Nonlinear diffusion filtering of images obtained by planar laser-induced fluorescence spectroscopy

Henrik Malm and Gunnar Sparr

Department of Mathematics, Lund Institute of Technology, Box 118, S-22100 Lund, Sweden

Johan Hult and Clemens F. Kaminski*

Division of Combustion Physics, Lund Institute of Technology, Box 118, S-22100 Lund, Sweden

Received January 24, 2000; revised manuscript received July 11, 2000; accepted August 10, 2000

The application of nonlinear anisotropic diffusion filtering to reduce noise and enhance contours in images obtained by two-dimensional planar laser-induced fluorescence (PLIF) spectroscopy is presented. In this process the diffusion coefficient is locally adapted, becoming negligible as object boundaries are approached. Noise is efficiently removed, and object contours are strongly enhanced. The technique is demonstrated with PLIF images obtained from the OH radical recorded in turbulent flames. We show that nonlinear diffusion is suitable as a preprocessing step, before image segmentation becomes possible, and we demonstrate how the technique is applied for the quantitative extraction of flame reaction boundaries from PLIF data. © 2000 Optical Society of America [S0740-3232(00)03612-7]

OCIS codes: 100.2000, 300.2530, 280.1740.

1. INTRODUCTION

Two-dimensional laser spectroscopic imaging has become one of the most valuable tools in the study of the physics and chemistry of turbulent reactive flows.[1] A variety of spectroscopic techniques are available for two-dimensional imaging of relevant scalars in combustive flows: Raman spectroscopy,[2] Mie scattering, Rayleigh scattering,[3] and laser-induced fluorescence[3–5] have all been applied in a two-dimensional fashion in these environments. These techniques have in common that a suitable, high-power laser beam is shaped into a planar sheet that cuts through the measurement volume of interest. The laser light is scattered by particles present in the measurement volume and is captured at right angles by a suitable camera (see Fig. 1). By a suitable choice of laser wavelength and other parameters, one is able to measure the concentrations of chemical species present in the flame environment, the flame temperature, the flow characteristics, etc. By use of high-power laser sources with short pulse lengths (of the order of 10^{-8} s), this information can be obtained practically instantaneously on the chemical and physical time scales prevailing in combustion systems. Such data are valuable for improving the understanding of combustion processes and in the quest to design cleaner, more efficient combustion devices. An overview of spectroscopic applications to combustion diagnostics is given in Ref. 6.

Laser-induced fluorescence is of particular importance in this context, since it is the only technique with the sensitivity required for making concentration measurements of radical species that are produced *in situ* during the combustion process and are only present at trace levels. Of particular importance in this context is the OH radical, since it occurs in all systems in which air is used as an oxidizer. The local OH concentration is a good indicator of the efficiency of the combustion process, since it is sensitively dependent on the local flow field properties.[7] Recently the evolution of OH concentration fields in a flame with shear flow-generated turbulence was visualized in real time in a filmlike manner[8] by use of planar laser-induced fluorescence (PLIF). The effects of convective stress on the flame front and local fuel–air mixing on the flame could be tracked in real time by these measurements. The PLIF spectrometer used for this work is unique in terms of the temporal and spatial resolution attainable. However, its dynamic range is limited to an 8-bit resolution, and in a typical experimental situation the signal-to-noise ratio and contrast-to-noise ratio are therefore limited. For example, Fig. 2B shows a low-quality raw image obtained by OH PLIF in a turbulent flame. For quantitative comparison with results obtained from theoretical calculations of the combustion system, some form of post-image-processing for quantitative data extraction is necessary. Usually this involves local averaging over noise or the enhancement of image contours. Simple smoothing techniques work by low-pass filtering, which removes the noisy high-frequency components from the image. Unfortunately, this results in image blurring and degrading of edge information. This can be partly overcome by nonlinear filters (median filtering), which preserve edges but compromise fine-scale resolution.

The present paper reports the application of a local anisotropic diffusion filter for the enhancement of PLIF images. In the present case the application was developed to allow the quantitative extraction of flame contour data from OH images obtained by the 8-bit detector system described, but the method is generally applicable to all two-dimensional laser imaging techniques. It is adapted from an approach originally formulated by Perona and

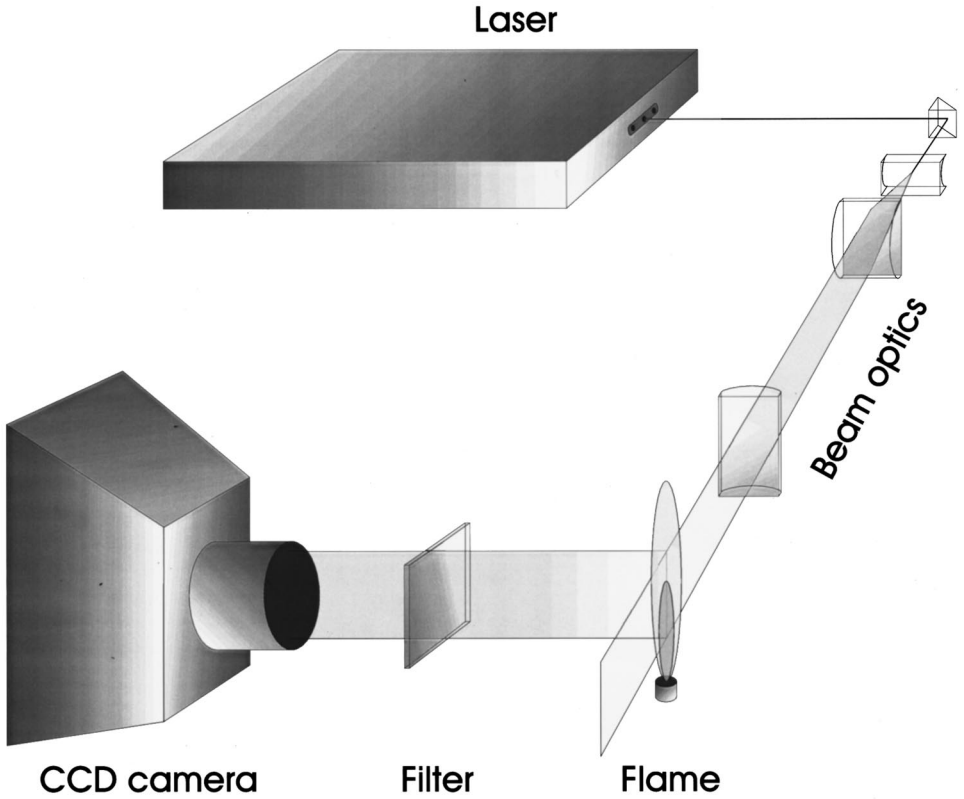

Fig. 1. Schematic setup of planar laser light-scattering techniques.

