SANDIA REPORT

SAND2010-8890

Unlimited Release

Printed September 2010

Uncertainty Quantification of Cinematic Imaging for Development of Predictive Simulations of Turbulent Combustion

Jonathan Frank, Matthew Lawson, Khachik Sargsyan, Bert Debusschere, Habib Najm

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>

Uncertainty Quantification of Cinematic Imaging for Development of Predictive Simulations of Turbulent Combustion

Jonathan Frank, Matthew Lawson, Khachik Sargsyan, Bert Debusschere, Habib Najm
Sandia National Laboratories, P.O. Box 969, Livermore, CA, 94551-0969
{jhfrank,mlawson,ksargsy,bjdebus,hnnajm}@sandia.gov

Abstract

Recent advances in high frame rate complementary metal-oxide-semiconductor (CMOS) cameras coupled with high repetition rate lasers have enabled laser-based imaging measurements of the temporal evolution of turbulent reacting flows. This measurement capability provides new opportunities for understanding the dynamics of turbulence-chemistry interactions, which is necessary for developing predictive simulations of turbulent combustion. However, quantitative imaging measurements using high frame rate CMOS cameras require careful characterization of their noise, non-linear response, and variations in this response from pixel to pixel. We develop a noise model and calibration tools to mitigate these problems and to enable quantitative use of CMOS cameras. We have demonstrated proof of principle for image de-noising using both wavelet methods and Bayesian inference. The results offer new approaches for quantitative interpretation of imaging measurements from noisy data acquired with non-linear detectors. These approaches are potentially useful in many areas of scientific research that rely on quantitative imaging measurements.

Acknowledgment

This work was supported by Sandia National Laboratories LDRD (Laboratory Directed Research and Development) program. Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Company, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Contents

1	Introduction	9
2	Camera calibration: linearity and noise measurements	10
3	Flame measurements	15
4	Analysis	18
4.1	Bayesian inference	18
4.2	Karhunen-Loève expansion	19
4.3	Summary of Bayesian and KL methods	22
4.4	Wavelet noise suppression	23
5	Summary	29
	References	30

List of Figures

1	Reciprocity verification for the intensified CMOS camera using two different methods to measure the detector response. In one method, the irradiance is varied using neutral density filters while maintaining a constant intensifier gate width. In the other method, the irradiance is kept constant and the gate width of the intensifier is varied.	11
2	Comparison of the nonlinearity of unintensified and intensified CMOS cameras. Response curve measurements are averaged over the entire detector. The solid line is a linear fit to the lowest twenty points.	11
3	Response curves of intensified and unintensified CMOS and CCD cameras.	12
4	(a) Average signal calculated from 632 images at each exposure time from $1\ \mu\text{s}$ to $97\ \mu\text{s}$ with a step of $2\ \mu\text{s}$. (b) Mean and standard deviation of 100 individual pixels.	12
5	Samples used for calibration of one pixel at location (200, 100), together with the cubic polynomial fits for the average of the samples and the standard deviation of the samples.	13
6	Covariance matrices from the intensified CMOS camera: (a) 100 pixels sampled uniformly across the image (b) 10×10 contiguous subimage for which 10 spaces in the covariance matrix corresponds to adjacent pixels. For the contiguous region, there are significant non-diagonal terms, indicating the correlation between adjacent pixels.	13
7	Covariance matrix from the unintensified CMOS camera: (a) 100 pixels chosen uniformly across the whole image (b) 10×10 subimage. The non-diagonal terms are negligible, indicating the lack of noise correlation between neighboring pixels.	14
8	(a) Probability distribution functions for 100 pixels at an exposure time of $57\ \mu\text{s}$. (b) The same PDFs, shifted and rescaled to a zero mean and unit variance, compared with the standard normal PDF. Clearly, there is a skewness, and a Gaussian assumption should in principle be replaced to incorporate the skewness.	14
9	Experimental configuration for high-repetition rate OH LIF imaging in flames. The laser beam from a 10 kHz tunable dye laser is formed into a sheet using a cylindrical lens. OH fluorescence emission is measured using an intensified CMOS camera that is oriented orthogonal to the laser sheet.	16

