), and global equivalence ratio ($\phi = 0.7$) were held constant, while four other parameters were varied. The staging nozzle (the nozzle through which staging fuel is injected) was varied, in order to understand the relative efficacy of different flames; results of these tests and description of the nozzle nomenclature are discussed in more detail in Culler et al. (2018a). The other variations concerned the fuel-staging transients by changing the transient timescale (τ_{tran}), the transient amplitude (ϕ_{stage}), and the transient direction. The transient timescale is the time scale over which the proportional control valve is actuated from the initial to final flow rates, and included times of 1 ms, 16 ms, 4 s, and 10 s. Infrared absorption measurements were used to quantify the actual fuel delivery time. For example, for an actual valve opening time of 1 ms, the fuel delivery time to the flame was between 10 ms (for $\phi = 0.85$) to 30 ms (for $\phi = 0.75$); further details of this measurement are provided in the supplementary material and Table 2 of

Table 1. Test matrix.

τ_{tran} (ms)	ϕ_{stage}	Staging nozzle
1	0.8	Center
1	0.8	Outer 2
1	0.8	Outer 4
1	0.85	Center
1	0.85	Nozzle 1
1	0.85	Nozzle 2
1	0.85	Nozzle 4
1	0.9	Center
1	0.9	Nozzle 1
1	0.9	Nozzle 2
1	0.9	Nozzle 3
1	0.9	Nozzle 4
16	0.8	Center
16	0.85	Center
4000	0.8	Center
4000	0.85	Center
10000	0.8	Center
10000	0.85	Center
10000	0.9	Center

Culler et al. (2018b). The transient amplitude is the equivalence ratio of the staging nozzle, and included values of $\phi_{stage} = 0.8$, $\phi_{stage} = 0.85$, and $\phi_{stage} = 0.9$. The transient direction refers to whether the system transitioned from stable to unstable (staged to unstaged operation) or unstable to stable (unstaged to staged operation) thermoacoustic conditions. Cases where the system moves from stable to unstable operation are said to be the “onset” cases and the change from unstable to stable operation are “decay” cases. Results from variations in the transient characteristics are reported in Culler et al. (2018b).

Figure 3 shows representative pressure time series for both stable and unstable operations and also reports their frequency characteristics in FFT graphs. The top left plot shows the low-amplitude pressure trace seen during a stable operation along with the corresponding zoomed section of the signal (middle left), while the FFT plot at the bottom left indicates the frequency information of the same. Similarly, the top right plot shows the high-amplitude pressure time series for an unstable operation along with the corresponding zoomed section of the signal (middle right), while the bottom right shows the corresponding FFT. In these reported FFT plots, the frequency cover used is 10–3,000 Hz as described further in sub-section III-A. It is seen that the unstable operation is characterized by a very strong peak at a particular frequency (~ 531 Hz in the reported time-series), while the stable operation has a more broadband spectrum with a mild peak (about 10 orders of magnitude lower than unstable) near the resonant frequency (~ 497 Hz in the reported time-series). This nature of the FFTs is easily explained by observing the zoomed in time series plots in the middle row of Figure 3, which shows ordered sinusoidal signals in an unstable operation and noisier signals during a stable operation.

Examples of the transient time-series pressure data are shown in the four plates of Figure 4. For all of the four cases, the center nozzle is staged with an equivalence ratio of $\phi_{stage} = 0.85$. All pressure time series have been high-pass filtered at 10 Hz to remove drift from DC-coupling in the pressure transducer electrical system. In the $\tau_{stage} = 1$ ms cases, the staging valve was actuated at 4 s, resulting in very different responses between the onset and decay

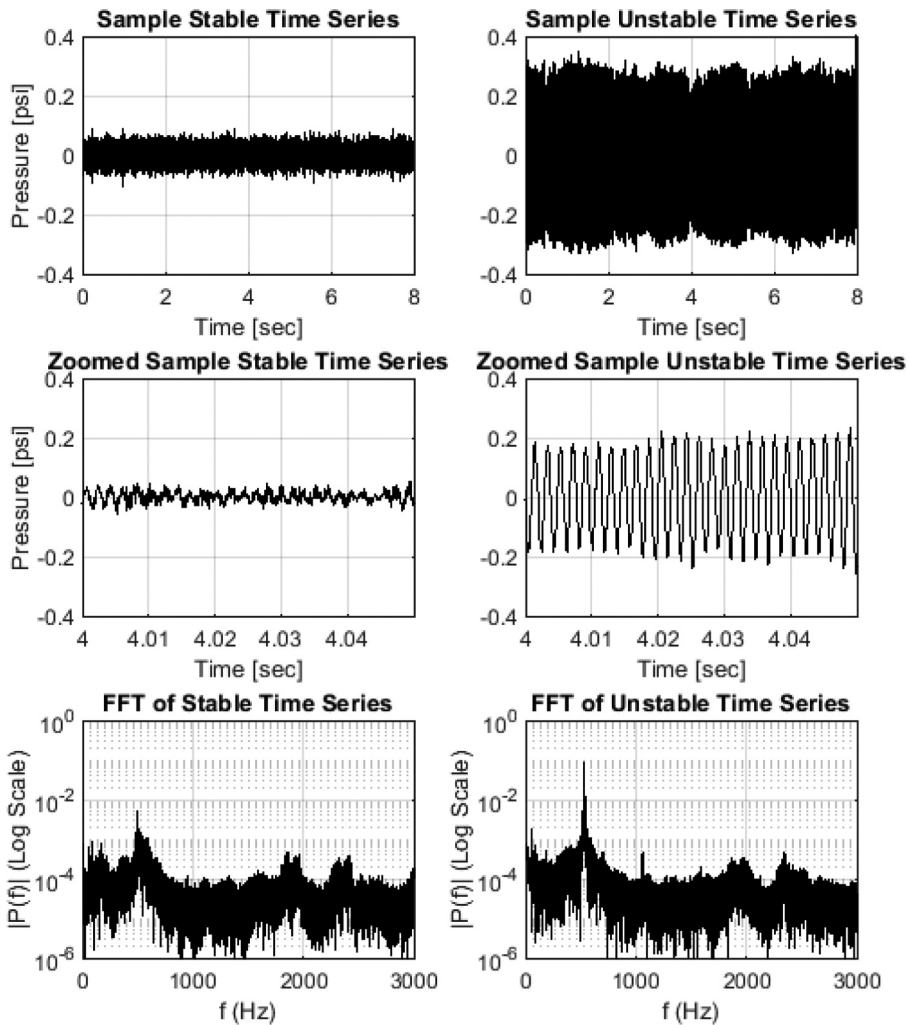


Figure 3. Samples of pressure signals for stable (top left) and unstable (top right) operation along with corresponding zoomed sections of the signals (middle left and right) and respective FFT plots (bottom left and right).

cases. In the decay case, the system responded quickly, resulting in instability suppression in less than 0.25 s after valve actuation. In the onset case, however, instability onset is more gradual, taking approximately 1 s to reach the limit-cycle amplitude. Although these data show two examples out of hundreds, the trends are similar for almost all of the 1 ms cases obtained. The valve in the $\tau_{stage} = 4$ s cases is actuated at 4 s as well, but actuation lasts until 8 s. Here, the onset of instability occurs faster after actuation than the decay, but the instability is highly intermittent during the valve actuation time. Again, these results are representative of those found across test cases.

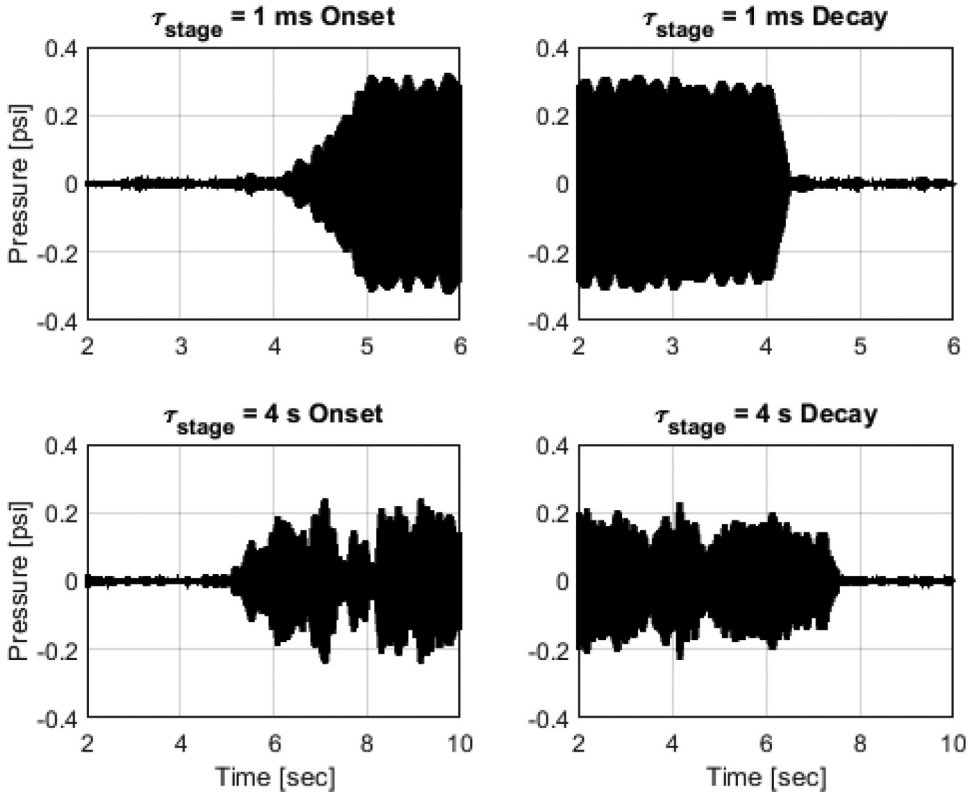


Figure 4. Pressure time series for four transient cases with different values of τ_{stage} .

Background mathematical information

This section addresses the underlying mathematical theory and intuition behind the three data-driven detection methods, namely fast Fourier transform (FFT), symbolic time series analysis (STSA), and hidden Markov modeling (HMM). The associated decision metric (i.e., the threshold method that is used to make the binary classification) and algorithm are also presented for each of these three methods. The mathematical theories for these three methods are extensively reported in open literature; however, for completeness and ease of readability, these concepts are very succinctly presented in this section.

Fast Fourier transform (FFT)

The FFT algorithm uses the knowledge of the physical process dynamics to distinguish between the stable and unstable regimes (Bhattacharya et al. 2020a; Mondal et al. 2017). During the stable operation of the combustor, the signal energy is generally broadband and so no strong frequency components can be clearly identified. However, when the system undergoes thermoacoustic instability (TAI), the frequency corresponding to the unstable mode becomes much stronger and contains most of the energy of the signal. In a plot of the FFT amplitude, this manifests as a large-magnitude spike at the resonant frequency during the TAI.

