

Analyzing Rocket Plume Spectral Data with Neural Networks for Condition Monitoring

Kevin W. Whitaker

Associate Dean for Academic Programs
College of Engineering
The University of Alabama
P.O. Box 870200
Tuscaloosa, Alabama 35487-0154
USA

Abstract

Space Shuttle Main Engine fault detection systems typically rely on sensor data analysis via redundant rule-based expert systems along with visual observations for the real-time assessment of engine health. A novel alternative to the traditional health monitoring approach is predicated upon the acquisition and subsequent neural network processing of electromagnetic plume emissions. Spectrometric examination of an emission spectrum provides a means for the identification and quantification of metallic species indigenous to the main engine plume flow. Knowledge of the metallic species eroding could pinpoint the specific location of component degradation within the engine as well as identify serious component failures at an early stage. Such an approach is advantageous because it allows for the detection of numerous internal failures that would otherwise go unnoticed by traditional monitoring methods. This paper details a radial basis function neural network architecture that is capable of inferring metallic state from a given plume spectrum. Specifically, a comprehensive discussion of the methodologies necessary for the development and implementation of the neural network approach are provided. The resulting neural networks are validated with actual test-stand data from an actual Space Shuttle Main Engine at NASA's Stennis Space Center.

Nomenclature

- y = desired function for neural approximation
 x = neural network input feature pattern
 g_j = Gaussian kernel function
 μ_j = kernel function center position
 σ_j = kernel function spread constant
 w = optimal array of network weighting coefficients
 G = array of Gaussian kernel responses to input patterns
 I_i = vector of spectral intensities, ($\text{Watts} \cdot \text{cm}^{-2} \cdot \text{str}^{-1} \cdot \text{nm}^{-1}$)

Introduction

Current Space Shuttle Main Engine (SSME) operating health assessment procedures employ an expert-rule based platform with redlining capability. This approach is very effective at identifying potentially catastrophic engine failures but it cannot quantify the degree of internal engine component degradation experienced during an engine's cycling period. For example, the degree of erosion of a pre-burner faceplate or nozzle side-wall is assessed by a costly post-test, or post-flight, visual inspection. In many instances, SSME inspectors do not have prior knowledge of what engine components to check and isolate as potentially dangerous for future

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 01 JUN 2003	2. REPORT TYPE N/A	3. DATES COVERED -	
4. TITLE AND SUBTITLE Analyzing Rocket Plume Spectral Data with Neural Networks for Condition Monitoring		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Associate Dean for Academic Programs College of Engineering The University of Alabama P.O. Box 870200 Tuscaloosa, Alabama 35487-0154 USA		8. PERFORMING ORGANIZATION REPORT NUMBER	
		10. SPONSOR/MONITOR'S ACRONYM(S)	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
		12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited	
13. SUPPLEMENTARY NOTES See also ADM001519. RTO-EN-022, The original document contains color images.			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	UU
			18. NUMBER OF PAGES 16
			19a. NAME OF RESPONSIBLE PERSON

engine operations. This paper offers an attractive alternative to the traditional health monitoring approach. It is a method that will allow for the isolation of SSME problem areas and also provide crucial real-time information that can be used to assess the internal health of the SSME. Specifically, radial basis function neural networks are used to approximate the underlying functional mapping between plume emission spectra and metallic species content. The resulting neural platform is capable of providing the metallic concentrations contained in the SSME plume.

For years, researchers at NASA's Marshall Space Flight Center (MSFC) have been filming developmental tests of the Space Shuttle Main Engine. Subsequent analysis of film that involved major engine failure incidents revealed a consistent feature common to most of the breakdowns. In 8 of 27 failure events observed, a visible discharge of some substance was seen in the plume prior to the engine failure.¹⁻³ These discharges ranged from extreme flashes to small regional streaks. The discovery suggested an interesting approach to SSME health assessment, namely, if it is possible to visually detect anomalous events with the naked eye, then perhaps it would be possible to discover anomalous events below the "visible threshold" (infrared, ultraviolet, etc.) that would be indicative of an imminent failure.^{1,2} This would be tantamount to real-time engine health monitoring via the quantitative analysis of the SSME exhaust plume spectral data.

The idea of extracting chemical data from the analysis of the electromagnetic (EM) spectrum is not new. Holding a copper wire in a sufficiently hot flame produces a characteristic green region in the flame. The copper atoms are excited to such a high energy state that they emit electromagnetic radiation at several wavelengths, with green light being dominant. The atomic structure determines the wavelengths of EM emitted, and since all elements are unique, no two elements will emit EM at exactly the same wavelengths. Thus, each element has its own unique spectral signature. For example, in the same flame, the element nickel will emit EM at wavelengths different than copper. A plot of radiant intensity versus wavelength is called the electromagnetic spectrum and an example of a typical spectrum for nickel is shown in Fig. 1. Since this spectrum is unique, a spectrometric detector some distance from the flame would allow a user to determine the presence of nickel, or any other material for that matter, if its spectrum was known.

