Photo-Consistent 3D Fire by Flame-Sheet Decomposition

Samuel W. Hasinoff Kiriakos N. Kutulakos

Dept. of Computer Science
University of Toronto
Toronto, Canada M5S 3G4
{hasinoff,kyros}@cs.toronto.edu

Abstract

This paper considers the problem of reconstructing visually realistic 3D models of fire from a very small set of simultaneous views (even two). By modeling fire as a semi-transparent 3D density field, we show that fire reconstruction is equivalent to a severely under-constrained computerized tomography problem, for which traditional methods break down. Our approach is based on the observation that every pair of photographs of a semi-transparent scene defines a unique density field, called a Flame Sheet, that (1) concentrates all its density on one connected, semi-transparent surface, (2) reproduces the two photos exactly, and (3) is the most spatially-coherent density field that does so. From this observation, we reduce fire reconstruction to the convex combination of sheet-like density fields, each of which is derived from the Flame Sheet of two input photos. Experimental results suggest that this method enables high-quality view extrapolation without over-fitting artifacts.

1. Introduction

The computational modeling of physical phenomena such as fire and smoke has received significant attention in computer graphics [1–6] as well as in other fields of experimental science [7–13]. While photographs provide a great deal of information about such phenomena, very little is known about how 3D models can be extracted from images. Extracting such information could open up new opportunities for creating better visual models [1–3]; developing image-based representations of fire [14]; permitting the manipulation of scene appearance in photographs that contain such phenomena; and developing new, dynamic simulation systems that use real-world data as input.

Toward this end, this paper considers the problem of reconstructing 3D fire densities from a set of simultaneously-captured images or videos. Using optical models of fire developed in the combustion literature [8, 9, 15], we model fire as a semi-transparent density field whose appearance at a projected pixel is a linear function of the transparencies along the corresponding ray.

We believe that any practical method for the image-based modeling of fire must satisfy three criteria:

- **Small viewpoint coverage:** Fire is an inherently dynamic phenomenon and, as such, requires simultaneous image capture. While multi-view systems with tens of video cameras do exist [8, 16], the method should not break down for small datasets.
- **Photo-consistency:** To preserve visual realism, reconstructed densities must reproduce the input images.
- **Good view extrapolation:** All views should be rendered with high quality, without "over-fitting" the limited number of input views.

Even though these criteria have received much attention in the case of opaque objects [16], satisfying them for fire or other semi-transparent scenes is not well understood.

At the heart of our approach lies the observation that every pair of photos of a transparent density field uniquely determines two special fields, called *Flame Sheets*, each of which reproduces the photographs exactly and concentrates all its density along a single connected surface. Using this observation as a starting point, we show that: (1) Flame Sheets can be decomposed into a family of solutions involving further sheet-like structures, (2) this family is a basis of the space of density fields photo-consistent with two or more views, and (3) the space of photo-consistent density fields is linear and convex. These results lead to a simple and efficient algorithm, called *Flame-Sheet Decomposition*, that computes the fire density field as a convex combination of sheet-like densities derived from pairs of input views.

While not a topic of previous computer vision research, several methods in the combustion literature have modeled 3D flames and their physical properties. Common approaches use lasers [12], complex optical systems [8], and thermography devices [7, 9] to extract a fire's physical properties. The most closely related work involves the use of computerized tomography methods for flame reconstruc-

The authors gratefully acknowledge the support of the National Science Foundation under Grant No. IRI-9875628, of the Natural Sciences and Engineering Research Council of Canada under the RGPIN and PGS-A programs, and of the Alfred P. Sloan Foundation.

tion. Because classic back-projection or algebraic methods [17] require tens or hundreds of frames to obtain accurate results, these methods have not been applied to fire datasets. Even tomography methods specialized to the sparse-view case, typically favoring local smoothness in a statistical framework [18, 19], do not generate 3D reconstructions adequate for view synthesis. Instead, previous methods of flame reconstruction capture images with multiple-second exposures [7, 9, 10] or assume a "stationary" fire [8] to reconstruct rough approximations to the density field. As a result, these methods are not appropriate for view extrapolation and cannot be used to model flickering or complex flames (e.g., Figure 9), whose structures can change dramatically from one instant to the next.

Unlike existing methods, Flame-Sheet Decomposition is specifically designed to capture models that are both photo-consistent and have good extrapolation capabilities. Intuitively, Flame Sheets represent the most spatially coherent interpretation of the input views. Hence, even though the problem of computing densities from a few images is inherently under-constrained, Flame Sheets provide a strong bias toward solutions that are both photo-consistent and spatially coherent. Moreover, unlike tomographic methods, which are ill-posed for a small number of views, Flame Sheets are unique and lead to a reconstruction method that is well-posed and easy to implement. This allows us to reconstruct flames from as few as two input views, while also being able to incorporate more views if they are available.

Our approach offers five contributions over the existing state of the art. First, unlike methods where photo-consistency is the only objective, our approach establishes both photo-consistency and spatial coherence as objectives for visual reconstruction. Second, it leads to a well-posed reconstruction algorithm that can handle any number of input views. Third, it introduces Flame Sheets as a complete basis for the space of photo-consistent density fields, enabling their potential use for reconstructing and rendering semi-transparent scenes other than fire (whenever the linear image formation model is valid). Fourth, the sheet-like structure of the Flame-Sheet Decomposition enables use of simple warp-based methods [20] to create photo-realistic views of the reconstructed fire. Fifth, our results show that the algorithm is able to render detailed reconstructions of complex flames without over-fitting artifacts.