10	Single-shot raw OH LIF signal in a steady laminar jet flame. The dashed red line indicates the row of pixels that was used for testing signal analysis approaches in 1-D.	17
11	Two frame sequences of raw data from 10 kHz OH LIF imaging in a turbulent piloted DME jet flame. The sequence was acquired from left to right with $100\ \mu\text{s}$ between each frame. The dashed red lines indicate the rows of pixels that were used for demonstrating signal analysis approaches in 1-D.	17
12	Application of the Bayesian methodology to the OH LIF signal in a row of pixels from the laminar flame data.	19
13	Two sample shots of OH LIF signal along a row of pixels together with their KL representations with 5, 10, 20 and 50 KL modes.	21
14	The average of $N = 10^4$ sample turbulent images, as well as the first 50 scaled KL modes $g_n(x)$	21
15	Result of a single inference run based on a parametrization by $L = 10$ KL modes. (a) Transformed back to the signal values to illustrate the posterior predictive distribution as a grey-shaded region, (b) The MAP estimate of the inferred intensity field. The data are from a single row of pixels in the top sequence of OH LIF images in Fig. 11	22
16	OH LIF signal from a single row of pixels in the laminar flame. The data is padded on both sides using mirror images of the signal. The central 640 pixels between the red dashed lines are the original signal. Zero padding is added such that the total number of pixels equals 2048. The edges of the padded profile are attenuated by a windowing function.	24
17	Eight level wavelet decomposition of two radial profiles of the OH LIF signal recorded in a steady laminar flame using the intensified CMOS camera. The two profiles are from rows of pixels that are separated by two pixels. The top plot (A8) shows the approximation coefficients, and the subsequent plots (D8-1) are the detail coefficients from each level. In each plot, the results from the two different rows of pixels are shown in green and blue.	25
18	Comparison of filtered and raw OH LIF signal profiles for the laminar flame. Local wavelet filtering function is the (a) magnitude of the correlation coefficient and (b) the square of the correlation coefficient between the two rows of pixels.	26
19	Comparison of filtered and raw OH LIF signal profiles for the laminar flame after corrections for background and detector nonlinearities.	26

20	Radial profiles of OH LIF signal recorded in a turbulent jet flame using the intensified CMOS camera. Each plot shows the signal along a single row of pixels from different single-shot OH LIF measurements. The data are padded on both sides using mirror images of the signal. The central 640 pixels between the red dashed lines are the original signal. The padded profile is multiplied by a windowing function that attenuates signal at the edges, and zero padding is added to the right hand side such that the total number of pixels equals 2048.	27
21	Eight level wavelet decomposition of OH LIF signal profiles in Fig. 20. The two profiles in each are from rows of pixels that are separated by two pixels. The top plot (A8) is the approximation coefficient, and the subsequent plots (D8-1) are the detail coefficients from each level. In each plot, the results from the two rows are shown in green and blue.	27
22	Comparison of filtered and raw OH LIF signal profiles for the turbulent jet flame from two single shot measurements. Local wavelet filtering function is the square of the correlation coefficient between two rows of pixels.	28
23	Comparison of filtered and raw OH LIF signal profiles for the laminar flame after corrections for background and detector nonlinearities.	28

1 Introduction

Accurate predictive simulations of turbulent flames are required for the development of advanced combustion technologies with high efficiency and low emissions. The development of high-fidelity numerical simulations of turbulent combustion and the associated submodels requires improved understanding of turbulence-chemistry interactions and detailed validation by experiments. The central goal of this work is to facilitate quantitative measurements of the spatial and temporal evolution of turbulent flames in order to advance the development of turbulent combustion simulations. To date, model validation experiments have primarily used single-point and line measurements, which do not capture the full dynamics of turbulent flames. We are developing approaches for comparing large-eddy simulations with 2-D and 3-D laser-based imaging measurements of turbulent flames. Central to this effort is the development of a high-repetition rate imaging capability for measuring the spatial and temporal evolution of turbulent flames. Recent advances in complementary metal oxide semiconductor (CMOS) detector technology and diode pumped solid-state lasers enable us to perform 2-D laser-induced fluorescence (LIF) measurements of combustion intermediates in turbulent flames at rates on the order of 10 kHz. However, quantitative measurements with high-speed CMOS cameras present significant challenges because of detector noise and nonlinear response. Many applications of high-speed CMOS cameras are focused on visualization of dynamic processes, such as automobile crash tests, and do not require the same level of quantitative signal analysis that is needed for understanding turbulence-chemistry interactions. In order to provide measurements that can be used for developing predictive simulations, the noise and nonlinearity issues must be addressed .

We combine detailed experimental characterization of high-speed CMOS cameras with uncertainty quantification analysis to design noise filters and provide uncertainty estimates for high-speed imaging measurements. Calibration experiments characterize the noise and nonlinearity of high-speed CMOS cameras with and without image intensifiers. Image intensifiers are necessary for high-speed laser-induced fluorescence measurements in flames but add another level of complexity to quantitative detection due to the noise and transfer function of the amplification process. Experimental results are used to develop noise reduction strategies that will enable quantitative cinematic imaging of turbulent combustion dynamics with proper accounting for uncertainty. We develop adaptive wavelet based noise filtering methods that capitalize on redundancy of information in measurements that are oversampled in space or time. We evaluate the utility of low dimensional representations of noisy spatio-temporal processes using Karhunen-Loève expansions along with Bayesian constructions. These methods are demonstrated using 1-D analysis of high-repetition rate OH laser-induced fluorescence imaging measurements in a steady laminar jet flame and a turbulent jet flame. The results establish methodologies for refining high-speed imaging data such that quantitative high-speed imaging measurements could be coupled with numerical simulations of the temporal evolution of turbulent flames.