Individual spectra for other elements vary in complexity, some having a few atomic transitions (peaks), others having many. Three germane points of importance result from considering Fig. 1: 1) *Every element has its own "spectral signature,"* 2) *The emission will contain atomic transitions at wavelengths which may not be part of the visible spectrum,* and 3) *The **intensity** of the emission is a function of the **quantity** of emitting matter present in addition to the system temperature and other quantum variables.*

Rocket plumes are emissive events subject to the same physics (with more complications of course) as burning nickel or copper in an open flame. The Optical Plume Anomaly Detection (or OPAD) program was initiated by researchers at MSFC as an effort to take advantage of the wealth of information contained in the exhaust plume of a rocket engine. The initial idea was to identify anomalous spectral events which were consistent with known mechanical failures and then use them as templates in the health monitoring of future engine tests, ground or in-flight. These spectral templates could then be coupled with the anomalous events found in the vibrational and other sensor data to determine the overall state of the engine.¹

The Technology Test Bed (TTB) at MSFC and the A1 Test Stand at NASA's Stennis Space Center (SSC) were the first to receive an OPAD instrument pack developed by combustion phenomenologists at the U. S. Air Force Arnold Engineering Development Center (AEDC). Soon thereafter, the process of building a cumulative database of the spectral templates began. Researchers wanted to catalogue the various spectral forms associated with changes in SSME operation (*i.e.*, changes in oxygen/fuel ratio, engine startup, etc.) so that a baseline of "expected" spectral signatures would be established. The "template idea," however, soon gave way to even more ambitious goals as a result of some initial findings in the TTB experimental program.^{1,3} Specifically, health monitoring now involved the simultaneous tasks of anomaly detection and metallic species quantification (see Fig. 2). This meant that the free atom densities of all the metals of interest within the engine would have to be inferred for every spectral scan taken by the instruments. The quantification process

would essentially give metal concentration versus time and thereby highlight any anomalous events indicative of a major failure.

Simply stated, the capability to detect anomalies and quantify metal erosions would be an enormous benefit to the propulsion engineering community. It could reduce costly engine servicing by specifying only those times that it would be needed. Moreover, it could pinpoint the specific engine component that needs attention (anomaly isolation). For example, post-test analysis of OPAD spectral data from one particular TTB firing revealed the erosion of a metal indigenous to the pre-burner faceplate; indeed, a subsequent analysis of the pre-burner by the engine group at MSFC found that the faceplate was in serious need of replacement.¹ Finally, the OPAD system could assist in the real-time health monitoring of launch vehicles. If an engine is operating in an “off nominal” condition, potentially jeopardizing the success of the mission or safety of the crew, the OPAD detection system could order an engine shutdown or mission termination.

The cursory overview of the OPAD capabilities given above says nothing of the complexity of the subsystems necessary to implement it. Specifically, metal quantification from an electromagnetic spectrum requires a computational model of the exhaust plume emission physics. Spectrometer hardware for the test stand, as well as the in-flight environment, had to be developed for the spectra acquisition. Lastly, methods for implementing the spectral model to allow metal determination would need to be integrated as well as validated. The central theme of the work presented here, however, will be the development of artificial intelligent systems that will perform real-time determination of the metallic species concentrations present in the exhaust plume. Specific objectives include: 1) Establishing appropriate techniques for feature extraction from a spectrum, 2) Identifying robust data transformation methods for accurate neural network performance, 3) Creating a neural architecture capable of quantifying the metallic content of the plume flow, and 4) Creating a neural platform that can perform the spectral mappings in “real-time.”

Spectral Investigation of Emissions

Any spectrum obtained with the OPAD instrumentation is composed of three components: 1) a dominant OH component which arises from the burning of dissociated hydrogen radicals, 2) a background noise component caused by the scattering of background light, and 3) a metallic component, if indeed there is one, which would be indicative of a metal erosion. Thus, the quantification of a metal erosion and the subsequent identification of any anomalies requires a spectral “cleaning” procedure followed by an evaluation of the plume metallic state. In other words, methods for removing the OH and background components of the spectrum would need to be employed so that the underlying metallic component could be seen. Then the metal quantities would have to be ascertained from the remaining metallic component in the spectrum.

For a given spectrum, ascertaining the metallic quantity could only be done through two methods: 1) by comparing the spectrum to past spectra obtained from plume seeding tests, or 2) using a theoretical model that emulates the emissive nature of the plume. The first option is plagued by an inability to precisely measure the erosion and survivability rates of the inserted species. Thus, there would exist plume spectra to compare to, but the metallic content associated with the spectra would be in error. Moreover, it would not be cost effective to run the SSME through all the possible metallic seeding combinations. For these reasons, option two was selected.

Spectral models have been in use since the late 1960’s at AEDC. Most of these models measured flow-field species concentrations using techniques such as emission resonance absorption, electron beam impact excitation, and most recently, Laser Induced Fluorescence (LIF). There are many excellent validated radiative transfer codes available for specific molecular species, but none have been developed that can handle the comprehensive analysis of atomic spectral data. To meet this need, researchers at Vanderbilt University and AEDC developed a theoretical plume model known as SPECTRA.^{4,5} The SPECTRA model can provide quantitative plume spectral analysis for emission, absorption, and LIF configurations installed at the rocket nozzle exit plane. Validation of this code is still ongoing, but carefully controlled SSME plume seedings at the TTB have shown the model to be consistent.⁴

The forward operation of the SPECTRA code involves the calculation of a theoretical plume spectrum from a pre-defined set of metal concentrations and flow parameters. The reverse operation of SPECTRA, namely, obtaining the metallic components which made up the spectrum, cannot be written in a numerically convenient form because of the insurmountable mathematics involved. This, therefore, mandates that the SPECTRA code be applied in an iterative manner until it converges with the spectrum obtained from the OPAD instrumentation. The set of SPECTRA input parameters which produced this convergence would then specify the current metallic state of the plume. Accordingly, higher order numerical fitting routines have been developed for such a purpose and the resulting spectral fits obtained using these codes are very accurate.^{6,7} However, the time necessary to process a single spectral scan is on the order of hours while the demands of real-time OPAD require processing a scan every half second.⁷ For this reason, neural network techniques have been investigated which can provide an initial estimate of metal concentrations in the rocket plume. These estimates could then be used for evaluative purposes in a real-time environment or as a starting point in post-test analysis for more accurate fitting routines, thereby significantly reducing the amount of computational effort.

