

Photo-Consistent Reconstruction of Semitransparent Scenes by Density-Sheet Decomposition

Samuel W. Hasinoff, *Student Member, IEEE*, and Kiriakos N. Kutulakos, *Member, IEEE*

Abstract—This paper considers the problem of reconstructing visually realistic 3D models of dynamic semitransparent scenes, such as fire, from a very small set of simultaneous views (even two). We show that this problem is equivalent to a severely underconstrained computerized tomography problem, for which traditional methods break down. Our approach is based on the observation that every pair of photographs of a semitransparent scene defines a unique density field, called a Density Sheet, that 1) concentrates all its density on one connected, semitransparent surface, 2) reproduces the two photos exactly, and 3) is the most spatially compact density field that does so. From this observation, we reduce reconstruction to the convex combination of sheet-like density fields, each of which is derived from the Density Sheet of two input views. We have applied this method specifically to the problem of reconstructing 3D models of fire. Experimental results suggest that this method enables high-quality view synthesis without overfitting artifacts.

Index Terms—Semitransparent scenes, fire, volumetric reconstruction, 3D reconstruction, computerized tomography, view synthesis, image-based modeling, image-based rendering.

1 INTRODUCTION

THE computational modeling of physical phenomena such as fire and smoke has received significant attention in computer graphics [1], [2], [3], [4], [5], [6], [7], [8] as well as in other fields of experimental science [9], [10], [11], [12], [13], [14], [15], [16]. 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], [2], [4], [6]; developing image-based representations of such phenomena [17]; 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 semitransparent scenes from a set of simultaneously captured images or videos. We take fire as our central example, and show that by using optical models developed in the combustion literature [10], [11], [18], we can model fire as a semitransparent 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 modeling dynamic semitransparent scenes from multiple viewpoints must satisfy three criteria:

- **Small viewpoint coverage.** Dynamic phenomena require simultaneous image capture. While multiview

- *The authors are with the Department of Computer Science, University of Toronto, 10 King's College Road, Room 3302, Toronto, Ontario, Canada M5S 3G4. E-mail: {hasinoff, kyros}@cs.toronto.edu.*

Manuscript received 31 May 2005; revised 27 Jan. 2006; accepted 11 July 2006; published online 18 Jan. 2007.

Recommended for acceptance by H. Shum.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0281-0505. Digital Object Identifier no. 10.1109/TPAMI.2007.1056.

systems with tens of video cameras do exist [11], [19], the method should not break down for small data sets.

- **Photo-consistency.** To preserve visual realism, reconstructed density fields must reproduce the input images.
- **Good view synthesis.** All views should be rendered with high quality, without “overfitting” the limited number of input views.

Even though these criteria have received much attention in the case of opaque objects [19], satisfying them for semitransparent scenes is not well understood.

At the heart of our approach lies the observation that every pair of photos of a semitransparent scene uniquely determines two special density fields, called *Density 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 the following: 1) Density 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 that are 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 *Density-Sheet Decomposition*, that computes the density field as a convex combination of sheet-like density fields derived from pairs of input views.

Unlike existing methods for reconstructing semitransparent scenes, Density-Sheet Decomposition is specifically designed to capture models that are both photo-consistent and have good view synthesis capabilities. Intuitively, Density Sheets represent the most spatially-compact interpretation of the input views. Moreover, Density Sheets are unique and lead to a reconstruction method that is well-posed and easy to implement. This allows us to reconstruct semitransparent scenes 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 and spatial smoothness are the only objectives [20], [21], [22], [23], our approach establishes spatial compactness as another objective for visual reconstruction. Second, it leads to a well-posed reconstruction algorithm that can handle any number of input views. Third, it introduces Density Sheets as a complete basis for the space of photo-consistent density fields, enabling their potential use for reconstructing and rendering any semitransparent scene where the linear image formation model is valid. Fourth, the sheet-like structure of the Density-Sheet Decomposition algorithm enables use of simple warp-based methods [24] to create photo-realistic views of the reconstructed scene. Fifth, our results on fire data sets show that the algorithm is able to render detailed reconstructions of complex flames without overfitting artifacts.

We begin this paper by discussing the relation of this work to current approaches for reconstructing semitransparent scenes in general and fire in particular. In Section 3, we present a simplified imaging model for fire, which establishes an equivalence between fire reconstruction and computerized tomography. Because of this equivalence, the rest of our analysis applies to the general case of semitransparent scenes (Section 4). In Section 5, we present an analysis of the space of density fields photo-consistent with two views, and introduce the Density Sheet solution. We then generalize the analysis to multiple views, and describe several properties that allow us to reduce the problem to the two-view case (Section 6). In Section 7, we describe a new family of photo-consistent density fields built from sheet-like structures and show that this family can be used to represent arbitrary density fields. This leads to an algorithm for M -view photo-consistent reconstruction that models density fields as superpositions of sheets. In Section 8, we present experimental results using real images of fire. We conclude with a discussion of our approach and some proposed directions for future work.

2 RELATED WORK

The 3D modeling of fire and other semitransparent scenes has received significant attention in computer graphics. Both physics-based [1], [2], [3], [5] and procedural [6], [7], [8] models have been used to produce visually realistic simulations of torches, gaseous jets of flame, and smoke. For scenes consisting of strongly scattering particles such as smoke, high-fidelity 3D capture of real scenes has been shown using laser illumination and high-speed photography [4]. But, the question of how we can capture models of general semitransparent scenes, particularly from standard images, remains largely open.

In addressing this question, we touch upon four distinct lines of work: computerized tomography, volumetric reconstruction in computer vision, combustion measurement, and 2D image-based modeling for dynamic scenes.

2.1 Reconstructing Semitransparent Scenes

2.1.1 Computerized Tomography

The problem of reconstructing semitransparent scenes has received the most attention in the context of computerized tomography. Our interest is in using very few views (two or more), which is closest to the “sparse-view tomography”

problem. Some sparse-view methods produce a binary segmentation rather than a full reconstruction [25], [26], [27] and, in the context of tomography, “sparse” can actually mean eight or more views [21], [23], [25]. On the other hand, classic methods such as filtered back-projection and algebraic techniques (e.g., [20], [28]) require tens or hundreds of input views for accurate results.

Tomography methods that are specialized to the ill-posed nature of the sparse-view case use various techniques to regularize the problem. This includes favoring local smoothness in a statistical framework [21], [27], assuming prior shape models [25], [26], [27], and coarsely discretizing the density levels [21]. Unfortunately, these methods still break down when the number of views is extremely limited and fail to generate 3D reconstructions adequate for photo-realistic view synthesis.

2.1.2 Volumetric Reconstruction in Computer Vision

In the stereo vision literature, several volumetric methods have attempted to recover transparencies along with voxel color information [22], [29], [30]. While some of these methods are more directed at modeling occlusion uncertainty [30] or mixed pixels at object boundaries [22], several of these methods have addressed the full tomography problem [22], [29]. For example, the Roxel method [29] is related to algebraic methods from computerized tomography [20] but its imaging model has an unclear physical basis since it confounds transparency and uncertainty. These methods do not guarantee photo-consistency nor do they produce results suitable for view synthesis, particularly when the number of views is limited.

Along similar lines, algebraic methods from tomography have also been applied to image-based rendering for trees with dense foliage, from tens of views [31]. Because a significant fraction of pixels in this context are opaque, this approach is closer to view-dependent texture mapping using the visual hull and is less suitable for complex semitransparent scenes that consist of multiple layers.

Although methods for reconstructing semitransparent scenes from a sparse number of views have attempted to address the ill-posed nature of the problem, none of the above methods explicitly characterize the fundamental ambiguities in the space of photo-consistent solutions. For this reason, they are ill-equipped for handling the extreme cases (i.e., two views). To our knowledge, the only previous work addressing tomographic reconstruction with two views involves binary density fields for segmentation [27].

2.2 Reconstructing Fire

2.2.1 Combustion Measurement

In the combustion literature, methods for the 3D measurement of fire have mainly been directed toward the measurement of specific physical properties (e.g., temperature) and the qualitative recovery of 3D shape [9], [10], [11], [12], [13], [14], [15], [16]. Most methods do not restrict themselves to visible light images from cameras, but rather employ a wide variety of sensors. Common approaches involve lasers [11], [16], special optical systems to exaggerate refraction [11], [12], [13], [14], [15], and thermography devices [9], [10]. None of these methods provide instantaneous 3D reconstruction at resolutions high enough for image-based rendering. In fact, many previous methods capture images with multiple-second exposures [9], [10], [12], [15] or assume a “stationary”

fire [11] to reconstruct rough approximations to the density field. As a result, these methods are not appropriate for view synthesis and cannot be used to model complex flickering flames (e.g., Fig. 11), whose structures can change dramatically from instant to instant.