To estimate how much of the total signal energy is contained in an unstable mode, a ratio of the energy of the unstable mode to that contained in the entire signal is defined. In this work, the total energy is defined as the total area under the square of FFT amplitude over a frequency range of 10–3,000 Hz; and the energy of the unstable mode is similarly defined as the area under the frequency range having the highest magnitude, using a bandwidth of 4 Hz (± 2 Hz). The choice of the frequency range of 10–3,000 Hz can be justified as follows. A lower bound of 10 Hz is selected as it is desirable to remove low frequency noise associated with DC effects. Similarly, the higher limit is chosen to be 3,000 Hz as that is higher (with a safe margin) than the frequencies of interest (400–700 Hz) as described in sub-section II-A and demonstrated in Figure 3. Furthermore, a sampling rate of 16,384 Hz sets the upper limit of the FFT (by Nyquist criterion) to 8,192 Hz. If the ratio of the above two areas is higher than a specified threshold (η_{FFT}), it implies that a significant part of the energy is contained in the unstable mode (i.e., it is unstable). During the testing phase, there is no specific training because the FFT method is largely physics-based and not data-driven. A window of specified length is observed, and the FFT ratio as defined above is computed. The final decision is made as:

$$\frac{\text{Area under highest magnitude frequency}}{\text{Area under total FFT}} \underset{\text{Stable Regime}}{\overset{\text{Unstable Regime}}{\gtrless}} \eta_{FFT}$$

Symbolic time series analysis (STSA)

The symbolic time series Analysis (STSA) method makes use of the concept of probabilistic finite state automata (PFSA) (Mukherjee and Ray 2014; Rajagopalan and Ray 2006; Ray 2004; Subbu and Ray 2008), to convert the given pressure time series into quantized sequence of alphabets and then generate features that are used for classification. The following sub-sub-sections III-B1 and III-B2 provide the mathematical details of STSA and the computation of D-Markov machines.

Theory of probabilistic finite state automata (PFSA)

The time series of a measured signal is converted to a symbol string. In this process, the signal space is partitioned into a finite number of cells, where the number of cells is identically equal to the user-chosen cardinality $|\Sigma|$ of the (symbol) alphabet Σ ; and a symbol from the alphabet Σ is assigned to each (signal) value corresponding to the cell to which it belongs (Rajagopalan and Ray 2006; Subbu and Ray 2008); details are reported in Mukherjee and Ray (2014). Thus, a symbol is associated with a data point at a given instant of time when the value of that data point is located in the particular cell corresponding to that symbol. The following definitions, which are available in standard literature (e.g., Ray (2004); Mukherjee and Ray (2014)), are recalled for completeness of the paper:

Definition 1: A finite state automaton (FSA) G , having a deterministic algebraic structure, is a triplet (Σ, Q, δ) where:

- Σ is a (nonempty) finite alphabet, i.e., its cardinality $|\Sigma|$ is a positive integer.
- Q is a (nonempty) finite set of states, i.e., its cardinality $|Q|$ is a positive integer.
- $\delta : Q \times \Sigma \rightarrow Q$ is the (deterministic) state transition map.

Definition 2: A symbol block, also called a word, is a finite-length string of symbols belonging to the alphabet Σ , where the length of a word $w \triangleq s_1 s_2 \cdots s_\ell$ with every $s_i \in \Sigma$ is denoted as $|w| = \ell$, and the length of the empty word ε is $|\varepsilon| = 0$. The parameters of FSA are extended as:

- The set of all words, constructed from symbols in Σ and including the empty word ε , is denoted as Σ^* .
- The set of all words, whose suffix (respectively, prefix) is the word w , is denoted as $\Sigma^* w$ (respectively, $w \Sigma^*$).
- The set of all words of a given (finite) length ℓ , where ℓ is a positive integer, is denoted as Σ^ℓ .

Remark 3: A symbol string (or word) is generated from a (finite-length) time series by symbolization.

Definition 4: A probabilistic finite state automaton (PFSA) K is a pair (G, π) , where:

- The deterministic FSA G is called the *underlying FSA* of the PFSA K .
- The probability map $\pi : \mathcal{Q} \times \Sigma \rightarrow [0, 1]$ is called the morph function (also known as symbol generation probability function) that satisfies the condition: $\sum_{\sigma \in \Sigma} \pi(q, \sigma) = 1$ for all $q \in \mathcal{Q}$.

Equivalently, a PFSA is a quadruple $\mathcal{K} = (\Sigma, \mathcal{Q}, \delta, \pi)$. The state transition probability mass function $\kappa : \mathcal{Q} \times \mathcal{Q} \times [0, 1]$ is constructed by combining δ and π , which can be structured as a $|\mathcal{Q}| \times |\mathcal{Q}|$ state transition probability matrix Π . In that case, the PFSA can also be described as the triple $K = (\Sigma, \mathcal{Q}, \Pi)$.

Theory of *D*-Markov machines

The PFSA model of a *D*-Markov machine generates symbol strings $\{s_1 s_2 \cdots s_\ell : \ell \in \mathbb{N}^+ \triangleq \{1, 2, 3, \dots\} \text{ and } s_j \in \Sigma\}$ on the underlying Markov process. When constructing a *D*-Markov machine, it is assumed that the generation of the next symbol depends only on the most recent *finite* history of at most *D* consecutive symbols, i.e., a symbol block of length not exceeding length *D*. A *D*-Markov machine (Mukherjee and Ray 2014) is defined as follows.

Definition 5: A *D*-Markov machine is a PFSA in the sense of Definition 4 and it generates symbols that solely depend on the most recent history of at most *D* consecutive symbols, where the positive integer *D* is called the depth of the machine. Equivalently, a *D*-Markov machine is a statistically stationary stochastic process $S = \cdots s_{-1} s_0 s_1 \cdots$, where the probability of occurrence of a new symbol depends only on the last consecutive (at most) *D* symbols, i.e.,

$$P[s_n | \cdots s_{n-D} \cdots s_{n-1}] = P[s_n | s_{n-D} \cdots s_{n-1}]$$

Consequently, for $w \in \Sigma^D$ (see Definition 2), the equivalence class $\Sigma^* w$ of all (finite-length) words, whose suffix is w , is qualified to be a *D*-Markov state that is denoted as w .

In this paper, the maximum entropy partitioning (MEP) scheme (Rajagopalan and Ray 2006; Subbu and Ray 2008) has been used to quantize the time series. A PFSA (or D -Markov machine) (Mukherjee and Ray 2014) is constructed from the symbol string that is generated from the partitioned time series and the state transition probability matrix Π is extracted as the feature (Π indicates the probability of observing a particular state in the next instant given the present state). In this paper, the depth of the D -Markov machine is taken to be unity (i.e., $D = 1$), which makes Σ and \mathcal{Q} identical.

During the training phase, the algorithm trains on the time-series data corresponding to fully stable or fully unstable regimes (i.e., no transience). The algorithm trains on a certain portion (80% has been chosen in this paper) of the available data. Two trained feature matrices, namely, one for the stable regime and the other for the unstable regime (called Π_{Stable} and $\Pi_{Unstable}$, respectively) are constructed and used as the trained features for the pattern classification that follows.

During the testing phase, the algorithm takes a data window from the unknown transient time series. The specified window length and downsampling rates now determine the number of data points to be selected. This windowing is done at equal intervals of 0.02 s (i.e., 50 Hz rate), which implies that, for window lengths of 10 ms and 20 ms, there is no overlap; however, for higher window lengths, there is overlap. Each window is studied individually, and the state transition probability matrix (Π_{Test}) corresponding to the window is generated using the selected data points. Then, the following ratio is computed, which depending on a threshold (η_{PFSA}), decides whether the given window corresponds to data belonging to the stable regime or the unstable regime:

$$\log\left(\frac{\|\Pi_{Test} - \Pi_{Stable}\|_{\infty}}{\|\Pi_{Test} - \Pi_{Unstable}\|_{\infty}}\right) \underset{\text{Stable Regime}}{\overset{\text{Unstable Regime}}{\gtrless}} \eta_{PFSA}$$

where $\|\bullet\|_{\infty}$ is the induced infinity norm of the matrix \bullet .

Figure 5 shows two samples of state transition matrices, one each for stable (Π_{Stable}) and unstable ($\Pi_{Unstable}$) regimes, at downsampling $DS = 2$, window length = 100 ms with alphabet size $|\Sigma| = 4$, and PFSA depth $D = 1$. The difference between the two is leveraged to discriminate between these two regimes. In this case, since $D = 1$, the number of states ($|\mathcal{Q}|$) is equal to the alphabet size ($|\Sigma|$), which is 4 in this case, as seen in Figure 5.

Π_{Stable}					$\Pi_{Unstable}$				
	q_1	q_2	q_3	q_4		q_1	q_2	q_3	q_4
q_1	0.4238	0.2750	0.1883	0.1129	q_1	0.7993	0.1922	0.0049	0.0036
q_2	0.2755	0.2813	0.2523	0.1909	q_2	0.3931	0.2068	0.3906	0.0095
q_3	0.1871	0.2530	0.2814	0.2785	q_3	0.0096	0.3870	0.2140	0.3894
q_4	0.1120	0.1911	0.2779	0.4191	q_4	0.0036	0.0047	0.1906	0.8012

Figure 5. Sample state transition matrices for stable and unstable regimes where $|\mathcal{Q}| = |\Sigma| = 4$, $D = 1$, $DS = 2$, window length = 100 ms.

Hidden Markov modeling (HMM)

Since the details of hidden Markov models (HMM) are extensively reported in technical literature (e.g., Rabiner (1989); Hajek (2015)), only the key concepts are introduced here for completeness of the paper. For further mathematical details, the readers are requested to follow the references.

HMMs have been used to represent long-range dependencies between observations, where the underlying models are assumed to be probabilistic functions of the hidden states (Murphy 2012). Considering a discrete-time representation of a data string $\mathbf{Y} = \{y_1, y_2, \dots, y_T\}$ of T continuous (real-valued) observations, and assuming a first-order Markov property (Bishop 2007) over the observations, the joint probability density function of \mathbf{Y} is obtained as:

$$p(\mathbf{Y}) = p(y_1) \prod_{t=1}^{T-1} p(y_{t+1}|y_t)$$

Although the above conditions have been widely used in practice, it may not always be valid because of long-range correlations among the observations, which are seen in certain systems (Bishop 2007; Murphy 2012). This has led to HMMs being used for applications primarily in speech recognition (Najkar, Razzazi, Sameti 2010), time series classification (Oates, Firoiu, Cohen 2000) and even in image classification (Ali and Ghani 2014). More recently it has shown good accuracy in the classification of chaotic data (Bhattacharya and Ray 2020a). In all of these problems, the HMM method has shown high classification accuracy.

Essentially, HMMs belong to a class of doubly-embedded stochastic processes, with a latent stochastic process of *hidden* state evolution. Although not directly observed, this evolution can be inferred by observing another stochastic process that produces the sequence of observations (Rabiner and Juang 1993), which capture the long-range dependencies among observations and enables the usage of HMMs as black-box density models on observation sequences. The major difference between the HMM and the standard Markov model is that the HMM does not directly assume the Markov property (i.e., conditional dependence on the states and being independent of each other) for the observations themselves. Instead, the hidden state sequence $\mathbf{Z} = \{z_1, z_2, \dots, z_T\}$ is assumed to follow Markovian dynamics. That is, given the current state z_t , the future state z_{t+1} is independent of all the states prior to time instant t .