The OPAD Neural Network Approach

A classic use of neural networks is function approximation. Specifically, for the OPAD case, the function would be the inverse of the theoretical SPECTRA model. Using large amounts of spectral data generated from SPECTRA, it should be possible to train a neural network to learn this inverse mapping. Subsequently, when presented with a spectrometer scan, the neural network could estimate the metallic content of the plume in real-time. (Predictions in under 0.5 sec was set as a goal for this study.)

Figure 3 supplies an overview of the multiple neural network architecture developed in this investigation. Notice that there is one neural network for each metallic species being monitored. With this setup, the steps in assessing the SSME internal health are as follows: 1) Acquire a spectrum from the plume emission, 2) Extract information about the spectral signatures of the individual metallic species being monitored, 3) Give this information to the appropriate metal neural network for a prediction of the metal's number density, 4) Monitor all the metallic erosions in real-time and identify anomalous emissions, and finally 5) From the list of eroding metals, identify the failing engine component and severity of the erosion.

There can be found numerous types of neural network architectures in the technical literature. For this problem, however, the radial basis function neural network (RBFNN) was chosen because of its ability to approximate highly nonlinear functions (precisely what the inverse of SPECTRA is) very quickly. The means by which these RBFNN's are trained and developed are the topics of the next section.

Radial Basis Function Neural Networks

The RBFNN can estimate a function $y(x)$ after training with a set of representative input/output pairings. The pairings in the present case would be sets of spectra and their corresponding metallic states. From a neural network point of view, this training data could come from either the SPECTRA code or emissions obtained during a test-stand firing.

With examples of input/output patterns, the RBFNN can adjust its internal parameters so that it approximates the functional mapping to within some degree of error. Qualitatively, the RBFNN does this by forming *localized* "bumps" or response regions within the input space. The superposition of these local response regions forms a response surface that spans the space covered by the input training patterns.

By definition, a radial basis function (RBF) is one which decreases (or increases) monotonically away from a central point thereby giving it an inherent "bump" form. Classic functions that exhibit this propensity are the Gaussian, Cauchy, and the Inverse Multiquadric.⁸ As an example, consider the Gaussian function given by Eq. (1):

$$g_j(\vec{x}) = \exp\left[\frac{-(\vec{x} - \vec{\mu}_j)^2}{2\sigma_j^2}\right] \quad (1)$$

The RBFNN positions a collection of these RBFs throughout the space covered by the input training patterns. The parameter μ_j specifies the location of a single RBF within the input space (μ_j has the same dimension as the input vector \vec{x}) and the parameter σ_j determines the width of the local function. Thus, a given RBF will be centered at μ_j within the input space and have a “receptive field” which is proportional to σ_j . Moreover, it will give a maximum response for input vectors, \vec{x} , which are nearest the RBF center, μ_j .

With a developed approximation surface, the RBFNN estimates an output for an incoming input case by first evaluating each of the RBFs (in other words, determining where the input vector lies on the approximation surface) and then forming a weighted linear summation of their responses. The difficulty arises not from the logical evaluation of an input but rather the establishment of the network parameters, namely center positions (μ_j), RBF widths (σ_j), and output layer weighting coefficients.

The construction of an RBFNN is accomplished in a two part learning scheme known as hybrid learning (see Fig. 4). The initial forward connections of the network contain the RBF centers, μ_j , obtained through unsupervised assimilation, followed by an output layer of weighting parameters, w_j , formed through supervised instruction. Training in the unsupervised mode is done without a pre-defined learning goal; input categorization and learning must be done using correlations within the input training data in contrast to feedback from a teacher or critic. For the RBFNN, the learning scheme essentially clusters the inputs and specifies where to position the RBF centers so that the desired response coverage is obtained. Thus, via unsupervised learning, the RBF center positions (the μ_j) are chosen *a priori* and remain fixed throughout the output layer training. Methods for determining these positions as well as the RBF widths have been previously investigated.⁹ The rearward connections, composing the output layer of the network in Fig. 4, specify the weighting (or regression) coefficients which are trained in a supervised fashion. Supervised means that learning is based on comparison of the network output with the known “correct” answers.