Because the straightforward application of classic tomography methods [20], [28] to fire reconstruction would require an impractical number of synchronized cameras, results in this vein have been severely limited in both resolution and accuracy [9], [13], [32]. One exception is a 20-camera apparatus designed to reconstruct quasi-static or nonvisible flame [11]. More recently, methods from sparse-view tomography have been applied to fire reconstruction as well using eight synchronized cameras [23]. While this approach yields a plausible density distribution at a scale coarser than a pixel, the high-frequency detail present in the original images is not reconstructed. For an extremely limited number of views (e.g., 2-5), even such sparse-view tomography methods fail to produce results appropriate for view synthesis [21], [23], [25], [26], [27].

2.2.2 Image-Based Modeling for Dynamic Scenes

An alternative approach is to avoid almost all reconstruction and model fire in strictly image-based terms. Image-based methods such as video textures [17], [33] and linear dynamical systems [34], [35], have already been applied to fit simple 2D dynamic models of fire to video sequences. Although these methods can generate realistic continuations of flickering flames, and even allow basic attributes such as speed [35] and direction of flow [36] to be manipulated, they do not allow us to vary viewpoint nor to model nonstationary effects such as materials consumed by fire.

3 LINEAR OPTICAL MODEL OF FIRE

Fire is typically defined as an oxygen-fueled chemical decomposition that releases heat to the environment [18]. In this work, we are strictly concerned with flames, i.e., the visible luminous reaction products of fire.¹ We rely on a first-order model of flame appearance, proposed in the combustion measurement literature, called the *soot volume model* [18]. Under this model, the appearance of flame is generated by a continuum of luminous carbonaceous soot particles, where in a given infinitesimal volume, both intensity and absorption are proportional to soot particle density. Following the approach of Stam [2], we ignore refraction² and derive this model as a specialization of a more general model for radiative transfer [37].

At a given image point, the image irradiance, \hat{I} , can be expressed as the sum of two terms—a term that integrates radiance from luminous fire material along the ray through that pixel and a term that incorporates background radiance (Fig. 1):

$$\hat{I} = \int_0^L D(t) \tau(t) J(t) dt + \hat{I}_{bg} \tau(L), \quad (1)$$

1. In general, not all fire burns with visible flames. Some gases that burn cleanly are visually detectable only by their secondary effects, for example, the refractive distortion of the atmosphere [11]. We do not deal with this type of fire here.

2. The combustion literature supports this assumption, as the refraction in visible flame is generally small, with deflections less than 0.1 degree [11], [14].




Fig. 1. Optical model of fire. The image irradiance, \hat{I} , at a pixel depends on the integral of luminous fire density along a ray through the center of projection, c . The density field, $D(t)$, along the ray is parameterized by the distance, $t \in [0, L]$, from the image plane. The background intensity, \hat{I}_{bg} , also contributes to the intensity, \hat{I} .

where $D(t)$ is the fire's density field along that ray, 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 nonzero, \hat{I}_{bg} is the radiance of the background, and τ models transparency. The transparency can be thought of as the degree by which the fire products between the pixel and a position x along the ray permit the radiance from further away to reach that pixel,

$$\tau(x) = \exp\left(-\sigma \int_0^x D(t) dt\right), \quad (2)$$

where σ is a positive medium-dependent constant known as the extinction cross-section [18], and is assumed to be constant throughout the fire.

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

- **Negligible scattering.** This is a good approximation for fire not substantially obscured by smoke because emission from luminous soot particles dominates radiance [2], [10], [11], [14]. In this case, the function $J(t)$ in (1) becomes a pure self-emission term.
- **Constant self-emission.** This assumption models fires whose brightness depends only on the density of luminous soot particles [18], allowing us to assume that self-emission is constant per unit mass, i.e., that $J(t)$ is a constant, J_0 .

Together with (1) and (2), these assumptions lead to an expression for \hat{I} that depends only on the fire's transparency along the ray:

$$\hat{I} = J_0(1 - \tau(L)) + \hat{I}_{bg}\tau(L). \quad (3)$$

In practice, the short image exposure times we use to prevent pixel saturation leads to an effectively invisible background, i.e., $\hat{I}_{bg} = 0$. This obviates the need to measure the background radiance in the absence of fire and also permits measurement in the presence of dynamic backgrounds.

A key feature of (3) is that it allows us, through a simple transformation, to establish a linear relation between the image irradiance and the density field along the corresponding ray. To see this, note that the image irradiance, \hat{I} , is bounded from above by J_0 , since taking the limit of infinite fire density causes the transparency to approach zero, $\lim_{D \rightarrow \infty} \tau(L) = 0$. Therefore, when the camera's settings do not allow pixel saturation, so that J_0 is mapped to the maximum intensity, we can linearize (3) using the following invertible transformation:

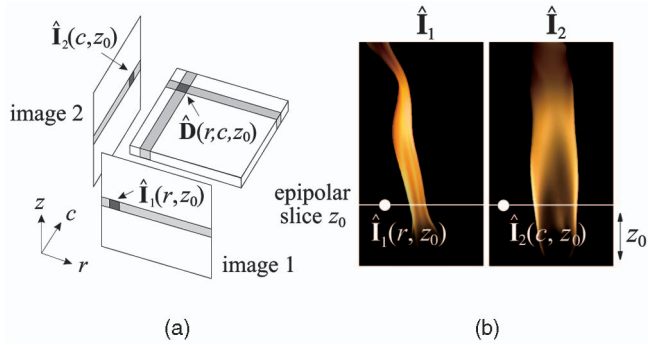


Fig. 2. Viewing geometry. (a) Two orthographic images corresponding to a 90 degree rotation about the vertical z -axis, along with an epipolar plane. (b) Two simultaneous images of a flame, with white horizontal lines indicating the corresponding epipolar lines at height z_0 .

$$I = -\frac{1}{\sigma} \log \left(\frac{\hat{I} - J_0}{\hat{I}_{bg} - J_0} \right), \quad (4)$$

to obtain

$$I = \int_0^L D(t) dt. \quad (5)$$

We further assume a linear camera response so that the image irradiance \hat{I} is proportional to the image intensity. The final unspecified aspect of our imaging model is the unknown constant σ , which is conflated with density by the transformation in (4). Since our goal is to model the appearance of fire, we only need to model the density up to scale, so both of these proportionality constants can be safely ignored.

4 PROBLEM STATEMENT

Consider a 3D semitransparent density field, $\hat{\mathbf{D}}(r, c, z)$, viewed under orthographic projection by a set of viewing directions perpendicular to the vertical z -axis (Fig. 2). Our goal is to obtain a photo-consistent reconstruction of the 3D density field. We assume the linear imaging model of (5), where a pixel's intensity is the integral of the density along its associated ray.

The reconstruction of the 3D density field can be reduced to a sequence of 2D problems whose goal is to reconstruct a single epipolar slice, $\hat{\mathbf{D}}(r, c, z_0)$, of the 3D field from corresponding epipolar lines in the input views. In practice, the pixels along corresponding epipolar lines are discrete and can be thought of as 1D column vectors. This leads to the following general problem:

Definition 1 (Photo-Consistent Density Reconstruction).

Given the N -dimensional column vectors $\mathbf{I}_1, \dots, \mathbf{I}_M$ corresponding to M orthographic views of a 2D density field represented by a nonnegative $N \times N$ matrix $\hat{\mathbf{D}}$, compute a matrix \mathbf{D} such that the following photo-consistency equation is satisfied for all pixels p (Fig. 3):

$$\mathbf{I}_m(p) = \sum_{(r,c) \in F(p)} w(r, c, p) \mathbf{D}(r, c), \quad (6)$$

where $F(p)$ collects the elements of \mathbf{D} that contribute to pixel p and $w(r, c, p)$ is a nonnegative weight that describes the contribution of element $\mathbf{D}(r, c)$ to the intensity of pixel p .




Fig. 3. Finding a photo-consistent density \mathbf{D} entails satisfying the discrete line integrals represented by (6) for all pixels. For a given view, the density element $\mathbf{D}(r, c)$ contributes to pixel $\mathbf{I}(p)$ according to the weight $w(r, c, p)$. We define this weight, represented by the lightly shaded region, as the fractional area of $\mathbf{D}(r, c)$ overlapping the cone corresponding to $\mathbf{I}(p)$.




Fig. 4. 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. We call the rightmost density field the “multiplication solution”. Note that view \mathbf{I}_3' has double the number of nonzero pixels compared to \mathbf{I}_3 and \mathbf{I}_3' .

In the following, we assume without loss of generality that the density field is normalized, i.e., $\sum_{r,c} \hat{\mathbf{D}}(r, c) = 1$. As in computerized tomography [20], this reconstruction problem is a linear system in the N^2 unknown density elements. This system is also subject to the nonnegativity of the density elements to ensure that the density field is physically meaningful. For M input images, the system has MN equations and is clearly underdetermined for the sparse-view case, when $M \ll N$. In general, $\mathbf{D} \neq \hat{\mathbf{D}}$, therefore photo-consistent density reconstruction is “weaker” than the well-posed problem examined by computerized tomography methods where $O(N^2)$ views are available.