2. Image Formation Model

Consider a fire viewed by a set of orthographic cameras lying on a single viewing plane (Figure 1). We represent this fire as a 3D density field, $\mathbf{D}(r, c, z)$, of luminous fire reaction products, with the z-axis perpendicular to the viewing plane. Reconstruction of this density field can therefore be expressed as a sequence of 2D reconstruction problems, each of which reconstructs a single slice, $\mathbf{D}(r, c, z_0)$, of the field. Our goal is to reconstruct such a slice from corresponding epipolar lines in the input views. To achieve

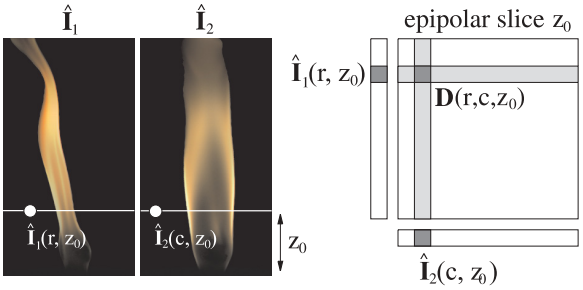

Figure 1. Viewing geometry. Two simultaneous orthographic images of a flame corresponding to a 90° rotation about the vertical z-axis. White horizontal lines indicate the corresponding epipolar lines at height z_0. A top view of their epipolar plane is shown on the right.

this, we first relate the 2D field, $\mathbf{D}(r, c, z_0)$, to the one-dimensional images it generates. For notational simplicity, we drop the z coordinate for the remainder of the paper.

The image formation model of fire can be expressed as a special case of the general radiative transfer model. Given an image $\hat{\mathbf{I}}$ aligned with the c-axis, the intensity of pixel $\hat{\mathbf{I}}(c)$ is a sum of two terms—a term that integrates radiance from luminous fire material on the ray through that pixel and a term that incorporates background radiance [21]:

$$\hat{\mathbf{I}}(c) \;=\; \int_0^L \mathbf{D}(r, c)\, \tau(r) J(r)\, \mathrm{d}r \;\; + \;\; \hat{\mathbf{I}}_{bg} \tau(L), \qquad (1)$$

where $\mathbf{D}$ is the fire's density field; J models the total effect of emission and in-scatter less any out-scatter per unit mass; L defines the interval $[0, L]$ along the ray where the field is non-zero; $\hat{\mathbf{I}}_{bg}$ is the radiance of the background; and τ models transparency, i.e., the degree by which the fire products between the pixel and a position x along the ray block radiance further away from reaching that pixel,

$$\tau(x) = \exp\left(-\sigma_t \int_0^x \mathbf{D}(r, c)\, \mathrm{d}r\right), \qquad (2)$$

where σ_t is a positive constant dependent on the medium.

To simplify this model further, we rely on two assumptions used in existing literature on combustion analysis:

- **Negligible scattering:** This is a good approximation for fire not substantially obscured by smoke because emission from luminous soot particles dominates radiance [3, 8, 9, 13]. In this case, the function $J(r)$ becomes a pure self-emission term.
- **Constant self-emission:** This assumption models fires whose brightness depends only on the density of luminous soot particles [15], allowing us to assume that self-emission is constant per unit mass, i.e., that $J(r)$ is a constant, J_0.

Together, these assumptions lead to an expression for $\hat{\mathbf{I}}(c)$ that depends only on the fire's transparency along the ray:

$$\hat{\mathbf{I}}(c) = J_0\,(1 - \tau(L)) \;\; + \;\; \hat{\mathbf{I}}_{bg} \tau(L). \qquad (3)$$

This expression is especially important because it allows us, through a simple intensity transformation, to establish a linear relation between a pixel's intensity and the density field along the ray through that pixel. In particular, when the camera's settings do not allow pixel saturation, we have

$$\mathbf{I}(c) = -\log\left(\frac{\hat{\mathbf{I}}(c) - J_0}{\hat{\mathbf{I}}_{bg} - J_0}\right) = \sigma_t \int_0^L \mathbf{D}(r, c)\,\mathrm{d}r. \qquad (4)$$

In practice, images are discrete and the factor σ_t, which is constant across all pixels and input views, is unknown. Moreover, the linearity of Eq. (4) tells us that the sum of all transformed pixel intensities in an image does not change with viewpoint. Therefore, without loss of generality, we represent each image as an N-dimensional column vector of transformed pixel intensities whose rows sum to one,[1] and represent the fire density field as an $N \times N$ matrix $\mathbf{D}$ that can be recovered only up to an unknown scale factor. This leads to the following 3D fire reconstruction problem:

Definition 1 (Photo-Consistent Fire Reconstruction). Given column vectors $\mathbf{I}_1, \ldots, \mathbf{I}_V$ representing V views of the fire density field, reconstruct a matrix $\mathbf{D}$ that is *photo-consistent* with these views, i.e., every pixel $\mathbf{I}_v(c)$ satisfies the discrete form of Eq. (4).

3. The Flame Sheet

For two views of a fire density field, the reconstruction problem is extremely under-determined (Figure 2). In fact, given images $\mathbf{I}_1$ and $\mathbf{I}_2$ corresponding to row and column sums of the field respectively, it is trivial to show that

$$\mathbf{D} = \mathbf{I}_1 \mathbf{I}_2^T$$

is a solution, which we call the *multiplication solution*.[2] While photo-consistent, the multiplication solution does not generalize to more views and causes significant artifacts during view synthesis. For instance, Figure 2 shows that double images are created in intermediate views, resulting in blurring (see also Figure 9, rows 1-2). Intuitively, the multiplication solution represents the most spread-out and least-coherent solution.[3]

To avoid these problems while preserving photo-consistency, we define an alternative, maximally-coherent solution. Intuitively, the construction of this solution can be visualized as follows (Figure 3a,b). Imagine "pushing" the densities accumulated in image $\mathbf{I}$ upward along the columns of the density field, concentrating them on a monotonic curve, so that the field becomes photo-consistent with the other input image. This pushing and spreading procedure will always be possible because both images sum to one. We describe a method for constructing two such curves of opposite diagonal orientation, called *Flame Sheets*, in the course of an inductive proof for their existence.

[1] This normalization implicitly assigns a value to the unknown σ_t.

[2] Note that any two images can be reduced to this orthogonal-view case by a known and invertible 2D warp [20] of the epipolar plane.