2 Camera calibration: linearity and noise measurements

High repetition rate cameras commonly require image intensifiers to amplify the signal and to provide fast gating for suppression of interference. It is thus important to characterize the linearity and noise characteristics of intensified CMOS cameras. In order to isolate the characteristics of the CMOS detector from those of the image intensifier, two CMOS cameras (Vision Research, Phantom 7.3) were used in this study. One camera was unintensified and the other had a fiber optically coupled Gen II image intensifier. To assess the nonlinearity and noise characteristics of these cameras, we uniformly illuminated each detector with a cw illumination source. We removed all lenses and filters from the camera and placed the detector near the output port of an 8-in. diameter integrating sphere (Oriel Instruments) that was illuminated by a xenon lamp. For convenience, the detector response was measured using a constant illumination intensity and varying the camera exposure time. The exposure times of the intensified and unintensified cameras were varied via the intensifier gate time and the CMOS detector integration time, respectively. The reciprocity of the cameras was verified by comparing these response curves to those recorded using a fixed exposure time and a variable illumination intensity. For the latter case, the illumination intensity was varied by inserting combinations of neutral density (ND) filters between the xenon lamp and the integrating sphere. Figure 1 shows a comparison of the response curves of the intensified CMOS camera recorded using these different methods. The intensifier gate time was varied from 1 μ s to 97 μ s in 2 μ s increments. The x-axis indicates the percent of full illumination. For both methods, the 100% illumination case was acquired under identical conditions with a 97 μ s intensifier gate and no ND filters. For each ND filter combination, the effective irradiance incident on the detector was determined from the product of the measured spectral transmission curves of the filters, the spectrum of the xenon arc lamp, and the quantum efficiency curve of the intensifier. The agreement between the two response curves in Fig. 1 indicates excellent reciprocity of the image intensifier. A separate set of measurements showed a similar level of reciprocity for the unintensified camera.

The nonlinear response shown in Fig. 1 results from a combination of the CMOS detector and image intensifier responses. In order to isolate these effects, we compare the responses of the unintensified and intensified cameras in Fig. 2. The solid line is plotted as a guide to indicate the linear response. The CMOS detector response begins to deviate from linear near the middle of its dynamic range, resulting in approximately 10% deviation at full scale. The addition of the image intensifier increases the nonlinearity of the detection system in the upper half of the dynamic range. At full scale, the response of the intensified CMOS detector deviates from linear by approximately 20%.

In Fig. 3, we compare the responses of the ICMOS and CMOS cameras to those of intensified and unintensified charge coupled device (CCD) cameras, which are commonly used for low repetition rate imaging measurements. The unintensified CCD camera (Cooke Sensicam) displays no perceptible nonlinearity, whereas the ICCD camera (Andor iStar with Gen II intensifier) behaves similarly to the unintensified CMOS camera, except at the top end of its dynamic range. The ICMOS camera has the largest nonlinearity of these four cameras.

We consider the pixel-to-pixel variations in the response of the ICMOS camera and frame-to-frame fluctuations using sets of 632 images acquired at each exposure time from 1 μ s to 97 μ s.

Figure 1. Reciprocity verification for the intensified CMOS camera using two different methods to measure the detector response. In one method, the irradiance is varied using neutral density filters while maintaining a constant intensifier gate width. In the other method, the irradiance is kept constant and the gate width of the intensifier is varied.

Figure 2. Comparison of the nonlinearity of unintensified and intensified CMOS cameras. Response curve measurements are averaged over the entire detector. The solid line is a linear fit to the lowest twenty points.

Each image is 640 x 480 pixels. Figure 4(a) shows the average signals at each gate time calculated over the entire image. The same shape of the curve is seen for individual pixels, as illustrated in Figure 4(b). It shows the mean and standard deviation of signals for 100 pixels sampled uniformly across the image. Each pixel has a distinctive nonlinear response. In order to correct for these nonlinearities, a separate response curve must be measured for each pixel. We have obtained cubic polynomial fits as calibration curves for means and standard deviations of each pixel separately.

Figure 3. Response curves of intensified and unintensified CMOS and CCD cameras.

Figure 4. (a) Average signal calculated from 632 images at each exposure time from $1 \mu\text{s}$ to $97 \mu\text{s}$ with a step of $2 \mu\text{s}$. (b) Mean and standard deviation of 100 individual pixels.

Figure 5 shows an example of the calibration data, as well as the polynomial fits for the average and the standard deviation, for a single pixel at location (200, 100).

We consider the degree of correlation between noise in neighboring pixels. Figure 6 shows the covariance structure of two different 10×10 pixel sub-images: a) the matrix is constructed from pixels that are uniformly sampled across the detector, i.e. every 48-th pixel in the vertical direction and every 64-th pixel in the horizontal direction, and b) a local, contiguous sub-image of size 10×10 pixels. The covariance matrix in both cases is a 100×100 matrix of highly diagonal nature. The large off-diagonal terms in the contiguous sub-image correspond to the neighboring pixels. The correlation decays very quickly with distance however. It decays by a factor of 3-4 from the diagonal to the ± 1 off-diagonal pixels. Accordingly, we will model the detector response

Figure 5. Samples used for calibration of one pixel at location (200, 100), together with the cubic polynomial fits for the average of the samples and the standard deviation of the samples.