While any density field \mathbf{D} satisfying the linear system can reproduce the input images, not all of them are equally good in generating novel views that are “similar” to the input views. We explore this issue further in the next section.

5 THE DENSITY SHEET

The photo-consistent reconstruction problem has many solutions when only two views of a density field are available (Fig. 4). In fact, given images \mathbf{I}_1 and \mathbf{I}_2 corresponding to row and column sums of the density field, respectively, it is trivial to show that the matrix

$$\mathbf{D} = \mathbf{I}_1 \mathbf{I}_2^T \quad (7)$$

is a solution, i.e., satisfies (6) for all pixels. We call this solution the *multiplication solution*. Note that any two images can be reduced to this orthogonal-view case by a known and invertible 2D warp [24] of the epipolar plane.




Fig. 5. Creating Density Sheets. (a) Moving from top to bottom in image I_1 , the density of each pixel in I_1 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.

While the multiplication solution is photo-consistent, it does not generalize to more than two views and it leads to significant artifacts during view synthesis. For example, Fig. 4c shows that double images are created in intermediate views, resulting in a blurred appearance. Intuitively, the multiplication solution represents the most spread-out, spatially incoherent solution to the two-view tomographic reconstruction problem.³

The key idea of our approach is to identify an alternative pair of photo-consistent solutions, called *Density Sheets*, which represent the conceptual opposite of the multiplication solution: rather than spread the density as much as possible, Density Sheets concentrate it along a monotonic curve, resulting in density fields that have maximal spatial coherence.

Theorem 1 (Density Sheet Theorem). *Every pair of nonnegative vectors, I_1 and I_2 , with equal sum define two unique nonnegative matrices, D , D' , which have the following properties:*

- For each matrix, the row and column sums are equal to I_1 and I_2 , respectively.
- D is nonzero only along a discrete monotonic curve connecting $D(1, 1)$ to $D(N, N)$.
- D' is nonzero only along a discrete monotonic curve connecting $D'(1, N)$ to $D'(N, 1)$.

We prove Theorem 1 constructively, by giving an algorithm that constructs Density Sheets from pairs of input views and then proving that it is correct.⁴ To gain some intuition, imagine “pushing” the intensities in image I_1 rightward along the rows of the density field, concentrating them on a monotonic curve, until the field becomes photo-consistent with image I_2 (Fig. 5a). This “pushing” and “spreading” procedure will always be possible because both images sum to one. In general, there are two such curves of opposite diagonal orientation, corresponding to the graphs of a monotonically decreasing (Fig. 5a) and monotonically increasing (Fig. 5b) function.

More concretely, we use the following iterative algorithm to construct Density Sheet D of Theorem 1 (Fig. 6); the

3. It is possible to show that the multiplication solution is the maximum entropy solution: since I_1 and I_2 are normalized and nonnegative, they can be viewed as marginals of a joint distribution, $p(I_1, I_2)$. Their joint entropy is maximized when I_1 and I_2 are independent, namely, when $p(I_1, I_2) = p(I_1)p(I_2)$, corresponding directly to the case $D = I_1 I_2^T$.

4. See Appendix A for the full proof of the algorithm’s correctness.




Fig. 6. Visualizing a step in the Density Sheet Construction Algorithm. Pixel intensities in the light gray regions of I_1 , I_2 are fully explained by a field that concentrates its nonzero density along a “partial” Density Sheet (dark gray elements) and has zero density elsewhere (light gray region of D). Depending on the amount of unexplained intensity at the projections of element (r, c) , the partial Density Sheet is expanded in one of the two directions indicated by the arrows. A sequence of six such expansions was used to create the Density Sheet in Fig. 5b.

second Density Sheet D' can be constructed analogously. The algorithm has an $O(N)$ time complexity.

Algorithm 1. Density Sheet Construction Algorithm

- 1) Set $D = 0$.
- 2) **(Current element).** Set $(r, c) = (1, 1)$ and set $D(r, c) = \min\{I_1(1), I_2(1)\}$.
- 3) **(Monotonic expansion).** Expand the curve in one of two directions:
 - (Rightward)** If $\sum_c D(r, c) < I_1(r)$ or $r = N$,
 - set $D(r, c + 1) = \min\{I_2(c + 1), I_1(r) - \sum_c D(r, c)\}$.
 - set $c = c + 1$.
 - (Downward)** Otherwise,
 - set $D(r + 1, c) = \min\{I_1(r + 1), I_2(c) - \sum_r D(r, c)\}$.
 - set $r = r + 1$.
- 4) Repeat Step 3) until $(r, c) = (N, N)$.

The Density Sheet Construction Algorithm is similar in spirit to stereo matching algorithms that use the cumulative magnitude of the intensity gradient [38]. Those algorithms match pixels via a monotonic correspondence curve, with invariance to global changes in contrast. Unlike the case of stereo, where the correspondence curve is unique and models both occlusion and corresponding pixels in images of an opaque surface, there are two such curves in the semitransparent case, both of which are consistent with our more general imaging model.

6 SPACE OF PHOTO-CONSISTENT DENSITY FIELDS

While two-view Density Sheets may be adequate for modeling simple scenes, they cannot model complex density fields with large appearance variations across viewpoint. For example, a Density Sheet cannot generate the “parallax” which characterizes the appearance of density distributed over multiple layers. Particularly for more complex flames, the structure of a Density Sheet may actually be perceived as a simple semitransparent surface (see Fig. 11, third row).

To overcome these limitations, we generalize our analysis to multiple views and complex density fields. To achieve this, we use the following three properties of the space of photo-consistent density fields. These properties give us a way to combine multiple two-view solutions




Fig. 7. Visualizing the space of density fields. A point in the figure represents an $N \times N$ density field D . The space of density fields photo-consistent with $M - 1$ views is a region in this space (outermost polygon). The Nesting Property tells us that as the number of input views increases to M , this region shrinks (dark gray polygon). Moreover, the Convexity Property implies that these regions are always convex polytopes. In general, the convex polytope defined by a set of “basis $(M - 1)$ -view density fields” (light gray polygon) may not span the entire M -view polytope.

without working directly in the prohibitively high-dimensional space of $N \times N$ density fields:

Property 1 (Nesting Property). *The space of density fields that are photo-consistent with M input images is a subset of the space of density fields photo-consistent with $M - 1$ of those images.*

Proof. Follows directly from the definition of photo-consistency (Definition 1). \square

The Nesting Property suggests that every density field photo-consistent with M 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 for a particular pixel in an input image, given by (6), 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 nonnegative). \square

A key consequence of the Convexity Property is that the space of photo-consistent density fields is a *convex polytope* (Fig. 7). This follows because the only constraints we impose on the space of density fields, i.e., nonnegativity and photo-consistency, are finite in number and linear in the density elements, implying a piecewise-linear boundary. Although only the photo-consistent density fields on the extreme corners of the polytope span the space of M -view solutions with their convex combination, a sparser sampling on the boundary may still permit a good approximation.

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

Proof. Follows directly from the linearity of image formation given by (6). \square

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

it allows us to operate exclusively on images of dimension N , rather than density fields of dimension N^2 .

7 DENSITY-SHEET DECOMPOSITION

The analysis of Section 6 tells us that if we have, for every pair of input views, a set of “basis density fields” spanning the associated 2-view solution space, we can express every density field that is photo-consistent with M 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 M 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.

7.1 Generating Basis Density Fields

By considering the views pairwise and applying the Density Sheet Construction Algorithm of Section 5, we can trivially generate up to $2 \cdot \binom{M}{2} = M(M - 1)$ Density Sheets from M input views. Unfortunately, this basis is not expressive enough to span the space of photo-consistent density field and, hence, cannot represent arbitrarily complex density fields. We therefore generalize the Density Sheet Construction Algorithm to compute basis density fields that provably span this space. We call the resulting bases *Decomposed Density Sheets*.

Recall that the Density Sheet Construction Algorithm takes into account all pixels in a pair of corresponding epipolar lines. To define the new Decomposed Density Sheets, we combine smaller Density Sheets built from *parts* of the input images that have equal sums. Specifically, given images I_1 and I_2 , we apply the following algorithm (Fig. 8).

Algorithm 2. Density Sheet Decomposition Algorithm

- 1) **(Weight assignment).** Choose a weight, w , with $0 \leq w \leq 1$.
- 2) **(Image decomposition).** Decompose each image into two parts, governed by weight w

$$\begin{aligned} I_1 &= wI'_1 + (1 - w)I''_1 \\ I_2 &= wI'_2 + (1 - w)I''_2, \end{aligned} \quad (8)$$

where all images are nonnegative and have equal sum.