[3] It is possible to show that it is the maximum entropy solution.

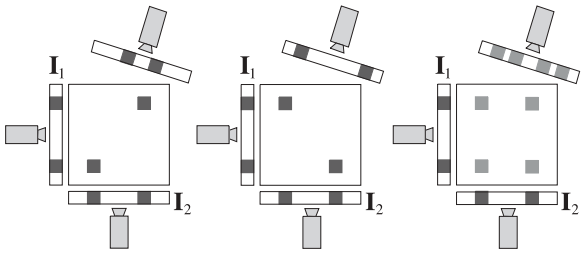

Figure 2. Ambiguities in two-view fire reconstruction. Dark gray squares represent twice the density of light gray ones. All three density fields are photo-consistent with the two orthogonal views, $\mathbf{I}_1, \mathbf{I}_2$, but their images differ along the third. The right-most field is the multiplication solution.

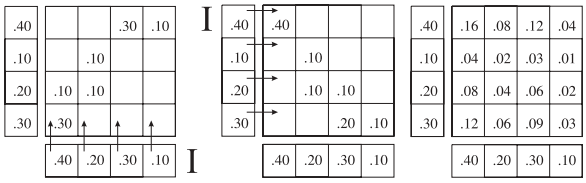

Figure 3. Creating Flame Sheets. (a) Moving from left to right in image $\mathbf{I}$, the density of each pixel in $\mathbf{I}$ is "pushed" in the direction of the arrows to ensure the field's photo-consistency with the second image. (b) A second sheet is defined by switching the roles of the input images. (c) The corresponding multiplication solution.

Theorem 1 (Flame Sheet Theorem). Every pair of non-negative vectors, $\mathbf{I}_1, \mathbf{I}_2$, with equal column sums define two unique matrices, $\mathbf{D}, \mathbf{D}'$, with the following properties:

- Their row and column sums are equal to $\mathbf{I}_1, \mathbf{I}_2$ respectively.
- $\mathbf{D} = 0$ everywhere except a discrete monotonic curve from $\mathbf{D}(1,1)$ to $\mathbf{D}(N,N)$.
- $\mathbf{D}' = 0$ everywhere except a discrete monotonic curve from $\mathbf{D}'(1,N)$ to $\mathbf{D}'(N,1)$.

Proof. (By induction) Denote the marginals $\sum_c \mathbf{D}(r,c)$ and $\sum_r \mathbf{D}(r,c)$ as $\mathbf{D}(r)$ and $\mathbf{D}(c)$ respectively. We maintain the following induction hypothesis for a moving "frontier element," (r,c), that traces a Flame Sheet in the density field (Figure 4):

1. The sub-vectors $\mathbf{I}_1(1,\ldots,r-1)$ and $\mathbf{I}_2(1,\ldots,c-1)$ define a Flame Sheet solution.
2. Intensities at the projection of (r,c) are unaccounted for in at most one vector. Specifically, either $\mathbf{D}(r) = \mathbf{I}_1(r)$ and $\mathbf{D}(c) \le \mathbf{I}_2(c)$, or $\mathbf{D}(r) \le \mathbf{I}_1(r)$ and $\mathbf{D}(c) = \mathbf{I}_2(c)$.

For the base case, set $\mathbf{D}$ to zero everywhere except $\mathbf{D}(1,1) = \min(\mathbf{I}_1(1), \mathbf{I}_2(1))$. There are now three cases to consider, each corresponding to a move of the frontier element that expands the curve monotonically while maintaining our induction hypothesis:

- $(\rightarrow)$ $\mathbf{D}(r) < \mathbf{I}_1(r)$
  Set $\mathbf{D}(r, c+1) = \min(\mathbf{I}_1(r) - \mathbf{D}(r), \mathbf{I}_2(c+1))$.
- $(\downarrow)$ $\mathbf{D}(c) < \mathbf{I}_2(c)$
  Set $\mathbf{D}(r+1, c) = \min(\mathbf{I}_1(r+1), \mathbf{I}_2(c) - \mathbf{D}(c))$.
- $(\searrow)$ $\mathbf{D}(r) = \mathbf{I}_1(r)$ and $\mathbf{D}(c) = \mathbf{I}_2(c)$
  Set $\mathbf{D}(r+1, c+1) = \min(\mathbf{I}_1(r+1), \mathbf{I}_2(c+1))$.

Because both vectors have the same sum, the induction will always terminate with no unaccounted density at (N,N). Furthermore,

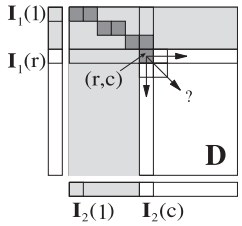

Figure 4. Visualizing the inductive step in the proof of Theorem 1. Intensities in the light gray sub-vectors are fully explained by a field that concentrates non-zero densities along a "partial" Flame Sheet curve (dark gray elements) and has zero density everywhere else (light gray region of $\mathbf{D}$). Depending on the unexplained intensities at the projections of (r, c), the partial Flame Sheet is expanded in one of the three directions shown. A sequence of five such expansions was used to create the Flame Sheet in Figure 3b.

from $\mathbf{D}(r, N)$ the only possible move is $(\downarrow)$, and from $\mathbf{D}(N, c)$, the only possible move is $(\rightarrow)$. The proof for $\mathbf{D}'$ is analogous. □

4. Space of Photo-Consistent Density Fields

While two-view Flame Sheets are adequate for modeling simple flames, they cannot model complex flames with large appearance variations across viewpoint and may actually be perceived as transparent surfaces. To overcome this limitation, we generalize our analysis to multiple views. The following properties of the space of photo-consistent densities give a way to combine multiple 2-view solutions without working directly in the prohibitively high-dimensional space of density fields (Figure 5):

Property 1 (Nesting Property). The space of density fields photo-consistent with k input views is a subset of the space of density fields photo-consistent with $k-1$ of those views.

The Nesting Property suggests that every density field photo-consistent with k views must lie in the intersection of all 2-view solution spaces, for all pairs of input views.