To formalize the mathematical structure, let a string of observations $\{y_1, \dots, y_T\}$ be assumed to be generated by a hidden state sequence $\{z_1, \dots, z_T\}$. A HMM is then constructed as a triplet $\lambda = \{A, B, \pi\}$ (Rabiner 1989), where:

- a) $A \triangleq [a_{ij}]$ is the $|N| \times |N|$ state-transition probability matrix, where $|N|$ is the finite number of hidden states belonging to the set N of hidden states:

$$a_{ij} = p(z_{t+1} = q_j | z_t = q_i) : q_i, q_j \in N$$

where $\sum a_{ij} = 1$ i and $a_{ij} \geq 0$ i, j . In this paper, for comparison of HMM with STSA, the cardinality of the hidden state set is made equal to that of the PFSA state set (see subsection III-B).

- b) $B \triangleq [b_j(y_t)]$ is the probability density of the observation given the state:

$$b_j(y_t) = p(y_t|z_t = q_j)$$

c) $\pi \triangleq [\pi_i]$ is the probability distribution of the initial state z_1 : $\pi_i = p(z_1 = q_i)$, where π is a $1 \times |N|$ vector with $\sum_{i=1}^{|N|} \pi_i = 1$ and $\pi_i \geq 0 \quad i$.

Following a model λ , the corresponding joint probability distribution of states and observations has the form:

$$p(\mathbf{Y}, \mathbf{Z}) = p(z_{1:T})p(y_{1:T}|z_{1:T}) = [p(z_1) \prod_{t=1}^{T-1} p(z_{t+1}|z_t)] [\prod_{t=1}^T p(y_t|z_t)]$$

In the training phase of a HMM (Rabiner 1989), time series data from each regime (i.e., classes $k = 1, \dots, \mathcal{K}$), are used to learn the HMM. Subsequently, in the testing phase, the learned HMMs are compared to the HMM constructed from the data belonging to an unknown regime. A succinct description of HMM is outlined below.

During training, a commonly used expectation maximization (EM) procedure, called the *Baum-Welch algorithm* (Hajek 2015; Rabiner 1989), is used to learn each of the HMM models belonging to \mathcal{K} classes. For each class (regime), an ensemble of time-series windows is obtained from the available time-series data. Then, the *Baum-Welch algorithm* Rabiner (1989); Hajek (2015) is applied to train the HMM λ_k which is a triplet $\lambda^k = \{A^k, B^k, \pi^k\}$ Rabiner (1989), $k = 1, \dots, \mathcal{K}$, where A^k, B^k , and π^k are defined similarly as A, B , and π for each class k . The procedure is repeated for each of the \mathcal{K} classes.

During the testing phase, data windows of the signal from an unknown regime are provided as inputs to the algorithm. Given this observational sequence and the HMM models for each of the \mathcal{K} classes, $\lambda^k = \{A^k, B^k, \pi^k\}$, the problem is to find the probability of the entire observation sequence being associated with the most likely model. This is expressed as:

$$p(\mathbf{Y}|\lambda^k) = \sum_{\mathbf{Z}} p(\mathbf{Y}|\mathbf{Z}, \lambda^k)p(\mathbf{Z}|\lambda^k) = \sum_{z_1, z_2, \dots, z_T} \pi_{z_1} b_{z_1}(y_1) a_{z_1 z_2} b_{z_2}(y_2) \dots a_{z_{T-1} z_T} b_{z_T}(y_T)$$

which is obtained by using the *Forward Procedure* (Hajek 2015; Rabiner 1989) to compute the log likelihood (L^k) of the given window of unknown time series data belonging to each of the \mathcal{K} classes. The final decision, as to which class the unknown data belongs, is made by selecting the class with the largest log likelihood as follows:

$$\text{Selected Class} = * \operatorname{argmax}_{k \in \{1, 2, \dots, \mathcal{K}\}} L^k$$

In this paper, a continuous HMM formulation has been used, where the emission is assumed to follow a Gaussian mixture model with M Gaussian components and N hidden states. The algorithms and theory for all the above HMM procedures are available in literature (e.g., Rabiner (1989); Hajek (2015); Murphy (2012)).

In the context of this paper, during the training phase, two HMM models are trained using the Baum-Welch algorithm, one for the stable regime (λ^{Stable}) and another for the unstable regime ($\lambda^{Unstable}$). Once again, similar to the STSA method, these models are trained using 80% of the data available from the fully stable and fully unstable regimes. During the testing phase, a similar method to the STSA method is adopted where a data window of a specified window length is taken and downsampled. Once again, the

windowing is done at equal intervals of 0.02 s (50 Hz rate), i.e., the raw pressure data used to make the decision is the same for both the STSA and HMM methods. Then using the forward procedure, the log likelihoods of the observed data to be from the stable and unstable regimes are computed (called L^{Stable} and $L^{Unstable}$ respectively). The final decision is once again made using a threshold (η_{HMM}) as:

$$L^{Unstable} - L^{Stable} \underset{StableRegime}{\overset{UnstableRegime}{>}} \eta_{HMM}$$

Computation of the threshold parameters

In order to investigate the performance of the binary classification problem (i.e., accurate detection of the stable or unstable regimes), it is best to study it in the form of receiver operating characteristics (ROC) curves (Bishop 2007). The ROC curves are created by plotting the true positive rate (i.e., the rate at which the stable condition is correctly classified to be stable) against the false positive rate (rate at which the stable condition is falsely classified to be unstable) for a range of threshold values. The area under the curve (AUC) in a ROC curve is an indicator of how accurate the algorithm is, with a value of 1 being perfect classification and lower values implying poorer classification.

In the computation of these ROC curves, the optimal thresholds (η_{FFT} , η_{PFSA} and η_{HMM}) have also been identified as those yielding the least errors, i.e., least number of mis-classified data windows. The optimal thresholds are seen to be functions of the window length and downsampling parameters as well as the alphabet size and number of mixture components for the STSA and HMM methods. Details of the ROCs are given in sub-section IV-A.

Results of experimental data analysis and discussions

This section discusses the comparison of the above three methods, namely, FFT, STSA, and HMM, based on the same set of experimental data obtained from the combustion apparatus, described in Section II. As mentioned earlier, in order to allow for a fair comparison of these three methods, the respective parameters in the algorithms have been kept similar, as explained below.

The root mean square (RMS) value of the pressure signal is defined as the ground truth that has been obtained using a moving window formulation, applied on the the experimental data. The signal is considered stable when the RMS value of the pressure is found to be below 0.07 psi (0.483 kPa), which is 0.5% of the mean combustor pressure (Culler et al. 2018b). As mentioned earlier in Section II, the available experimental time-series data consisted of completely stable signals, completely unstable signals and transient signals, where the combustor went from a stable state to an unstable state or vice-versa. For training both the HMM and STSA methods, data-driven models in the algorithms used 80% of the available completely stable and completely unstable data. For testing each of the three algorithms, only the transient data were used.

For each of the three algorithms, six window lengths of 10 ms, 20 ms, 50 ms, 100 ms, 500 ms, and 1000 ms duration have been investigated. The analysis is performed every 0.02 s (i.e., 50 Hz), yielding non-overlapping windows for data lengths of 10 ms and 20 ms, and

overlapping windows for the rest. However, each window is processed individually with no bias, making the comparison fair. The windowed data from the transient time-series are downsampled to examine the algorithm performance. Downsampling parameters of 1 (i.e., no downsampling), 2, 5, and 10 have been investigated. The FFT method uses the non-downsampled data to assure that signals at all pertinent frequency bands are captured and to avoid the risk of having too few points for generation of a sensible FFT graph; however, downsampling is applied to the data used for STSA and HMM to mitigate self-looping in the finite-state graph models. The final parameter under investigation is the alphabet size $|\Sigma|$ in the STSA method (see sub-section III-B) and equivalently, $|N|$, the number of hidden states, as seen in sub-section III-C; the values of $|N|$ that have been studied are 2, 4, 6, and 8. In all of the analyses, the number, M , of output Gaussian mixtures in the emissions of the HMM model has been kept fixed (e.g., $M = 2$ was seen to be sufficient for initial analysis). Detailed analyses are made in the following sub-sections, where Table 2 tabulates the number of data-points used for analyzing FFT, STSA, and HMM algorithms as a function of the window-length and the downsampling parameter. This table will be referred to subsequently to show how the methods work for ultra-short windows of time series.

Comparison of FFT, STSA, and HMM performance

As mentioned in sub-section III-D, ROC curves are computed to obtain the optimal thresholds. The ROC curves can also be used to study the comparative accuracy of the various methods for different values of hyper-parameters. A total of 96 ROC curves have been generated (e.g., 6 window lengths \times 4 downsampling parameters \times 4 alphabet sizes/number of hidden states). Given below are 8 selected sets of ROC curves in Figures 6 and 7 to accurately describe the effectiveness of the methods without cluttering the paper. These correspond to the window length of 10 ms (Figure 6) with downsampling parameter of 1 (i.e. no downsampling) and that of window length 20 ms (Figure 7) with downsampling parameter of 5.

These ROC curves are indicative of several very interesting observations. It is seen in Figures 6 and 7 that the HMM method consistently yields the highest accuracy (i.e., highest AUC). It is concluded from Figure 6 that the STSA method is the second best, with the AUC nearly matching that of the HMM method when the data is sufficiently partitioned (i.e.,

Table 2. Different window lengths and downsampling parameters (DS) for FFT, STSA, and HMM.

Window length (ms)	FFT window data points	STSA\HMM window data points			
		DS = 1	DS = 2	DS = 5	DS = 10
10	164	164	82	33	17
20	328	328	164	66	33
50	819	819	410	164	82
100	1639	1639	819	328	164
500	8192	8192	4096	1639	820
1000	16384	16384	8192	3277	1639

NOTE: Total accuracy, as reported in the subsequent sub-sections, is the accuracy in classifying only the transient test set, expressed as; the total percentage of correct classification (i.e., stable as stable and unstable as unstable) obtained.

NOTE: All the computations for this paper were conducted using in-house MATLAB codes for the STSA (Available at: <https://github.com/Chandrachur92/PFSA>) and using Murphy's (Murphy 2012) Hidden Markov Model (HMM) Toolbox for MATLAB (Available at: <https://www.cs.ubc.ca/~murphyk/Software/HMM/hmm.html>). The computations are conducted serially on a single core processor of a DELL Precision Tower 7910 Workstation running an Intel® Xeon® E5-2670 CPU.