For an RBFNN with a single layer of Gaussians, given that the RBF centers (μ_j) are fixed, the optimal weight array for the output connections which gives the best functional mapping can be found using the least squares normal equation developed in multiple linear regression theory. Details of this procedure can be found in Ref. 8; the results are simply stated here. For a set of training input vectors, \vec{x} , with corresponding RBF centers μ_j , there will be an array of Gaussian neuron responses, \mathbf{G} . Given this, the optimal weight array can be stated as:

$$\mathbf{w} = [\mathbf{G}^T \mathbf{G}]^{-1} \mathbf{G}^T \mathbf{y} \quad (2)$$

With the weights, centers, and widths set, the fundamental mapping can then be written as:

$$f(\vec{x}) \approx \sum_{i=1}^m w_i g_i(\vec{x}) \quad (3)$$

Thus, for an input vector \vec{x} , the solution $f(\vec{x})$ is a weighted linear summation of each RBF’s response to \vec{x} . RBFs that are within the region of \vec{x} will give the largest responses, whereas those farthest away will give negligible contributions to the series solution. Moreover, the RBF neuron responses, $g_i(\vec{x})$ will be bounded between 0 and 1 with the assigned weights (w_i) specifying the neurons heights.

A quick glance at Eq. (3) reveals that the RBFNN is a weighted summation of basis functions that are “tuned” using the training data. This is a technique which is frequently used in asymptotics and series representations of functions.¹⁰

Extracting Features From the Plume Spectrum

The overall goal here is to quantify, in real-time, the amount of metal erosion as well as the plume combustion temperature. The only available information given to the RBFNN is a spectral plot procured from the OPAD instrumentation. Given this, what should the neural network inputs be? In other words, what features within the data supply the most information about the state of a particular metal? The answer to such a question is one of the fundamental discriminants that ultimately determines the success of this neural network application.

Developing neural network inputs most often requires the researcher's in-depth knowledge of both the systems involved and the underlying physical phenomena occurring. There are times, however, when a bad choice of input can lead to a new form of insight into the detailed processes of the physical problem. For example, some inputs may not be as sensitive to changes in the output as was once thought, thereby forcing an inquiry as to why and ultimately leading to some form of insight. This has certainly happened throughout the years of the OPAD neural network development.⁹

Choosing network inputs starts with a firm understanding of all the possible error sources contained within the data. For the OPAD system, the effects of the background light and OH continuum on the observed emission are two such examples. Additional errors could also be caused by spectral interference; that is, some metals may emit at nearly identical wavelengths causing interference in the actual intensities obtained at those wavelengths. It is important to note that these are external error sources and they do not include the internal errors incurred from the failure of the linear regression model to fit all the RBFNN training points. The proper choice of an input structure can never eliminate the external errors, but it can help alleviate their effects on the RBFNN performance.

With the foregoing thoughts in mind, a study was initiated to find input data transformations that would assist in countering the effects of errors associated with the instrumentation and external radiation. Consider a typical OPAD spectrum shown in Fig. 5. The training of a neural network to predict properties of element "A" is desired. Assume that all the peaks (and regions) of element "A" that are not overshadowed or significantly affected by other elements and OH emissions have been located. Further suppose that a total of 200 sampling points which comprise these regions can be extracted and formed into vectors (like \vec{I}_c for example). Based on the foregoing discussions, three things need to be accomplished: 1) use the smallest number of inputs to reduce computational cost, 2) apply a transformation to the inputs that will reduce sensitivity to errors, and 3) scale the data so that it has relatively the same variance.

Step 1 is accomplished by selecting a few emission regions for a metal which are most responsive to changes in concentration, temperature, etc., and relatively non-interfered with by other metals or the OH continuum. Traditionally, step 3 is handled with a logarithmic transformation and this has been found to be sufficient for the analysis here as well. For step 2, three transformation methods were selected for use in pilot studies. These transformations are given in Table 1. The appropriateness, and ultimate selection for use, of the three transformation methods was determined by a detailed empirical study.^{9,11} The study isolated a metal and trained three RBFNNs to analyze the metal's spectrum. Each of the three networks trained used a different data transformation. The results found that the subinterval energy method gave the best performing network in light of all the external error sources.

Results and Discussion

The first step in the development of a RBFNN involves the creation of n -sets of input/output pairings. To do this, the three parameters for a chosen metal are randomly varied over specified ranges within the SPECTRA code and the resulting spectra are stored. The parameter ranges were set by combustion phenomenologists as the most likely to be observed during a firing of the SSME.⁵ Table 2 details these parameters and their typical ranges appropriate for the SSME.

Transformation Method	Operator
Area Input	$\int(\vec{I}_n) d\lambda$
Maximum Intensity	$\max(\vec{I}_n)$
Subinterval Energy Method	$\ \vec{I}_n\ $

Table 1 Possible transformations for a given input vector.

Flow Field Parameter	Range
Mach Disc Temperature [K]	2500 - 3200
Number Density	1.0E09 - 5.0E14
Broadening Parameter	0.0 - 1.2

Table 2 SPECTRA flow field parameter ranges.

Keeping track of these combinations and the associated spectra produced provides the RBFNN with the necessary sequence of input/output pairings. The number of training sets to generate varies depending on the specific application. For the OPAD project, RBFNNs were trained with 200, 400, and 800 training sets. During training, it was observed that significant increases (on the order of hours) in training time and network complexity accompanied the use of large training sets. The problem was further augmented by the SPECTRA costs incurred: the creation and storage of training data. The constraints of the OPAD project demand real-time, low computational cost software that can be developed in a short time. Furthermore, adding additional training sets (such as the 600 and 800 runs) did not increase the accuracy of the predictions. For these reasons, only moderately large (on the order of 200) training sets were used in the development of RBFNNs. If these are uniformly generated across the parameter ranges, then enough information is provided to allow the RBFNNs to interpolate.