Property 2 (Convexity Property). Every convex combination of photo-consistent density fields is photo-consistent.

Proof. By definition, the discrete line integral corresponding to a particular pixel in an input image has the same value for all photo-consistent density fields. Hence, by the linearity of discrete integration, every convex combination of these integrals has the same value. Thus, convex combinations of density fields are themselves photo-consistent (and non-negative). □

A key consequence of the Convexity Property is that the space of photo-consistent density fields is a *simplex*. This follows because the only two constraints we impose on the space of density fields, i.e., photo-consistency and non-negativity, imply a piecewise-linear boundary. Although only the $O(N^2)$ photo-consistent density fields on the extreme corners of the simplex span the V-view solutions with their convex combination, a sparser sampling on the boundary may still permit a good approximation (Figure 5).

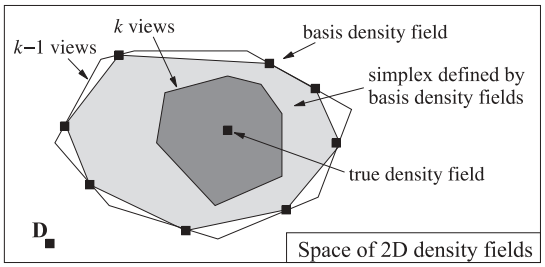

Figure 5. Visualizing the space of density fields. A point in the figure represents a 2D density field $\mathbf{D}$. The space of density fields photo-consistent with $k-1$ views is a region in this space (outer-most polygon). The Nesting Property tells us that as the number of input views increases to k, this region shrinks (dark gray polygon). Moreover, the Convexity Property implies that these regions are always simplices. In general, the simplex defined by a set of "basis $(k-1)$-view density fields" (light gray polygon) may not span the entire V-view solution simplex.

Property 3 (Image Linearity Property). The image of a convex combination of photo-consistent density fields is the convex combination of their respective images.

Proof. Follows directly from linearity of image formation. □

The Image Linearity Property implies that if we represent the space of photo-consistent fields as the convex combination of "basis fields" on its boundary, we can represent photo-consistency constraints as linear functions on the *images* of these fields, rather than the fields themselves. This is especially important from a computational standpoint.

5. Flame-Sheet Decomposition

The solution-space analysis of Section 4 tells us that if we have, for every pair of input views, a set of "basis density fields" whose simplex spans the associated 2-view solution simplex, we can express every density field photo-consistent with V views as a convex combination of all these basis fields. Moreover, we can always find such a convex combination using image-space computations, i.e., by finding a convex combination of the images of these basis fields that reproduces all input views. To turn these ideas into an algorithm we answer two questions: (1) how do we efficiently generate an appropriate set of basis fields from the input views and (2) how do we find convex combinations of these fields that maximize photo-consistency with all views? We consider each question below.

5.1. Generating Basis Density Fields

The Flame Sheet creation process outlined in Section 3 and in Figures 3-4 takes into account all pixels in two corresponding epipolar lines. To generate a whole family of 2-view photo-consistent solutions, we simply apply the same procedure to portions of the input views. The resulting solutions, which we call *Decomposed Flame Sheets*, retain

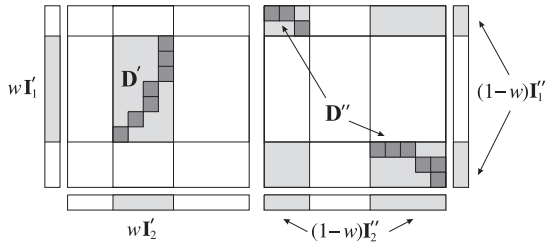

Figure 6. Breaking $\mathbf{I}_1$ and $\mathbf{I}_2$ into component parts. $\mathbf{I}_1$ is decomposed into a single interval of pixels $\mathbf{I}_1'$, whose intensities sum to w, and a sub-image $\mathbf{I}_1''$ that contains the remaining mass of $\mathbf{I}_1$ (i.e., $1-w$). $\mathbf{I}_2$ is decomposed in a similar way, with the restriction that the intensities of $\mathbf{I}_2'$ also sum to w. Note that, in general, $\mathbf{I}_1'$ and $\mathbf{I}_1''$ need not cover disjoint sets of pixels. Also shown are the two flame sheets, $\mathbf{D}'$ and $\mathbf{D}''$, computed from the partial images. These flame sheets have opposite diagonal orientations (Section 3).

the photo-consistency and sheet-like structure of the original Flame Sheets.

Given two input images $\mathbf{I}_1$ and $\mathbf{I}_2$, we construct a Decomposed Flame Sheet by (1) splitting each image into two parts and (2) using these parts to construct two new Flame Sheets, $\mathbf{D}'$ and $\mathbf{D}''$. The precise splitting process, described in Figure 6, is of the general form

$$\begin{array}{rcl} \mathbf{I}_1 & = & w\mathbf{I}_1' + (1-w)\mathbf{I}_1'' \\ \mathbf{I}_2 & = & w\mathbf{I}_2' + (1-w)\mathbf{I}_2'' \end{array}, \tag{5}$$

where $0 < w < 1$, and all images sum to one. The new Flame Sheets can now be combined to create a Decomposed Flame Sheet that explains the complete set of input pixels:

$$\mathbf{D} = w\mathbf{D}' + (1-w)\mathbf{D}''. \tag{6}$$

It follows from the Convexity and Image Linearity Properties that the field $\mathbf{D}$ is photo-consistent with $\mathbf{I}_1$ and $\mathbf{I}_2$.

Observe that any particular solution in the family of Decomposed Flame Sheets can be parameterized by the weight w, the diagonal orientation of $\mathbf{D}'$, and the position of $\mathbf{I}_1'$ and $\mathbf{I}_2'$ within the original images. Therefore, given S different weights and T distinct positions, we can create $2ST^2$ Decomposed Flame Sheets.

