Early Detection of Combustion Instability from Hi-speed Flame Images via Deep Learning and Symbolic Time Series Analysis

Soumalya Sarkar[1], Kin G Lore[2], and Soumik Sarkar[3], Vikram Ramanan[4], Satyanarayanan R Chakravarthy[5], Shashi Phoha[6], Asok Ray[7]

[1,6,7] *Pennsylvania State University, University Park, Pennsylvania, USA*
svs5464@psu.edu, sxp26@arl.psu.edu, axr2@psu.edu

[2,3] *Iowa State University, Ames, Iowa, USA*
kglore@iastate.edu, soumiks@iastate.edu

[4,5] *Indian Institute of Technology Madras, Chennai, India*
vikrambest@yahoo.co.in, src@ae.iitm.ac.in

ABSTRACT

Combustion instability, characterized by self-sustained, large-amplitude pressure oscillations and periodic shedding of coherent vortex structures at varied spatial scales, has many detrimental effects on flight-propulsion dynamics and structural integrity of gas turbine engines. Hence, its early detection is one of the important tasks in engine health monitoring and prognostics. This paper proposes a dynamic data-driven approach, where a large volume of sequential hi-speed (greyscale) images is used to analyze the temporal evolution of coherent structures in combustion chamber for early detection of combustion instability at various operating conditions. The proposed hierarchical approach involves extracting low-dimensional semantic features from images using Deep Neural Networks followed by capturing the temporal evolution of the extracted features using Symbolic Time Series Analysis (STSA). Extensive experimental data have been collected in a swirl-stabilized dump combustor at various operating conditions for validation of the proposed approach. Intermediate layer visualization of deep learning reveals that meaningful shape-features from the flame images are extracted, which enables the temporal modeling layer to enhance the class separability between stable and unstable regions. At the same time, the semantic nature of intermediate features enables expert-guided data exploration that can lead to better understanding of the underlying physics. To the best of the authors knowledge, this paper presents one of the early applications of the recently reported Deep Learning tools in the area of prognostics and health management (PHM).

1. INTRODUCTION

Combustion instability is a very undesirable phenomenon characterized by high-amplitude flame oscillations at discrete frequencies. These frequencies typically represent the natural duct/resonator acoustic modes. Combustion instability, in its most basic form arises when there is a positive coupling between the heat release rate oscillations and the pressure oscillations, provided this driving force is higher than the damping present in the system. The mechanisms of pressure-heat release rate coupling are system dependent and thus, the problem of combustion instability becomes very system specific.

The underlying principle of heat release rate oscillations, that drives the pressure oscillations-which result in velocity oscillations and in turn modulate heat release rate oscillations-all in a turbulent background in case of actual gas turbine combustors pose significant complexities in determining the mechanisms of combustion instability. Crocco (Mcmanus, Poinsot, & Candel, 1993) modeled unsteady heat release rate as a function of unsteady velocity to determine stability of a ducted zero-mean flow flame. Subsequently, a whole class of reduced order modeling-Flame Transfer/Describing functions were theoretically and experimentally (Palies, Schuller, Durox, & Candel, 2011; Noiray, Durox, Schuller, & Candel, 2008; Bellows, Bobba, Forte, Seitzman, & Lieuwen, 2007) formulated to understand the stability of the system by means of solving the dispersion relation. In addition, flame oscillation saturation mechanisms were also experimentally diagnosed which in addition to experiments based on turbulent non reacting and reacting flows led to the universal feature of combustion instability- heat release rate oscillations driven by coherent structures.

Coherent structures are fluid mechanical structures associ-

ated with coherent phase of vorticity, high levels of vorticity among other definitions (Hussain, 1983). These structures, whose generation mechanisms vary system wise, cause large scale velocity oscillations and overall flame shape oscillations by curling and stretching. These structures can be caused to shed/generated at the duct acoustic modes when the forcing (pressure) amplitudes are high. The interesting case of the natural shedding frequency of these structures, causing acoustic oscillations, has been observed by Chakravarthy et al. (Chakravarthy, Shreenivasan, Bhm, Dreizler, & Janicka, 2007).

Recently, a swirl combustor has been characterized and a wide range of experiments relating swirl flows and coherent structures associated with swirl flows has been reported (Syred, 2006; Paschereit, Gutmark, , & Weisenstein, 1998). The presence of Precessing vortex core as the dominant coherent structure has been reported and non linear interactions between heat release rate oscillations and PVC as the cause of superposed frequencies in time series data has also been reported (Moeck, Bourgouin, Durox, Schuller, & Candel, 2012). Much of the literature is dedicated to detection and correlation of these coherent structures to heat release rate and unsteady pressure. The popular methods resorted for detection of coherent structures are proper orthogonal decomposition (POD) (Berkooz, Holmes, & Lumley, 1993) and dynamic mode decomposition (DMD) (Schmid, 2010), which use tools from spectral theory to derive spatial coherent structure modes. DMD has been used to estimate the growth rates and frequencies from experimental data and thus offered to perform stability analysis on experimental data.

This paper proposes a data-driven hierarchical framework for early detection of thermo-acoustic instability from hi-speed greyscale images. In the lower layer, large volume of hi-speed sequential images are used to train a deep neural network model that extracts hierarchical features from the training data (G. E. Hinton & Salakhutdinov, 2006) through the use of multiple layers of latent variables. An unsupervised pre-training approach with deep-belief networks (DBN) (G. E. Hinton, 2009) is used in particular to automatically learn the coherent structures while reducing the dimension of the images for temporal modeling at the top layer (Erhan, Bengio, et al., 2010). Symbolic time series analysis (STSA) (Ray, 2004), a fast probabilistic graphical model is placed at the top layer to extract temporal feature from the output of deep learning model. The concept of STSA has been used for anomaly detection in physical systems as reported in (Ray, 2004; Rao, Ray, Sarkar, & Yasar, 2009; Sarkar, Jin, & Ray, August, 2011). Recently, STSA is applied on pressure and chemiluminescence time series for early detection of Lean-blow out (Mukhopadhyay, Chaudhari, Paul, Sen, & Ray, 2013; Sarkar, Ray, Mukhopadhyay, Chaudhari, & Sen, 2014) and thermo-acoustic instability (Ramanan, Chakravarthy, Sarkar, & Ray, 2014).