Figure 6. Covariance matrices from the intensified CMOS camera: (a) 100 pixels sampled uniformly across the image (b) 10×10 contiguous subimage for which 10 spaces in the covariance matrix corresponds to adjacent pixels. For the contiguous region, there are significant non-diagonal terms, indicating the correlation between adjacent pixels.

employing uncorrelated pixels. In contrast to the intensified camera, the covariance matrices for the unintensified camera in Fig. 7 are the same for contiguous and non-contiguous regions. Neither covariance matrix has significant off-diagonal terms. Evidently, the image intensifier is responsible for the significant off-diagonal terms. These results indicate that the point spread function of the image intensifier spans multiple pixels of the CMOS detector thereby producing a correlation

Figure 7. Covariance matrix from the unintensified CMOS camera: (a) 100 pixels chosen uniformly across the whole image (b) 10×10 subimage. The non-diagonal terms are negligible, indicating the lack of noise correlation between neighboring pixels.

Figure 8. (a) Probability distribution functions for 100 pixels at an exposure time of $57 \mu s$. (b) The same PDFs, shifted and rescaled to a zero mean and unit variance, compared with the standard normal PDF. Clearly, there is a skewness, and a Gaussian assumption should in principle be replaced to incorporate the skewness.

between noise in adjacent pixels.

To evaluate the noise characteristics of the intensified CMOS camera, we consider the frame-to-frame signal fluctuations of individual pixels during constant illumination by the xenon lamp via the integrating sphere. Figure 8a shows the PDFs of 100 different pixels at the illumination with an intensifier gate time of $57 \mu s$. In Fig. 8b, the PDFs are rescaled to a mean of 0.0 and a standard deviation of 1.0 in order to compare the shapes of the PDFs from different pixels. The rescaled

PDFs collapse relatively well indicating the similarity of their shapes. In the analysis section that follows, we use a Gaussian noise distribution to approximate the shape of the PDFs. Figure 8b shows that the PDFs exhibit some skewness that is not captured by a Gaussian distribution. In principle, however, the methodology that we outline below, can be readily generalized for any parametric noise distribution.

3 Flame measurements

Laser-induced fluorescence (LIF) imaging measurements of the hydroxyl radical (OH) were performed in steady laminar jet flames and turbulent jet flames. The noise and nonlinearity characteristics of the intensified CMOS camera were used to test the feasibility of uncertainty quantification and noise reduction of OH LIF measurements. Figure 9 shows the experimental configuration for OH LIF imaging. OH LIF is excited at 285 nm by the second harmonic of a dye laser, which is pumped by a 40W 10 kHz repetition rate diode-pumped Nd:YAG laser ($\lambda = 532$ nm). A cylindrical lens forms the laser beam into a vertical sheet that intersects the jet axis. The OH LIF emission is collected at 90 deg. to the laser sheet using an f/1.8 fused silica camera lens. The OH LIF signal is recorded on the high-speed intensified CMOS camera. The laser and camera are synchronized to an external 10 kHz clock. The camera has on-board memory for recording approximately 3 second sequences, or 30,000 frames. The data is then downloaded to a computer before acquisition of another sequence.

The partially premixed DME jet flames are stabilized on a 7 mm diameter nozzle with a low-speed air coflow. The fuel stream mixture consists of 20% DME and 80% air, by volume, which results in a stoichiometric mixture fraction of 0.353. For the turbulent jet flames, an annular pilot is used to anchor the flames to the nozzle. The pilot is a mixture of acetylene, hydrogen, air, carbon dioxide, and nitrogen with the same enthalpy and equilibrium composition as a lean premixed DME/air flame.

Figure 9. Experimental configuration for high-repetition rate OH LIF imaging in flames. The laser beam from a 10 kHz tunable dye laser is formed into a sheet using a cylindrical lens. OH fluorescence emission is measured using an intensified CMOS camera that is oriented orthogonal to the laser sheet.

Figure 10. Single-shot raw OH LIF signal in a steady laminar jet flame. The dashed red line indicates the row of pixels that was used for testing signal analysis approaches in 1-D.

Figure 11. Two frame sequences of raw data from 10 kHz OH LIF imaging in a turbulent piloted DME jet flame. The sequence was acquired from left to right with $100 \mu\text{s}$ between each frame. The dashed red lines indicate the rows of pixels that were used for demonstrating signal analysis approaches in 1-D.

Measurements were performed in both steady laminar flames and turbulent jet flames. We focused on DME flames because of their importance in understanding turbulent combustion of oxygenated fuels for reduced pollution emissions in practical combustion devices. Figure 10 shows a single frame of raw OH LIF signal from a steady laminar DME jet flame. No corrections have been performed for laser beam profile, background level, or detector nonlinearity. Figure 11 shows two sequences of raw OH LIF images acquired at 10 kHz in a turbulent piloted DME jet flame. The OH distribution marks the high-temperature reaction zone, and the temporal evolution of the turbulent reaction zone is evident from the changes in flame morphology as the turbulent structures are convected downstream. In the analysis that follows, we use the rows of pixels that are indicated by the dashed red lines for proof-of-principle demonstrations of signal analysis in 1-D.