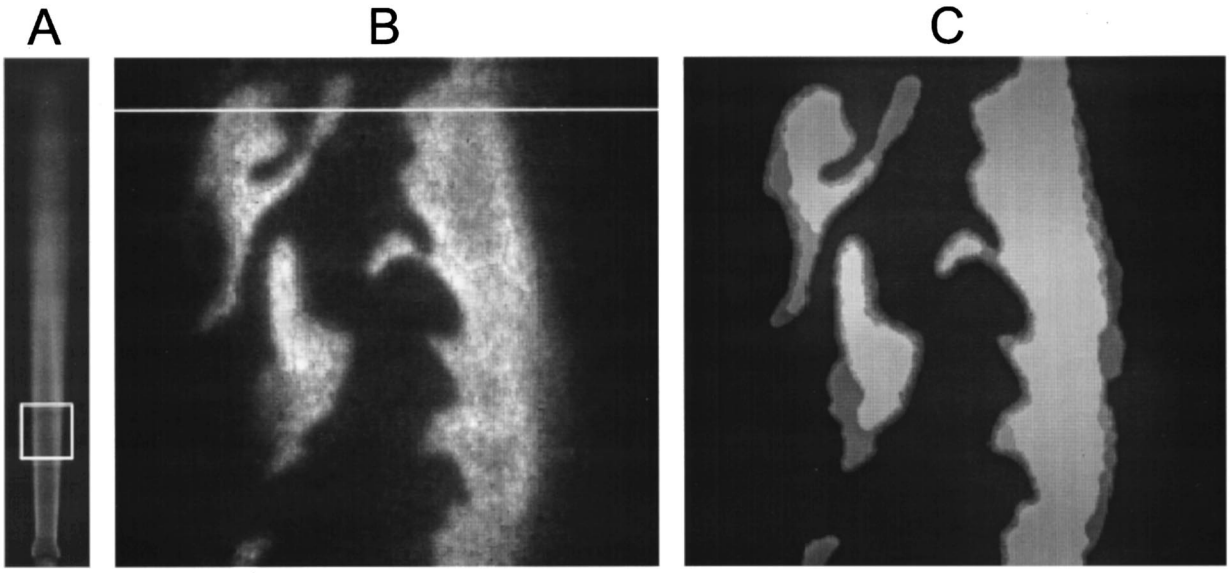

Fig. 2. A, Photograph of line-of-sight flame emission from a turbulent nonpremixed flame. The area imaged by the PLIF experiments is indicated by the white box, corresponding to 24 mm × 23 mm, starting 40 mm above the burner nozzle B, PLIF raw image showing two-dimensional OH-radical concentrations in a cut through the flame. C, PLIF image after nonlinear diffusion filtering. The position of the flame front is strongly enhanced.

Malik.[9] The method has several advantages: Noise is smoothed locally within regions defined by object boundaries, whereas little or no smoothing occurs *between* image objects. Local edges are enhanced because discontinuities, such as boundaries, are amplified. Mathematically one treats the problem as a diffusion process, where the diffusion coefficient is adapted locally to the effect that diffusion stops as soon as an object boundary is reached. Local noise is efficiently removed, and object boundaries are enhanced.

The paper is organized as follows: In Section 2 a brief description of the technique and its computational implementation is given. In Section 3 examples are given to visualize how the method is optimized for a given appli-

cation, and the technique is then applied to find flame contours from two-dimensional PLIF images obtained in a turbulent diffusion flame. The particular objective was to find flame boundaries and to calculate contour lengths and areas from OH radical distributions. Practical aspects concerning the technique in the context of the present application will be discussed. Finally, Section 4 summarizes the paper and highlights advantages of and problems with the technique.

2. THEORY

A. Nonlinear Diffusion

We define an image as a function u on $\mathbf{R}^2$, where $u(x)$ represents the gray level or intensity at point x. A traditional way to smooth an image in a controlled way is to convolve it with a Gaussian kernel:

$$G_\sigma(x) := \frac{1}{2\pi\sigma^2}\exp(-|x|^2/2\sigma^2), \tag{1}$$

where σ determines the width of the kernel. By "controlled," one means that the smoothing, or low-pass filtering, fulfills some natural requirements, such as invariance under rotations, translations, and constant shifts in the gray scale. Another requirement is that the smoothing not create false structure in the image; that is, the process should not create new local extrema. It is a classical fact that the solution of the linear diffusion equation

$$\partial_t u = \mathrm{div}(d\nabla u), \tag{2}$$

where d is a scalar constant, for time t is exactly the same operation as convolving the image u with a Gaussian kernel of width $\sqrt{2t}$.

In the current application we are interested in finding a good approximation of the contour of an object in a noisy image. Smoothing with a Gaussian kernel will help to suppress noise, but it will also blur and dislocate edges and contours. In Ref. 9 Perona and Malik proposed to exchange the scalar diffusion constant d in Eq. (2) for a scalar-valued function g of the gradient ∇u of the gray levels in the image. The diffusion equation then reads as

$$\partial_t u = \mathrm{div}[\, g(|\nabla u|)\nabla u\,]. \tag{3}$$

The length of the gradient, $|\nabla u|$, has been empirically proved to be a good measure of the edge strength of the current location and expresses a dependence on the differential structure of the image. This dependence, of course, makes the diffusion process nonlinear. The function g is required to be smooth and nonincreasing with $g(0) = 1$ and $g(s) \to 0$ for $s \to \infty$.

Using the alternative so-called gauge coordinates η and ξ, where η is the normalized direction parallel to the gradient and ξ is the normalized direction orthogonal to the gradient, the right-hand side of Eq. (3) can be rewritten as

$$\mathrm{div}[\, g(|\nabla u|)\nabla u\,] = g(|\nabla u|)u_{\xi\xi} + \Phi'(|\nabla u|)u_{\eta\eta}, \tag{4}$$

where $\Phi(|\nabla u|) = |\nabla u| g(|\nabla u|)$ and $u_{\xi\xi}$ and $u_{\eta\eta}$ are the second derivatives in the directions ξ and η, respectively. For a detailed derivation of relation (4), see Ref. 10. In this new form it can be seen that the coefficient $\Phi'(|\nabla u|)$ before the term $u_{\eta\eta}$ governs diffusion across edges.