From the above perspectives major contributions of the paper are delineated below.

- A novel data-driven framework, with DBN at lower layer and STSA at upper layer, is proposed for early detection of thermo-acoustic instability from hi-speed images.
- In the above framework, the DBN layers extract meaningful shape-features to represent the coherent structures in the flame images. This phenomenon enables STSA at the temporal modeling layer to enhance the class separability between stable and unstable modes of combustion, which implies higher precision for early detection of the onset of combustion instability.
- The proposed theory and the associated algorithms have been experimentally validated at multiple operating conditions in a swirl-stabilized combustor by characterizing the stable and unstable states of combustion.
- Training and testing of the proposed framework have been performed on different operating conditions (e.g., Reynolds number (Re), fuel flow rate, and air-fuel premixing level) of the combustion process to test the transferability of the approach. Performance of the proposed framework (DBN+STSA) have been evaluated by comparison with that of a framework, where DBN is replaced by another extensively used dimensionality reduction tool, principal component analysis (PCA) (Bishop, 2006).

The paper is organized in five sections, including the present one. Section 2 describes a laboratory-scale swirl-stabilized combustor, which serves as a test apparatus for experimental validation of the proposed architecture for early detection of thermo-acoustic instability. Section 3 describes the proposed framework along with its building blocks via explaining the concepts of DBN and STSA. Section 4 presents the capability and advantages of the proposed approach along with the feature visualization at intermediate layers of DBN. Finally, the paper is summarized and concluded in Section 5 with selected recommendations for future research.

2. EXPERIMENTAL SETUP

The swirl combustor test bed used in this study has a swirler of diameter 30 mm with 60 degree vane angles, thus yielding a geometric swirl number of 1.28. Air to the combustor is fed through a settling chamber of diameter 280 mm with a sudden contraction leading to a square cross section of side 60 mm. This provides an area ratio of around 17, which thus acts as an acoustically open condition at the contraction. A mesh and honeycomb are mounted in immediate downstream of the contraction to provide uniform flow to the swirler. The combustor, shown in figure 1(a) consists of an inlet section of length 200 mm, an inlet optical access module(IOAM) of length 100 mm to provide optical access to the fuel tube, a primary combustion chamber of length 370 mm, and secondary

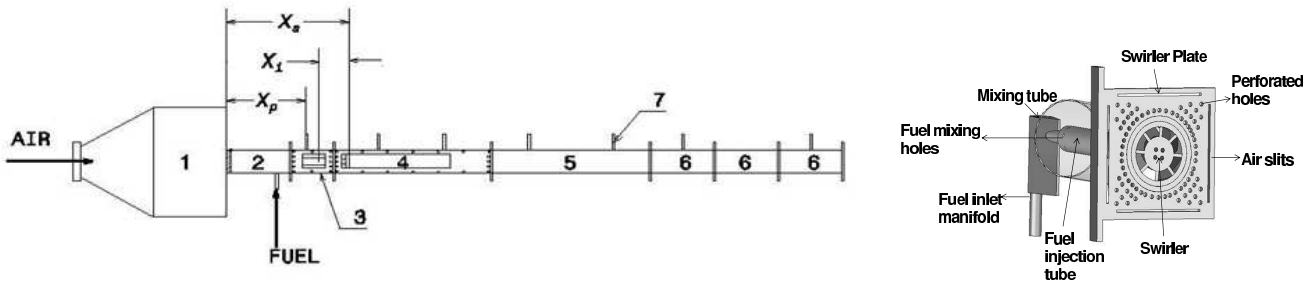

(a) (b)

Figure 1. (a)Schematic of the experimental setup. 1 - settling chamber, 2 - inlet duct, 3 - IOAM, 4 - test section, 5 - big extension duct, 6 - small extension ducts, 7 - pressure transducers, X_s - swirler location measured downstream from settling chamber exit, X_p - transducer port location measured downstream from settling chamber exit, X_i - fuel injection location measured upstream from swirler exit, (b) Swirler assembly used in the combustor

duct of the same length. Extension ducts of the same cross section are added to provide length flexibility. The overall length of the constant area ducts was chosen to be 1340 mm.

The fuel injection is done by injecting it coaxially with the air in a fuel injection tube with slots on the surface as shown in Figure 1(b). The fuel injection tube is coaxial to a mixing tube which has the same diameter as that of the swirler. The bypass air that does not enter the mixing tube passes through slots on the swirl plate. The slots on the fuel injection tube are drilled at designated distance upstream of the swirler. The larger this distance, more fuel mixes with the primary air in the mixing tube thus leading to more premixedness. Two upstream distances of $X_1 = 90mm$ and $X_2 = 120mm$ were chosen for this work. The upstream distance of 120 mm provides for full premixing of the fuel with the air thus henceforth, it will be referred to as the premixed case. The 90 mm upstream injection case causes partial premixing of the fuel with air and thus will be referred to as the partially premixed case. The images were acquired at 3 kHz using Photron High speed star with a spatial resolution of 1024×1024 pixels. The data acquisition was triggered simultaneously using NI card and taken for a duration of $3s$ yielding in a sequence of 9,000 images for every operating condition.