- 3) **(Density Sheet construction).** Build Density Sheets D' and D'' from image pairs I'_1, I'_2 and I''_1, I''_2 , respectively.
- 4) **(Density Sheet combination).** Combine Density Sheets D' and D'' to obtain the Decomposed Density Sheet

$$D = wD' + (1 - w)D''. \quad (9)$$

Note that the resulting Decomposed Density Sheet, D , will always be photo-consistent with the input images I_1, I_2 . This is a direct consequence of the Convexity and Image Linearity Properties of Section 6.

To fully specify the algorithm we must describe how to decompose the images in Step 2, given some choice of the weight w . To do this, we split I_1 and I_2 into a central interval of




Fig. 8. Decomposing \mathbf{I}_1 and \mathbf{I}_2 into component parts according to w . \mathbf{I}_1 is decomposed into a central interval of pixels, $w\mathbf{I}'_1$, whose fraction of the total sum is w , and the remainder, $(1-w)\mathbf{I}'_1$, whose fraction of the sum is $1-w$. \mathbf{I}_2 is decomposed analogously, according to w as well. The central intervals of \mathbf{I}_1 , \mathbf{I}_2 are offset from the first pixel in the images by t_1 and t_2 , respectively. Also shown are the two Density Sheets with opposite diagonal orientations, \mathbf{D}' and \mathbf{D}'' , computed from corresponding pairs of parts (Section 5).

pixels, whose fraction of the total sum is w , and the remainder (Fig. 8). The resulting density field therefore retains the sheet-like structure of the basic Density Sheet in a piecewise fashion. Note that we construct \mathbf{D}' and \mathbf{D}'' with opposite diagonal orientation, otherwise their combination would simply be one of the Density Sheets for the input images.

The family of Decomposed Density Sheets can thus be parameterized by the diagonal orientation of \mathbf{D}' , the weight of the central interval w , and the offsets of the central intervals \mathbf{I}'_1 and \mathbf{I}'_2 from the first pixel in the images. Therefore, given W distinct weights and T distinct offsets, we can create $2WT^2$ Decomposed Density Sheets using this method.

Intuitively, when we are restricted to a limited number of these bases, we would like to choose them so that the central Density Sheet “covers” all areas of the density field. As we change the offset parameters, the rectangular region containing $w\mathbf{D}'$ will translate over the $N \times N$ field, motivating a simple scheme where the T offsets for a central interval of weight $w \geq 1/T$ are distributed linearly over the whole range. Specifically, we define the set of offsets for each input image as $t_1, t_2 \in \{\frac{i}{T-1}(1-w)\}_{i=0}^{T-1}$. By varying the diagonal orientation of the central Density Sheet and selecting the offset for each input image, we thus obtain $2T^2$ bases for each choice of the weight. Likewise, we adopt a simple method to distribute the W different weights linearly, according to $w \in \{\frac{1}{T} + \frac{T-1}{T} \frac{k}{W}\}_{k=0}^{W-1}$.

7.2 Completeness Theorem

A fundamental property of Decomposed Density Sheets is that even though they are easy to generate, the subspace they span contains all density fields that are photo-consistent with the input views. We derive this result by analyzing the space of photo-consistent density fields and showing that we can generate enough Decomposed Density Sheets to span this space. Specifically, we know from Section 6 that the space of photo-consistent density fields is a convex polytope. First, we show that this polytope lies in an $(N-1)^2$ -dimensional hyperplane:

Theorem 2. *The space of $N \times N$ density fields that are photo-consistent with two views is an $(N-1)^2$ -dimensional convex polytope in the space of $N \times N$ matrices.*

Proof. See Appendix B for a constructive proof. \square

Next, we show that the Decomposed Density Sheets form a complete basis for the space of photo-consistent density fields, by constructing a family of Decomposed Density Sheets that spans the $(N-1)^2$ -dimensional hyperplane. In other words, the parameters of the Decomposed Density Sheets (diagonal orientation, weight, and offsets in the input

images) give us enough degrees of freedom to generate bases that span the space:

Theorem 3 (Density Sheet Completeness). *The family of Decomposed Density Sheets forms a complete basis for the space of M -view photo-consistent density fields.*

Proof: (Sketch). We first show how to construct a linearly independent basis of Decomposed Density Sheets for the special case of density fields that are photo-consistent with two constant input images, $\mathbf{I}_1 = \mathbf{I}_2 = \alpha[1 \dots 1]^T$. We then derive transformations that let us reduce the general problem to this case. See Appendix C for a full proof. \square

7.3 Decomposed Density Sheet Reconstruction

With the procedure for generating Decomposed Density Sheets (Section 7.1) as a starting point, we rely on the following algorithm to compute an epipolar slice of the density field that maximizes photo-consistency with M views:

Algorithm 3. Decomposed Density Sheet Reconstruction Algorithm

- 1) **(Basis construction).** For each of $P = M(M-1)/2$ pairs of views, generate a family of B Decomposed Density Sheets, $\{\mathbf{D}_{p1}, \mathbf{D}_{p2}, \dots, \mathbf{D}_{pB}\}$ using Algorithm 2. Warp these into a common global coordinate system.
- 2) **(Basis projection).** Project the Decomposed Density Sheets to all input viewpoints, and stack their images into MP blocks, each of size $N \times B$:

$$\mathbf{F}_{mp} = [\mathbf{I}_m(\mathbf{D}_{p1}) \quad \mathbf{I}_m(\mathbf{D}_{p2}) \quad \dots \quad \mathbf{I}_m(\mathbf{D}_{pB})],$$

where the notation $\mathbf{I}_m(\mathbf{D})$ refers to the projection of density field \mathbf{D} to image \mathbf{I}_m using discrete integration (6).

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

$$\mathcal{F} = \begin{bmatrix} \mathbf{F}_{11} & \mathbf{F}_{12} & \dots & \mathbf{F}_{1P} \\ \mathbf{F}_{21} & \mathbf{F}_{22} & \dots & \mathbf{F}_{2P} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{F}_{M1} & \mathbf{F}_{M2} & \dots & \mathbf{F}_{MP} \end{bmatrix}, \quad \mathcal{I} = \begin{bmatrix} \mathbf{I}_1 \\ \mathbf{I}_2 \\ \vdots \\ \mathbf{I}_M \end{bmatrix},$$

to form the following quadratic programming problem:

$$\begin{aligned} & \text{minimize} && \|\mathcal{F}\mathbf{x} - \mathcal{I}\|^2 \\ & \text{subject to} && \sum \mathbf{x} = 1, \quad \mathbf{x} \geq 0. \end{aligned}$$

- 4) **(QP solution).** Solve this quadratic optimization problem by any standard method (e.g., [39]). The resulting optimal \mathbf{x} gives the weights needed to express the density field as a convex combination of the Decomposed Density Sheets computed in Step 1.

Because the quadratic programming problem is convex, the solver will find a global minimum of the objective function. This solution is not necessarily unique because the minimum $\|\mathcal{F}\mathbf{x} - \mathcal{I}\|^2 = 0$ will be attained by all density fields of this form that are photo-consistent.

7.4 Discussion

The result of this algorithm is a reconstruction of the density field as a superposition of piecewise Density Sheets defined by pairs of views, namely a convex combination of the various Decomposed Density Sheet bases. As with the basic Density Sheet solution, the sheet-like structures are well-suited to view synthesis because they allow us to effectively warp pieces of the original images in order to render them.

If the $P \cdot B$ Decomposed Density Sheets computed by a run of Algorithm 3 are such that a convex combination is photo-consistent with all views, the algorithm will attain the optimal lower bound, $\|\mathcal{F}\mathbf{x} - \mathcal{I}\|^2 = 0$. Otherwise, it will return the convex combination of bases that minimizes inconsistency with the input images in a least-squared sense.

Note that Step 3 of the algorithm is suboptimal because a combination of Decomposed Density Sheets constrained to be convex may not, in general, span all M -view photo-consistent solutions (Fig. 7). While optimizing over linear, rather than convex, combinations of sufficiently many sheets would avoid this problem, this would incur a significant computational cost. This is because explicit enforcement of nonnegativity constraints for arbitrary linear combinations requires working directly in the N^2 -dimensional space of density fields. In practice, we enforce the convex combination constraint exclusively.

As described in Section 4, this algorithm is applied independently to a collection of epipolar planes slicing through the 3D volume. The coherence between adjacent epipolar lines in the input images, however, implies that adjacent slices of the volume will have similar decompositions. In practice, this tends to generate a superposition of coherent 2D surfaces in the 3D volume.⁵

8 EXPERIMENTAL RESULTS

We performed experiments on a variety of scenes of fire, some simulated and some real. Here, we show results from three scenes containing real fire that try to convey a range of different flame structures.⁶ We show that the basic Density Sheet produces excellent results for simple to moderately complex flames and that the Density-Sheet Decomposition algorithm is useful for reconstructing more complex scenes of fire.

5. Formally, it is easy to show that the function that maps input images to Density Sheets is smooth. This is why Density Sheets from adjacent scanlines have similar shape when the intensity variation across these scanlines is small (e.g., Fig. 9, rows 3-4).