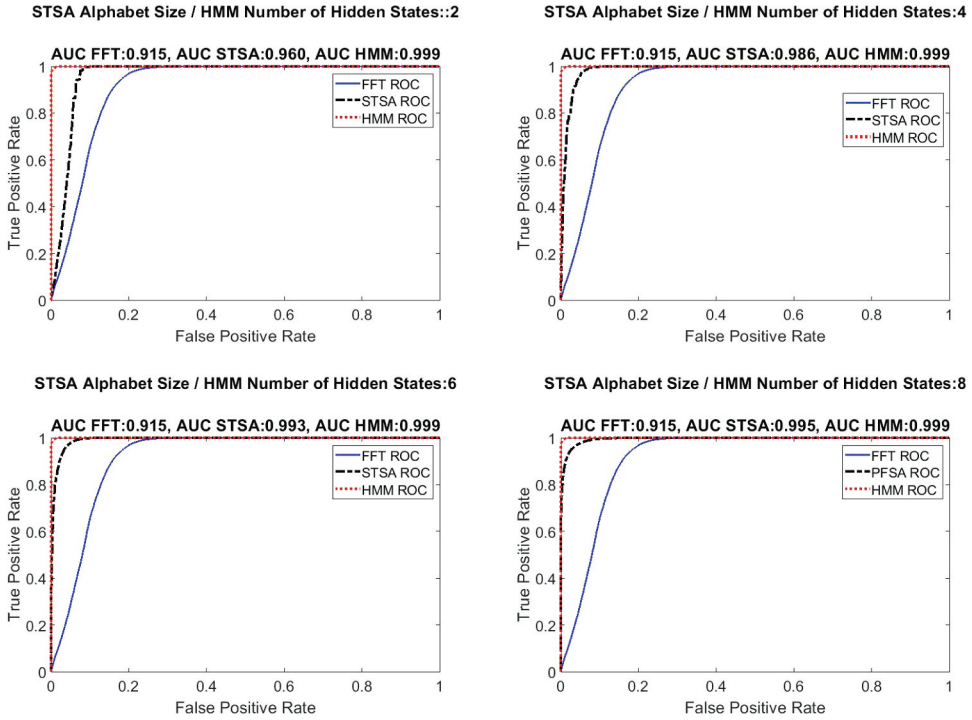


Figure 6. ROC Curves of the three methods with window length of 10 ms and no downsampling, i.e., $DS = 1$ (yielding window sizes of 164 data points for the FFT, STSA, and HMM methods) for different values of PFSA alphabet size and number of HMM hidden states.

more discrimination power). The FFT method consistently has the poorest performance, although it still has a rather high AUC of about 0.91; there are no changes in the FFT responses as the data for the FFT method is not downsampled. Figure 7 shows that downsampling degrades the performance of STSA when the number of symbols is 2 because it cannot sufficiently capture the system dynamics, but the performance recovers when the number of symbols is increased with $|\Sigma| \geq 4$. In the following sub-sections, the effects of each of the model parameters are explained in detail.

Figure 8 demonstrates the working of the threshold-based metrics for FFT, STSA, and HMM classification, where the optimal thresholds are obtained as elaborated in sub-section III-D. Each column indicates the representative result for a particular time series. For each, the top-most panel shows the pressure trace and the true regime based on the pressure RMS ground truth. The second row shows the FFT ratio computed in a windowed fashion superimposed with the computed threshold (see Table 6) and the corresponding classification based on the FFT method. Similarly the third and fourth rows show the computed log ratio, threshold and classification for the STSA and HMM methods. It is noted that the FFT method produces the most noisy and hence the lowest accuracy in classification. Both STSA and HMM produce similar classification that closely matches the ground truth. These findings are further elaborated and discussed in the remainder of the paper.

Model Parameters: $|\Sigma| = N = 4$, $M = 2$, $D = 1$, $DS = 2$, Window-length = 100 ms

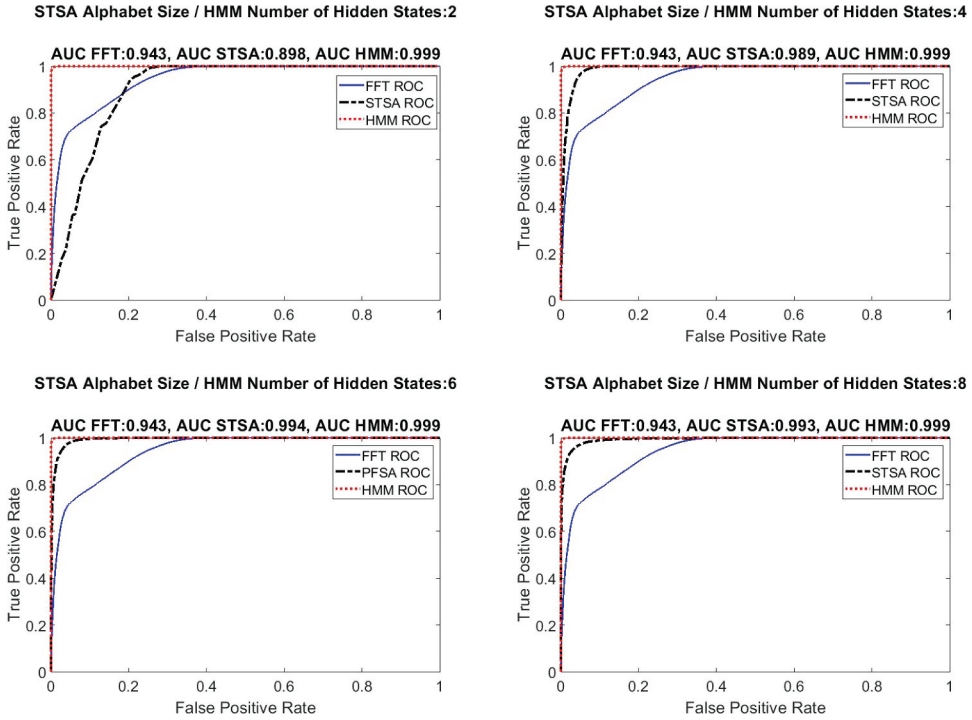


Figure 7. ROC Curves of the three methods with window length of 20 ms and downsampling of 5 for only the STSA/HMM methods, i.e., $DS = 5$ (yielding window sizes of 328 data points for the FFT method while the STSA and HMM methods have window sizes of 66 data points) for different values of PFSA alphabet size and number of HMM hidden states.

Effects of window length on performance of STSA, HMM, and FFT

This section investigates the effects of data window length on accuracy of FFT, STSA, and HMM. The following figures plot the maximum total accuracy of each of the methods under a specified choice of model parameters. Maximum accuracy is expressed as the total percentage of correct classification (i.e., stable as stable and unstable as unstable) obtained at the optimal threshold. The optimal threshold is identified during the generation of the ROCs, mentioned earlier, as the threshold that yields the least errors, which implies the least number of mis-classified data windows. For the purpose of online detection, it is essential to be able to make a fast decision based on a short-length time series; thus, it is an important parameter to be decided by the user.

It follows from [Figure 9](#) that the accuracy of the FFT method is highest when the window length is about 50–100 ms, which yields a sufficient number of data points to accurately generate the FFT, as seen in [Table 2](#); smaller window lengths yield less accurate results. However, even higher window lengths perform poorly, because the FFT is an averaging method and is not able to capture the unstable mode amplitude spike over a large window, especially near change-points.

[Figure 10](#) indicates that STSA yields the best performance for a window length of 100 ms at a downsampling parameter of 2; if the downsampling parameter is increased to 5, the best performing window length is 50 ms. However, in general, it is observed that window length

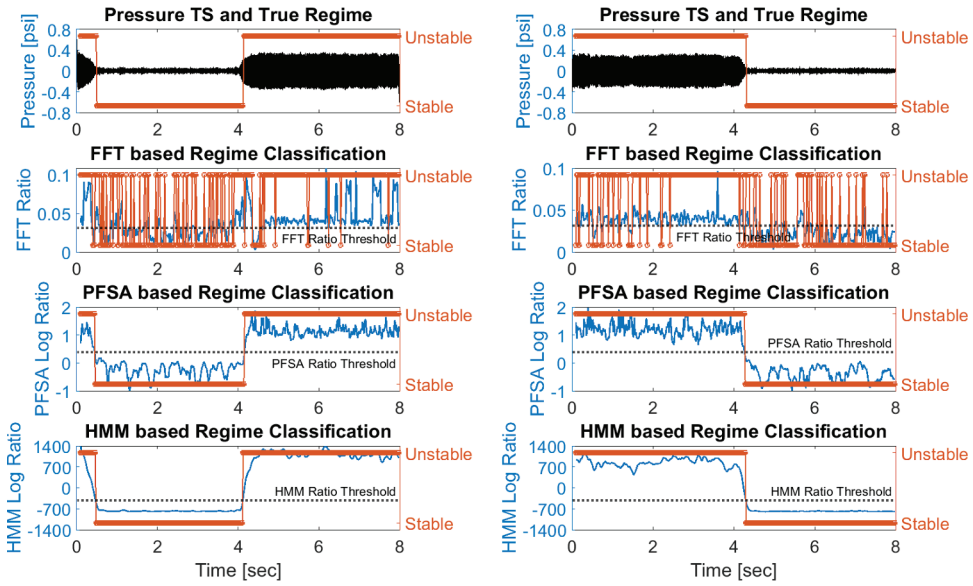


Figure 8. Two sample pressure time series: first row – left and right (true regime superimused); second row – left and right (FFT-based ratio metric for threshold and regime classification); third row – left and right (PFSA-based ratio metric for threshold and regime classification); and fourth row – left and right (HMM-based ratio metric for threshold and regime classification).

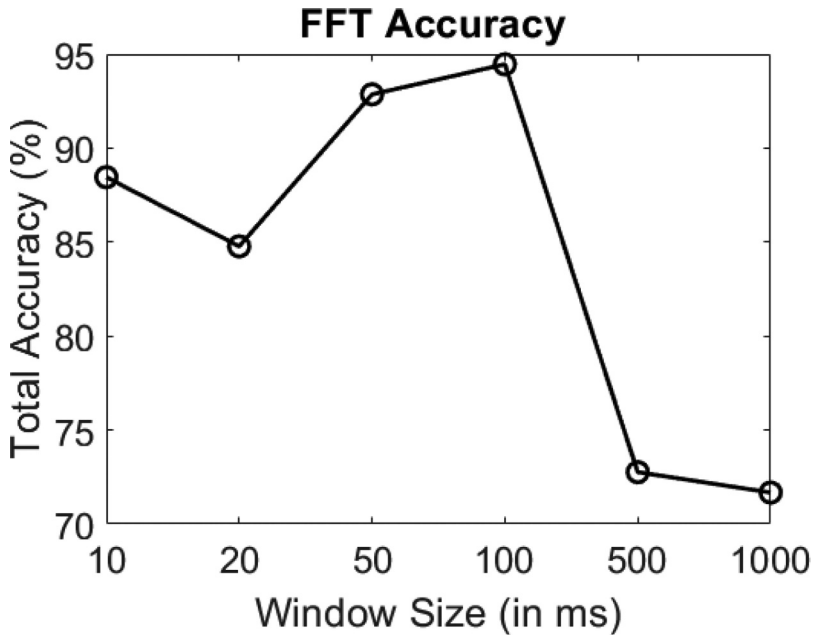


Figure 9. Detection accuracy of the FFT Method for different window lengths.

of 50 ms and 100 ms produced the most accurate detection for STSA. Even then, a smaller window length does not degrade the performance significantly (e.g., accuracy of ~95%). It