Two inlet Reynolds numbers (Re), based on the swirler diameter were chosen, the lower Re having stable combustion behavior and higher Re having exhibiting unstable behavior. The Re's were chosen to be 7,971 and the higher Re being 15,942 for a fuel flow rate (FFR) of 0.495 g/s. Another protocol followed was keeping the inlet Re constant at 10,628 and having two different fuel flow rates. The higher FFRs exhibited stable combustion, whereas the leaner configuration was unstable. The two FFRs were chosen to be 0.66 g/s and 0.308 g/s. These corresponded to equivalence ratios of 0.955 and 0.445 respectively. Besides these conditions, 3 seconds of images are also collected for $Re = 1,771$ and $FFR = 0.083$ at relatively stable state of combustion. The details of the operating conditions along with their ground truth (e.g., stable or unstable) are presented in table 1.

Table 1. Description of operating conditions along with respective ground truth (stable or unstable) for hi-speed image data collection. $3s$ of greyscale image sequence at $3kHz$ is collected for each condition

Premixing	FFR (g/s)	Re	Ground truth
Partial ($X_1 = 90mm$)	0.495	7,971	Stable
		15,942	Unstable
	0.308	10,628	Unstable
	0.66		Stable
Full ($X_2 = 120mm$)	0.495	7,971	Stable
		15,942	Unstable
	0.308	10,628	Unstable
	0.66		Stable
	0.083	1,771	Relatively stable

Figure 2 presents sequences of images of dimension 392×1000 pixels for both stable ($Re = 7,971$, $FFR = 0.495g/s$ and full premixing) and unstable ($Re = 15,942$, $FFR = 0.495g/s$ and full premixing) states. The flame inlet is on the right side of each image and the flame flows downstream to the left. It can be observed that the flame does not have any prominent coherent structure when the combustion is stable. While the combustion is unstable, vortex shedding along the flow is observed. Bottom segment of the figure 2 shows formation of mushroom-shaped vortex at $t = 0, 0.001s$ and the shedding of that towards downstream from $t = 0.002s$ to $t = 0.004s$.

3. DECISION FRAMEWORK AND TOOLS

This section describes the proposed architecture for early detection of thermo-acoustic instability in a combustor via analyzing a sequence of hi-speed images. Figure 3 presents the schematics of the framework where a deep belief net-

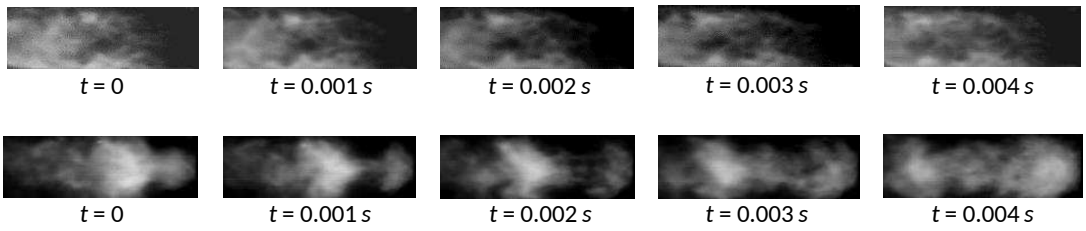

Figure 2. Top: greyscale images at $Re = 7,971$ and full premixing for a fuel flow rate of 0.495 g/s, bottom: greyscale images at $Re = 15,942$ and full premixing for a fuel flow rate of 0.495 g/s

work (DBN) is stacked with symbolic time series analysis (STSA). In the training phase, images (or a segment of the images) from both stable and unstable states for various operating conditions are used as the visible layer V of a DBN. Multiple hidden layers (i.e., h_1 to h_n) with reducing dimensions (G. E. Hinton & Salakhutdinov, 2006) are stacked after the visible layer.The weights (i.e., W_1 to W_n), connecting adjacent layers, are learned first via greedy layer-wise pretraining (G. E. Hinton & Salakhutdinov, 2009) and then they are fine-tuned in a supervised manner. In this paper, unsupervised pre-training step is emphasized more for capturing the coherent structures in flame images at unstable state. The vector of activation probabilities of the hidden units at the topmost hidden layer is used as input to the STSA module.

While testing, sequence of images are passed through the learned DBN and a time series of l_2 norm (equivalent to signal energy) of the activation probability vectors is obtained. In STSA module, the time-series is symbolized via partitioning the signal space and a symbol sequence is created as shown in the figure 3. A probabilistic finite state automata (PFSA) (Ray, 2004) is constructed from the symbol sequence, which models the transition from one state to another as state transition matrix. State transition matrix is the extracted feature which represents the sequence of images, essentially capturing the temporal evolution of coherent structures in the flame. DBN and STSA are explained in detail later in this section.

3.1. Deep Learning techniques

Deep Learning is an emerging branch of machine learning with a strong emphasis on modeling multiple levels of abstraction (from low-level features to higher-order representations, i.e., features of features) from data (Deng & Dong, 2014; Bengio, Courville, & Vincent, 2013). For example, in a typical image processing application while low-level features can be partial edges and corners, high-level features may be combination of edges and corners to form parts of an image.

Among various deep learning techniques, Deep Belief Networks (DBNs) have become an attractive option for data dimensionality reduction (G. E. Hinton & Salakhutdinov, 2006), collaborative filtering (Salakhutdinov, Mnih, & Hinton, 2007), feature learning (Coates, Ng, & Lee, 2011), topic modeling (G. E. Hinton & Salakhutdinov, 2009), and solving classification problems (Larochelle & Bengio, 2008). Several other deep learning architectures such as Convolutional Neural Networks, Stacked Denoising Autoencoders, and Deep Recurrent Neural Networks have also gained immense traction recently as they have been shown to outperform all other state-of-the-art machine learning tools for handling very large dimensional data spaces to learn features in order to perform detection, classification and prediction. The basic building block of DBN is the Restricted Boltzmann Machine (RBM), where multiple RBMs are stacked on top of another to form a deep network. An RBM is essentially a generative probabilistic graphical model that is capable of learning a probability distribution over the inputs to best explain the observed data. Individual RBMs consists of visible units (the inputs) which are connected to latent variables in the hidden units. Note that connections exist only between the visible layer and the hidden layer but not among visible units and hidden units– hence termed *Restricted*. While a single layer of RBM is already quite powerful to represent complex distributions, increasing the number of hidden layers greatly improves modeling capacity where the output of one hidden layer becomes the input of another placed over it.