Once the RBFNNs are trained for the respective metals, the question then becomes, how good are the networks at accomplishing the task of predicting the number density, broadening parameter, and temperature of a particular metal from a given test-stand spectrum? Additional sets of test data can be created with the SPECTRA code and used to validate the RBFNNs. The data is easily obtained and can be quite useful for discovering trends and testing new ideas. As a validation agent, however, it does not provide a definitive statement as to the performance of a network within the “real world.” To accomplish this, the SPECTRA code would require a complete modeling of the absorption phenomena, the OH continuum, variable wavelength dispersion effects and many other unforeseen physical influences. To use actual test-stand data for validation, however, would require knowledge of what is in the flow, (something that is unknown), so that comparisons can be made with the network’s predictions. However, test-stand data can provide validation by other means as will be seen.

In January of 1996, the engine group at SSC was testing the new Pratt and Whitney SSME turbo-pump when an observer noticed a visible object being ejected from the engine accompanied by numerous significant emissive events. Immediate shutdown was ordered to preclude a possible engine failure. Post-test analysis found that the turbo-pump had failed. Further analysis of all the video tape footage, vibrational data, and other instrumentation identified five specific moments during the test where major anomalous events had occurred. Had these anomalies been detected ahead of time, the engine could have been shutdown and the destroyed engine components salvaged.

With this information, OPAD researchers at MSFC attempted to identify the same anomalies within the spectral data. Using the SPECTRA code, the researchers *iteratively* fit many of the spectral scans obtained from the SSC test stand and found major anomalous events in the spectra at precisely the same points in time as indicated by the other data sources.⁷ It took the SPECTRA fitting routines two hours to fit a *single* spectral scan taken in time. There were over a thousand scans for the 500 second SSME test, thereby making the job extremely arduous.

As a validation of the RBFNNs, seven metals were selected by analysts at MSFC as the most important to monitor in terms of engine component makeup and erosive materials. They are chromium, aluminum, nickel, manganese, silver, cobalt, and iron. All the scans taken during the SSME failure at SSC were then given to RBFNNs trained to monitor the selected metals. It took the RBFNNs less than two minutes on a Silicon Graphics Indigo to make predictions on a thousand scans for the seven metals. The improvement of computational time was significant in comparison to the two hours required for the iterative fitting of a single scan using the SPECTRA routines.

Figure 6 gives the RBFNN predicted number density trace for each metal. For comparison, the major anomalous events found by the engine group test groups at MSFC and the Stennis Space Center via the alternate data sources (primarily vibrational and visual data) occurred at the test times given in Table 3. Figure 6 demonstrates that at every one of the points listed in Table 3, the RBFNN discovered an anomalous erosive event. (These events are annotated on Fig. 6). The 531 sec event marked the total failure of the SSME turbo-pump and the metallic emissions were much higher than those found in Fig. 6. Therefore, the plots were truncated prior to 531 seconds so the anomalies leading up to the failure could be viewed.

Event 1	Event 2	Event 3	Event 4	Turbo-pump Failure
130-131 [sec]	276 [sec]	283 [sec]	404-405 [sec]	531-554 [sec]

Table 3 Test times where the most erosive anomalous events occurred for the January 1996 test of an SSME at the Stennis Space Center.

The predictions of the RBFNNs duplicated previous conclusions drawn from all the other data sources and this was accomplished in under two minutes. Recall that SSME health monitoring via plume spectral assessment is predicated upon the following tripartite evaluation scheme: anomaly identification, quantification, and isolation. The results illustrated by Fig. 6 explicitly *validate* the neural network architecture developed herein as a real-time anomaly identifier.

The question of the RBFNN's ability to *quantify* the amount of metal in the spectrum will be addressed momentarily. As for anomaly isolation, such as identifying a failed component, this task would not be one for the RBFNN alone. Isolation requires the simultaneous use of the network predicted quantities and knowledge

of the distribution of the various metals in the components throughout the SSME. To date, a database such as this has yet to be established.

In an attempt to validate the anomaly quantification ability of the RBFNNs, a formal investigation was undertaken. There are three main sources of error in the RBFNN estimations: 1) Uncertainties in the calibration of the spectrometers, 2) Uncertainties caused by the inability of the SPECTRA model to completely account for all the flow physics, and 3) Uncertainties in the inability of the RBFNN model to completely approximate the underlying physical function. The study is ongoing, but the initial results indicate approximate errors of 2% for the temperature prediction, 8% for the number density prediction, and 35% for the broadening parameter prediction.¹²

Further qualitative conclusions can be drawn by examining the spectral “fits” produced by the RBFNN predictions. To illustrate this, predictions for the seven metals of interest were taken from the 130 sec and 276 sec spectral events. There is not a preference of one event time over another; these two were merely selected to provide a qualitative measure of the results. Indeed, the results given here are representative of spectral fits found at other erosion events. RBFNN predictions were fed back into the SPECTRA code to obtain all the individual metal contributions to the spectrum. The resulting spectra were then compared against the *original* test-stand data scans and the wavelength regions indigenous to the various metals were isolated. Typical results of these actions are shown in Figs. 7-10. Additional spectral fits can be found in Ref. 9.