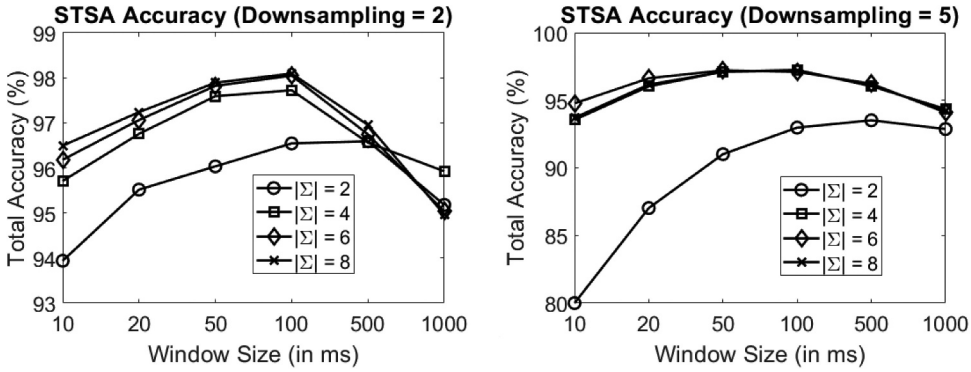


Figure 10. Detection accuracy of the STSA method for different window lengths at downsampling parameters of 2 and 5, for four different values of alphabet size $|\Sigma|$.

is observed that an alphabet size of 2 is probably insufficient for capturing the system dynamics. Typically, for proper control of a dynamic process, it is needed to be able to predict/detect changes using short data (i.e. at a time scale shorter than the time-scale the physical process needs in order to evolve into the anomalous state). Typically the growth of combustion instability takes to the order of 1 s (Bhattacharya et al. 2020b; Culler et al. 2018b), so detection using time-windows of 100 ms (0.1 s) at a rate of 50 Hz allows for speedy diagnosis and control.

A similar trend is observed in Figure 11 for the HMM, which shows that the window length of 100 ms yields the best results for both downsampling parameters of 2 and 5; smaller window lengths show comparable accuracies. In general, it is observed that the accuracy of HMM is higher than that of STSA. For HMM, the results for each of the selected number of hidden states are largely similar with minor improvements when using larger number of hidden states, almost similar to STSA, however, the training and testing time increases quite a bit when increasing the number of hidden states. This is further discussed in sub-section IV-D.

It is to be noted that, the primary frequencies of interest are in the range of 400–700 Hz (see sub-section II-A and Figure 3). Thus, a 1-s-long data clip has about 400–700 cycles of

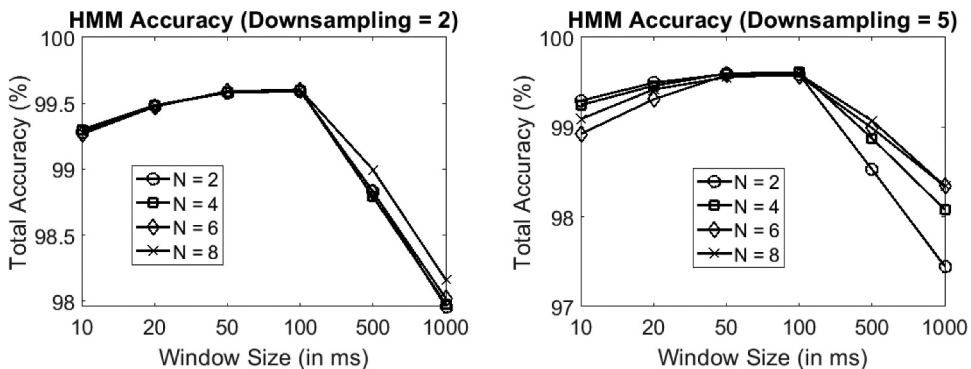


Figure 11. Detection accuracy of the HMM Method for different window lengths at downsampling parameters of 2 and 5, for the four values of number of hidden states.

pressure fluctuations. In tune with that 10 ms long windows have about 4–7 cycles, 100 ms windows correspondingly have about 40–70 cycles, with 1000 ms (i.e., 1 s) windows having 400–700. From a quick detection point of view, smaller observation windows are better, but due to them not capturing sufficient cycles, the algorithms are not able to learn the best models of the system regimes. About 1000 ms, being extremely long again degrades results for the transient data cases due to the averaging effect of both HMM and STSA algorithms which make it difficult to capture change points accurately. The 100 ms window is the best of both worlds, capturing sufficient signal data to make good assessments about the regime while at the same time being small enough to capture changes well.

Effects of downsampling on STSA and HMM performance

Downsampling is another important parameter to be selected for accurate execution of the algorithms. No downsampling implies that all data points are used, which leads to higher computational times (see sub-section IV-E). At the same time, no downsampling may cause high probabilities of self-looping in the graphs of STSA because many closely spaced consecutive points may lie in the same data segment of the partition (see sub-sub-section III-B1), which may reduce the discriminative power of STSA. On the other hand, a high downsampling parameter may discard some of the important dynamical information contained in the data, which could degrade the quality of learning and classification.

It is seen in Figure 12 that STSA yields similar accuracy for downsampling parameters of both 1 and 2 for a window length of 20 ms. In contrast, for a window length of 100 ms, the optimal downsampling parameter is 1 for alphabet sizes of 2 and 8 and the downsampling parameter of 2 for alphabet sizes 4 and 6. In general, the downsampling parameter of 2 yields similar or better results, and higher downsampling parameters degrade the accuracy for STSA.

Figure 13 shows that all of the downsampling parameters yield similar results (i.e., with accuracy differences in the order of 0.2%). Ignoring these differences, a modest amount of downsampling has no major bearing on the HMM performance as long as the training is also done with a data-set with the same downsampling parameter. In fact, slight downsampling of the data reduces computation time during testing and training. In general, it is observed that HMM slightly outperforms STSA.

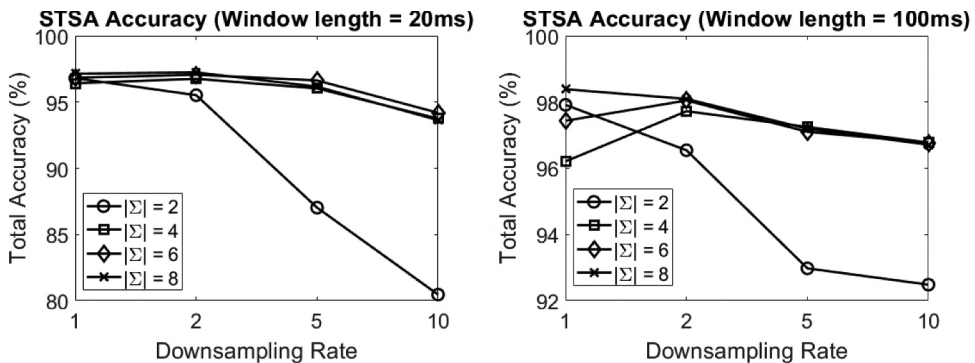


Figure 12. Detection accuracy of the STSA Method for different values of downsampling parameter DS with window lengths of 20 ms and 100 ms and for four different values of alphabet size $|\Sigma|$.

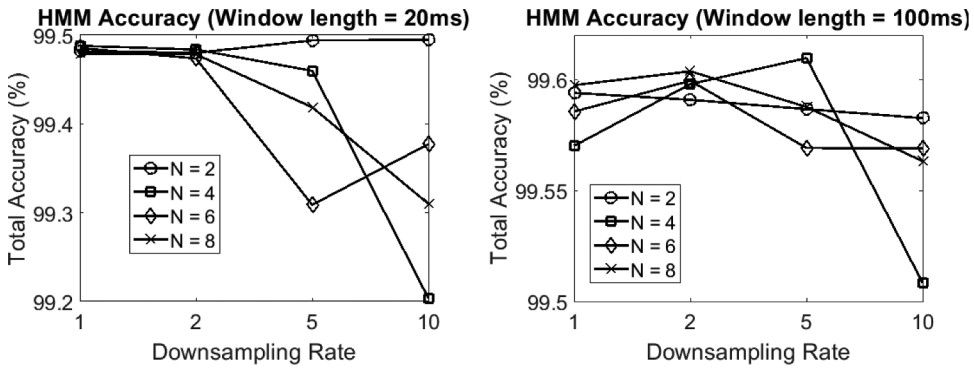


Figure 13. Detection accuracy of the HMM Method for different downsampling parameters, four different values of number of hidden states $|N|$, and window lengths of 20 ms and 100 ms.

The reason why downsampling emerges to be an important hyper-parameter for STSA-based methods is two-fold. Higher downsampling rates lead to poor detection since much of the signal information is not accounted for and is removed. However for STSA, if the signal is not sufficiently downsampled, a situation emerges where the state transition matrices for the STSA method have ‘heavy diagonals’ due to a large number of self-loops within the same state. This is a problem seen when there is oversampling. The fundamental frequencies of the signals are about the 400–700 Hz while the normal signal sampling rate is 16,384 Hz, which even using the maximum downsampling tested (DS = 10) yields a sampling rate of ~ 1640 Hz which is more than two times the highest frequency of interest, thus satisfying the Nyquist criterion. Other than this, there needs to be no mathematical linkage between the fundamental frequencies of the dynamical system and the sampling frequency. In other words, since the best sampling rate in this case is about 8192 Hz (16,384 Hz at DS = 2) when the frequencies of interest are in the range of 400–700 Hz, it means that a desirable sampling frequency would be about 20 times larger than the important frequencies.

Effects of alphabet size on STSA/PFSA performance and number of mixtures on HMM performance

The alphabet size of the STSA/PFSA method largely governs the accuracy of the method. Similarly, the number of hidden states changes the accuracy of the HMM method. Typically higher values of these yield better data separation. However having an extremely large number of alphabets/hidden states may degrade performance by needing more data to train the models, longer window lengths of data for good classification, and more training/testing times. Thus an optimal value of the alphabet size and number of hidden states must be arrived at.

It is observed from Figure 14 that the alphabet size of 2 is insufficient for capturing the system dynamics in the combustion pressure data. However, alphabet sizes of 4, 6, and 8 all yield largely similar results for each window length, which implies that an alphabet size of 4 is sufficient. The trend similar to that seen in sub-section IV-B, where the best performance is obtained for a window-length of 100 ms (closely followed by the 50 ms window-length)

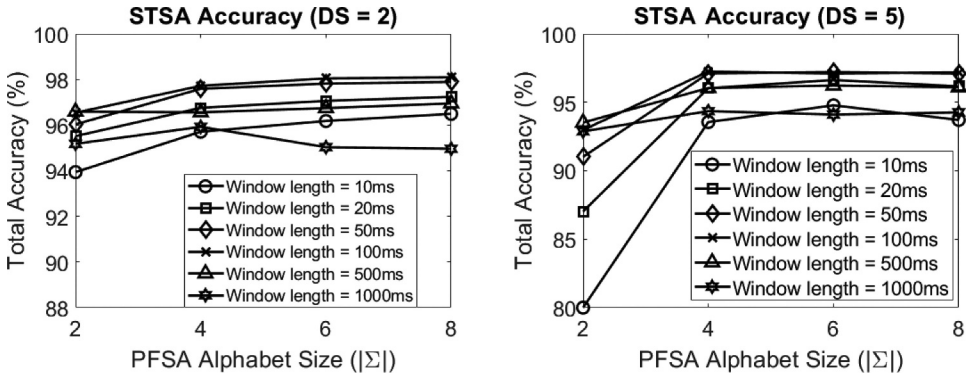


Figure 14. Detection accuracy of the STSA Method for different window lengths and for four different values of alphabet size $|\Sigma|$.