Deep Belief Networks can be trained in an unsupervised greedy layer-wise manner. In simpler terms, the first RBM layer is trained with the raw input as the visible layer. During training, the first layer acquires a representation of the input by updating its weights and biases between the visible and hidden layers (usually through computing the mean activations or by sampling) which in turn becomes the input of the second layer (G. Hinton, Osindero, & Teh, 2006). The objective during layer-wise training is to find the weight vector $\boldsymbol{W}$ (and biases for both visible and hidden units) that maximizes the expected log likelihood of the training data $\boldsymbol{V}$ (Fischer & Igel, 2014). More formally, the optimization problem can be represented (ignoring the biases) as:

$$\arg\max_{W} \mathbb{E}\left[\sum_{v \in V} \log P(v)\right]$$

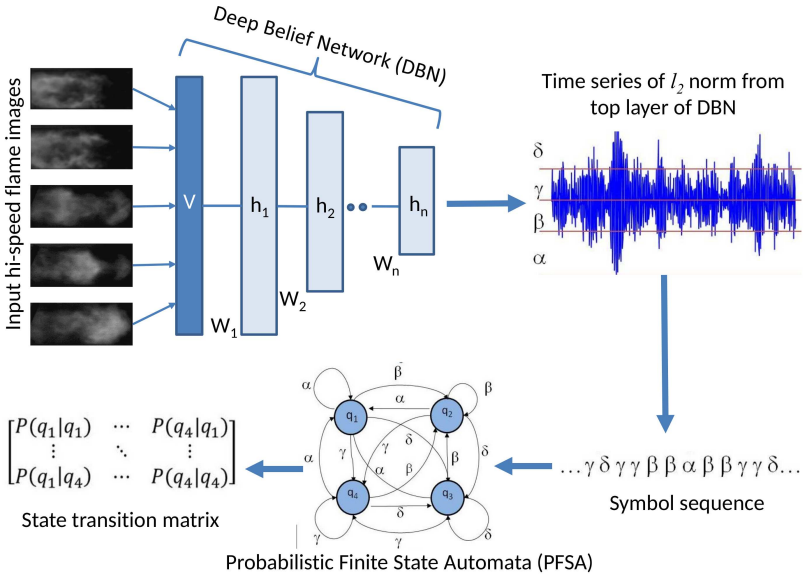

Figure 3. Framework for early detection of combustion instability from hi-speed flame images via semantic feature extraction using deep belief network (DBN) followed by symbolic time series analysis (STSA)

Typically, the optimization is solved in a gradient descent manner. Keeping the weights and biases of the first layer constant after it is trained, the transformed input from the layer is utilized to train the next layer. This process is repeated for the desired number of layers in the network with each iteration propagating either the samples or mean activations to higher levels. As training continues, the product of probabilities assigned to the input is maximized. Once all the layers are trained, the *pre-trained* model is finetuned via supervised backpropagation. It is important to note that layer-wise training helps with initializing weights and biases in the network prior to the actual supervised training. Taking classification as an example, a logistic classifier is used to classify the input based on the output of the final hidden layer of the DBN. A predefined error metric is computed between the class labels and the resultant output of the DBN (after applying the logistic classifier) and then the error is backpropagated down the network to further adjust and optimize the weights and biases.

Visualization of Learned Features

One of the main claim of a hierarchical semantic feature extraction tool such as DBN is that it learns meaningful patterns in the data that can signify the underlying characteristics of the process. Therefore, visualizing the learned features is crucial to both understand and verify the performance of the feature extractor. Furthermore, intermediate feature visualization may lead domain experts to scientific discoveries that are not easy to figure out via manual exploration of large volume of data.

For the lowest RBM layer, simply plotting the weight matrix may be sufficient to visualize the features learned by the first hidden layer. Since the dimensionality of the input and the weights are in the same order, the vectors of weights for each input can be reshaped into the dimension equal to the resolution of the input image. Thus, the visualizations are usually intelligible. Complexity arises for visualizing features learnt at deeper layers because they lie in a different space from the visible data space. At the same time, the dimension of weight matrix depends on the number of hidden units between the layer and the layer before. Thus, plotting the weight matrix will result in an incomprehensible visualization which typically resembles the appearance of white noise. To obtain filter-like representations of hidden units in the DBN, a recent technique known as Activation Maximization (AM) is used (Erhan, Courville, & Bengio, 2010). This technique seeks to find inputs that maximize the activation of a specific hidden unit in a particular layer and the technique is treated as an optimization problem. Let θ denote the parameters of the network (weights and biases) and $h_{ij}(\theta, \boldsymbol{x})$ be the value of the activation function $h_{ij}(\cdot)$ (usually the logistic sigmoid function) of hidden unit i in layer j on input $\boldsymbol{x}$. Assuming the network has been trained, θ remains constant. Therefore, the optimization process aims to find

$$\boldsymbol{x}^* = \underset{\boldsymbol{x} s.t. ||x||=\rho}{\operatorname{argmax}} \; h_{ij}(\theta, \boldsymbol{x})$$

where $\boldsymbol{x}^*$ denotes the inputs that maximizes the hidden unit activation. Although the problem is a non-convex optimization problem, it is still useful to find the local optimum by

performing a simple gradient ascent along the gradient of $h_{ij}(\theta, \boldsymbol{x})$ because in many cases, the solutions after convergence are able to visualize the patterns of the inputs that are being learned by the hidden units.