4 Analysis

In the following, we present a preliminary analysis of the OH LIF measurements from the intensified CMOS camera, and of their utilization for inference of intensity field images with quantified uncertainty. The overall goal is to use a data analysis method that accounts for the noise in the detection system, and provides estimates of quantities of interest with quantified uncertainty. While this can in principle be done with least squares fitting (LSF) methods, there are numerous simplifying assumptions involved in the LSF approach for quantification of uncertainty in estimated quantities. In order to provide full probabilistic uncertainty quantification (UQ) in estimated parameters, we rely rather on a Bayesian inference formalism. The Bayesian formalism provide a general means of handling information employing probability theory. Moreover, it provides the means of seamlessly merging subjective prior information and empirical data in probabilistic inference, formal and non-adhoc handling of *nuisance* parameters, and is a sound foundation for sequential learning, *e.g.* from subsequent experiments.

We begin by introducing the Bayesian formalism, outlining its implementation for inference with laminar flame data. We then address the turbulent flame measurements, describing a general random field construction for the intensity image, which allows the Bayesian inference of arbitrary uncertain image structures.

4.1 Bayesian inference

We use a Bayesian inference methodology [1] to infer the true intensity field $I(x)$, where x denotes the spatial coordinates, *given* the measured signal $S_i(x)$, $i = 1, \dots, R$, and the results from Section 2. For an introduction to Bayesian techniques for image analysis and reconstruction, see [2].

Bayes' formula states

$$P(\mathcal{M}|\mathcal{D}) \propto P(\mathcal{D}|\mathcal{M})P(\mathcal{M}), \quad (1)$$

relating the prior probability $P(\mathcal{M})$ of a model $\mathcal{M}$ to the posterior probability $P(\mathcal{M}|\mathcal{D})$ using the likelihood function $P(\mathcal{D}|\mathcal{M})$, which is the probability of having the given dataset $\mathcal{D}$ if the data was drawn from the model $\mathcal{M}$. In the present context, the model $\mathcal{M}$ consists of a parameter vector c that defines the intensity field $I(x; c)$, while data are the signals in a single image, $S(x)$. The object is to use Bayes formula (1) to estimate the parameters c . The likelihood function $L(\mathcal{M}) = P(\mathcal{D}|\mathcal{M})$ is key. The logarithm of the likelihood function takes the following form given the Gaussian noise assumption,

$$\log L(\mathcal{M}) = -\frac{1}{2} \sum_x \log (2\pi\sigma_x^2(I(x; c))) - \sum_x \frac{(S(x) - \mu_x(I(x; c)))^2}{2\sigma_x^2(I(x; c))}, \quad (2)$$

where $\mu_x(I)$ and $\sigma_x(I)$ are the calibration curves for the mean and standard deviation corresponding to the pixel x , obtained in Section 2. The log-likelihood function (2) relies on the assumption of the Gaussian noise in the signal. The methodology can clearly be adapted to implement any other, perhaps more realistic, parametric distribution, *e.g.* Poisson or lognormal.

We illustrate the methodology using a single row of 640 pixels, at pixel height 200, from the laminar flame data in Fig. 10. Based on the form of the signal, let us assume a smooth model for the intensity field that depends on the distance r from the 300-th pixel, i.e. $r = |x - 300|$ and

$$I(x; c) = I_c(r) = \exp(c_0 + c_1 r + c_2 r^2 + c_3 r^3). \quad (3)$$

The polynomial coefficients $c = (c_0, c_1, c_2, c_3)$ then become the objects of inference in the Bayesian formulation. We use an adaptive Markov chain Monte Carlo (MCMC) algorithm to sample from the posterior distribution of c . Figure 12 shows the calibrated intensity data from a single row of pixels, their average values, as well as the intensities obtained from fitting the data through the bimodal-exponential model (3) via Bayesian inference, outlined above. The methodology provides the full probability distribution of the uncertain inferred model parameters c representing the image. However, for now we only present the *maximum a posteriori* (MAP) estimate. Namely, the parameter value c_{MAP} that maximizes the posterior distribution.

Figure 12. Application of the Bayesian methodology to the OH LIF signal in a row of pixels from the laminar flame data.

For the laminar flame data, the parametrization (3) seems to capture the dominant features of the OH LIF signal profile, although it does not fully capture the peaks signals. A more accurate representation of the OH LIF signal is possible with larger orders of the polynomial. However, for much more complicated structures, such as the turbulent flame data, a representation basis is extremely difficult to choose up-front. For that reason, we suggest using Karhunen-Loève representation, outlined below.

4.2 Karhunen-Loève expansion

We assume the intensity field $I(x)$ to be a stochastic process with R samples of the field available from the measurements. Denote the mean of this process over sample realizations by $\bar{I}(x)$.

Separating the mean, one can write

$$I(x) = \bar{I}(x) + I_0(x), \quad (4)$$

where $I_0(x)$ is a *centered* stochastic process, i.e. its mean vanishes for all pixels x .