As diffusivity function g we use a function originally proposed by Weickert[11] in a similar setting. In the present case this function was found to be optimally suited to the properties of the PLIF images recorded by our camera. Weickert's diffusivity function reads as

$$g(s) = 1 - \exp\left[\frac{-C_m}{(s/\lambda)^m}\right]. \tag{5}$$

Here, m is a positive integer, $C_m > 0$, and $\lambda > 0$. The constant C_m is chosen in such a way that the flux function $\Phi(s)$ is increasing for $s < \lambda$ and decreasing for $s > \lambda$. That is, $\Phi'(s) > 0$ when $s < \lambda$, and $\Phi'(s) < 0$ when $s > \lambda$. In this way the diffusion in the direction across the edges will be a forward diffusion when $|\nabla u| < \lambda$ and a backward diffusion when $|\nabla u| > \lambda$, and the slopes at the edges will increase. The parameter λ acts as a contrast parameter that separates regions where forward diffusion occurs from those where backward diffusion takes place.

Function (3) is apparently rather unstable and sensitive to noise because of the backward-diffusion properties. This makes the stability of the solutions of the equation strongly dependent on the regularizing effect of the discretization and numerical implementation. Therefore it seems natural to introduce a controlled regularization directly into the continuous process. Catté *et al.* in Ref. 12 proposed to smooth the gradient inside the diffusivity and replace $g(|\nabla u|)$ with the estimate $g(|\nabla(G_\sigma * u)|)$, where G_σ is a Gaussian kernel of width σ. That is, they propose the following diffusion equation:

$$\partial_t u = \mathrm{div}[\, g(|\nabla(G_\sigma * u)|)\nabla u\,]. \tag{6}$$

The model of Catté *et al.*[12] has the advantage that it is insensitive to noise at a scale that is determined by σ. Therefore it does not misinterpret noise as edges that should be enhanced. This makes Eq. (6) much more stable than Eq. (3).

B. Numerical Implementation

Good results can be obtained with rather simple discretizations of Eq. (6). The discretization used for the examples shown in this paper is based on central finite difference approximations of the spatial derivatives with reflecting boundary conditions. The step sizes Δx and Δy are both set equal to 1. For the differentials in the x and the y directions, respectively, one has

$$\partial_x u_{i,j} = \frac{u_{i+1,j} - u_{i-1,j}}{2}, \tag{7}$$

$$\partial_y u_{i,j} = \frac{u_{i,j+1} - u_{i,j-1}}{2} \tag{8}$$

The difference calculations can be implemented by multiplication with sparse matrices, offering a substantial improvement in computational speed. One has $\partial_x u = uX$ and $\partial_y u = Yu$, where X and Y are the sparse forms of

$$X = \frac{1}{2}\begin{bmatrix} 0 & -1 & 0 & \dots & 0 & 0 \\ 1 & 0 & -1 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & -1 \\ 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}_{N\times N},$$

$$Y = \frac{1}{2}\begin{bmatrix} 0 & 1 & 0 & \dots & 0 & 0 \\ -1 & 0 & 1 & \dots & 0 & 0 \\ 0 & -1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 1 \\ 0 & 0 & 0 & \dots & -1 & 0 \end{bmatrix}_{M\times M}. \quad (9)$$

Here N represents the number of discrete points (pixels) of u along the x axis and M is the number of pixels along y.

To iterate 1 time step, one starts by calculating $G_\sigma * u$ and proceeds by differentiating the result. An approximate solution of $|\nabla(G_\sigma * u)|$ is thereby obtained. The function $g[|\nabla(G_\sigma * u)|]$ is then evaluated at all discrete points, and finally u is updated according to

$$u^n = u^{n-1} + \Delta t[(g \cdot \partial_x u^{n-1})X + Y(g \cdot \partial_y u^{n-1})]. \quad (10)$$

Here $g \cdot \partial_x u^{n-1}$ denotes elementwise multiplication of the matrices $g[|\nabla(G_\sigma * u^{n-1})|]$ and $\partial_x u^{n-1}$.

The time step Δt has to be chosen carefully for the process to be stable. Calculations of an upper bound for stability in similar settings[13] suggest $\Delta t = 0.2$ as an appropriate choice, but simulations using the current discretization have shown that in practice we can obtain stable solutions by using time steps as large as $\Delta t = 0.8$.

The explicit scheme explained here is, of course, very basic, and because of the restriction in the size of the time step it can be rather time consuming. More advanced numerical methods should be well suited for improving the efficiency of the current filtering scheme. Among these are the use of adaptive grid finite elements,[14] semi-implicit schemes based on additive operator splitting,[15] and parallel implementations.[16]

3. RESULTS AND DISCUSSION

A. Planar Laser-Induced Fluorescence Imaging of OH

Figure 2B corresponds to an example of a low-quality OH PLIF image obtained in the turbulent diffusion flame displayed in Fig. 2A. It represents a two-dimensional cut through the region of the flame indicated by the rectangular box. Similar measurements have been used by the present group to study turbulence chemistry interactions and have been the subject of other publications.[8,17] For comparisons with model calculations of such flames it is necessary to extract information such as contour lengths and areas from such data. Nonlinear diffusion was used to enhance the images before extraction of these parameters.

The PLIF image shown in Fig. 2B is an 8-bit image with a size of 261 pixels × 241 pixels, corresponding to a region of 24 mm × 23 mm in the flame, starting 40 mm above the burner nozzle. The PLIF image shown in Fig. 2B corresponds to raw data from which a nearly constant background level has been subtracted and which has been compensated for laser beam profile variations. Details of the beam profile compensation scheme are given in Ref. 18. Apart from this the image is unprocessed.

A main goal of the current research is to extract information about the position of the so-called flame front, which is where chemical reactions in the flame are taking place. As a good approximation, the flame front lies close to the region where the maximum gradient of the OH concentration lies.[17]

In the PLIF image the OH concentration gradients are not perfectly sharp, owing to a limited detector resolution caused by imperfect imaging and cross talk between pixels and also owing to physical phenomena such as convective and diffusive mixing of OH with ambient gases. Since the production terms for OH are highly nonlinear in nature and are strongly temperature dependent, there are also large gray-scale variations (and associated gradients) *within* regions where OH occurs. Also, there are high levels of noise present, which can be attributed to several sources: scattered stray laser light that enters the detector, noise stemming from the signal amplification process in the image intensifiers of the camera, and noise associated with the CCD devices in the camera. These properties make it difficult to identify the flame front directly with other conventional methods, such as thresholding or Canny–Deriche edge detection.