3.2. Symbolic Time Series Analysis (STSA)

STSA is a fast time series feature extraction tool that models the temporal evolution of a quasi-stationary time series via symbolization (Ray, 2004).The algorithms of STSA are formulated via symbolization of the time series generated from dynamical systems along with subsequent state machine construction. First, the time series data are partitioned by maximum-entropy partitioning (MEP) (Rajagopalan & Ray, 2006) to construct the symbol alphabet Σ for generating symbol sequences. MEP maximizes the Shannon entropy (Cover & Thomas, 2006) of the symbol sequence via generating more partitions at the information-dense zones in the range domain than information-sparse zones. Once the partitions are obtained, each data point of the time series is assigned a symbol $s_i \in \Sigma$ same as the partition it belongs to.Then, a D-Markov machine, based on the algebraic structure of probabilistic finite state automata (PFSA) (Ray, 2004), is constructed from the symbol sequence. D-Markov machine is defined as follows.

Definition 3.1 *(Ray, 2004; Sarkar et al., 2014) (D-Markov) A D-Markov machine is a 4-tuple PFSA ($K = (\Sigma, Q, \delta, \pi)$), in which each state is represented by a finite history of D symbols as defined by:*

- *Σ is a non-empty finite set, called the symbol alphabet, with cardinality $|\Sigma| < \infty$;*
- *Q is the finite set of states with cardinality $|Q| \leq |\Sigma|^D$, i.e., the states are represented by equivalence classes of symbol strings of maximum length D where each symbol belongs to the alphabet Σ; D is the depth of the Markov machine;*
- *$\delta : Q \times \Sigma \rightarrow Q$ is the state transition function that satisfies the following condition if $|Q| = |\Sigma|^D$, then there exist $\alpha, \beta \in \Sigma$ and $x \in \Sigma^\star$ such that $\delta(\alpha x, \beta) = x\beta$ and $\alpha x, x\beta \in Q$.*
- *$\tilde{\pi} : Q \times \Sigma \rightarrow [0, 1]$ is the symbol generation function (also called probability morph matrix) that satisfies the condition $\sum_{\sigma \in \Sigma} \tilde{\pi}(q, \sigma) = 1 \ \ \forall q \in Q$, and π_{ij} is the probability of occurrence of a symbol $\sigma_j \in \Sigma$ at the state $q_i \in Q$.*

State transition matrix, denoted by Π ($\Pi \triangleq [\pi_{ij}], \ i = 1, 2, \cdots, |Q|, \ j = 1, 2, \cdots, |Q|$), is obtained via combining $\tilde{\pi}$ and δ. Each element of Π, π_{ij} is the probability of moving from state q_i to q_j upon occurrence of a symbol at the next time step.In this paper, depth of the D-Markov machine is chosen to be one and it results in the equality of state transition matix (Π) and probability morph matrix $\tilde{\pi}$. Depth greater than one can also be chosen via applying generalized D-Markov machine construction (Sarkar et al., 2014; Mukherjee & Ray, 2014). Π is considered as the output feature of the D-Markov machine, which represents the time-series in reduced dimension. More details on STSA can be found in (Ray, 2004; Sarkar et al., 2014).

4. RESULTS AND DISCUSSIONS

The DBN used for the study is comprised of three hidden layers with 1000, 100, and 10 hidden units for the first, second, and third hidden layer respectively. The input image has a dimension of 56×98 pixels flattened to a 1×5488 row vector. The input image segments are taken from respective images at the flame entry (right end of the images) zone after scaling the original images down by 4 times.

4.1. DBN feature visualization

For visualization, the training set consists of 54,000 training images containing 6,000 images each from 9 conditions, 9,000 validation images containing 1,000 images each from 9 conditions, and 18,000 test images containing 2,000 images each from 9 conditions. A learning rate of 0.01 is used for the gradient descent algorithm for both pre-training and supervised finetuning. Pre-training is performed in batches of 50 samples and each layer undergoes 30 complete iterations of pre-training before moving onto the next layer. During supervised finetuning, classification errors on the validation images is compared against the errors from training set as a measure to prevent overtraining the network and consequently overfitting the data. The optimized model is obtained prior to the point when the validation error becomes consistently higher than the training error in subsequent training iterations.

Figure 4 (d) shows the visualization of weights from the first layer with each tile representing a hidden unit in the layer immediately after pre-training. Values of weights connecting from all visible units to this single hidden unit are represented as pixel intensities. Panels (c), (b), and (a) visualize the input that maximizes the activation of the hidden units in the first, second, and third hidden layers respectively. As expected, the weights and the inputs that maximizes the activation of the first hidden layer are similar except that the pixel intensities are inverted. For higher layers, the network is able to capture the whole mushroom-shaped features from the input images. However, visualization for the third hidden layer (with only 10 hidden units) is not as clear due to the activation maximization algorithm converging to a non-ideal local optimum. A faint mushroom shape is still visible, however. In general, the pretrained model acquires a good representation of the input. Prominent features serving as the key to distinguishing between stable and unstable flames can clearly be seen in the visualized weight matrices.