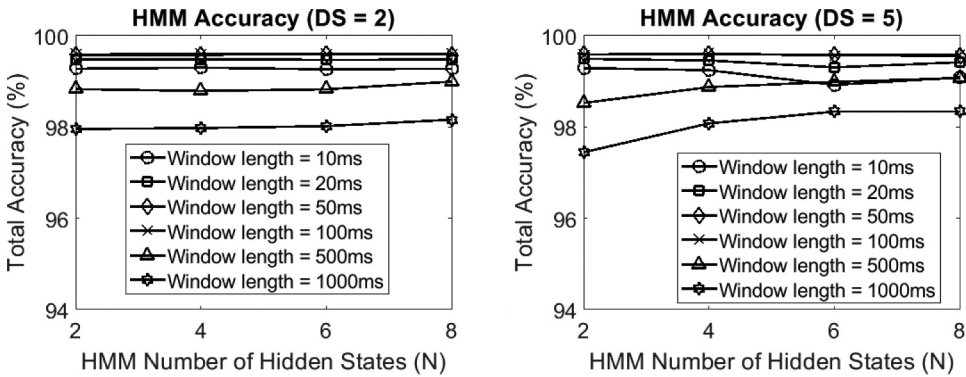


Figure 15. Detection accuracy of the HMM Method for different window lengths and four different values of number of hidden states $|N|$.

for both downsampling parameters of 2 and 5, with the performance dropping off when reducing or increasing the window-size.

Figure 15 shows that selecting the number of hidden states to be 2, 4, 6, and 8, causes the HMM method to yields similar results. This implies (as discussed in sub-section IV-B) that selecting the number of hidden states as $N = 2$ is sufficient and higher values of N are unnecessary, which increase the time complexity (see sub-section IV-E). Once again, window lengths of 50 ms and 100 ms produce the best results.

Time complexity

Table 3 lists the average testing time taken to analyze a single window of data for six different window lengths. Since no downsampling or symbolization of data is considered in

Table 3. Testing time (per window) using the FFT method for the different window lengths.

Window length (ms)	10	20	50	100	500	1000
Testing Time (in ms)	0.037	0.059	0.096	0.127	0.402	0.695

the FFT method, there is no impact on the downsampling parameter or alphabet size \number of mixture components. It is apparent that the testing time increases as the window-length increases, as is expected. It is noted that the highest testing time is much less than 1 ms, i.e., the FFT method is the fastest of the three data-driven methods compared. However, as seen previously, this is indeed the most inaccurate method under the present formulation. Therefore, the time complexities of the STSA and HMM methods are truly of major interest and are listed in Tables 4 and 5, respectively, in tabular forms for ease of comparison and visualization.

Table 4 shows that, in STSA for each downsampling parameter, the higher window-lengths require longer time to train or test. As the downsampling parameter is increased, the training and testing time reduce, because the algorithm needs to handle fewer data points (see Table 2). It is interesting to note that the training and testing times (per window) are very close in value for STSA. The testing time is slightly higher because the algorithm has to compute the Π matrix of the unknown data and compare it with two trained state transition matrices (i.e., stable and unstable), and then to use a threshold to make the final decision (see sub-section III-B).

A different observation can be made for HMM from Table 5. First, the training time of HMM is much higher for any given downsampling parameter or window length. In fact, the training time of HMM is about 2 orders of magnitude higher than that of STSA. Also, as HMM is an iterative method, there is no monotonic behavior with respect to window length

Table 4. Training and testing times (per window) of the STSA method for different window lengths and four different values of downsampling parameters with $|\Sigma| = 4$.

Window length (ms)	STSA training time (ms)				STSA testing time (ms)			
	DS = 1	DS = 2	DS = 5	DS = 10	DS = 1	DS = 2	DS = 5	DS = 10
10	0.055	0.048	0.045	0.044	0.082	0.074	0.071	0.073
20	0.062	0.065	0.059	0.056	0.093	0.084	0.078	0.074
50	0.090	0.068	0.054	0.061	0.12	0.11	0.082	0.084
100	0.13	0.10	0.063	0.054	0.17	0.12	0.10	0.084
500	0.57	0.33	0.13	0.087	0.59	0.36	0.17	0.12
1000	0.91	0.51	0.25	0.12	1.00	0.57	0.30	0.17

Table 5. Training and testing times (per window) of the HMM method for different window lengths and four different downsampling parameters with $N = 4$.

Window length (ms)	HMM training time (ms)				HMM testing time (ms)			
	DS = 1	DS = 2	DS = 5	DS = 10	DS = 1	DS = 2	DS = 5	DS = 10
10	17.29	8.59	3.98	1.91	1.013	0.70	0.52	0.49
20	21.42	10.16	4.58	2.44	1.85	1.08	0.72	0.58
50	16.76	8.15	4.45	2.28	3.46	2.13	1.01	0.80
100	20.01	7.92	3.64	2.29	6.16	3.42	1.76	1.01
500	19.43	9.77	4.39	2.02	27.97	14.54	6.08	3.30
1000	16.29	7.02	3.37	1.86	55.11	26.48	11.39	6.12

Table 6. Parameters and optimal thresholds for comparison of efficacy in early detection.

	window length (in ms)	DS	$ \Sigma $	M	Optimal η_{FFT}	Optimal η_{STSA}	Optimal η_{HMM}
Case I	100	2	4	2	0.031	0.371	(-)440
Case II	10	2	4	2	0.176	0.346	(-)45.9

for the same downsampling parameter. On the other hand, a simple algebraic structure of STSA allows a strong correlation between the number of data points and the training time. Similar to STSA, the training time of HMM reduces as the downsampling parameter is increased, but even then, on the average, the training time of HMM is still about 2 orders of magnitude higher than that of STSA. For the testing phase, the time complexity of HMM increases monotonically with window length and decreases monotonically as the downsampling parameter is increased. However, the forward method used to obtain the log-likelihood ratio (see sub-section III-C) is more complex than the simple distance norm used for STSA. Thus, the testing time of HMM is about 1 order of magnitude larger than that of STSA for the lower window lengths and creeps up to 2 orders of magnitude higher for longer window lengths. Thus, from a time complexity perspective, STSA is far superior to HMM for both training as well as testing.

It is also seen that both training and testing times of STSA increase as the alphabet size $|\Sigma|$ is increased. Similarly, for HMM, both training and testing times increase as the number of hidden states N is increased. The relative trend of time complexity is largely similar for STSA and HMM, although the magnitude for HMM train and test times is several times larger than STSA with a very minor drop in accuracy.

Thresholds

For each method there are a variety of parameters that need to be tuned, namely, window length for all of the three methods (i.e., FFT, STSA, and HMM), downsampling parameters for both STSA & HMM, and alphabet size, $|\Sigma|$, for STSA and number, N , of hidden states for HMM. Furthermore, there is an optimal threshold for each parameter combination, which yields the best performance. As mentioned earlier, the optimal threshold is computed as the threshold yielding the least errors (i.e., least number of misclassified data windows from the ROC curves). Thus, there is no “global threshold” that work across all the parameter sets, and instead the optimal threshold values is a strong function of the parameters, which must be determined from a training set before implementation of the algorithm. Optimal values of the three thresholds, η_{FFT} , η_{STSA} , and η_{HMM} , are obtained from the respective ROC plots of FFT, STSA, and HMM, as listed in Table 6. In the following subsection, results corresponding to a particular set of parameters and the corresponding optimal threshold have been discussed that elucidates this point further.

Classification versus early detection

This paper focuses on two important aspects of thermoacoustic instability (TAI) analysis of pressure time-series in combustors. The first is accurate classification of a given time series window into thermoacoustic stability or instability. The second aspect, which has not yet been addressed in this paper, is the need for efficient and prompt control of the thermoacoustic instabilities, namely, early detection of instability precursors and ability to predict the onset of a forthcoming instability as early as possible.

To study this second requirement, a set of parameters must be chosen and then the different methods (i.e., FFT, STSA, and HMM) need to be compared for evaluation of their performance and efficacy. As seen earlier, the window length of 100 ms yields the best

performance for all three methods. Similarly, the optimal alphabet size for STSA is 4 for this specific application, while the optimal number of hidden states for HMM is seen to be 2. A downsampling parameter of 1 or 2 is similar, but the downsampling parameter of 2 requires a lower training and testing time for both STSA and HMM. Thus, this set of parameters should yield the best performance. However, since detection from short-length time series is a thrust of this paper, it is necessary to compare the efficacy of the methods using the same set of parameters but replacing the window length with the (smallest) window length of 10 ms. Table 6 lists the parameters and optimal thresholds used in the 2 cases to compare the efficacy of the 3 methods when it comes to early detection of TAI. The windowing has been done at 50 Hz (i.e., once every 0.02 s) as discussed at the beginning of this section, which has been kept consistent throughout the paper.

It has been seen that the FFT method yielded poor results (as consistent with earlier findings). In order to reduce confusion, only the classification done by the STSA and HMM methods have been shown in the following images. In each of the images, the left hand plate shows the entire pressure time series with the true regime (based on the RMS value of the pressure) superimposed. The regime predicted by the STSA and HMM methods have also been superimposed. The right hand plate shows the same time series but zoomed into the portion where the regime changes from stable to unstable to appreciate how the algorithms are performing at the “change point.” Figures 16 and 17 show the classification on 2 sample time-series using a window length of 100 ms, while Figures 18 and 19 show the classification on the exact same 2 sample time-series but using a window length of 10 ms. the remaining method parameters are those listed in Table 6.

It is seen from Figures 16 and 17 that, for the 100 ms window, the HMM method goes closest to the the true regime, with the HMM method exactly matching the true state in Time Series 1 and with a very minor offset in Time Series 2. It can be seen from the zoomed plates that, in Time Series 1, STSA detects the change slightly late (by 0.02 s) while, for Time Series 2, it captures the impending change in regime far ahead of the change itself (0.08 s earlier) and also ahead of the HMM method (which detects it 0.02 s in advance).