B. Properties of Nonlinear Diffusion

In this section some of the properties of the nonlinear diffusion approach described here will be demonstrated by application to a number of artificially generated one-dimensional test curves. These test curves simulate characteristics expected in real experimental situations using the PLIF system mentioned above. The behavior of the filtering technique for images containing gradient

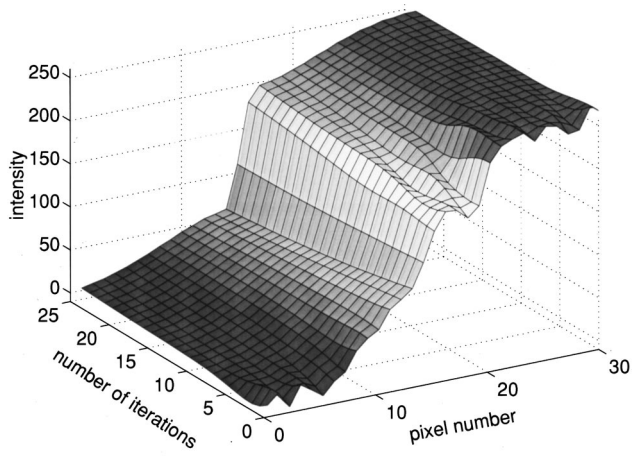

Fig. 3. Sharpening of an edge by nonlinear anisotropic diffusion. The plot shows the evolution of a noisy, blurred step function during 25 computational iterations. Clearly the edge is strongly enhanced, while noise fluctuations are effectively smoothed out by the process.

variations as well as variations in signal-to-noise ratios can be visualized in this way. This part of the paper also explains the issues regarding the optimal adaptation of the technique to a given experimental situation.

The basic properties of the process are exhibited in Fig. 3. It shows the progressive filtering of a blurred, noisy edge, which was generated by an error function and to which white noise was added. The edge intensity varies between 0 and ~250, corresponding to the dynamic range of the present camera system. The standard deviation of the noise was 7.5% of the peak intensity of the error function. Weickert's diffusivity function [Eq. (5)] was used in Eq. (6), with $m = 8$, $C_m = 3.315$, $\lambda = 15$, and $\sigma = 1$. After only 25 computational iterations the edge is strongly enhanced, although its position is unaffected. At the same time the noise is effectively removed.

Physically one can picture the diffusion process as approaching the one-dimensional edge from both directions without actually crossing it. In the process noise is averaged out, and, given an appropriate diffusivity function, the edge may be sharpened at the same time. Because in the present application the position of the edge contour is of interest, this edge enhancement is a highly desirable feature, and Eq. (5) has proved to be very effective to this end. Note, however, that other functions may be chosen for $g(s)$ if this edge sharpening effect is not desirable.

To see which edges are preserved during the process, it is useful to define a flow $\Phi(\nabla u) = g(|\nabla u|)\nabla u$ and to plot Φ against the gradient[13] ∇u. In Fig. 4 this is shown for three different diffusivity functions. The two dotted curves correspond to functions originally proposed by Perona and Malik,[9] whereas the solid curve corresponds to Eq. (5) for $m = 8$. All three curves are plotted in terms of λ and are normalized to effect maximum flow at $\lambda = 1$. Areas are normalized to unity. The principal behavior is the same for the three functions: Diffusion is rising to a maximum near $\lambda = 1$ and then falling off again for gradients $\nabla u > 1$. In a given situation one must choose a parameter set for which the flow $\Phi(\nabla u)$ is a maximum for the gradients exhibited by noisy features and a minimum for the edges that must be preserved. Weickert's function, for $m = 8$, is the most selective of the three, peaked sharply around $\lambda = 1$. The flow of Weickert's diffusivity function for lower m values is less sharply peaked and more closely resembles the ones proposed by Perona and Malik. In the present case Weickert's function for $m = 8$ was found to match the noise characteristics of the detector in an ideal way (the noise spectrum resembled a normal distribution in the present case). This means that good results can be achieved after relatively few iterations (typically 100 iterations or less for a given PLIF image; see subsection 3.C). A drawback is that the approach becomes more sensitive to the proper choice of λ and is numerically less stable than the other two functions shown in Fig. 4. Much of this instability is prevented by use of the smoothed kernel as shown in Eq. (6). In practice we find that good results are achieved if λ is set to ~70% of the minimum edge strength that we wish to preserve and a σ within the interval 0.5–1 is used. Figure 5 shows the behavior of different terms appearing in Eq. (6) when a noise-free but blurred edge is used. Figure 5 is an aid to qualitative understanding of the edge-sharpening process. Figure 5A shows the edge function u after n equal to 1, 10, and 100 computational iterations. The gradual steepening as n increases is clearly exhibited. In Fig. 5B the gradient ∇u is plotted. This shows quantitatively how the gradient in u becomes steeper in time. Finally, Fig. 5C displays the rate of change of the gradient, $\partial_t \nabla u$. Initially one can see that the largest rate of change occurs at the point where the gradient ∇u is largest. $\partial_t \nabla u$ is positive here, so the edge increases its sharpness. Simultaneously there are two negative lobes in $\partial_t \nabla u$ to the left and to the right of the center peak, giving the effect that the slope of u is flattened at these positions, which is reflected in Fig. 5A. At $n = 10$, $\partial_t \nabla u$ has reached a maximum near the center. In contrast, the sidelobes are decreasing again in magnitude, indicating that edge diffusion is slower there than

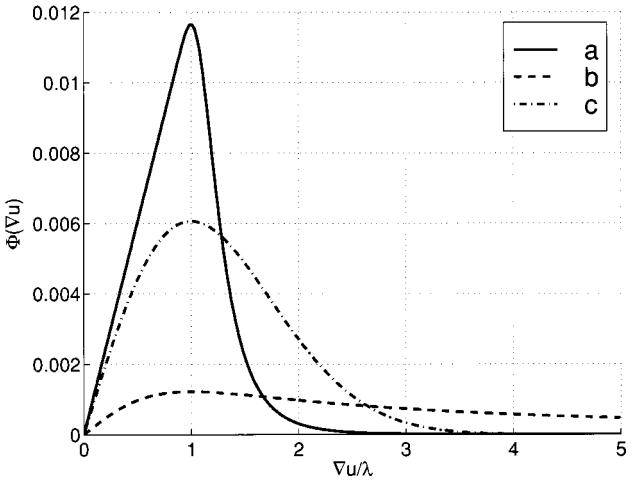

Fig. 4. Graph displaying the diffusive flow $g(|\nabla u|)\nabla u$ for three different diffusivity functions g as a function of the edge strengths $\nabla u/\lambda$ in the image. The functions are, a, the function used in this paper [Eq. (5)], $g(s) = 1 - \exp[-C_m/(s/\lambda)^8]$ and, b and c, two functions originally proposed by Perona and Malik, $g(s) = 1/[1 + (|s|/\lambda)^2]$ and $g(s) = \exp[-(|s|/\lambda)^2]$, respectively. All three curves are normalized to unity area with maximum flow occurring at $\nabla u = \lambda$.