5.2. Flame-Sheet Decomposition Algorithm

Using the procedure for generating Decomposed Flame Sheets as a starting point, we rely on the following algorithm to compute an epipolar slice of the fire density field that maximizes photo-consistency with V views:

1. **(Basis construction).** For each of the $P = V(V-1)/2$ pairs of views, generate a family of B Decomposed Flame Sheets, $\{\mathbf{D}_{p\,1}, \mathbf{D}_{p\,2}, \ldots, \mathbf{D}_{p\,B}\}$, according to Section 5.1. Warp these into a common global coordinate system.
2. **(Basis projection).** Project the Decomposed Flame Sheets to all input viewpoints, stacking their images into VP blocks of size $N \times B$:

$$\mathbf{F}_{v\,p} = \left[\begin{array}{cccc} \mathbf{I}_v(\mathbf{D}_{p\,1}) & \mathbf{I}_v(\mathbf{D}_{p\,2}) & \cdots & \mathbf{I}_v(\mathbf{D}_{p\,B}) \end{array} \right].$$

3. **(QP formation).** Stack the blocks from Step 2 and the input images together,

$$\mathcal{F} = \left[\begin{array}{cccc} \mathbf{F}_{1\,1} & \mathbf{F}_{1\,2} & \cdots & \mathbf{F}_{1\,P} \\ \mathbf{F}_{2\,1} & \mathbf{F}_{2\,2} & \cdots & \mathbf{F}_{2\,P} \\ & \vdots & & \\ \mathbf{F}_{V\,1} & \mathbf{F}_{V\,2} & \cdots & \mathbf{F}_{V\,P} \end{array} \right], \quad \mathcal{I} = \left[\begin{array}{c} \mathbf{I}_1 \\ \mathbf{I}_2 \\ \vdots \\ \mathbf{I}_V \end{array} \right],$$

and form the following quadratic programming problem:

$$\begin{array}{ll} \text{minimize} & ||\mathcal{F}\mathbf{x} - \mathcal{I}||_2 \\ \text{subject to} & \sum \mathbf{x} = 1, \mathbf{x} \geq 0. \end{array}$$

4. **(QP solution).** Solve this quadratic optimization program by any standard method [22]. The resulting optimal **x** gives the weights needed to express the density field as a convex combination of the Decomposed Flame Sheets from Step 1.

Note that Step 3 of the algorithm is suboptimal. This is because the convex combination of sheets computed in Step 1 may not, in general, intersect the space of V-view photo-consistent solutions (Figure 5). While an optimization over arbitrary linear combinations of sufficiently many sheets would avoid this problem (since Section 5.3 shows that our basis is complete), the computational requirements would be increased substantially. This is because explicit enforcement of non-negativity constraints for arbitrary linear combinations requires working in the N^2-dimensional space of density fields. In practice, we therefore use the convex approach exclusively.

5.3. Completeness Theorem

Even though Decomposed Flame Sheets are easy to generate, a sufficiently large number of them spans the space of V-view photo-consistent densities. We derive this result by first analyzing the space of 2-view photo-consistent densities through the notion of a *switching matrix*.

Lemma 1. Given two N-column vectors, all non-negative $N \times N$ matrices that generate these vectors via their row and column sums lie in an $(N-1)^2$-dimensional hyperplane.

Proof. Given any particular matrix $\mathbf{D}_0$ that generates these vectors, we can represent the matrices in this space by the hyperplane $\mathbf{D}_0 + \mathbf{S}$, where $\mathbf{S}$ is a matrix with two properties:

- *Row and column sums preserved:* We require that $\mathbf{1}^T\mathbf{S} = 0$ and $\mathbf{S1} = 0$ so that the vectors are still reproduced.
- *Non-negativity preserved:* For all matrix elements (i,j), we require that $\mathbf{S}(i,j) \geq -\mathbf{D}_0(i,j)$.

Note that fixing the $(N-1) \times (N-1)$ upper-left block of $\mathbf{S}$ uniquely defines the remaining elements, giving us an upper bound on the dimensionality of the space. □

By explicitly constructing an $(N-1)^2$-dimensional basis for the matrices photo-consistent with two views, we show that the $(N-1)^2$ dimensionality bound in Lemma 1 is tight.

Lemma 2. The hyperplane described in Lemma 1 can be spanned by a canonical basis of $(N-1)^2$ *switching matrices*.

Proof. Extending the analysis in [19] to the case of continuous-valued densities, we define a *switching matrix* $N \times N$ to be a matrix with only four non-zero entries in a rectangle aligned with the rows and columns as follows:

$$\mathbf{S} = \begin{pmatrix} & \vdots & & \vdots & \\ \cdots & +\alpha & \cdots & -\alpha & \cdots \\ & \vdots & & \vdots & \\ \cdots & -\alpha & \cdots & +\alpha & \cdots \\ & \vdots & & \vdots & \end{pmatrix}. \quad (7)$$

Adding a switching matrix to a density field has the effect of shifting mass from two diagonal corners of the rectangle to the other two corners.

We define a canonical set of switching matrices $\{\mathbf{S}_{ij}\}$ by fixing $\mathbf{S}_{ij}(N, N) = +1$, and moving another $+1$ over all positions in the upper-left block of $(N-1) \times (N-1)$ elements. Then, the positions of the corresponding -1's follow immediately from the rectangular form of the switching matrix. By expressing these matrices as 1D vectors we see that they are clearly linearly independent, thus forming a basis for the $(N-1)^2$-hyperplane. This hyperplane can also be expressed as a linear combination of the 2-view photo-consistent matrices $\{\mathbf{D}_0 + \mathbf{S}_{ij}\}$. □

Lemmas 1-2 and the Nesting Property imply that any linearly independent set of $(N-1)^2$ densities that are photo-consistent with two views spans the space of V-view photo-consistent densities.[4]

Theorem 2 (Flame Sheet Completeness). The family of Decomposed Flame Sheets forms a complete basis for the space of V-view photo-consistent density fields.