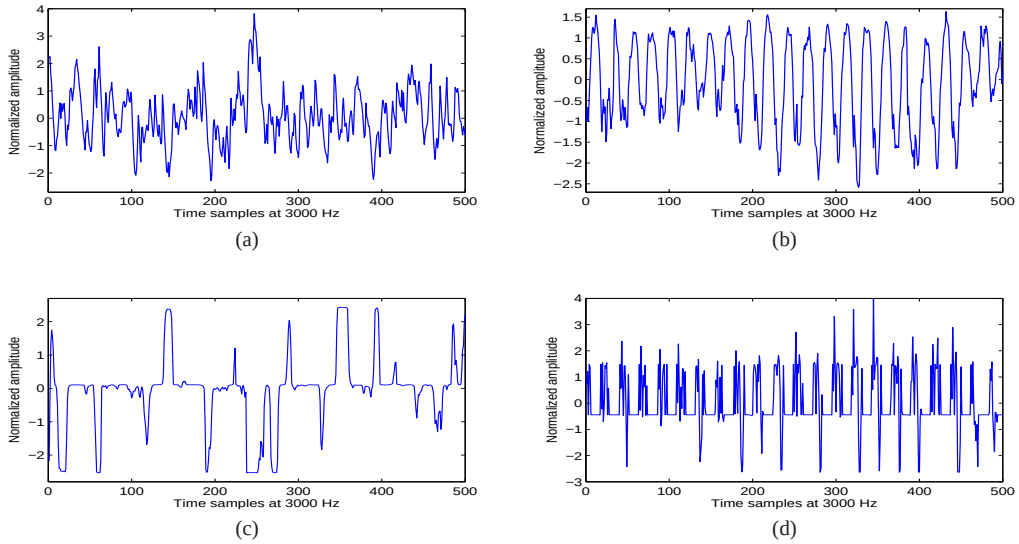

Figure 6. $0.2s$ long time series of l_2 norms of (i) 10 largest variance components of PCA performed on images at (a) stable and (b) unstable states and (ii) activation probabilities of last hidden layer after pre-training a DBN on images at (c) stable and (d) unstable states

In Figure 5, visualization of weights from the first layer and inputs that maximizes activations for all hidden layers after supervised finetuning are shown. An immediate difference can be clearly observed: visualized weights are now less noisy, whereas the third hidden layer is able to produce a visualization with more clarity compared to the weights prior to finetuning.

For both cases, the learning rate used in the AM algorithm is 0.01. Results have also indicated that depending on the initial value of the input vector, the resulting visualization from solving the optimization problem will be very different in terms of clarity. Thus, initial values of the input vectors are manually tuned by trial-and-error in order to obtain the best result. However, random initialization of the input vectors over a uniform distribution yielded undesirable results most of the time, showing images that are completely noisy without any perceivable features. Even if the results do converge, there are no significant differences between the solution from random initialization compared to the solution from tuning the initial values manually.

Remark: It is observed from the feature visualization that, though the DBN is trained on both stable and unstable flame images, the features gravitate more towards the coherent structure which is a characteristic of thermo-acoustic instability. An expert can use this feature visualization as an important tool to choose templates for unstable combustion, especially from the higher layer features. Those templates can be applied in post-processing of images to calculate the extent of instability via appropriate metrics that can effectively replace the age-old need for hand-crafted visual feature.

4.2. Performance of STSA module

In this subsection, DBN is pre-trained with 36,000 training images coming from 4 different operating conditions (see table 1) at partial premixing. Half of the training data is collected during stable combustion and other half during unstable combustion. Two sequences of images, consisting of one at stable ($Re = 7,971$, $FFR = 0.495g/s$ and full premixing) and another at unstable ($Re = 15,942$, $FFR = 0.495g/s$ and full premixing) combustion states, are reduced dimensionally via DBN with the parameters learned at pre-training phase. It is to be noted that, pre-training and testing of DBN are done on data at different levels of premixing to test the transferability of the proposed architecture.

Time series of l_2 norm of 10 dimensional activation probability vectors from each image are obtained as shown in figure 6(c) and (d). For comparison, l_2 norm of 10 largest variance components of those images, based on principal component analysis (PCA) (Bishop, 2006) coefficients learned on same training images, are constructed as presented in the top half of the figure 6. It is observed that the difference in textures of the l_2 time series between stable and unstable combustion is amplified in the case of DBN feature learning.

STSA is performed with increasing alphabet size on the l_2 time series that are mentioned above. Time series for stable and unstable combustion are partitioned separately via MEP and respective state transition matrices are calculated by the method explained in subsection 3.2. Euclidean distance between state transition matrices of stable and unstable combustion is a measure of class separability between those. The more the class separability is the more would be the precision

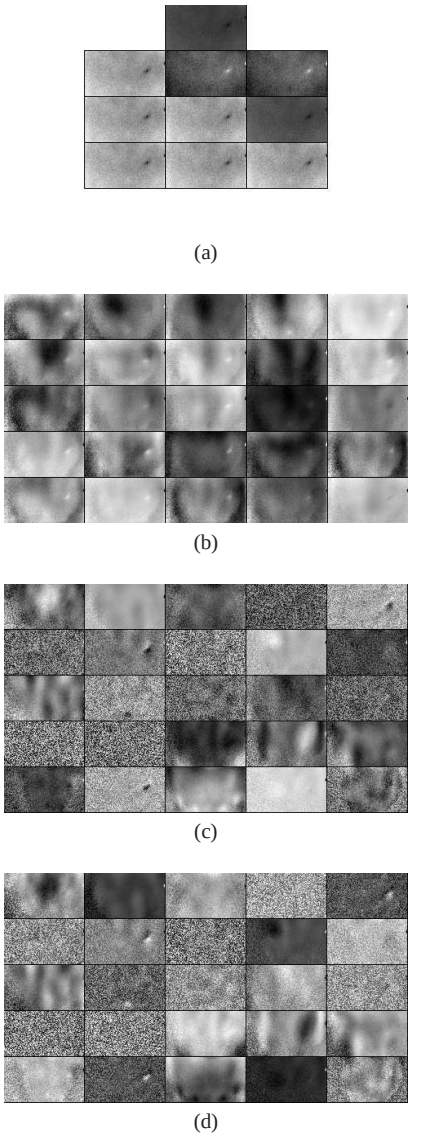

Figure 4. (d) Visualization of weights from the first layer and inputs that maximizes the hidden unit activations for the (c) 1st layer, (b) 2nd layer, and (a) 3rd layer after pre-training and prior to supervised finetuning.