6. See <http://www.cs.toronto.edu/~hasinoff/fire> for color images and video clips, as well as more results.

8.1 “Torch” Data Set

The first scene (“torch”) was a citronella patio torch, burning with a flickering flame about 10 cm tall. Two synchronized progressive-scan Sony DXC-9000 cameras, roughly 90 degrees apart, were used to acquire videos of the flame at 640×480 resolution (Fig. 2). While ground truth for this scene was not available, the structure of the flame is simple, and it appears that epipolar slices actually contain a single elongated blob of density.

The cameras were calibrated using Bouguet’s Camera Calibration Toolbox for Matlab [40] to an accuracy of about 0.5 pixels. The input views were then rectified so that corresponding horizontal scanlines define the epipolar slices. For each of the 25 frames in the video sequence, each epipolar slice was reconstructed independently.

We compared three different reconstruction methods with respect to the quality of synthesized views interpolating between the two cameras (Fig. 9). First, the multiplication solution shows typical blurring and doubling artifacts. Interpolated views of this solution do not contain the same high frequency content as the input images and suggest that the viewpoints were “accidentally” aligned so as to hide significant structures in the scene. Second, an algebraic method based on fitting Gaussian blobs to the density field [41] overfits the two input images and produces a badly mottled appearance for synthesized views. This confirms that sparse-view tomography methods are not suitable when the number of viewpoints is extremely limited. Third, the Density Sheet reconstruction produces very realistic views that appear indistinguishable in quality from the input views.

If the viewpoint is varied while time is fixed, the true nature of the Density Sheet reconstruction as a transparent surface can be perceived. The addition of temporal dynamics, however, enhances photo-realism. For simple flames like the “torch” scene, a two-view reconstruction consisting of a single Density Sheet serves as a good impostor for the true scene.

To render the Density Sheet solution, we first compute the geometry of the Density Sheet and then use this sheet to warp the input images to the interpolated view by backward mapping [24]. We then blend the warped input images according to a simple view-dependent linear interpolation. To render the other solutions, we use standard volume rendering which applies discrete integration to the reconstructed 3D density field.

8.2 “Burner” Data Set

The second scene (“burner”), courtesy of Ihrke and Magnor [23], [42], consists of a colorful turbulent flame emanating from a gas burner. This scene was captured for 348 frames using eight synchronized video cameras with 320×240 resolution. The cameras are roughly equidistant from the flame and distributed over the viewing hemisphere. Ground truth is not available for this scene either, but an algebraic tomography method restricted by the visual hull produces a reasonable approximation to the true density field [23].

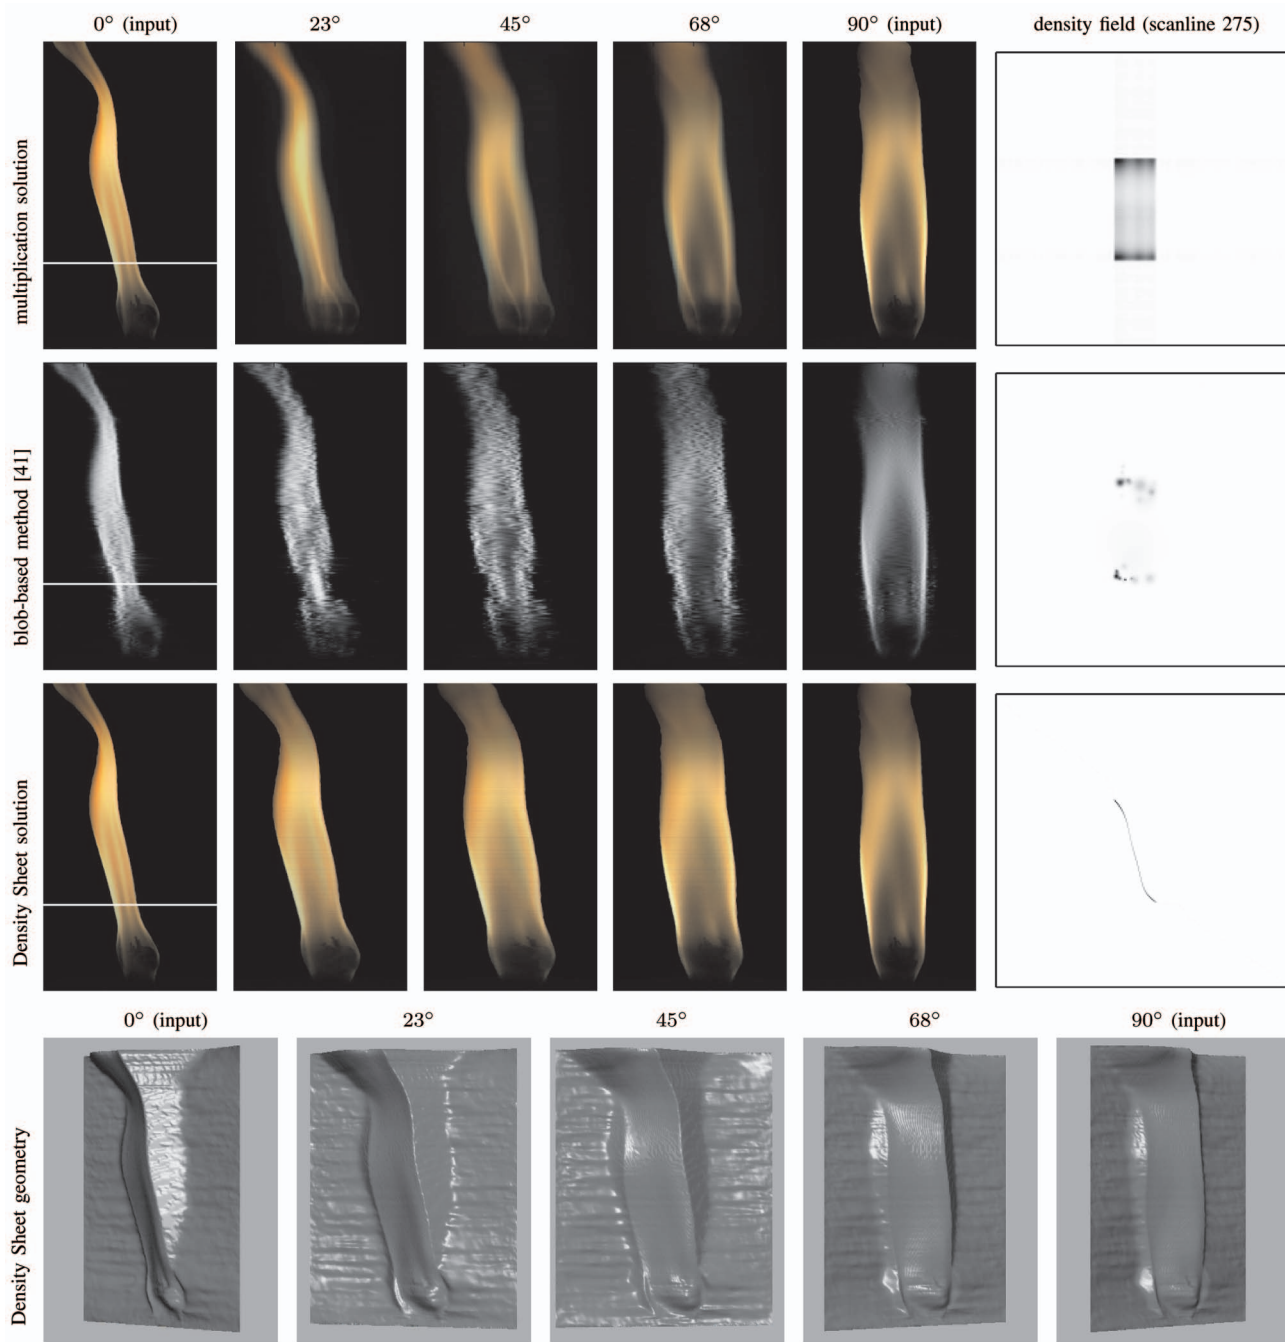


Fig. 9. Three different reconstructions of the “torch” data set (frame 8) from two input images. Labels on the left indicate the method used. On the right, we show a top view of a reconstructed epipolar slice, corresponding to the white horizontal line in the 0° images (dark regions denote areas with concentrated fire density). Note that the multiplication and Density Sheet solutions reproduce the input images exactly. The bottom row illustrates the 3D geometry of the Density Sheet solution, rendered as a specular surface.

For the “burner” data set, we applied our algorithms to just two of the input cameras, spaced about 60 degrees apart. We used the calibration provided to rectify the images from this pair of cameras and then computed the Density Sheet from corresponding scanlines. Finally, as in view morphing [43], we used a homography to warp the images synthesized in the rectified space back to the interpolated camera (Fig. 10).

Although the “burner” scene is significantly more complex than the previous scene, two-view reconstruction using a single Density Sheet can still produce realistic

intermediate views (Fig. 10, rows 1–2). As with the “torch” data set, the view interpolation appears even more plausible when the dynamics of the flame are added.