Proof. The Decomposed Flame Sheets allow us to define a matrix analogous to a switching matrix. Specifically, let $\mathbf{D}$ be a Decomposed Flame Sheet defined by images $\mathbf{I}_1, \mathbf{I}_2$ in Eq. (6) and let $\hat{\mathbf{D}}$ be one of the two basic Flame Sheets defined by these images. Then, the difference $\hat{\mathbf{D}} - \mathbf{D}$ can be thought of as a generalization of a switching matrix. By varying the parameters of $\mathbf{D}$ (weight, position, diagonal orientation), we can generate $(N-1)^2$ linearly independent matrices of this form that span the space. □

Note that Theorem 2 does not imply that a *convex* combination of Decomposed Flame Sheets spans the space of photo-consistent density fields. It does suggest, however, that these sheets comprise an expressive family of solutions.

6. Experimental Results

We performed experiments with images captured from two scenes of fire.[5] The first scene ("torch") was a citronella patio torch, burning with a flickering flame about 10cm tall. Two synchronized progressive-scan Sony DXC-9000 cameras, roughly 90° apart, were used to acquire videos of the flame (Figure 1). Calibration accuracy was about 0.5 pixels.

For simple flames like the "torch" scene, a two-view reconstruction consisting of one Flame Sheet produces very realistic views (Figure 7). By comparison, the multiplication solution shows typical blurring and doubling artifacts. An algebraic method based on fitting Gaussian blobs to the density field [23] over-fits the two input images and produces a badly mottled appearance for extrapolated views.

The second scene ("jet") consisted of a complex flame emerging from a gaseous jet (Figure 9, first row). This sequence contained 47 frames from a promotional video for a commercial 3D freeze-frame system [24], roughly corresponding to inward-facing views arranged around a quarter-circle. Since no explicit calibration was available for this sequence, we assumed that the views were evenly spaced.

To test the extrapolation capabilities of our approach, we used a subset of the 47 frames as input and used the rest to evaluate agreement with the synthesized views (Figure 9). Results from the multiplication and single Flame Sheet solutions, which are restricted to only two input views, suggest that these methods cannot capture the flame's appearance and structure accurately because of its complexity and its large appearance variation across viewpoint. To incorporate more views, we applied the algebraic blob method in [23] and the Flame Sheet Decomposition Algorithm with the 45° view as a third image. In the latter algorithm, we generated $B = 150$ Decomposed Flame Sheets for each of the $P = 3$ pairs of input views in Step 1, giving rise to a total of 450 basis fields. To further explore the benefit of using multiple views, we optimized the convex combination of these fields in Steps 2-4 in two ways: (1) by maximizing photo-consistency with the three input images and (2) by maximizing photo-consistency with four additional images from the sequence. While the results in Figure 9 and Table 1 suggest that the Flame Sheet Decomposition Algorithm is superior to the other methods, they also show that increasing the number of images during optimization has a clear benefit. On the other hand, increasing the number of basis fields has little effect (Figure 8).

Three observations can be made from our experiments. First, the convex combination of Decomposed Flame Sheets can capture the structure of complex flames because the sheets sweep out regions of the density field, explaining fire density at multiple depths. Second, the resulting algorithm can render complex flames from a small number of input views, without the doubling artifacts of the multiplication solution or the over-fitting artifacts of the blob-based method. This is because the family of Decomposed Flame Sheets is expressive and induces a bias toward coherent solutions. Third, the sheets' surface-like structure enables photo-realistic rendering of complex flames by warping and blending the partial images in Eq. (5). Thus, the fire models we extract are suitable for interactive graphics applications.

[4] Another consequence of Lemma 2 is that every Decomposed Flame Sheet derived from two views lies on the boundary of the space of density fields photo-consistent with those views. This is because the sheets' structure ensures that adding an infinitesimally-small switching matrix will make the density field negative.

[5] See the ICCV 2003 CD-ROM for color images and video clips, and visit http://www.cs.toronto.edu/∼hasinoff/fire for more results.

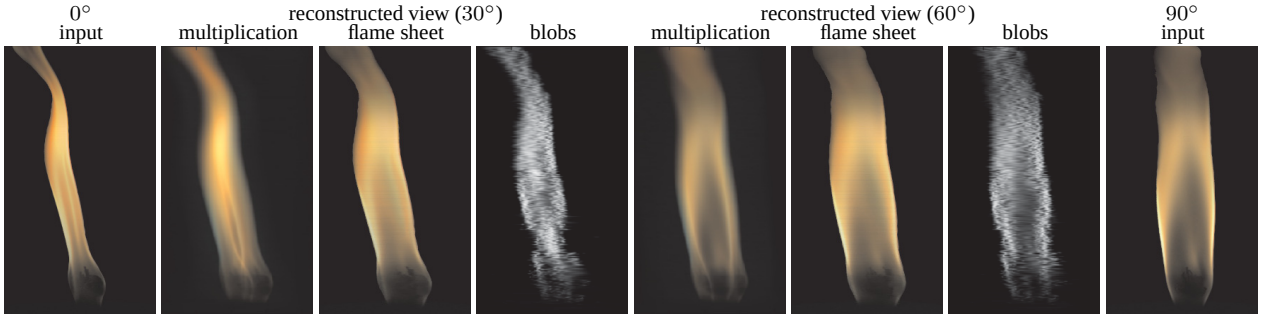

Figure 7. Views of three different reconstructions of the "torch" dataset. The same two input images were used for all methods. Since all methods reconstruct the input views, only novel views of the reconstructed flame are shown.

reconstruction method (views)	RMS error	
	input	all
flame sheet solution (2)	0	26.6
multiplication solution (2)	0	21.5
blob-based method (3)	13.3	19.0
flame sheet decomposition (3)	10.1	18.4
flame sheet decomposition (3/7)	11.5	15.8

Table 1. Per-pixel RMS image reconstruction error for different algorithms applied to the "jet" dataset. For the last entry, three input views were used for basis generation and seven for optimizing photo-consistency. Note that the first two reconstructions reproduce the input views exactly.