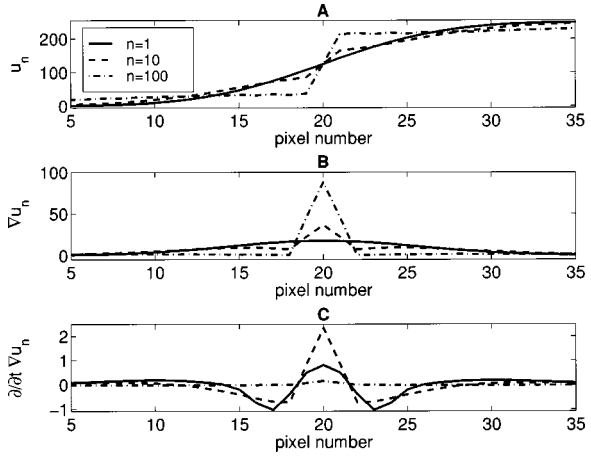

Fig. 5. Behavior of an edge propagating through the diffusion filtering process. The evolutions of A, u; B, ∇u; and C, $\partial_t \nabla u$ are shown for n equal to 1, 10, and 100 iteration steps, respectively.

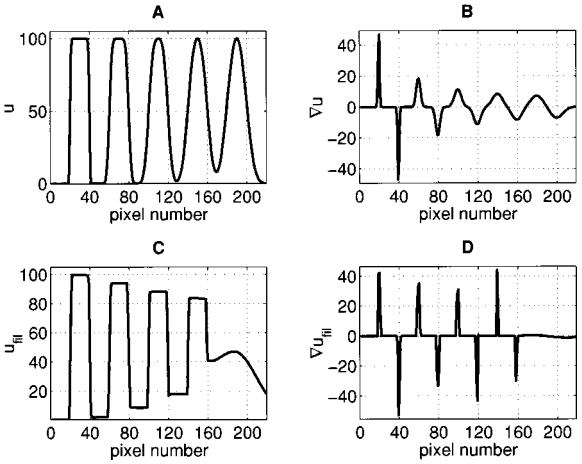

Fig. 6. Illustration of the gradient-selective behavior of the diffusion process. A, Unprocessed periodic function u; B, its gradient ∇u. Function u contains a large range of gradients, successively decreasing for each period. C, Filtered function u_{fil}; D, its gradient ∇u_{fil}. It can be seen that the peak with the lowest gradient (the rightmost peak on in graph A) is not retained in the filtered image. All other gradients are strongly enhanced.

initially. Finally, at $n = 100$, diffusion has virtually come to a standstill over the entire region shown, although a minute amount of steepening is still active near the center of the graph. From this time on nothing much is gained by a continuation of the process. In practice, therefore, one should balance the degree of smoothing and sharpening that one wishes to achieve with the associated computational cost (number of iterations required to achieve the task). We found that $n = 150$ was sufficient for all measurement situations that we encountered. Note, however, that there is no rule, in principle, governing the maximum number of iterations, although numerical rounding errors will ultimately lead to a uniform intensity distribution in the limit $n \to \infty$.

Figure 6 displays the gradient selectivity of the method. In Fig. 6A a function u is shown, consisting of a series of steps with equal periods (40 pixels) but decreasing edge gradients. The gradients corresponding to u are shown in Fig. 6B. Figure 6C shows u_{fil}, which is the same function after filtering for 25 iterations; in Fig. 6D the gradients corresponding to the filtered function are shown. In this example, all gradients are strongly enhanced except for the last feature shown, the gradient of which is too low to be preserved by the filtering process. The sizes of gradients preserved are governed by the shape of g in Eq. (6), and are strongly dependent on the value of the contrast parameter λ. In the present case $\lambda = 6.4$ was chosen. A smaller λ would have enhanced all features in the process. In a practical situation one has to adapt λ to the physical situation prevailing, and it is clear that the method will fail if the gradients of the structures that one wishes to preserve are similar to a significant part of the gradients present in the noise spectrum.

The performance of the method in the presence of various degrees of noise is seen in Fig. 7. Again a series of structures is shown with a period of 40 pixels and a nominal amplitude of 100 counts (on top of a dc level of 100 counts). The standard deviation of the noise added to each peak (from left to right) is 0, 10%, 20%, 30% and 40% of the maximum signal amplitude. It can be seen that the maximum gradients in the original function are almost completely restored, despite the severe noise levels present.

In a given experimental situation the best choice of λ can be obtained by a spectral analysis of the gradients present in the noise and by adapting $g(\nabla u, \lambda)$ to maximize diffusion over the noise components while preserving the frequency components corresponding to edges that one wishes to preserve. An example of this is shown in Fig. 8, which corresponds to a function of the type displayed in Fig. 7. Figure 8A shows the gradient of the original, noise-free periodic function, and Fig. 8B shows its gradient spectrum, upon which the flow function $g(|\nabla u|)\nabla u$ is superimposed. Figure 8C shows the gradient of the line image corrupted by increasing amounts of

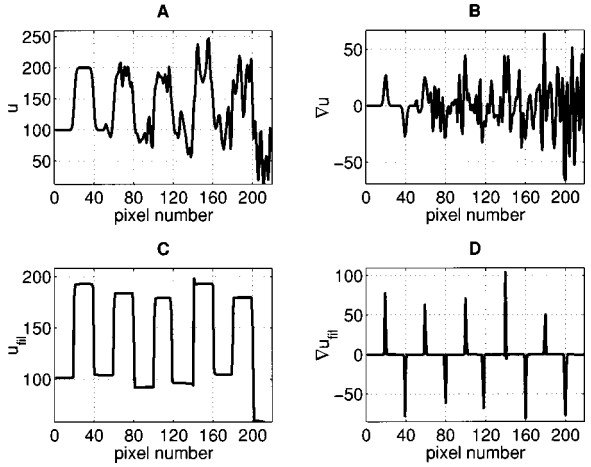

Fig. 7. Filtering of periodic structures with increasing degrees of white noise added and consequently decreasing signal-to-noise ratios. A, Unfiltered function u; B, its gradient. C, After filtering, the features of the original periodic function are restored and enhanced (u_{fil}); D, the corresponding gradient.