Looking at Figs. 7-10, it can be concluded that the RBFNNs are capable of quantifying metallic content in the plume. (Note that an “exact fit” would be when the predicted and test-stand data curves are identical.) The fits for nickel, iron, and chromium seem to be the best. Cobalt is nestled in-between a series of nickel peaks, but with closer examination it can be seen that the RBFNN’s predictions matched very well with most of the cobalt peaks within the scan. The nickel prediction was slightly higher than that observed in the test-stand data; this could probably best be explained by the corruption of the nickel peaks by the neighboring cobalt lines.

A statement that can be made with certainty here is that the RBFNNs accomplished the anomaly *detection* task. As for the *quantification* task, initial uncertainty analysis along with many spectral fits, similar to the foregoing fits, have demonstrated an initial capacity of the RBFNNs to correctly predict metal properties in the SSME exhaust plume.

Conclusions

SSME health monitoring via plume spectral assessment is predicated upon the following tripartite evaluation scheme: anomaly identification, quantification, and isolation. Recognizing anomalies in the plume spectral scans and determining their severity in real-time situations can be computationally exhaustive. For this reason, neural network techniques have been investigated that accomplish these tasks in an expeditious manner.

The SPECTRA plume emulation code determines the specific plume spectral signature for a given metal composition and flow condition. The ultimate goal of the RBFNN architecture was to provide the “inverse” of this SPECTRA operation. More specifically, it was to discover a functional mapping allowing for the inference of the plume metallic state for a given spectrum. Using data sets created by the SPECTRA code, data mining procedures were incorporated that provided means for the selection of RBFNN inputs that compressed the computational space while still adequately representing the spectral information. Trained RBFNN architectures were then developed using this data. Results showed that the networks were capable of quantifying metals contained in the SSME plume.

To validate the RBFNN architecture developed, spectral data taken from the January 1996 SSME failure at NASA’s Stennis Space Center were used. The RBFNNs were able to quickly identify anomalous events. Their ability to quantify has also been verified through uncertainty analyses and spectral fits with the test stand data.

Overall, the use of RBFNNs as real-time anomaly identifiers has been substantiated. They provide quick detection of anomalous erosions, identify which metals are eroding, and quantify the degree of erosion. This information could be used to isolate failing internal components, determine engine life, or obviate an imminent engine failure on any propulsive platform that contains some type of EM emission.

References

¹ Powers, W. T., Cooper, A. E., and Wallace, T. L., "OPAD Status Report: Investigation of SSME Component Erosion," SAE Paper 92-1030, 1992.

² Cikanek, H. A., and Powers, W. T., "Analysis of UV-Visible Spectral Radiation from SSME Plume," *Advanced Earth to Orbit Propulsion Technology - 1988*, NASA CP-30120, Vol. 2, 1988, pp.595 - 611.

³ Madzsar, G. C., Bickford, R. L., and Duncan, D. B., "An Overview of In-Flight Plume Diagnostics for Rocket Engines," AIAA Paper 92-3785, 1992.

⁴ Powers, W. T., Cooper, A. E., and Wallace, T. L., "Validation of UV-VIS Atomic Spectral Model For Quantitative Prediction of Number Density, Temperature, and Broadening Parameter," *Proceedings of the 1995 JANNAF Propulsion Systems Hazards Subcommittee*, Published by the Chemical Propulsion Information Agency, Columbia, MD, 1995.

⁵ Wallace, T. L., Powers, W. T., and Cooper, A. E., "Simulation of UV Atomic Radiation for Application in Exhaust Plume Spectrometry," AIAA Paper 93-2512, 1993.

⁶ Buntine, W. L., "OPAD Data Analysis," *Advanced Earth to Orbit Propulsion Technology - 1994*, NASA CP-3282, Vol. 1, 1994, pp. 168-177.

⁷ Cooper, A. E., Powers, W. T., Wallace, T. L., Buntine, W., "Recent Results in the Analysis of Large Rocket Engine Anomalies Utilizing State-of-the-Art Spectral Modeling Algorithms," *Proceedings of the 1997 JANNAF Propulsion Systems Hazards Subcommittee*, Published by the Chemical Propulsion Information Agency, MD, 1997.

⁸ M. J. L. Orr, "Regularisation in the Selection of Radial Basis Function Centres", *Neural Computation*, Vol. 7, No. 3, 1995, pp. 606-623.

⁹ Benzing, D. A., "A Neural Network Approach to Space Shuttle Main Engine Health Monitoring Using Plume Spectra," Master of Science Thesis, Department of Aerospace Engineering and Mechanics, Univ. of Alabama, Tuscaloosa, AL, 1996.

¹⁰ Arfken, G. B., and Weber, H. J., *Mathematical Methods for Physicists*, Academic Press, New York, 1995, pp. 191-280.

¹¹ Benzing, D. A., Whitaker, K. W., and Moore, D. C., "A Neural Network Approach to Anomaly Detection in Spectra," AIAA Paper 97-0223, 1997.

¹² Hopkins, R. C., and Benzing, D. A., "Uncertainty in Calibration, Detection, and Estimation of Metal Concentrations in Engine Plumes Using OPAD," Final Report - NASA/ASEE Summer Faculty Fellowship Program, Marshall Space Flight Center, August 1997, pp. 126-132.




Fig. 3 Multiple neural network architecture.




Fig. 4 RBF neural network architecture.