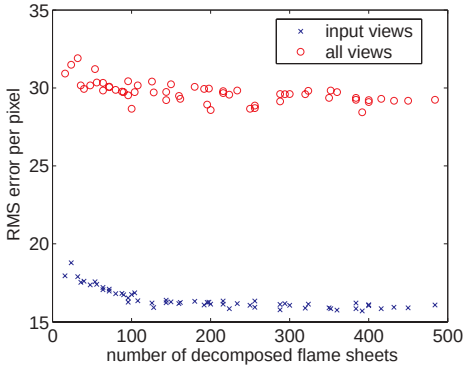

Figure 8. Dependence of image error on the number B of generated basis fields (Section 5.2, Step 1).

7. Concluding Remarks

A current limitation of our multi-view technique is the planar configuration of the input viewpoints. While planarity is used for efficiency reasons, it may be possible to derive an efficient reconstruction method that combines Flame Sheets from non-planar views.

More generally, the question of how best to capture the global 3D structure and dynamics of fire remains open. Toward this goal, we are investigating new spatio-temporal coherence constraints and are studying ways to integrate our approach with traditional fire simulation methods.

Finally, while our limited experiments suggest that Flame Sheets and the Flame-Sheet Decomposition Algorithm are useful for fire reconstruction, these tools may also prove useful in more general contexts. This includes sparse-view tomography problems in medical diagnostics [17] and accelerated image-based methods for volume rendering.

References

[1] D. Nguyen, R. Fedkiw, and H. Jensen, "Physically based modeling and animation of fire," in *Proc. SIGGRAPH*, pp. 721–728, 2002.

[2] A. Lamorlette and N. Foster, "Structural modeling of flames for a production environment," in *Proc. SIGGRAPH*, pp. 729–735, 2002.

[3] J. Stam and E. Fiume, "Depicting fire and other gaseous phenomena using diffusion processes," in *Proc. SIGGRAPH*, pp. 129–136, 1995.

[4] N. Chiba, K. Muraoka, H. Takahashi, and M. Miura, "Two-dimensional visual simulation of flames, smoke and the spread of fire," *J. Visualiz. and Comp. Anim.*, vol. 5, no. 1, pp. 37–54, 1994.

[5] R. Fedkiw, J. Stam, and H. W. Jensen, "Visual simulation of smoke," in *Proc. SIGGRAPH*, pp. 15–22, 2001.

[6] N. Foster and D. N. Metaxas, "Modeling the motion of a hot, turbulent gas," in *Proc. SIGGRAPH*, pp. 181–188, 1997.

[7] R. T. Baum, K. B. McGratten, and M. Nyden, "An examination of the applicability of computed tomography for the measurement of component concentrations in fire-generated plumes," *Combustion and Flame*, vol. 113, pp. 358–372, 1998.

[8] A. Schwarz, "Multi-tomographic flame analysis with a schlieren apparatus," *Measurement Science & Tech.*, vol. 7, pp. 406–413, 1996.

[9] D. P. Correia, P. Ferrao, and A. Caldeira-Pires, "Advanced 3D emission tomography flame temperature sensor," *Combustion Science and Technology*, vol. 163, pp. 1–24, 2001.

[10] W. Xue, W. Donglou, and P. Gongpei, "Use of Moire tomography to measure the temperature field of the flame of a pyrotechnical composition from its infrared radiation," *Combustion, Explosion, and Shock Waves*, vol. 37, no. 4, pp. 440–442, 2001.

[11] G. W. Faris and R. L. Byer, "Beam-deflection optical tomography," *Optics Letters*, vol. 12, no. 3, pp. 72–74, 1987.

[12] T. Otsuka and P. Wolanski, "Particle image velocimetry (PIV) analysis of flame structure," *Journal of Loss Prevention in the Process Industries*, vol. 14, no. 6, pp. 503–507, 2001.

[13] A. K. Agrawal, N. Butuk, S. R. Gollahalli, and D. Griffin, "Three-dimensional rainbow schlieren tomography of a temperature field in gas flows," *Applied Optics*, vol. 37, no. 3, pp. 479–485, 1998.

[14] A. Schödl, R. Szeliski, D. H. Salesin, and I. Essa, "Video textures," in *Proc. SIGGRAPH*, pp. 489–498, 2000.

[15] D. Drysdale, *An Introduction to Fire Dynamics*. Chichester: John Wiley and Sons, Second ed., 1998.

[16] P. J. Narayanan, P. W. Rander, and T. Kanade, "Constructing virtual worlds using dense stereo," in *Proc. ICCV*, pp. 3–10, June 1998.

[17] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. New York: IEEE Press, 1988.

[18] T. Frese, C. A. Bouman, and K. Sauer, "Adaptive wavelet graph model for Bayesian tomographic reconstruction," *IEEE Transactions on Image Processing*, vol. 11, pp. 756–770, July 2002.

[19] G. T. Herman and A. Kuba, eds., *Discrete Tomography: Foundations, Algorithms, and Applications*, (Boston), Birkhäuser, 1999.