As shown in the third row of Fig. 10, a significant artifact occurs when the true flame consists of multiple structures and when our imaging model (Section 3) is not satisfied exactly. In this example, the true scene and both input images consist of two major flame structures, but the images disagree on the total intensity (i.e., the total density) of each structure. For the Density Sheet solution, this discrepancy leads to a “tearing” effect where a spurious




Fig. 10. Density Sheet reconstruction of the “burner” data set [23], [42], interpolating between two cameras about 60 degrees apart. At the top, the first and last columns are input views; the middle columns are interpolated. Each row represents a different frame from the turbulent flame sequence. The highlighted region draws attention to a spurious thin structure present in interpolated views of the third row. This artifact is due to inexact correspondence between major flame structures because our imaging model is violated. At the bottom, we compare views interpolated using the Density Sheet, midway between the two cameras, alongside nearby views generated by an algebraic tomography method using all eight cameras [23].

thin vertical structure is used to account for the lack of correspondence. Note that this problem applies more generally to any technique assuming a linear image formation model for this particular data set.

We also compare interpolated views generated using the Density Sheet reconstruction and the algebraic tomography method of Ihrke and Magnor [23]. While the tomography method uses all eight cameras to reconstruct a coarse 96^3 voxel grid, the Density Sheet reconstruction explains the density field locally using a single transparent surface defined by the closest two cameras. Despite this imbalance, the simple Density Sheet reconstruction (Fig. 10, bottom) may still be preferable for view synthesis. The volume renderings produced by algebraic tomography are relatively noisier, lower resolution, and less “similar” to the sharp and smooth input images.

8.3 “Jet” Data Set

The third scene (“jet”) consists of a complex flame emerging from a gaseous jet (Fig. 11, first row). The data set consists of 47 synchronized views, roughly corresponding to inward-facing views arranged around a quarter-circle, captured from a promotional video for a commercial 3D freeze-frame system [44]. Since no explicit calibration was available for this sequence, we assumed that the views were evenly spaced.

To test the view synthesis 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 (Fig. 11). Rendering results from the multiplication solution and the Density Sheet solution (Algorithm 1) suggest that these solutions cannot accurately capture the flame’s appearance and structure, and are not well suited to more complex flames with large appearance variation across viewpoint.




Fig. 11. Reconstruction of the “jet” data set. 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 synthesized for the multiplication and Density Sheet solutions and is an input view for the other methods; 68° view, synthesized 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).

To incorporate more views, we applied the blob-based algebraic tomography method in [41] and the Density-Sheet Decomposition algorithm (Algorithm 3) with the 45 degree view as a third image. In the latter algorithm, we generated $B = 150$ Decomposed Density Sheets for each of the $P = 3$ pairs of input views in Step 1, giving rise to a total of 450 basis density 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, but using the same three-view basis.⁷ The results in Fig. 11 and Table 1 suggest that the Density-Sheet Decomposition algorithm produces rendered images that are superior to the other

7. The reason for using $M' > M$ views for photo-consistency and M for basis computation is that this leaves unchanged the number of unknowns in the quadratic programming problem. With a constant number of bases per view-pair (Section 7.1), the number of unknowns grows according to $O(M^2)$, not $O(M'^2)$, which makes the problem less computationally intensive.

TABLE 1
Per-Pixel RMS Reconstruction Error for the “Jet” Data Set

reconstruction method (views)	RMS error	
	input	all
density sheet solution (2) ^a	0	26.6
multiplication solution (2)	0	21.5
blob-based method [41] (3)	13.3	19.0
density sheet decomposition (3)	10.1	18.4
density sheet decomposition (3+4) ^b	11.5	15.8

^aThe first two reconstructions reproduce the input views exactly.

^bThree input views were used to generate the basis, and four additional views were used to evaluate photo-consistency only.

methods. They also show that increasing the number of images has a clear benefit.

Three observations can be made from these experiments: First, the convex combination of Decomposed Density 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 overfitting artifacts of the blob-based method. This is because the family of Decomposed Density Sheets is expressive and induces a bias toward compact solutions. Third, the sheets’ surface-like structure enables photo-realistic rendering of complex semitransparent scenes by simply warping and blending partial input images (8). Thus, the density models we extract may be suitable for interactive graphics applications.

9 DISCUSSION

One feature of semitransparency is that no part of the scene will be occluded as long as the camera exposure is set to prevent saturation. This benefit, however, comes at the cost that the space of photo-consistent density fields is fundamentally ambiguous when the number of input views is small. In this work, we have assumed spatial coherence in the density field in order to achieve good view synthesis.

It would be useful to understand how photo-consistency constraints from three or more views restrict the space of two-view and M -view solutions. In particular, a third view would be expected to eliminate much of the ambiguity that causes the doubling of sharp features found in the multiplication solution (Fig. 4). Such a view is less likely to be helpful for disambiguating smooth and diffuse regions because, as in standard stereo, the lack of high-frequency detail makes unambiguous reconstruction in these regions more difficult.

So far, we have assumed that the cameras are restricted to lie on the same plane, but most previous methods for semitransparent reconstruction extend readily to general viewing configurations (at least in theory). Unfortunately, in such configurations the basis coefficients from different 2D slices of the 3D density field become linked, leading to much larger quadratic programming problems. In the worst case, with three orthogonal views, the coefficients of all Decomposed Density Sheet bases from all slices are needed to evaluate photo-consistency for a single pixel. To enhance

computational efficiency, it might be possible to consider Decomposed Density Sheet bases defined in 3D and constructed from bands of epipolar lines in the input views.

From the point of view of fire reconstruction, simultaneous recovery of the 3D temperature, density, and pressure fields, commonly used in combustion engineering and computer graphics [1], [11] appears to be a daunting task. Consequently, simpler models of fire, like the one we describe in Section 3, dominate current reconstruction techniques. This suits our purpose, since we are more interested in modeling appearance variation than extracting particular physical properties of fire. In fact, the advantage of any image-based technique is that any flames we can observe can be reconstructed directly, without modeling or simulating their interaction with the environment [1], [6].

10 CONCLUDING REMARKS

In this paper, we introduced Density Sheets as a basic primitive for representing, reconstructing, and rendering semitransparent scenes from a small number of input views. This primitive was derived through a detailed analysis of the space of densities that are photo-consistent with two or more views and was designed with four key properties in mind—*photo-consistency*, *uniqueness*, *spatial compactness*, and *representational completeness*. We showed that these properties lead to simple reconstruction algorithms that produce stable solutions even for very limited numbers of input views (even two), despite the ill-posed nature of the volumetric reconstruction problem. Importantly, we showed that the behavior and solution space of these algorithms can be characterized and that the resulting reconstructions have good view interpolation and extrapolation capabilities.

The question of how best to capture the global 3D structure and dynamics of fire and other physical phenomena remains open. Toward this goal, we are investigating new spatio-temporal coherence constraints and ways to integrate our approach with traditional simulation methods. In theory, image-based reconstructions could be used to train or validate physical simulators, and these simulators could be used to evaluate reconstructions in terms of their physical plausibility.

While our experiments suggest that Density Sheets and the Density-Sheet Decomposition algorithm are useful for fire reconstruction, these tools may also prove useful in more general contexts. This includes the reconstruction of other volumetric phenomena such as smoke [4], sparse-view tomography problems in medical diagnostics [20], and accelerated image-based methods for volume rendering.

APPENDIX A

PROOF OF THEOREM 1

Proof (by induction). We show that Algorithm 1 generates a density field \mathbf{D} with the desired properties. The downward and rightward expansions in Step 3 guarantee that the elements of \mathbf{D} traced by the pair (r, c) define a monotonic four-connected curve. Therefore, it is sufficient to show that \mathbf{D} is nonnegative and photo-consistent. For ease of notation, we denote the row and column sums of \mathbf{D} by $\mathbf{s}_1 = \mathbf{D}\mathbf{1}$ and $\mathbf{s}_2 = \mathbf{D}^T\mathbf{1}$, respectively, where $\mathbf{1} = [1 \dots 1]^T$.

We show that properties **H1-H4**, corresponding to our induction hypothesis, hold throughout the course of the algorithm, as the element (r, c) traces a path from $(1, 1)$ to (N, N) :

- H1:** \mathbf{D} is photo-consistent with the subvectors $\mathbf{I}_1(1, \dots, r-1)$ and $\mathbf{I}_2(1, \dots, c-1)$;
- H2:** \mathbf{D} is photo-consistent with $\mathbf{I}_1(r)$ or $\mathbf{I}_2(c)$;
- H3:** \mathbf{D} satisfies $s_1 \leq \mathbf{I}_1$ and $s_2 \leq \mathbf{I}_2$; and
- H4:** $\mathbf{D} \geq 0$.

For the base case $(r, c) = (1, 1)$, **H1-H4** are trivially satisfied after Step 2 of the algorithm.