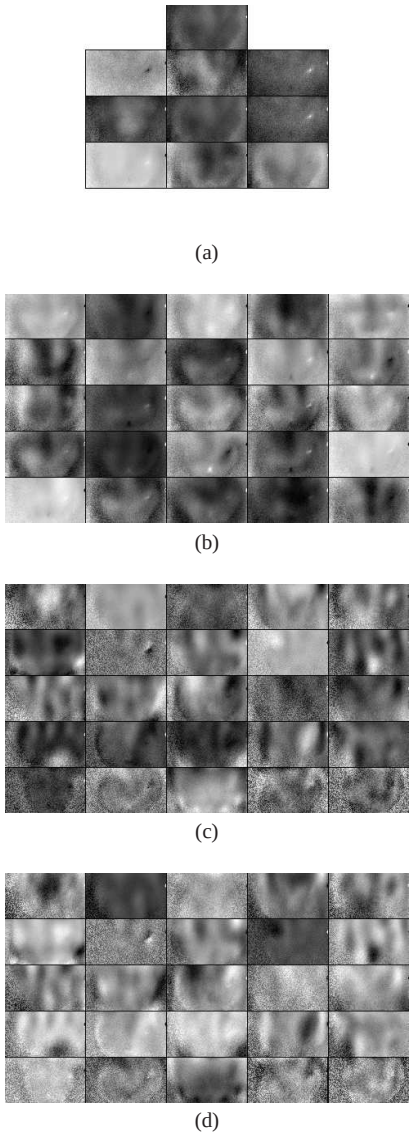

Figure 5. (d) Visualization of weights from 1st layer and inputs that maximizes the hidden unit activations for the (c) 1st layer, (b) 2nd layer, and (a) 3rd layer after supervised finetuning.

of detecting the intermediate states of the combustion while shifting from stable to unstable state. Therefore, this framework is better suited for early detection of onset of instability. It is presented in figure 7 that the class separability is much higher when STSA is applied on pre-trained DBN features than the PCA features. A probable rationale behind this observation is that, while PCA is averaging the image vector based on just maximum spatial variance, DBN is learning semantic features based on the coherent structures seen during unstable combustion. This rationale is also supported by the DBN feature visualizations that are shown in the subsection 4.1.

In a PHM context, the state transition matrix emerging from STSA module at the top can be used in supervised manner to detect instability from hi-speed image data. As the 'DBN+STSA' architecture provides a large class-separability between stable and unstable conditions, the state transition matrix can help in early prediction of thermo-acoustic instability. While the training of the proposed architecture is carried out offline in a GPU, the testing in a PHM application can be performed online with a processing power of a regular CPU. This is possible because the feed-forward computation of DBN along with STSA is feasible in real-time.

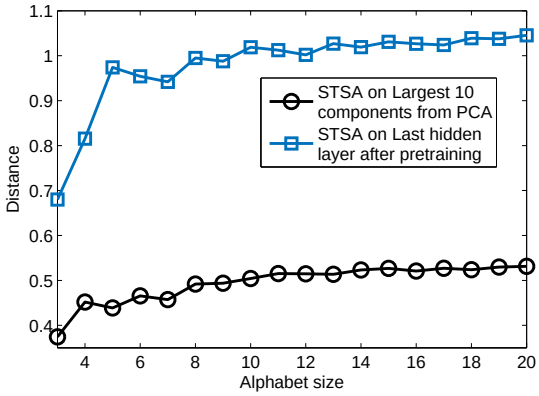

Figure 7. Variation of Euclidean distance between STSA features of image sequences from stable and unstable combustion as a function of alphabet size for STSA

5. CONCLUSION AND FUTURE WORK

The paper proposes a framework that synergistically combines the recently introduced concepts of DBN and STSA for early detection of thermo-acoustic instability in gas turbine engines. Extensive set of experiments have been conducted on a swirl-stabilized combustor for validation of the proposed method. Sequences of hi-speed greyscale images are fed into a multi-layered DBN to model the fluctuating coherent structures in the flame, which are dominant during unstable combustion. DBN hidden layers along with bottom layer weight matrix are visualized via activation maximization method and mushroom-shaped vortex are demonstrated by higher layers after, both, pre-training and finetuning stages. Although visualization after fine tuning is less noisy, it may lead to overfitting due to limitation of the data volume. Therefore, an ensemble of time series data is constructed from sequence of images based on the l_2 norm of the activation probability vectors of last hidden layer at the DBN. Then, STSA is applied on the time series that is generated from an image sequence and 'DBN pre-training+STSA' is found to exhibit more class separability with varying alphabet size than 'PCA+STSA'. More class separability between stable and unstable combustion implies more precision at detecting early onset of thermo-acoustic instability. In summary, while DBN captures the sematic features (i.e., coherent structures) of the combustion flames, STSA models the temporal fluctuation of those features at a reduced dimension.

One of the primary advantages of the proposed semantic dimensionality reduction (as opposed to abstract dimensionality reduction, e.g., using PCA) would be seamless involvement of domain experts into the data analytics framework for expert-guided data exploration activities. Developing novel use-cases in this context will be a key future work. Some other near-term research tasks are:

- Application of deep convolutional network on entire (large) flame images to model coherent structure at varying scales and orientations.
- Dynamically tracking multiple coherent structures in the flame to characterize the extent of instability.
- Multi-dimensional partitioning for direct usage of the last hidden layers for the sequence of images to the STSA module without converting it to time series of l_2 norm.