The covariance function $C(x_1, x_2) = \langle I_0(x_1)I_0(x_2) \rangle$ of a centered stochastic process is symmetric, bounded and positive definite. Hence, it can be expanded as a sum [3]

$$C(x_1, x_2) = \sum_{n=1}^{\infty} \lambda_n f_n(x_1) f_n(x_2) \quad (5)$$

with respect to its complete, orthonormal set of eigenfunctions $f_n(x)$ and real positive eigenvalues λ_n (in a descending order) that are defined as solutions of the integral equation

$$\int_{x_1} C(x_1, x_2) f_n(x_1) dx_1 = \lambda_n f_n(x_2). \quad (6)$$

The underlying stochastic process $I(x)$ then admits the *KL decomposition* [4, 5]

$$I(x) = \bar{I}(x) + \sum_{n=1}^{\infty} \xi_n \sqrt{\lambda_n} f_n(x) \quad (7)$$

in terms of the eigenfunctions $f_n(x)$ and random variables ξ_n that are uncorrelated, have zero mean and unit variance. Using the orthonormality of the eigenfunctions, one can recover these uncorrelated - but not independent - random coefficients ξ_n by projection of the sample trajectories onto the eigenfunction basis:

$$\xi_n = \frac{1}{\sqrt{\lambda_n}} \int_x I_0(x) f_n(x) dx. \quad (8)$$

In practice, the sum in Eq. (7) is truncated, leading to a finite KL decomposition in terms of L zero-mean, uncorrelated random variables $\{\xi_i\}_{i=1}^L$. The finite KL decomposition essentially represents the intensity field $I(x)$ in terms of a finite number of generally dependent random variables. The structures containing higher spatial frequencies are taken into account as more terms in the decomposition are considered.

Figure 13 illustrates the truncated KL representation of a single row of 546 pixels, at pixel height 238, for the turbulent flame data in Fig. 11. As more KL modes are retained, progressively rougher structures are captured by the expansion. Assuming the flame is in a statistical steady state, it can be represented as a random field, where the KL representation serves as a reduced order model of this random intensity field.

The KL modes will serve as a basis of expansion for the forward model $I(x; c)$ in the Bayesian inference methodology, outlined in the Section 4.1. Namely, we will take

$$I(x; c) = \bar{I}(x) + \sum_{n=1}^L c_n g_n(x), \quad (9)$$

Figure 13. Two sample shots of OH LIF signal along a row of pixels together with their KL representations with 5, 10, 20 and 50 KL modes.

Figure 14. The average of $N = 10^4$ sample turbulent images, as well as the first 50 scaled KL modes $g_n(x)$.

where $g_n(x) = \sqrt{\lambda_n} f_n(x)$ are the eigenvalue-scaled KL modes, shown in Figure 14. The results of Bayesian inference, based on the turbulent flame data, of a 10-parameter (10-term KL) expansion are presented in Figure 15. The shaded region in Figure 15(a) corresponds to the posterior predictive [6] (PP) uncertainty which takes all sources of uncertainty into account for a given 10-dimensional model, including both the uncertainty in the inferred parameters and the noise model information from the calibration data. The PP uncertainty, as expected, spans the distribution of the noisy data.

Figure 15. Result of a single inference run based on a parametrization by $L = 10$ KL modes. (a) Transformed back to the signal values to illustrate the posterior predictive distribution as a grey-shaded region, (b) The MAP estimate of the inferred intensity field. The data are from a single row of pixels in the top sequence of OH LIF images in Fig. 11

4.3 Summary of Bayesian and KL methods

We have outlined above a preliminary analysis of CMOS detector data, allowing construction of both laminar and turbulent flame intensity field images with quantified uncertainty. The approach is based broadly on probabilistic representation of information, and on the use of Bayesian inference for estimation of the uncertain image structure. We illustrated a range of complexity in model structure, allowing representation of arbitrary intensity fields from turbulent flame measurements, with quantified uncertainty.

The above illustrations have largely focused on proof-of concept and preliminary demonstrations employing single-line data. The extension to full 2D images is straightforward from a formulation viewpoint, although it is certainly more computationally intensive. Analysis of 2D images with flame fronts will require significant attention to issues of local filtering/smoothing, *e.g.* using wavelet noise reduction. The above KL model representation can be extended for application in such a local context, and it is of interest to explore the potential of using wavelet noise reduction coupled with efficient local random field KL representations.

Moreover, the above analysis did not address the handling of time-sequences of image data. Here again, in principle, the KL representation of the uncertain intensity field is extendable in the time axis. However, this has to be done carefully and with attention to computational complexity. Potentially, one approach is to estimate the KL eigenfunctions, and associated truncation, from the full data set, but to update the random basis in time using sequential Bayesian updating methods.