Part I: The algorithm terminates with \mathbf{D} photo-consistent with $\mathbf{I}_1, \mathbf{I}_2$. First, we show that if \mathbf{D} is photo-consistent with one of the vectors $\mathbf{I}_1, \mathbf{I}_2$, then **H1-H4** imply that \mathbf{D} must be photo-consistent with both of them. In particular, assume wlog that $(r, c) = (N, j)$ and $s_1 = \mathbf{I}_1$. From **H3**, we know that $s_2(j) \leq \mathbf{I}_2(j)$, and because $\sum \mathbf{I}_1 = \sum \mathbf{I}_2$, all density must be distributed within the first j columns. Thus, $\mathbf{I}_2(j+1, \dots, N) = 0$ and $s_2(j) = \mathbf{I}_2(j)$, implying that \mathbf{I}_2 is photo-consistent as well. Second, we note that the expansions prescribed by Step 3 of the algorithm are valid even at the boundary cases: If $r = N$, then the expansion will be rightward; if $c = N$, the expansion will be downward unless $s_1(r) < \mathbf{I}_1(r)$. But, this inequality is never satisfied: if $s_1(r) < \mathbf{I}_1(r)$, **H2-H3** imply that $s_2 = \mathbf{I}_2$ and, therefore, $s_1 = \mathbf{I}_1$, leading to a contradiction. Thus, the algorithm always terminates with $(r, c) = (N, N)$.

Part II: Every application of Step 3 preserves the induction hypothesis. If the expansion is rightward, then either $s_1(r) < \mathbf{I}_1(r)$ or $s_1 = \mathbf{I}_1$. In the first case, **H2-H3** imply that $s_2(c) = \mathbf{I}_2(c)$. Then, by construction, $\mathbf{D}(r, c+1)$ will be set to some nonnegative value so that either the new column is photo-consistent (i.e., $s_2(c+1) = \mathbf{I}_2(c+1)$), or the current row becomes photo-consistent (i.e., $s_1(r) = \mathbf{I}_1(r)$). Both outcomes preserve **H1-H4** for $(r, c+1)$. If $s_1 = \mathbf{I}_1$, it must also be the case that $s_2 = \mathbf{I}_2$ and setting $\mathbf{D}(r, c) = 0$ as prescribed by Step 3 of the algorithm preserves **H1-H4**. Analogously, every downward expansion also preserves the inductive hypothesis.

Therefore, the algorithm terminates with $(r, c) = (N, N)$ and **H1-H4** intact. Hence, $\mathbf{D} \geq 0$ and \mathbf{D} is photo-consistent with both vectors. The proof for \mathbf{D}' is analogous. \square

APPENDIX B

PROOF OF THEOREM 2

Proof. From Section 6, we know that the space of photo-consistent density fields is a convex polytope. We show that it is $(N-1)^2$ -dimensional in the two-view case.

Given two nonnegative, N -dimensional column vectors \mathbf{I}_1 and \mathbf{I}_2 , consider an $N \times N$ matrix \mathbf{D} whose row and column sums, respectively, are equal to these vectors, i.e., $\mathbf{D}\mathbf{1} = \mathbf{I}_1$ and $\mathbf{D}^T\mathbf{1} = \mathbf{I}_2$. Every such matrix can be expressed as $\mathbf{D} = \mathbf{D}_0 + \mathbf{S}$, where \mathbf{D}_0 is any photo-consistent density (e.g., the multiplication solution), and \mathbf{S} preserves row and column sums, i.e., $\mathbf{S}\mathbf{1} = \mathbf{0}$ and $\mathbf{S}^T\mathbf{1} = \mathbf{0}$. 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 of $(N-1)^2$ on the dimensionality of the polytope.




Fig. 12. The matrices $\{\mathbf{S}_{ij}\}_{i,j < N}$ form a basis for the space of $N \times N$ photo-consistent matrices. Each matrix contains only four nonzero elements arranged in a rectangle whose bottom-right corner is fixed at element (N, N) . \mathbf{S}_{ij} is zero in the gray-shaded regions. Clearly, $\mathbf{S}_{ij}\mathbf{1} = \mathbf{0}$ and $\mathbf{S}_{ij}^T\mathbf{1} = \mathbf{0}$.

We now show that the $(N-1)^2$ upper bound is tight by explicitly constructing a basis for the polytope's hyperplane that consists of $(N-1)^2$ linearly independent matrices. In particular, consider the family of $N \times N$ matrices $\{\mathbf{S}_{ij}\}_{i,j < N}$ defined in Fig. 12, where each matrix consists of four nonzero elements arranged in a rectangle. Adding a matrix \mathbf{S}_{ij} to \mathbf{D}_0 has the effect of shifting mass from two diagonal corners of the rectangle to the other two corners.⁸ The matrices $\{\mathbf{S}_{ij}\}_{i,j < N}$ are linearly independent since $\mathbf{S}_{ij}(i, j) = 1$ and $\mathbf{S}_{i'j'}(i, j) = 0$ for all other $i', j' < N$. Therefore, they span the $(N-1)^2$ -dimensional hyperplane.

The hyperplane can be expressed as a linear combination of the resulting two-view photo-consistent matrices, \mathbf{D}_0 and $\{\mathbf{D}_0 + \mathbf{S}_{ij}\}_{i,j < N}$. \square

APPENDIX C

PROOF OF THEOREM 3

Lemma 1. Let Σ_c be the set of nonnegative $N \times N$ matrices whose row and column sums are constant vectors, i.e., $\Sigma_c = \{\mathbf{D} \mid \mathbf{D}\mathbf{1} = \alpha\mathbf{1} \text{ and } \mathbf{D}^T\mathbf{1} = \alpha\mathbf{1}\}$ for some $\alpha \in \mathbb{R}$. The Decomposed Density Sheets form a basis spanning Σ_c .

Proof. Assume, without loss of generality, that $\alpha = 1$. Theorem 2 tells us that the set Σ_c is an $(N-1)^2$ -dimensional polytope. We use a constructive proof, specifying Decomposed Density Sheet matrices that span this polytope's hyperplane.

First, consider the family of $N \times N$ binary matrices $\{\mathbf{T}_{ij}\}_{i,j < N}$ defined in Figs. 13a and 13b. These matrices are zero except for three identity-matrix blocks and a 2×2 block at position (i, j) . Each matrix \mathbf{T}_{ij} satisfies $\mathbf{T}_{ij}\mathbf{1} = \mathbf{1}$ and $\mathbf{T}_{ij}^T\mathbf{1} = \mathbf{1}$ and, therefore, belongs to Σ_c . Moreover, \mathbf{T}_{ij} is a Decomposed Density Sheet corresponding to the output of Algorithm 2 with $\mathbf{I}_1 = \mathbf{I}_2 = \mathbf{1}$ and specified by the parameters $t_1 = \frac{i-1}{N}$, $t_2 = \frac{j-1}{N}$, and $w = \frac{2}{N}$ (see Section 7.1).

We now define matrix $\hat{\mathbf{T}}_{ij} = \mathbf{T}_{ij} - \mathbf{I}_N$ and show that the family $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N}$ is linearly independent for $i \neq j+1$. This provides us with a basis of dimension $(N-1)^2 - (N-2)$, each of whose members, $\hat{\mathbf{T}}_{ij}$, is the difference of two Decomposed Density Sheets (\mathbf{I}_N and \mathbf{T}_{ij}). Note that each

8. The matrices \mathbf{S}_{ij} can be thought of a generalization to continuous-valued density fields of the "switching matrices" described in [27].




Fig. 13. Matrix definitions for the proof of Theorem 3. (a) and (b) For a given $i, j < N$, the $N \times N$ matrix \mathbf{T}_{ij} contains a 2×2 -pixel central block at position (i, j) and three identity matrix blocks. \mathbf{I}_k denotes an identity matrix of size $k \times k$. The elements in gray-shaded regions are zero. (c) Definition of the $N \times N$ matrix \mathbf{U}_k .

$\hat{\mathbf{T}}_{ij}$ will have elements with values 1, 0, or -1. Furthermore, note the recursive structure of $\{\hat{\mathbf{T}}_{ij}\}$, where the upper-leftmost $(N-1) \times (N-1)$ block of a given $\hat{\mathbf{T}}_{ij}$ with $i, j < N-1$ corresponds directly to matrix $\hat{\mathbf{T}}_{ij}$ for a problem of size $N-1$.