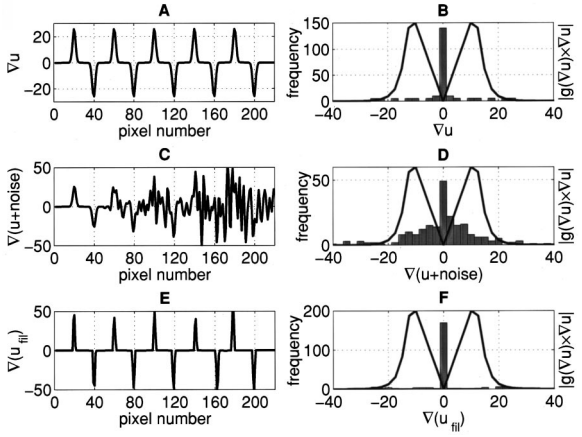

Fig. 8. A, Gradients ∇u of an unperturbed periodic structure u; C, gradients $\nabla(u + \text{noise})$ of the same structure with increasing levels of noise added (analogous to Fig. 7); E, gradients $\nabla(u_{\mathrm{fil}})$ of the filtered structure. B, D, and F, Corresponding histograms of the gradient distributions. The magnitude of the flow $|g(\nabla u)\nabla u|$ as a function of ∇u is superimposed on the histogram to help visualize which gradients are preserved by the process.

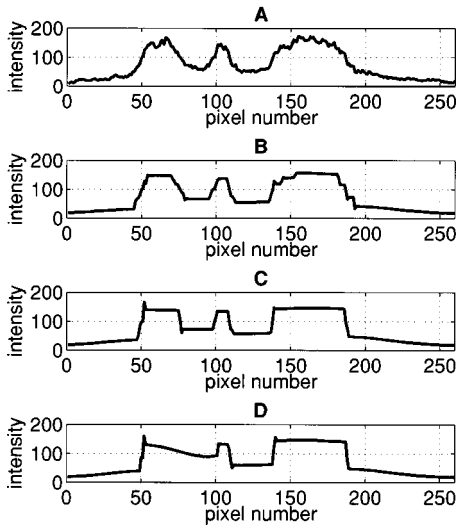

Fig. 9. Effects of the contrast parameter λ on the retainment of image structure during nonlinear diffusion filtering. The line profile shown corresponds to the pixel row indicated in white in Fig. 2B. A, Raw intensity profile; B, C, and D, filtered profile for λ equal to 3, 5, and 7, respectively.

white noise. Figure 8E shows the gradients in the noisy line image after filtering with $n = 100$. Comparing the histograms, one can clearly see how components corresponding to noise have effectively diffused away. Note that some of the moderate gradient components in the original also diffuse away, which leads to the gradient sharpening. In the presented case the optimal diffusivity function corresponds to $\lambda = 11$. For $\lambda > 11$ the original edges are not retained, but for $\lambda < 11$ some noise gradients may be amplified. In extreme situations such as this example, where the gradient spectrum of the noise overlaps significantly with that of the original function, the correct choice of λ is crucial. As is discussed in Ref. 13 in a similar context, the optimal λ should be chosen to be as small as possible so that the largest possible range of physical gradients can be preserved. In the context of the PLIF images that we present below, we found $\lambda = 1.2\sigma_n$, where σ_n is the standard deviation of the noise gradient distribution, to be a good choice for almost all experimental situations encountered.

C. Nonlinear Diffusion Filtering of OH Planar Laser-Induced Fluorescence Images

An example of an OH PLIF image before and after nonlinear filtering is shown in Fig. 2, where Fig. 2B is the unprocessed OH PLIF image and Fig. 2C shows the image after nonlinear diffusion filtering, by use of Weickert's diffusivity function [Eq. (5)], with $m = 8$. The image corresponds to 300 computational iterations with $\lambda = 4.5$ and $\sigma = 1$. Clearly the signal-to-noise and contrast-to-noise ratios are strongly enhanced, thus simplifying a successive segmentation of OH boundaries (see subsection 3.D). In Fig. 9 the effect of the contrast parameter λ on the filtered images is illustrated. Figure 9A shows the intensity along one pixel row, corresponding to the line indicated in white in Fig. 2B; in Figs. 9B–9D the filtered data is shown for λ equal to 3, 5, and 7. The gradient sharpening is evident from these cuts. Note the effect of the contrast parameter on the preservation of image structure: As λ is increased, small-scale structure is lost, and in Fig. 9D, for example, the middle peak of OH is lost. It is clear, therefore, that a proper choice of the contrast parameter is crucial. However, since the noise-gradient spectrum was similar in all experimental situations we encountered, a single λ value (obtained by the method outlined in Subsection 3.B) was usually sufficient to produce satisfactory results for nearly all images processed within an experimental run. The choice of the smoothing kernel, characterized by σ, and the number of time steps n is far less critical. Good results were achieved for $n = 150$ and $\sigma = 1$ in all cases.

Figure 10 shows three-dimensional plots of Figs. 2B and 2C. A comparison of the two images illustrates several attractive properties of the filtering process. Noise levels are drastically reduced without compromising overall intensities. Small structure is preserved, while OH boundaries are strongly enhanced.

D. Flame-Front Identification

One can easily identify the OH boundary (the flame front) from the diffusion filtered images by computing the gradient of the filtered image and thresholding and binarizing the resulting gradient image. The boundary of the

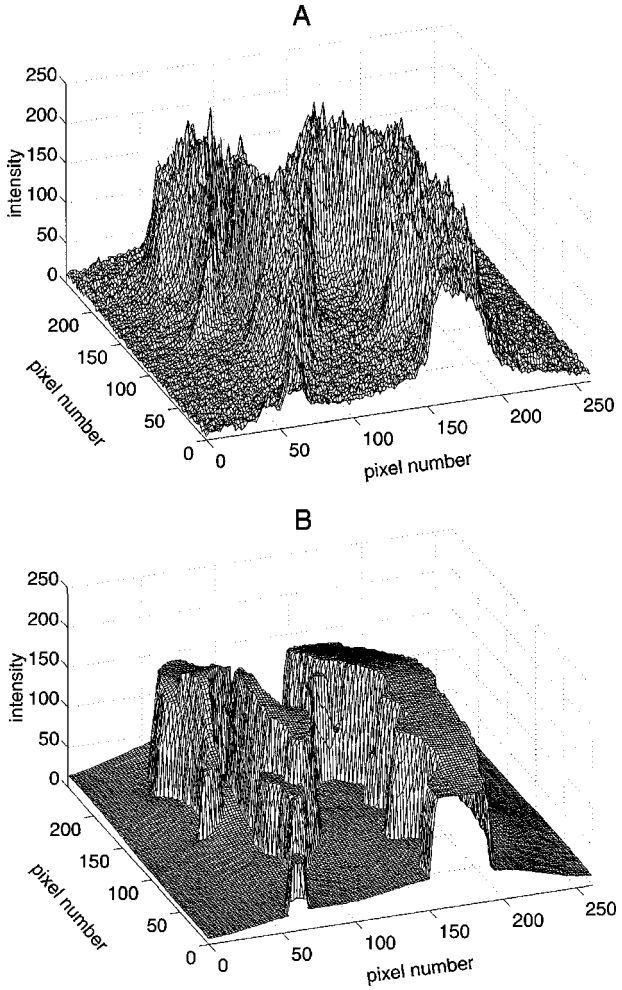

Fig. 10. Three-dimensional plots of, A, the original PLIF image and, B, the nonlinear-diffusion-filtered image.