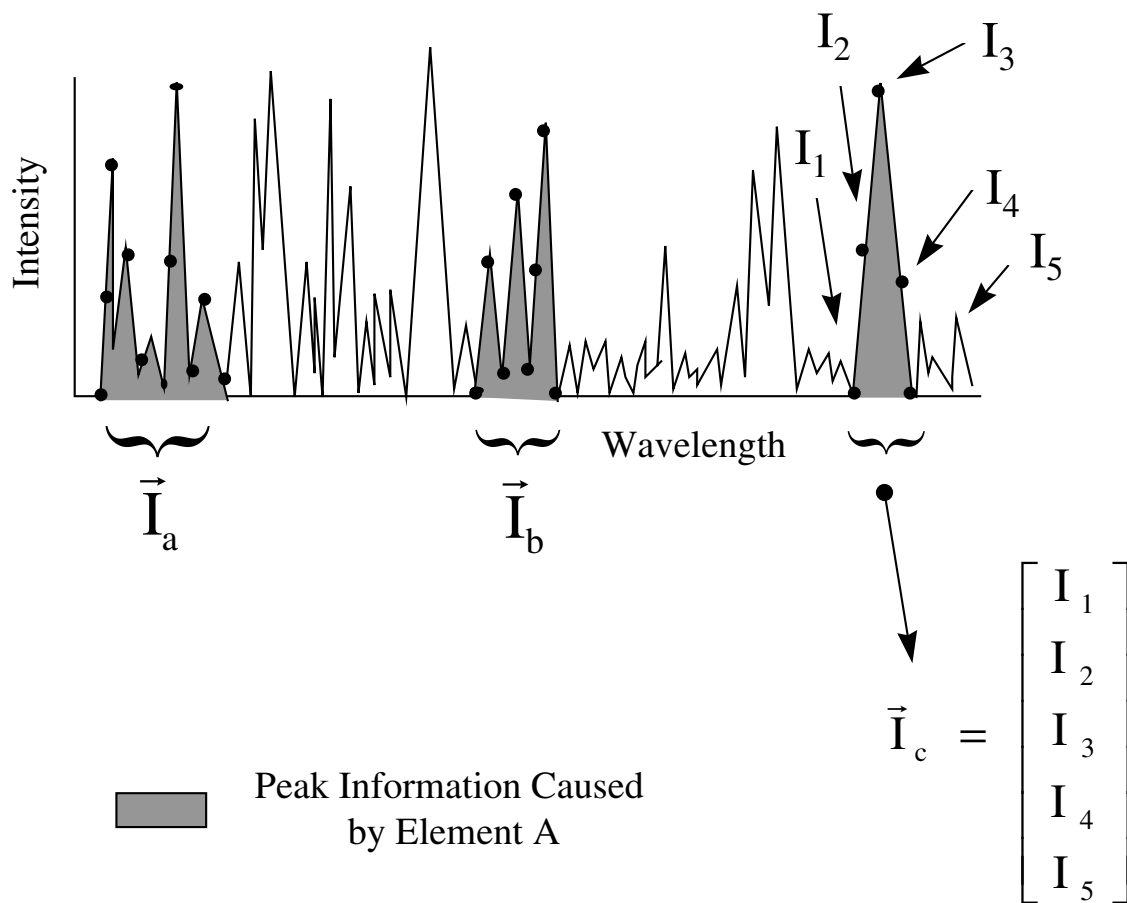


Fig. 5 Network input feature extraction.




Fig. 6 Number density predictions versus test time (actual SSME test-stand data shown).

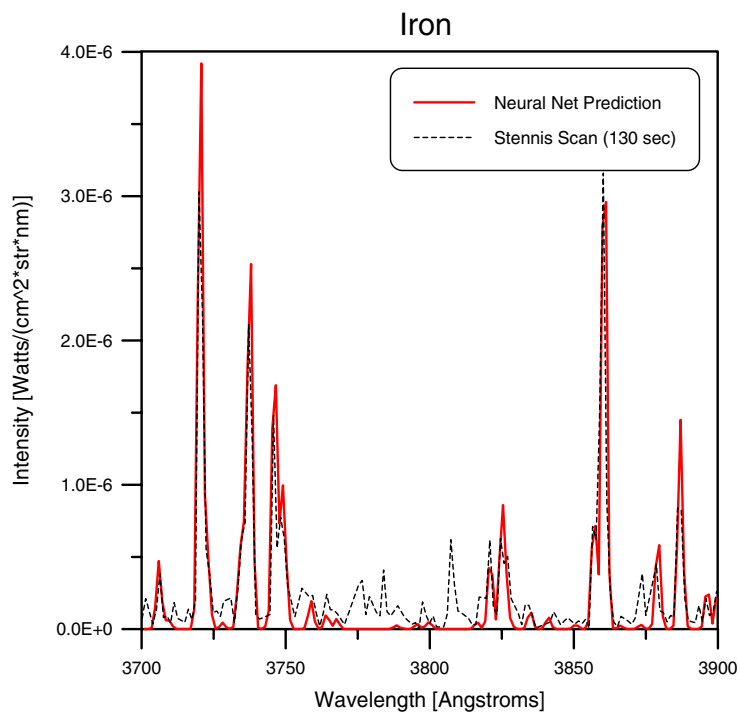


Fig. 7 Iron prediction for the 130 sec Stennis Space Center scan.