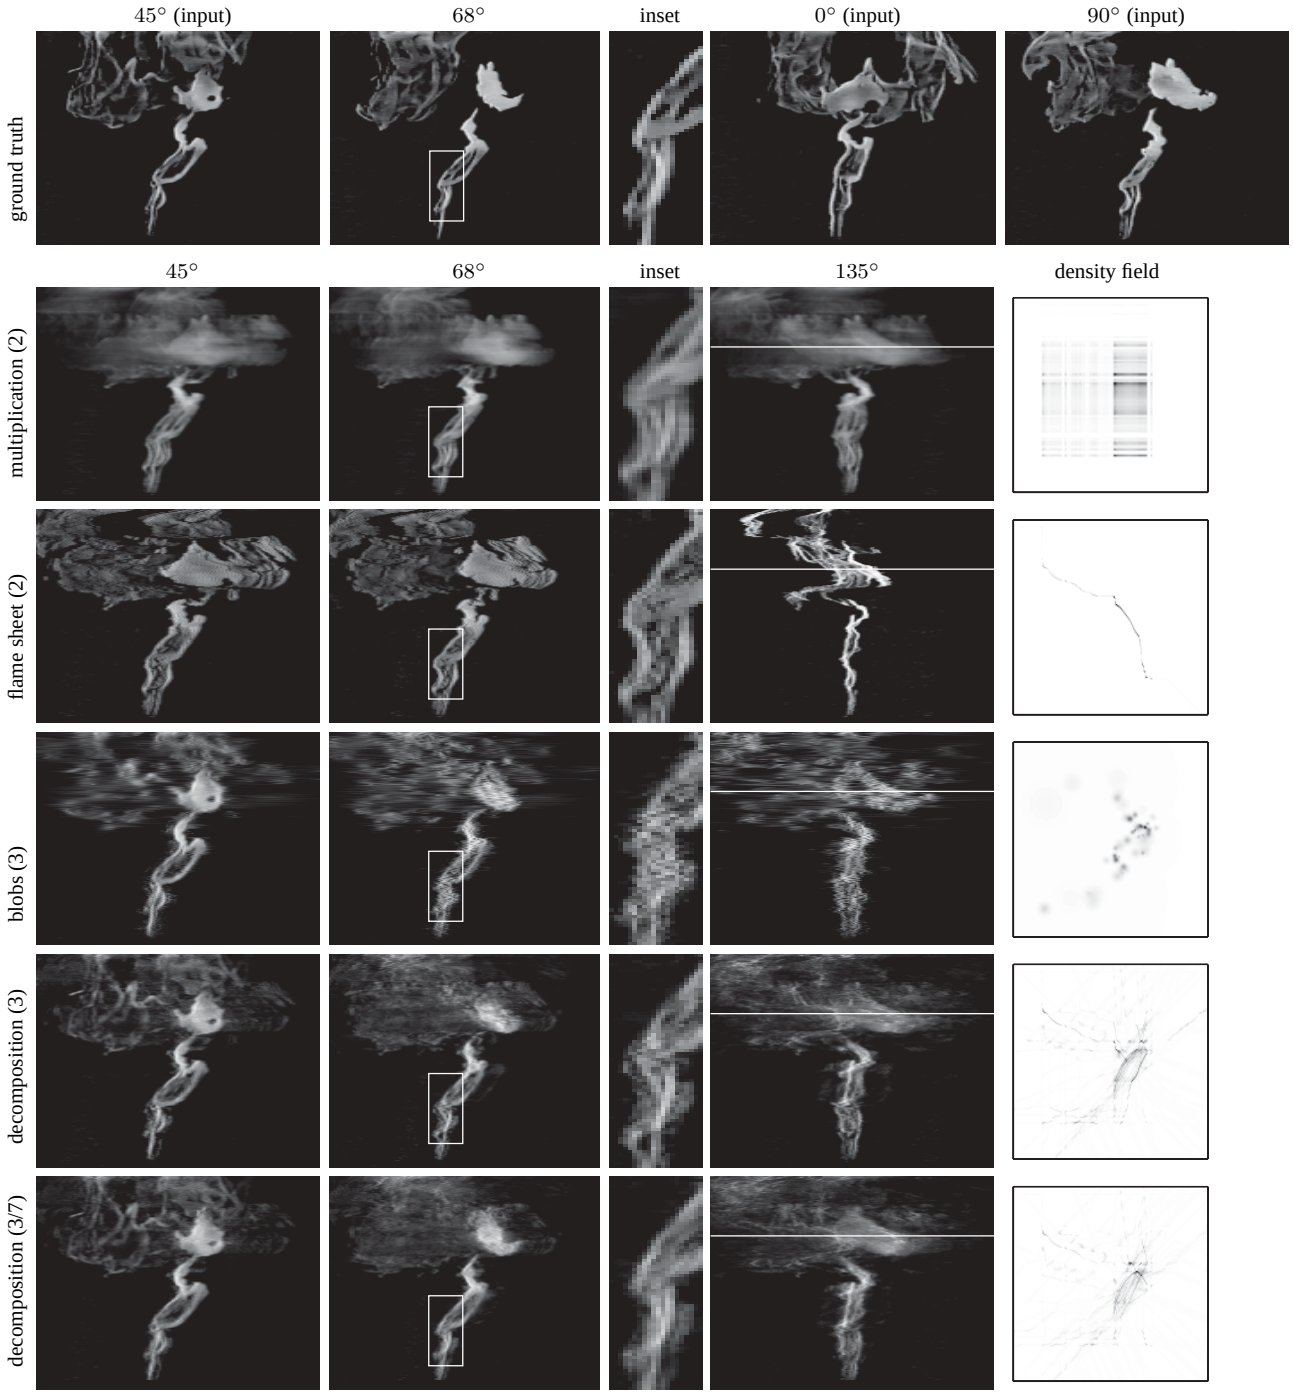

Figure 9. Reconstruction of the "jet" dataset. Ground truth images of the fire are shown in the top row; the remaining rows correspond to different reconstruction methods. Labels on the left indicate the method used, with the number of input views in brackets. From left to right: 45° view, which is extrapolated for the multiplication and Flame Sheet solutions and is an input view for the other methods; 68° view, extrapolated for all methods; zoomed-in view of the region marked in the 68° image; 135° view, which is a distant extrapolation; and top view of a reconstructed epipolar slice, corresponding to the white horizontal line in the 135° image (dark regions denote areas with concentrated fire density).

[20] G. Wolberg, *Digital Image Warping*. Los Alamitos, CA: IEEE Computer Society Press, 1990.

[21] J. Duderstadt and W. Martin, *Transport Theory*. New York, NY: John Wiley and Sons, 1979.

[22] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, 1999.

[23] S. Hasinoff, "Three-dimensional reconstruction of fire from images," Master's thesis, Univ. of Toronto, Dept. of Computer Science, 2002.

[24] Digital Air Incorporated. http://www.virtualcamera.com.