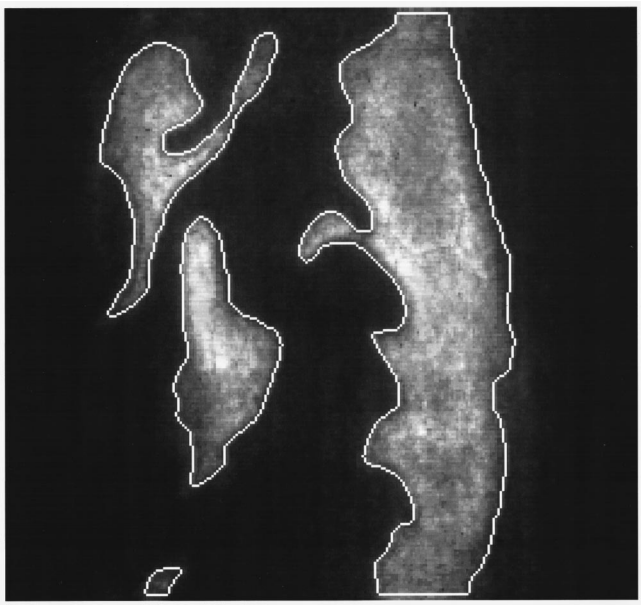

Fig. 11. Original PLIF raw image with the identified flame fronts superimposed as white outlines.

resulting image is then taken as the OH boundary. In Fig. 11 this procedure has been applied to Fig. 2C, and the detected boundary has been superimposed on the raw image (Fig. 2B). From the identified flame front, the circumference-to-area ratio, η, of the burned regions in the flame can be calculated, which is a parameter that is highly suitable for comparisons with theoretical flame calculations.[17]

An important issue in the present application is how the choice of parameters for the nonlinear diffusion filter influences η. To investigate this quantitatively, we evaluated the sensitivity of η to the choice of parameters in Eq. (6) and to the signal-to-noise level present in the raw image.

For this purpose the OH PLIF image in Fig. 2B was first nonlinear diffusion filtered with a range of values used for λ, σ, n, and m. The flame front was then extracted from the filtered images by the method outlined above, and η was evaluated. The variation of η in relation to these parameters is plotted in Figs. 12A–12D.

In Fig. 12A, η (normalized to 1 for the parameter values used in subsection 3.C) is shown as a function of the contrast parameter λ; η is seen to increase with increasing λ. This result can be explained, because small features with low gradients are not preserved by the filter at large λ (see Fig. 9D), and as a consequence the evaluated area decreases, thus making η larger. As can be seen, η increases by ~13% for when λ increases by nearly a factor of 2. As was discussed above, the critical dependence on the choice of λ (caused by the high selectivity of Weickert's diffusion kernel) is a weakness of the method, which may limit its use in some applications. For λ larger than 5.5, in the present example, structures start to merge or disappear, and some edges start to diffuse away after multiple iterations.

In Fig. 12B, η is plotted as a function of the smoothing kernel width, σ. Over the range $0.4 \leq \sigma \leq 1.2$, η stays nearly constant, and thus the choice of σ is not critical in the present application. For values of σ smaller than 0.4, the flame area is overestimated because noisy structures are preserved. Values of σ larger than 1.2 result in the same problem as the classical diffusion discussed in Section 2: Edges become blurred to a degree that physical information is lost.

In Fig. 12C the evolution of η with n, the number of computational iterations, is shown. It can be seen that η stabilizes after slightly more than 100 iterations, although there is actually a continuous small decrease caused by numerical rounding errors.

In Fig. 12D, η is plotted for the different diffusivity functions associated with m values ranging from 5 to 15. It is seen that η decreases only slightly with m over this range. It must be noted, however, that m is not interpretable as a diffusion filtering parameter like λ, σ, and n in Figs. 12A–12C, but rather as a number that defines the shape for the diffusivity function. In this context the optimal choice of λ to suppress noise and sharpen gradients might not be the same for the diffusivity functions associated with different m values (see also the discussion in subsection 3.B).

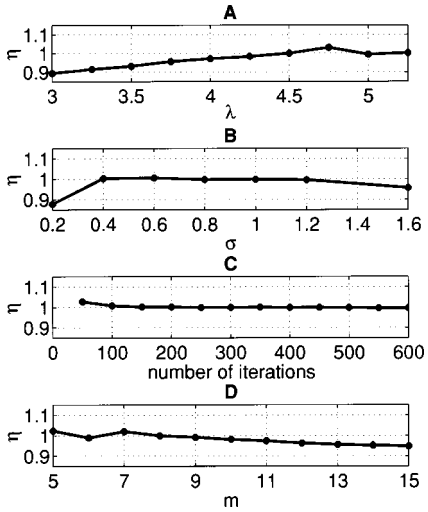

Fig. 12. Dependence of the circumference to area ratio, η, as a function of different filter parameters. A, Dependence on λ; B, on σ; C, on n; D on m.

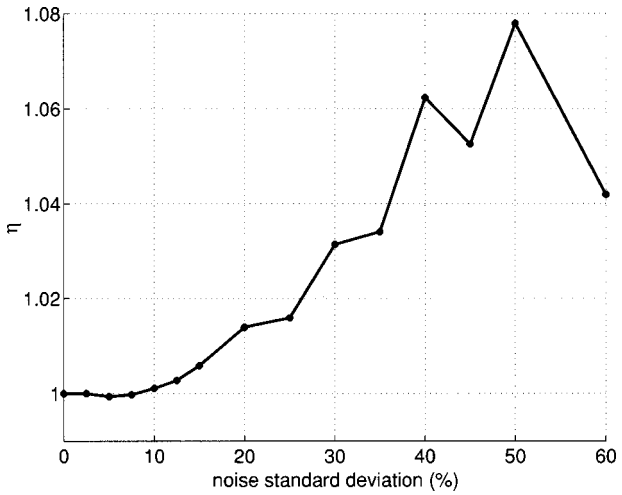

Fig. 13. Circumference to area ratio, η, as a function of the noise level in an artificially generated test image.