Figures 18 and 19 show that, the window-size 10 ms, the STSA oscillates slightly at a few places giving rise to false alarms (classifying a stable regime window as unstable) or mis-detections (classifying an unstable regime window as stable). However, these errors last for

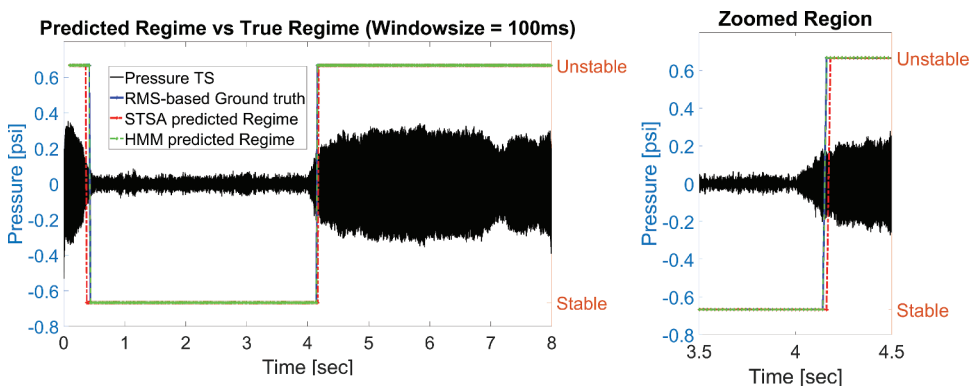


Figure 16. Sample time series 1 – Predicted vs actual regime with window length of 100 ms. Left: Entire time-series Right: Zoomed-in on the change point region

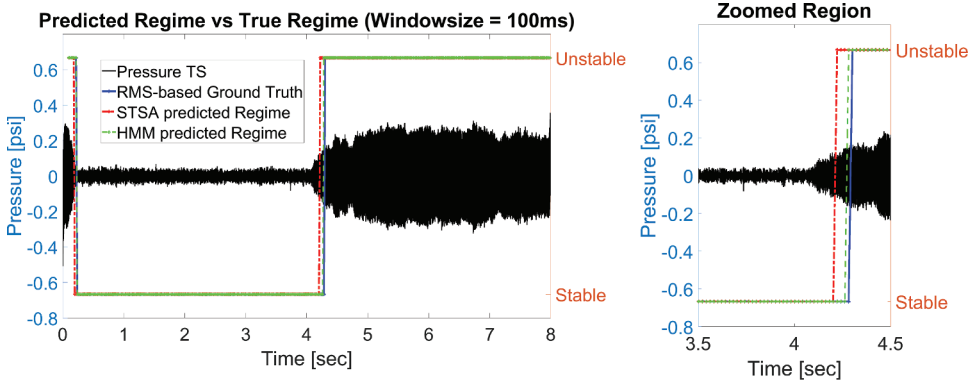


Figure 17. Sample time series 2 – Predicted vs actual regime with window length of 100 ms. Left: Entire time-series Right: Zoomed-in on the change-point region

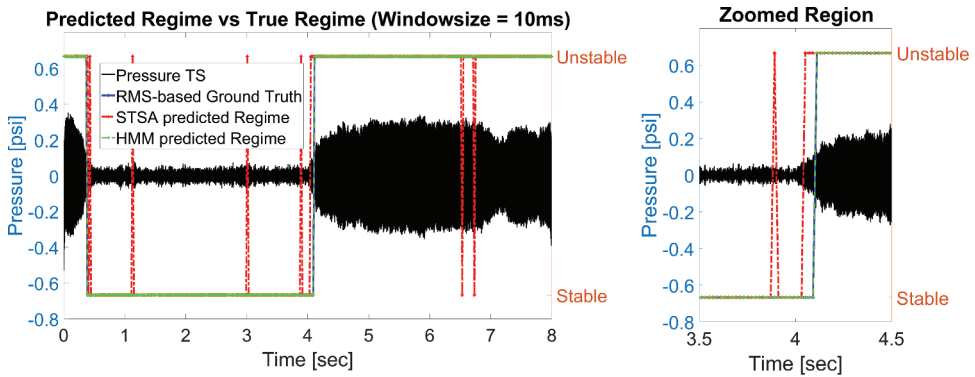


Figure 18. Sample time series 1 – Predicted vs actual regime with window length of 10 ms. Left: Entire time-series Right: Zoomed-in on the change-point region

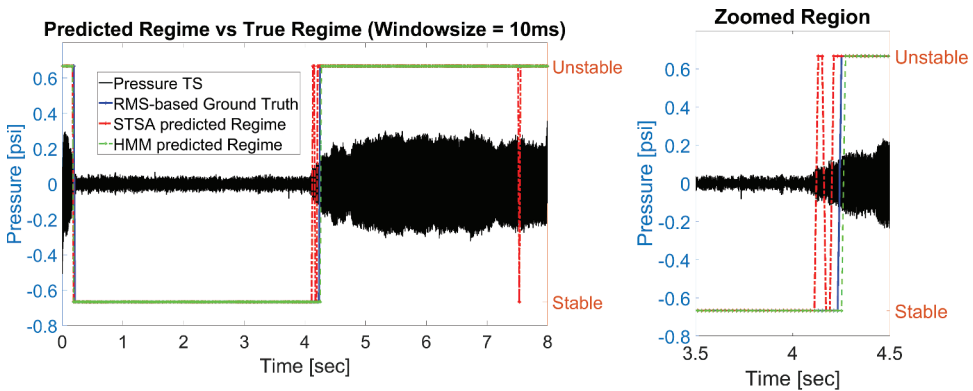


Figure 19. Sample time series 2 – Predicted vs actual regime with window length of 10 ms. Left: Entire time-series Right: Zoomed-in on the change-point region

a single window generally and then returns to a correct classification, so that would not cause a major issue of incorrect control. However, it is interesting to note that in both the time series shown, STSA detected the impending change to instability earlier than the RMS-based ground truth or the HMM method. In time series 1, the STSA method detects the change 0.06 s in advance, while the HMM method exactly traces the ground truth. In time series 2, STSA captures the impending change 0.04 s in advance and ahead of the HMM method, which detects it 0.02 s late.

Although only two pressure time series are reported here, it has been observed that, for smaller window lengths, STSA mainly detected the impending change in advance and, at most, coincided with the exact change in the system operation from a the stable to unstable regime. However, the HMM method could not detect the change so far in advance. For higher window lengths there was no clear winner between STSA and HMM for early detection.

In the experiments, the staging fuel was injected 3.75 inch (i.e., 95.25 mm) upstream of the flame. The bulk flow speed of the air, which would roughly be the same speed at which the control fuel would convect, was 26 m/s (Culler et al. 2018b). This implies that the time taken for the control fuel to convect from the control nozzle to the burning zone would be about 3.6 ms. The actuator delay was found to be about 6 ms. Allowing for a data assimilation delay of 1 ms, STSA data processing delay of 0.074 ms (Table 4), the total time taken for the control to be implemented and take effect in the burning zone is about 11 ms.

Thus, being able to detect the impending changes from the stable to the unstable regime ~ 40 ms in advance allows for control to be implemented before the instability starts growing to dangerously high amplitudes. For the purpose of preemptive control using short time series data, STSA would be preferable. The FFT method yields poor results in general for short time-series data and should be used only as a baseline method.

Summary, conclusions, and future research

In this research, three standard data-driven methods, namely, fast Fourier transform (FFT), symbolic time series analysis (STSA), and hidden Markov modeling (HMM), have been investigated for detection of thermoacoustic instability (TAI). These methods are compared for their efficacy and accuracy for fast detection and early detection of instability using short data lengths. The effects of various algorithm parameters (e.g., window length and downsampling parameter of time series), used in STSA and HMM, have been studied to identify an optimal set of algorithm parameters. The time complexity for training as well as testing a given data window has also been investigated. Finally a comparison has been made to find which method is best for predicting impending instability to be included in an (active) control policy. The following conclusions are drawn from this investigation:

- The HMM method reports highest detection accuracy, very closely followed by the STSA method. The FFT method is found to be most inaccurate for the windowed analysis. However, it must also be noted that the FFT method produces high accuracy ($\sim 95\%$) for 100 ms windows, which compares well to PFSA ($\sim 97\%$) and

HMM (~99%) for the same window length. This occurs due to the two primary drawbacks of the FFT methods traditionally used, (1) for accurate FFT plots, adequate length of signal is needed and (2) being an averaging method, it reduces the applicability of FFT-based methods for analysis of non-stationary signals using longer window lengths.

- The FFT method is the fastest for testing, with the STSA method having slightly higher, but comparable CPU execution time. The HMM method requires at least 1 order of magnitude larger CPU execution time for testing.
- For training, CPU execution time of STSA is about 2 orders of magnitude smaller than that of HMM. This makes STSA a clear winner in the computational speed of both training and testing as compared to HMM, with a very small loss of accuracy. This property of STSAs finds use in the concept of online regime discovery which may be used to detect additional sub-regimes as demonstrated by Bhattacharya and Ray (2020b) where the authors ‘discovered’ anomalous operation in a Rijke tube (Rijke 1859) in an online and partly supervised manner.
- For STSA, the alphabet size $|\Sigma|$ should be greater than or equal to 4 to be able to accurately capture the dynamics of the signal for the combustion TAI detection. For HMM, 2 hidden states (N) and a 2 mixture emission (M) are sufficient for achieving similar accuracy.
- The ideal downsampling parameter for both STSA and HMM is found to be 2.
- The window length of 100 ms yields the best detection performance for both STSA and HMM; but even a 10 ms window (i.e., 82 data points at a downsampling parameter of 2) also produces almost comparable results.
- For predictive analysis (i.e., detection of precursors to impending thermoacoustic instability (TAI)), the STSA is superior to both FFT and HMM. However, the results do not indicate a clear superiority for the larger window lengths. This early detection capability of STSA is achieved at an apparent cost of loss of TAI detection accuracy. It is to be noted that TAI detection accuracy is not sufficient in deciding the “best” method for early detection.
- If the computational speed is of prime importance, or if it is essential to have early detection of TAI (especially in an online setting), then STSA is recommended for use with carefully tuned model parameters. If, however, the goal is to have a classification algorithm that closely matches the ground truth, and if there is no strict time constraints in the training or testing phases, the HMM method appears to be preferable. The FFT method should be used as a baseline, because it does not act as a good online detector with short-length time series.

While there are many topics of both theoretical and experimental research, the authors suggest the following topics for research in the near future:

- a) **Topic 1:** Finding a method to have a global set of thresholds to work across a range of window lengths.
- b) **Topic 2:** Development of more powerful and accurate FFT based methods as it has the lowest computation time.
- c) **Topic 3:** Testing these algorithms in a three regime setting; stable, unstable and unstable (increasing and decreasing).

- d) **Topic 4:** Exploring other pre-processing techniques and study if they improve the accuracy of the PFSA and HMM methods mentioned here, e.g., using wavelet transforms for noise removal, etc.
- e) **Topic 5:** Study other genres of data-driven tools, namely neural networks, such as Recurrent Neural networks (RNNs) and Long Short-Term Memory (LSTM) networks.

Acknowledgments

The work reported here has been supported in part by the U.S. Air Force Office of Scientific Research (AFOSR) under Grant Nos. FA9550-15-1-0400 and FA9550-18-1-0135 in the area of dynamic data-driven application systems (DDDAS). The data collection was funded by the U.S. Department of Energy University Turbine Systems Research Program under grant DE-FE0025495 with program monitor Mark Freeman. The first author is also thankful to Indo-US Science and Technology Forum (IUSSTF) for granting the Research Internship for Science and Engineering (RISE) scholarship to him for collaboration between Pennsylvania State University and Jadavpur University. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

Funding

This work was supported by the Air Force Office of Scientific Research [FA9550-15-1-0400,FA9550-18-1-0135]; Indo-US Science and Technology Forum [Research Internship for Science and Engineering]; U.S. Department of Energy [DE-FE0025495].