4.4 Wavelet noise suppression

We investigated the feasibility of using wavelets as an approach to suppressing noise in measurements from the intensified CMOS camera. Wavelets provide a compact basis set for efficient decomposition of signals for a wide range of signal processing applications. For background on wavelet analysis, see for example Walker [7] and references therein. Noise suppression using wavelets involves the decomposition of the signal into multiple levels of wavelet coefficients, filtering of the wavelet coefficients, and reconstruction of the signal from the filtered coefficients. We use stationary wavelet transforms (SWT), also referred to as undecimated wavelet transforms [8]. An important advantage of the undecimated transform over the decimated transform is that the wavelet decomposition is translation-invariant. Wavelet noise suppression is traditionally implemented by attenuating the wavelet coefficients using two different schemes, referred to as hard and soft thresholding [9]. In hard thresholding, wavelet coefficients that have a magnitude less than a fixed threshold are set to zero, and only coefficients that are greater than the threshold are retained. One of the disadvantages of hard thresholding is the abrupt truncation of the wavelet coefficients, which can produce artifacts in the reconstructed signal. In soft thresholding, coefficients below a fixed threshold are set to zero and the values of the retained coefficients are reduced by the threshold value. This approach eliminates abrupt transitions to zero in the filtered coefficients but can still produce artifacts. In both of these thresholding approaches, only the magnitude of the correlation coefficients is used as a measure of the signal content. However, wavelet coefficients contain a combination of signal and noise, which are not so easily separated. An adaptive thresholding approach is needed to address these shortcomings. Adaptive filtering attenuates the wavelet coefficients using a locally adapted threshold value or attenuation coefficient. The use of a locally varying attenuation coefficient is preferable to avoid abrupt truncation of the wavelet coefficients.

In the present investigation, we developed an adaptive thresholding method in which the local signal and noise content of the wavelet coefficients are locally evaluated and the results are used to determine the attenuation coefficients. Our approach takes advantage of the redundancy of signal content in neighboring rows of pixels within an image. This method is ideally suited for an over-sampled imaging measurement or two independent measurements of the same quantity using two cameras. The former configuration is used in the present study. The basic concept is that the signals in neighboring rows of pixels are highly correlated, but the noise is uncorrelated. The correlation coefficient between the wavelet coefficients in neighboring rows is used to provide a measure of the local signal content, and the coefficients are attenuated accordingly. In regions where the wavelet coefficients are highly correlated, the wavelets have a high signal content, whereas low correlation indicates higher noise content and these wavelet coefficients are attenuated.

The success of this adaptive filtering approach requires that noise in neighboring rows of pixels is uncorrelated. The covariance matrix for the intensified CMOS camera in Fig. 6b shows the degree of noise correlation between neighboring rows of pixels. The finite point spread function of the image intensifier produces a noise correlation between adjacent pixels on the intensified camera. This correlation decays to negligible levels for separation distances of two or more pixels. Therefore, the adaptive filtering scheme for the intensified camera uses the correlation between rows that are separated by two pixels.

For proof of concept, we apply 1-D adaptive wavelet noise filtering to individual rows of pixels, but the approach is extendable to 2-D filtering. We first consider a radial profile of the OH LIF measurement in the steady laminar flame shown in Fig. 10. Prior to applying a wavelet transform to the measured signal, we reconfigure the data in order to minimize the impact of artifacts in the wavelet decomposition and reconstruction. The row of data is padded on either side using mirror images of the measured signal. Zero padding is then added at the edges such that the total length of the padded signal equals a power of two. In the present case, the padded signal length is 2048 pixels. The padded signal is multiplied by a windowing function to suppress edge effects. The resulting padded data for a single row is shown in Fig. 16. The center region between the red dashed lines is the original signal, and the padded regions are to the right and left of these lines.

Figure 16. OH LIF signal from a single row of pixels in the laminar flame. The data is padded on both sides using mirror images of the signal. The central 640 pixels between the red dashed lines are the original signal. Zero padding is added such that the total number of pixels equals 2048. The edges of the padded profile are attenuated by a windowing function.

Wavelet decomposition is performed using an eight level wavelet transform with 10th order Daubechies wavelets. The resulting approximation and detail coefficients for the row of pixels in Fig. 16 and a row separated by two pixels are displayed in Fig. 17. A comparison of the detail coefficients from the two rows of pixels shows that the D8 level detail coefficients are very highly correlated and represent the larger spatial scales within the signal. As the decomposition level decreases, the spatial scales and the magnitude of the detail coefficients decrease and the coefficients from the two rows become less correlated. At intermediate decomposition levels, D4-D6, there are regions of relatively high correlation and those with lower correlation. For D3-D1, the correlation is very low, indicating that the high spatial frequencies represented by those levels is dominated by noise.

In the adaptive filtering approach, we locally evaluate the correlation between the detail coefficients from the two rows of pixels and multiply the coefficients by the magnitude of the local correlation coefficient. As a result, the attenuation is greater for coefficients that are less correlated.

The filtered data is reconstructed using the attenuated coefficients. Figure 18a shows the original and reconstructed signals for the row of pixels from the laminar flame data. The degree of filtering can be varied using different schemes. For example, the attenuation coefficient can be constructed from raising the correlation coefficient to a power. Figure 18b shows the same data that has been filtered using the squared correlation coefficient as the filtering function. In both Figs. 18a and 18b, the filtered signal tracks the original signal but with significantly less noise. The higher level of filtering decreases the fluctuations. The filtered result can then be corrected for background and nonlinear response, and the result is shown in Fig. 19.

Figure 17. Eight level wavelet decomposition of two radial profiles of the OH LIF signal recorded in a steady laminar flame using the intensified CMOS camera. The two profiles are from rows of pixels that are separated by two pixels. The top plot (A8) shows the approximation coefficients, and the subsequent plots (D8-1) are the detail coefficients from each level. In each plot, the results from the two different rows of pixels are shown in green and blue.