To investigate the influence of the signal-to-noise ratio on η, we generated a series of test images with increasing degrees of noise superimposed. For this purpose a test image was generated from the OH-contour depicted in Fig. 11 with infinite contrast and gradients like those that naturally occur in such flames. Gaussian noise, with a standard deviation ranging from 0 to 60% of the signal level, was added to this test image. The ratio η was then calculated in the same manner as described above. The results are shown in Fig. 13. For noise levels up to 10%, corresponding to a signal-to-noise ratio of 10, η is constant to within less than a percent, highlighting the strength of the current method. For larger noise levels η is increasing because some artificial structure is being generated by the noise. Note, however, that even for noise levels as high as 35% (corresponding to a signal-to-noise ratio of only 3, a situation rarely present in an actual experiment) the increase in η is only $\sim$4%.

4. CONCLUSION

The present paper describes the adaptation and application of anisotropic nonlinear diffusion filtering to planar, laser-based, spectroscopic imaging methods. The filter is well suited to suppress noise and to enhance structural information in such data. We use nonlinear diffusion filtering to extract information on object contours, in particular, for the evaluation of reaction boundaries in turbulent flows. Advantages of the methods are its efficiency, even in the presence of severe noise levels, and its contour-enhancing properties, which greatly simplify subsequent object identification and image segmentation steps. Issues regarding the application of the method for given environments are discussed, as well as the influence of several critical parameters on the model's precision.

Although the technique has been applied in the particular context of laser-induced fluorescence imaging, it should be equally suited for other planar laser diagnostic techniques, for example, planar imaging by Mie, Rayleigh, or Raman spectroscopy. The technique is currently in use for data extraction from measurements performed in turbulent flames.[19,20]

Corresponding author Johan Hult's e-mail address is johan.hult@forbrf.lth.se

*Present address, Department of Chemical Engineering, Pembroke Street, Cambridge C32 3RA, UK.

REFERENCES

1. R. K. Hanson, "Combustion diagnostics: planar imaging techniques," in *Proceedings of the Combustion Institute* (Combustion Institute, Pittsburgh, Pa., 1986), Vol. 21, pp. 1677–1691.
2. B. Yip, J. K. Lam, M. Winter, and M. B. Long, "Time resolved three-dimensional concentration measurements in a gas jet," Science **235**, 1209–1211 (1987).
3. A. Orth, V. Sick, J. Wolfrum, R. Maly, and M. Zahn, "Simultaneous 2D single-shot imaging of OH concentrations and temperature fields in an SI engine simulator," in *Proceedings of the Combustion Institute* (Combustion Institute, Pittsburgh, Pa., 1994), Vol. 25, pp. 143–150.
4. I. van Cruyningen, A. Lozano, and R. K. Hanson, "Planar imaging of concentration by planar laser induced fluorescence," Exp. Fluids **10**, 41–49 (1990).
5. C. F. Kaminski, J. Engström, and M. Aldén, "Quasi instantaneous 2 dimensional temperature measurements in a spark ignition engine using 2-line atomic fluorescence," in *Proceedings of the Combustion Institute* (Combustion Institute, Pittsburgh, Pa., 1998), Vol. 27, pp. 85–93.
6. A. C. Eckbreth, *Laaser Diagnostics for Combustion Temperature and Species*, Vol. 3 of Combustion Science and Technology Book Series (Gordon & Breach, Amsterdam, 1996).
7. J. Warnatz, U. Maas, and R. W. Dibble, *Combustion* (Springer-Verlag, Berlin, 1996).
8. C. F. Kaminski, J. Hult, and M. Aldén, "High repetition rate planar laser induced fluorescence of OH in a turbulent non-premixed flame," Appl. Phys. B **68**, 757–760 (1999).
9. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," IEEE Trans. Pattern Anal. Mach. Intell. **12**, 629–639 (1990).
10. H. Malm, "Multiscale image analysis and image enhancement based on partial differential equations," Master's thesis (Lund Institute of Technology, Lund, 1998).
11. J. Weickert, "Anisotropic diffusion in image processing," Ph.D. dissertation (Universität Kaiserslautern, Kaiserslautern, 1996).
12. F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll, "Image selective smoothing and edge detection by nonlinear diffusion," SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **29**, 182–193 (1992).
13. G. Gerig, O. Klüber, R. Kikinis, and F. A. Jolesz, "Nonlinear anisotropic filtering of MRI data," IEEE Trans. Med. Imaging **11**, 221–232 (1992).
14. A. Handlovicova, K. Mikula, and F. Sgallari, "Numerical methods for nonlinear diffusion equations arising in image processing," Preprint No. 1 (Slovak University of Technology, Bratislava, 1999).
15. J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever, "Efficient and reliable schemes for nonlinear diffusion filtering," IEEE Trans. Image Process. **7**, 398–410 (1998).
16. J. Weickert, K. J. Zuiderveld, B. M. ter Haar Romeny, and W. J. Niessen, "Parallel implementation of AOS schemes: A fast way of nonlinear diffusion filtering," in *Proceedings of the 1997 International Conference on Image Processing* (Institute of Electrical and Electronics Engineers, New York, 1997), Vol. 3, pp. 396–399.
17. C. F. Kaminski, J. Hult, M. Aldén, S. Lindenmaier, A. Dreizler, and U. Maas, "Spark ignition of turbulent methane/air mixtures revealed by time resolved planar laser induced fluorescence and direct numerical simulations," in *Proceedings of the Combustion Institute* (Combustion Institute, Pittsburgh, Pa., 2000), Vol. 28.
18. A. Dreizler, S. Lindenmaier, U. Maas, J. Hult, M. Aldén, and C. F. Kaminski, "Characterization of a spark ignition system by planar laser induced fluorescence of OH at high repetition rates and comparison with chemical kinetic calculations," Appl. Phys. B **70**, 287–294 (2000).
19. T. Ding, Th. H. van der Meer, M. Versluis, M. Golombok, J. Hult, M. Aldén, and C. F. Kaminski, "Time-resolved PLIF measurements in turbulent diffusion flames," in *Third International Symposium on Turbulence, Heat and Mass Transfer*, Y. Nagano, K. Hanjalic, and T. Tsuji, eds. (Aichi Shuppan, Tokyo, 2000), pp. 857–864.
20. C. F. Kaminski, X. S. Bai, J. Hult, A. Dreizler, S. Lindenmaier, and L. Fuchs, "Flame growth and wrinkling in a turbulent flow," Appl. Phys. B (to be published).