References

- Ali, S. S., and M. U. Ghani. 2014. Handwritten digit recognition using dct and hmms. 2014 12th International Conference on Frontiers of Information Technology, Islamabad, Pakistan, 303–06.
- Bhattacharya, C., S. De, A. Mukhopadhyay, S. Sen, and A. Ray. 2020a. Detection and classification of lean blow-out and thermoacoustic instability in turbulent combustors. *Appl. Therm. Eng.* 180. 115808.
- Bhattacharya, C., S. Mondal, A. Ray, and A. Mukhopadhyay. 2020b. Reduced-order modelling of thermoacoustic instabilities in a two-heater Rijke tube. *Combust. Theor. Model.* 24 (3):530–48. doi:10.1080/13647830.2020.1714080.
- Bhattacharya, C., and A. Ray. 2020a. Data-driven detection and classification of regimes in chaotic systems via hidden markov modeling. *ASME Lett. Dyn. Syst. Control.* 1:1–6.
- Bhattacharya, C., and A. Ray. 2020b. Online discovery and classification of operational regimes from an ensemble of time series data. *J. Dyn. Syst. Meas. Control* 142 (11). doi:10.1115/1.4047449.
- Bishop, C. 2007. *Pattern recognition and machine learning*. New York, USA: Springer.
- Candel, S. 2002. Combustion dynamics and control: Progress and challenges. *Proc. Combust. Inst.* 29 (1):1–28. doi:10.1016/S1540-7489(02)80007-4.
- Chattopadhyay, P., S. Mondal, C. Bhattacharya, A. Mukhopadhyay, and A. Ray. 2017. Dynamic data-driven design of lean premixed combustors for thermoacoustically stable operations. *J. Mech. Des.* 139 (11):111419–1–111419–10. doi:10.1115/1.4037307.
- Culler, W., X. Chen, S. Peluso, D. Santavicca, J. O'Connor, and D. Noble. 2018a. Comparison of center nozzle staging to outer nozzle staging in a multi-flame combustor. In *Turbo expo: Power for land, sea, and air*, Vol. 51050, V04AT04A024. American Society of Mechanical Engineers.
- Culler, W., X. Chen, J. Samarasinghe, S. Peluso, D. Santavicca, and J. O'Connor. 2018b. The effect of variable fuel staging transients on self-excited instabilities in a multiple-nozzle combustor. *Combust. Flame* 194:472–84. doi:10.1016/j.combustflame.2018.04.025.

- Ghalyan, N. F., S. Mondal, D. J. Miller, and A. Ray. 2019. Hidden Markov modeling-based decision-making using short-length sensor time series. *J. Dyn. Syst. Meas. Control* 141 (10). doi:10.1115/1.4043428.
- Hajek, B. 2015. *Random processes for engineers*. 1st ed. Cambridge, UK: Cambridge University Press.
- Hauser, M., Y. Li, J. Li, and A. Ray. 2016. Real-time combustion state identification via image processing: A dynamic data-driven approach. 2016 American Control Conference (ACC), Boston, MA, 3316–21.
- Hochreiter, S., and J. Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9 (8):1735–80. doi:10.1162/neco.1997.9.8.1735.
- Howe, A., D. Doleiden, S. Peluso, and J. O'Connor. 2020. The effect of the degree of premixedness on self-excited combustion instability. ASME Turbo Expo, London, England.
- Husken, M., and P. Stagge. 2003. Recurrent neural networks for time series classification. *Neurocomputing* 50:223–35. doi:10.1016/S0925-2312(01)00706-8.
- Kabiraj, L., and R. I. Sujith. 2012. Nonlinear self-excited thermoacoustic oscillations: Intermittency and flame blowout. *J. Fluid Mech.* 713:376–97. doi:10.1017/jfm.2012.463.
- Lacasa, L., B. Luque, F. Ballesteros, J. Luque, and J. C. Nuño. 2008. From time series to complex networks: The visibility graph. *Proc. Natl. Acad. Sci.* 105 (13):4972–75. doi:10.1073/pnas.0709247105.
- Lee, J. G., and D. A. Santavicca. 2003. Experimental diagnostics for the study of combustion instabilities in lean premixed combustors. *J. Propul. Power* 19 (5):735–50. doi:10.2514/2.6191.
- Lieuwen, T. 2005. Online combustor stability margin assessment using dynamic pressure data. *J. Eng. Gas Turbines Power* 127 (3):478–82. doi:10.1115/1.1850493.
- Lieuwen, T., and V. Yang. 2005. *Combustion instabilities in gas turbine engines: Operational experience, fundamental mechanisms, and modeling*. Reston, VA: AIAA.
- Matveev, K. I. 2003. Thermoacoustic instabilities in the Rijke tube: Experiments and modeling. PhD thesis, California Institute of Technology, Pasadena, CA.
- Mondal, S., C. Bhattacharya, P. Chattopadhyay, A. Mukhopadhyay, and A. Ray. 2017. Prediction of thermoacoustic instabilities in a premixed combustor based on fft-based dynamic characterization. 53rd AIAA/SAE/ASEE Joint Propulsion Conference, Atlanta, GA.
- Mondal, S., N. F. Ghalyan, A. Ray, and A. Mukhopadhyay. 2019. Early detection of thermoacoustic instabilities using hidden markov models. *Combust. Sci. Technol.* 191 (8):1309–36. doi:10.1080/00102202.2018.1523900.
- Mukherjee, K., and A. Ray. 2014. State splitting and merging in probabilistic finite state automata for signal representation and analysis. *Signal Process.* 104:105–19. doi:10.1016/j.sigpro.2014.03.045.
- Murphy, K. 2012. *Machine learning: A probabilistic perspective*. 1st ed. Cambridge, MA: The MIT Press.
- Murugesan, M., and R. I. Sujith. 2015. Combustion noise is scale-free: Transition from scale-free to order at the onset of thermoacoustic instability. *J. Fluid Mech.* 772:225–45. doi:10.1017/jfm.2015.215.
- Nair, V., G. Thampi, S. Karuppusamy, S. Gopalan, and R. I. Sujith. 2013. Loss of chaos in combustion noise as a precursor of impending combustion instability. *Int. J. Spray Combust. Dyn.* 5 (4):273–90. doi:10.1260/1756-8277.5.4.273.
- Najkar, N., F. Razzazi, and H. Sameti. 2010. A novel approach to hmm-based speech recognition systems using particle swarm optimization. *Math. Comput. Model.* 52 (11):1910–20. The BIC-TA 2009 Special Issue.
- Noiray, N., and B. Schuermans. 2012. Theoretical and experimental investigations on damper performance for suppression of thermoacoustic oscillations. *J. Sound Vib.* 331 (12):2753–863. doi:10.1016/j.jsv.2012.02.005.
- O'Connor, J., V. Acharya, and T. Lieuwen. 2015. Transverse combustion instabilities: Acoustic, fluid mechanic, and flame processes. *Prog. Energy Combust. Sci.* 49:1–39. doi:10.1016/j.pecs.2015.01.001.
- Oates, T., L. Firoiu, and P. Cohen. 2000. Using dynamic time warping to bootstrap HMM-based clustering of time series. In *Sequence learning*, 35–52. New York, USA: Springer.
- Rabiner, L., and B.-H. Juang. 1993. *Fundamentals of speech recognition*. Upper Saddle River, NJ: Prentice-Hall, Inc.

- Rabiner, L. R. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77 (2):257–86. doi:10.1109/5.18626.
- Rajagopalan, V., and A. Ray. 2006. Symbolic time series analysis via wavelet-based partitioning. *Signal Process.* 86 (11):3309–20. doi:10.1016/j.sigpro.2006.01.014.
- Ray, A. 2004. Symbolic dynamic analysis of complex systems for anomaly detection. *Signal Process.* 84 (7):1115–30. doi:10.1016/j.sigpro.2004.03.011.
- Rayleigh, L. 1845. *The theory of sound*. Mineola, NY: Dover Publications.
- Richecoeur, F., S. Ducruix, P. Scoufflaire, and S. Candel. 2008. Experimental investigation of high-frequency combustion instabilities in liquid rocket engine. *Acta Astronaut.* 62 (1):18–27. doi:10.1016/j.actaastro.2006.12.034.
- Rijke, P. L. 1859. Notiz über eine neue art, die in einer an beiden enden offenen röhre enthaltene luft in schwingungen zu versetzen. *Annalen der Physik und Chemie* 183 (6):339–43. doi:10.1002/andp.18591830616.
- Samarasinghe, J., W. Culler, B. D. Quay, D. A. Santavicca, and J. O’connor. 2017. The effect of fuel staging on the structure and instability characteristics of swirl-stabilized flames in a lean premixed multinozzle can combustor. *J. Eng. Gas Turbines Power* 139 (12):121504.
- Sarkar, S., S. R. Chakravarthy, V. Ramanan, and A. Ray. 2016. Dynamic data-driven prediction of instability in a swirl-stabilized combustor. *Int. J. Spray Combust. Dyn.* 8 (4):235–53. doi:10.1177/1756827716642091.
- Sarkar, S., K. G. Lore, S. Sarkar, V. Ramanan, S. R. Chakravarthy, S. Phoha, and A. Ray. 2015. Early detection of combustion instability from hi-speed flame images via deep learning and symbolic time series analysis. Annual Conference of the Prognostics and Health Management Society, Coronado, CA.
- Sen, U., T. Gangopadhyay, C. Bhattacharya, A., . S. Misra, P. S. Karmakar, A. Mukhopadhyay, and S. Sen. 2016. Investigation of ducted inverse nonpremixed flame using dynamic systems approach. ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition, Seoul, South Korea, Vol. 4B.
- Sen, U., T. Gangopadhyay, C. Bhattacharya, A. Mukhopadhyay, and S. Sen. 2018. Dynamic characterization of a ducted inverse diffusion flame using recurrence analysis. *Combust. Sci. Technol.* 190 (1):32–56. doi:10.1080/00102202.2017.1374952.
- Subbu, A., and A. Ray. 2008. Space partitioning via hilbert transform for symbolic time series analysis. *Appl. Phys. Lett.* 92 (8):084107. doi:10.1063/1.2883958.
- Unni, V. R., A. Mukhopadhyay, and R. I. Sujith. 2015. Online detection of impending instability in a combustion system using tools from symbolic time series analysis. *Int. J. Spray Combust. Dyn.* 7 (3):243–55. doi:10.1260/1756-8277.7.3.243.
- Wang, Z., W. Yan, and T. Oates. 2017. Time series classification from scratch with deep neural networks: A strong baseline, 2017 International Joint Conference on Neural Networks (IJCNN), Anchorage, AK, 1578–85.
- Zhao, D. 2012. Transient growth of flow disturbances in triggering a Rijke tube combustion instability. *Combust. Flame* 159 (6):2126–37. doi:10.1016/j.combustflame.2012.02.002.
- Zhao, D., and Z. H. Chow. 2013. Thermoacoustic instability of a laminar premixed flame in Rijke tube with a hydrodynamic region. *J. Sound Vib.* 332 (14):3419–37. doi:10.1016/j.jsv.2013.01.031.