We demonstrated this adaptive wavelet noise suppression in the turbulent jet flames as well. Figure 20 shows the padded raw OH LIF signals from the rows of pixels in the two frames shown in Fig. 11. The eight wavelet decompositions for each of these signals is shown in Fig. 21. The filtered signals are compared with the raw signals in Fig. 22. The filtered signals capture the key features of the OH LIF signal while suppressing much of the noise. For these examples, the filtering function was the square of the local correlation coefficients. In Fig. 23, the filtered result is corrected for background and nonlinearities using the polynomial fits for the response of individual pixels. These results demonstrate the feasibility of using wavelets for noise suppression in single-shot imaging measurements of complex flame structures.

Figure 18. Comparison of filtered and raw OH LIF signal profiles for the laminar flame. Local wavelet filtering function is the (a) magnitude of the correlation coefficient and (b) the square of the correlation coefficient between the two rows of pixels.

Figure 19. Comparison of filtered and raw OH LIF signal profiles for the laminar flame after corrections for background and detector nonlinearities.

Figure 20. Radial profiles of OH LIF signal recorded in a turbulent jet flame using the intensified CMOS camera. Each plot shows the signal along a single row of pixels from different single-shot OH LIF measurements. The data are padded on both sides using mirror images of the signal. The central 640 pixels between the red dashed lines are the original signal. The padded profile is multiplied by a windowing function that attenuates signal at the edges, and zero padding is added to the right hand side such that the total number of pixels equals 2048.

Figure 21. Eight level wavelet decomposition of OH LIF signal profiles in Fig. 20. The two profiles in each are from rows of pixels that are separated by two pixels. The top plot (A8) is the approximation coefficient, and the subsequent plots (D8-1) are the detail coefficients from each level. In each plot, the results form the two rows are shown in green and blue.

Figure 22. Comparison of filtered and raw OH LIF signal profiles for the turbulent jet flame from two single shot measurements. Local wavelet filtering function is the square of the correlation coefficient between two rows of pixels.

Figure 23. Comparison of filtered and raw OH LIF signal profiles for the laminar flame after corrections for background and detector nonlinearities.

5 Summary

We combined detailed experimental characterization of high-speed CMOS cameras with uncertainty quantification analysis to develop noise filters and to provide uncertainty estimates for high-speed imaging measurements. The results establish the feasibility of coupling quantitative high-speed imaging measurements with numerical simulations of the temporal evolution of turbulent flames. We investigated wavelet based noise filtering methods and evaluated the utility of low dimensional representations of noisy using Bayesian constructions.

The proof-of-principle analysis presented here is implemented in 1-D. However, the methods are extendable to 2-D and may even be developed to use both spatial and temporal information. For example, the optimal filtering methods could be informed by available knowledge of the spatio-temporal structure of the field being measured. Such knowledge includes the spatial and temporal correlation structure of the flow coupled with image time history.

The adaptive wavelet filters could be further developed with specific regard to the local length and time scale distributions in the images. Adaptive filter sizes based on temporal autocorrelation functions or spatio-temporal covariance matrices have been used in CMOS and MRI sensor processing in a general context. However, their use in quantitative laser-induced fluorescence imaging of turbulent flow structures introduces specific challenges associated with fidelity in both quantitative signal levels and evolving length/time-scales. In a previous study, we demonstrated the use of adaptive Gaussian smoothing with kernel size scaling determined from local mean turbulence length scales. More local control over relevant time and length scales is possible with wavelet-based transformations, which are well suited to handling transients and local variations in noise levels.

The use of adaptive wavelet-based noise suppression and probabilistic Bayesian methods to properly characterize both the signal and its noise, taking into account all available information, will eventually allow quantitative validation of simulations against laser-induced fluorescence measurements that rely on high-speed CMOS cameras, thereby enabling predictive simulations of turbulent reacting flows.

This page intentionally left blank.

References

- [1] D.S. Sivia. *Data Analysis: A Bayesian Tutorial*. Oxford Science, 1996.
- [2] K.M. Hanson. Introduction to Bayesian image analysis. In M.H. Loew, editor, *Medical Imaging: Image Processing*, volume 1898, pages 716–731. Proc. SPIE, 1993.
- [3] R.G. Ghanem and P.D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer Verlag, New York, 1991.
- [4] K. Karhunen. Zur spektraltheorie stochastischer prozesse. *Ann. Acad. Sci. Fennicae*, 37, 1946.
- [5] M.M. Loève. *Probability Theory*. Van Nostrand, Princeton, NJ, 1955.
- [6] A. Gelman, X.-L. Meng, and H. Stern. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6:733–807, 1996.
- [7] J.S. Walker. *A Primer on Wavelets and Their Scientific Applications*. Chapman and Hall/CRC, second edition, 2008.
- [8] G.P. Nason and B.W. Silverman. The stationary wavelet transform and some statistical applications. In A. Antoniadis and G. Oppenheim, editors, *Wavelets and Statistics, Lecture Notes in Statistics*, pages 281–300. Springer Verlag, New York, 1995.
- [9] D. L. Donoho and I. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455, 1994.

This page intentionally left blank.