ACKNOWLEDGMENT

This work was supported in part by the Air Force Office of Scientific Research (AFOSR) under Grant No. FA9550-12-1-0270. 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 agency.

REFERENCES

Bellows, B., Bobba, M., Forte, A., Seitzman, J., & Lieuwen, T. (2007). Flame transfer function saturation mechanisms in a swirl stabilized combustor. *Proceedings of the combustion institute*, *31*(2), 3181-3188.

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, *35*(8), 1798–1828.

Berkooz, G., Holmes, P., & Lumley, J. L. (1993). The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, *25*(1), 539-575. doi: 10.1146/annurev.fl.25.010193.002543

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer, New York, NY, USA.

Chakravarthy, S. R., Shreenivasan, O. J., Bhm, B., Dreizler, A., & Janicka, J. (2007). Experimental characterization of onset of acoustic instability in a nonpremixed half-dump combustor. *Journal of the Acoustical Society of America*, *122*, 120127.

Coates, A., Ng, A. Y., & Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. In *International conference on artificial intelligence and statistics* (pp. 215–223).

Cover, T., & Thomas, J. (2006). *Elements of information theory, 2nd ed.* New York, NY, USA: Wiley.

Deng, L., & Dong, Y. (2014). Foundations and trends® in signal processing. *Signal Processing*, *7*, 3–4.

Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., & Bengio, S. (2010). Why does unsupervised pre-training help deep learning? *The Journal of Machine Learning Research*, *11*, 625–660.

Erhan, D., Courville, A., & Bengio, Y. (2010). Understanding representations learned in deep architectures. *Department d'Informatique et Recherche Operationnelle, University of Montreal, QC, Canada, Tech. Rep, 1355*.

Fischer, A., & Igel, C. (2014). Training restricted boltzmann machines: An introduction. *Pattern Recognition*, *47*(1), 25–39.

Hinton, G., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation*, *18*(7), 1527–1554.

Hinton, G. E. (2009). Deep belief networks. *Scholarpedia*, *4*(5), 5947.

Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, *313*(5786), 504–507.

Hinton, G. E., & Salakhutdinov, R. R. (2009). Replicated softmax: an undirected topic model. In *Advances in neural information processing systems* (pp. 1607–1614).

Hussain, A. K. M. F. (1983). Coherent structures - reality and myth. *Physics of Fluids*, *26*(10), 2816-2850.

Larochelle, H., & Bengio, Y. (2008). Classification using discriminative restricted boltzmann machines. In *Proceedings of the 25th international conference on machine learning* (pp. 536–543).

Mcmanus, K. R., Poinsot, T., & Candel, S. M. (1993, July). A review of active control of combustion instabilities. *Progress in energy combustion science*, *14*(1), 1-29.

Moeck, J. P., Bourgouin, J. F., Durox, D., Schuller, T., & Candel, S. (2012). Non-linear interaction between a precessing vortex core and acoustic oscillations in a turbulent swirling flame. *Combustion and Flame*, *159*(8), 26502668.

Mukherjee, K., & Ray, A. (2014, November). State splitting and state merging in probabilistic finite state automata for signal representation and analysis. *Signal Processing*, *104*, 105-119.

Mukhopadhyay, A., Chaudhari, R. R., Paul, T., Sen, S., & Ray, A. (2013). Lean blow-out prediction in gas turbine combustors using symbolic time series analysis. *Journal of Propulsion and Power*, *29*(4), 950–960.

Noiray, N., Durox, D., Schuller, T., & Candel, S. (2008). A unified framework for nonlinear combustion instability analysis based on flame describing function. *Journal of Fluid Mechanics*, *615*, 139-167.

Palies, P., Schuller, T., Durox, D., & Candel, S. (2011). Modelling of premixed swirling flame transfer functions. *Proceedings of the combustion institute*, *33*(2), 2967-2974.

Paschereit, C. O., Gutmark, E., , & Weisenstein, W. (1998). Structure and control of thermoacoustic instabilities in a gas turbine combustor. *Combustion Science and Technology*, *138*(1-6), 213232.

Rajagopalan, V., & Ray, A. (2006, November). Symbolic time series analysis via wavelet-based partitioning. *Signal Processing*, *86*(11), 3309-3320.

Ramanan, V., Chakravarthy, S. R., Sarkar, S., & Ray, A. (2014, December). Investigation of combustion instability in a swirl-stabilized combustor using symbolic time series analysis. In *Proc. asme gas turbine india conference, gtindia 2014, new delhi, india* (p. 1-6).

Rao, C., Ray, A., Sarkar, S., & Yasar, M. (2009). Review and comparative evaluation of symbolic dynamic filtering for detection of anomaly patterns. *Signal, Image and Video Processing*, *3*(2), 101–114.

Ray, A. (2004, July). Symbolic dynamic analysis of complex systems for anomaly detection. *Signal Processing*, *84*(7), 1115–1130.

Salakhutdinov, R., Mnih, A., & Hinton, G. (2007). Restricted boltzmann machines for collaborative filtering. In *Proceedings of the 24th international conference on machine learning* (pp. 791–798).

Sarkar, S., Jin, X., & Ray, A. (August, 2011). Data-driven fault detection in aircraft engines with noisy sensor measurements. *Journal of Engineering for Gas Turbines and Power-Transactions of the ASME*, *133*(8), 081602-081611.

Sarkar, S., Ray, A., Mukhopadhyay, A., Chaudhari, R. R., & Sen, S. (2014). Early detection of lean blow out (lbo) via generalized d-markov machine construction. In *American control conference (acc), 2014* (pp. 3041–3046).

Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, *656*, 5-28. doi: 10.1017/S0022112010001217

Syred, N. (2006). A review of oscillation mechanisms and the role of precessing vortex core (pvc) in swirl combustion systems. *Progress in Energy and Combustion Science*, *32*(2), 93161.