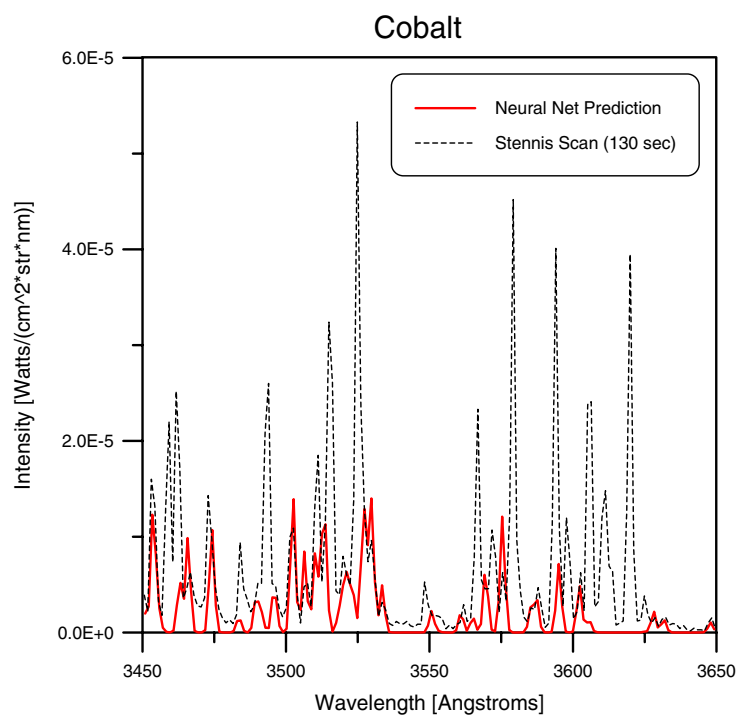


Fig. 8 Cobalt Prediction for the 130 sec Stennis Space Center scan.

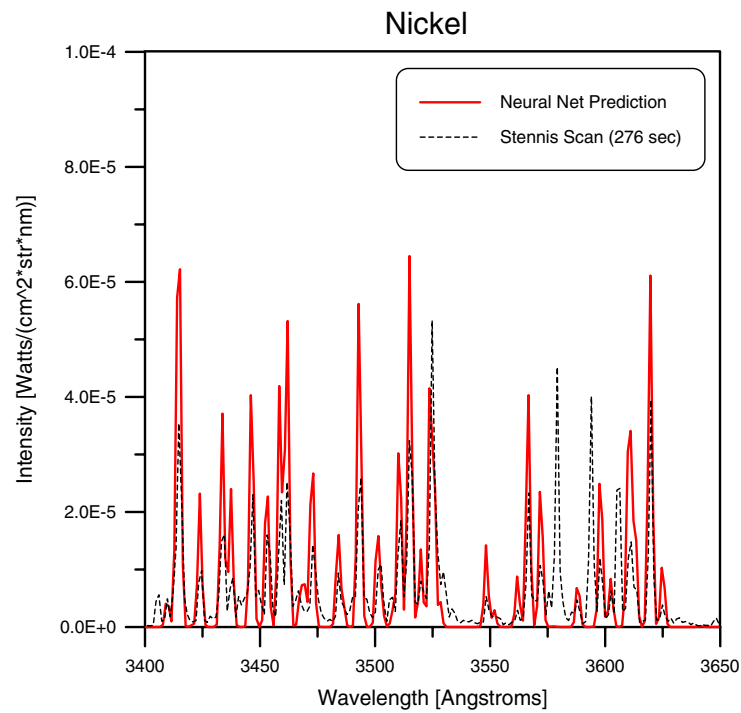


Fig. 9 Nickel Prediction for the 276 sec Stennis Space Center scan.

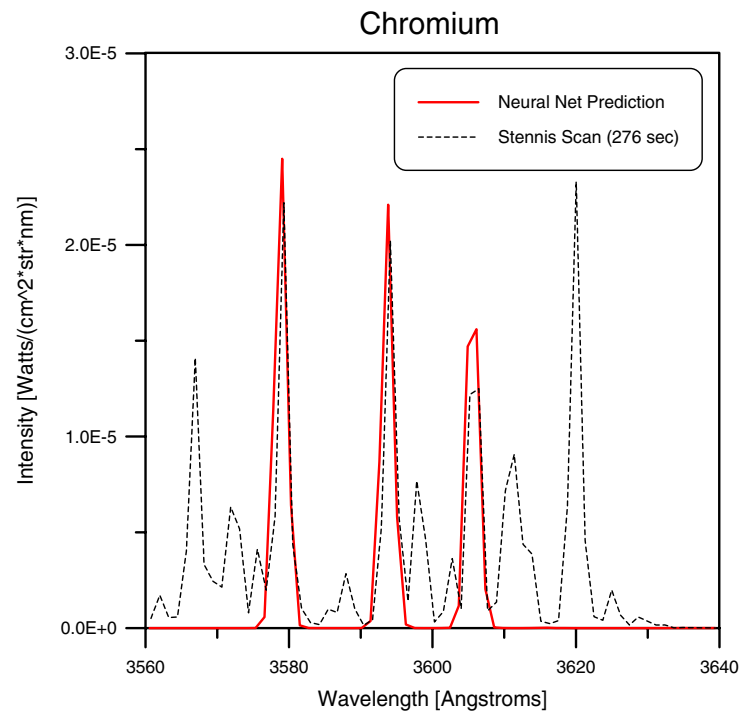


Fig. 10 Chromium Prediction for the 276 sec Stennis Space Center scan.