We use induction on N to obtain our linear independence result, where the induction hypothesis for $N-1$ implies that $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N-1}$ are linearly independent for $i \neq j+1$. The base cases for $N \leq 2$ are trivial. Given the induction hypothesis, it remains to show that the addition of the $2N-4$ matrices, $\{\hat{\mathbf{T}}_{N-1,j}\}_{j < N-2}$ and $\{\hat{\mathbf{T}}_{i,N-1}\}_{i < N}$, preserves linear independence. This is simple to show, since $\hat{\mathbf{T}}_{N-1,j}(N, j) = 1$ for $j < N-2$, but the element (N, j) is zero for all other $\hat{\mathbf{T}}_{ij}$. Similarly, $\hat{\mathbf{T}}_{i,N-1}(i, N) = 1$ for $i < N$, but apart from the case $i = N-2$, the element (i, N) is zero for all other $\hat{\mathbf{T}}_{ij}$. This leaves just $\hat{\mathbf{T}}_{N-2,N-1}$, for which element $(N-2, N)$ is nonzero in $\{\hat{\mathbf{T}}_{N-1,j}\}_{j < N-2}$ as well. However, none of these other matrices can participate in a linear combination reproducing $\hat{\mathbf{T}}_{N-2,N-1}$, since they comprise the only $\hat{\mathbf{T}}_{ij}$ for which any of the elements $\{(N, k)\}_{k=1}^{N-3}$ are nonzero.

The family $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N}$ gives us $(N-1)^2 - (N-2)$ linearly independent matrices. So, to obtain a basis for the $(N-1)^2$ dimensions of Σ_c we need $N-2$ additional linearly independent matrices. To construct them, we consider a second family of $N \times N$ matrices $\{\mathbf{U}_k\}_{k < N-1}$, shown in Fig. 13c. This family contains a total of $N-2$ members, each of which belongs to Σ_c . Moreover, \mathbf{U}_k is a Decomposed Density Sheet by construction, corresponding to the output of Algorithm 2 with $\mathbf{I}_1 = \mathbf{I}_2 = \mathbf{1}$ and parameters $t_1 = \frac{k+1}{N}$, $t_2 = 0$, and $w = \frac{1}{2N}$. Similarly, we define $\hat{\mathbf{U}}_k = \mathbf{U}_k - \mathbf{I}_N$.

The matrices in the family $\{\hat{\mathbf{U}}_k\}_{k < N-1}$ are linearly independent, and they are linearly independent of $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N}$ as well. In particular, the matrices $\{\hat{\mathbf{U}}_k\}_{k < N-1}$ are linearly independent because $\hat{\mathbf{U}}_k(k+2, 1) = 1$ and this element is zero for all $\hat{\mathbf{U}}_{k' \neq k}$. They are linearly independent of $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N}$ because $\hat{\mathbf{U}}_k(1, 2) = 1$ for all k , and the only member of $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N}$ for which this element is nonzero is $\hat{\mathbf{T}}_{11}$, which cannot participate in a linear combination reproducing $\hat{\mathbf{U}}_k$ (because $\hat{\mathbf{T}}_{11}$ is the only $\hat{\mathbf{T}}_{ij}$ for which element $(2, 1)$ is nonzero, whereas $\hat{\mathbf{U}}_k(2, 1) = 0$ for any k).

Therefore, the combined family $\{\hat{\mathbf{T}}_{ij}\}_{i,j < N} \cup \{\hat{\mathbf{U}}_k\}_{k < N-1}$ contains $(N-1)^2$ linearly independent matrices.

It follows that the polytope Σ_c can be expressed as a linear combination of the Decomposed Density Sheets $\{\mathbf{T}_{ij}\}_{i,j < N}$, $\{\mathbf{U}_k\}_{k < N-1}$, and \mathbf{I}_N . \square

Lemma 2. Let Σ_r be the set of nonnegative $N \times N$ matrices whose row and column sums are vectors of rational elements, i.e., $\Sigma_r = \{\mathbf{D} \mid \mathbf{D}\mathbf{1} = \frac{1}{m}\mathbf{n}_1 \text{ and } \mathbf{D}^T\mathbf{1} = \frac{1}{m}\mathbf{n}_2\}$ for $m > 0$, $m \in \mathbb{N}$ and $\mathbf{n}_1, \mathbf{n}_2 \in \mathbb{N}^N$. The Decomposed Density Sheets form a basis spanning Σ_r .

Proof. Given a particular $\mathbf{D} \in \Sigma_r$, we describe a mapping \mathcal{T} that reduces the problem to the case where row and columns are constant vectors. This lets us apply Lemma 1 to obtain a set of Decomposed Density Sheets spanning the transformed space. To show that \mathbf{D} is spanned by Decomposed Density Sheets it suffices to show that 1) the inverse mapping, \mathcal{T}^{-1} , is linear and onto, and 2) that the basis defined in Lemma 1, when transformed by \mathcal{T}^{-1} , is still a set of Decomposed Density Sheets.

Assume, without loss of generality, that $\sum \mathbf{n}_1 = \sum \mathbf{n}_2 = m$. We define the mapping $\mathcal{T} : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{m \times m}$ by expanding each matrix element $\mathbf{D}(i, j)$ to a matrix block of size $\mathbf{n}_1(i) \times \mathbf{n}_2(j)$, or collapsing it to an empty matrix if either $\mathbf{n}_1(i)$ or $\mathbf{n}_2(j)$ is zero. In the first case, each of the block elements is set to $\frac{\mathbf{D}(i,j)}{\mathbf{n}_1(i)\mathbf{n}_2(j)}$. Intuitively, we can think of matrix \mathbf{D} as being stretched nonuniformly, along the columns and then the rows, so that its row and column sums become constant vectors. The transformed matrix $\mathcal{T}(\mathbf{D})$ satisfies $\mathcal{T}(\mathbf{D})\mathbf{1} = \mathcal{T}(\mathbf{D})^T\mathbf{1} = \frac{1}{m}\mathbf{1}$.

The inverse mapping $\mathcal{T}^{-1} : \mathbb{R}^{m \times m} \rightarrow \mathbb{R}^{N \times N}$ is defined by taking sums of the matrix blocks defined by \mathcal{T} , where the “sum” of an empty matrix block is taken to be zero, so that $\mathcal{T}^{-1}(\mathcal{T}(\mathbf{D})) = \mathbf{D}$ and \mathcal{T}^{-1} is onto. Furthermore, because summation is a linear operator, \mathcal{T}^{-1} is also linear, i.e., $\mathcal{T}^{-1}(\mathbf{D}_1 + \alpha\mathbf{D}_2) = \mathcal{T}^{-1}(\mathbf{D}_1) + \alpha\mathcal{T}^{-1}(\mathbf{D}_2)$.

It remains to show that a Decomposed Density Sheet is still a Decomposed Density Sheet after transformation by the inverse mapping, \mathcal{T}^{-1} . But, this is indeed the case because block-wise summation preserves the Density Sheet structure, i.e., the properties of Theorem 1. \square

Theorem 3 (Density Sheet Completeness). *The family of Decomposed Density Sheets forms a complete basis for the space of M -view photo-consistent density fields.*

Proof. For any sufficiently small $\varepsilon > 0$, we reduce the 2-view problem to an ε -approximation of the case where row and column sums consist of rational elements. We then apply Lemma 2 to this rational approximation and show that the elements of the resulting Decomposed Density Sheet basis can be adjusted to span the original space. Finally, we apply the Nesting Property (Property 1, Section 6) to extend this completeness result to M views.

Specifically, without loss of generality consider an $N \times N$ nonnegative matrix \mathbf{D} satisfying $\sum \mathbf{D} = 1$. Define the matrix $\tilde{\mathbf{D}} = \lceil \frac{\mathbf{D}}{\varepsilon} \rceil / \lceil \frac{1}{\varepsilon} \rceil$ with rational elements such that $\tilde{\mathbf{D}} \in \Sigma_r$ and $\max |\tilde{\mathbf{D}} - \mathbf{D}| < \varepsilon$.

We apply the proofs of Lemmas 1-2 to construct a basis of Decomposed Density Sheets that spans the space of matrices photo-consistent with $\tilde{\mathbf{D}}\mathbf{1}$ and $\tilde{\mathbf{D}}^T\mathbf{1}$. Denote this basis by $\{\mathbf{B}_k\}_{k < K}$. Now, using the same parameters corresponding to a particular \mathbf{B}_k (i.e., the parameters t_1 , t_2 , and w used in Lemmas 1-2), we create a corresponding Decomposed Density Sheet matrix \mathbf{B}'_k that is photo-consistent with $\mathbf{D}\mathbf{1}$ and $\mathbf{D}^T\mathbf{1}$. By construction $\max |\mathbf{B}'_k - \mathbf{B}_k| < N^2\varepsilon$, since in the worst case all N^2 perturbations of magnitude ε will be concentrated on a single element.

It now suffices to show that the family of Decomposed Density Sheets $\{\mathbf{B}'_k\}_{k < K}$ consists of linearly independent matrices and, therefore, spans \mathbf{D} . Since the group of nonsingular matrices is open, there exists some $\delta > 0$ such that any family $\{\mathbf{B}''_k\}_{k < K}$ with $\max |\mathbf{B}''_k - \mathbf{B}_k| < \delta$ consists of linearly independent matrices (see [45], for example). We can therefore guarantee linear independence by choosing $\varepsilon < \delta/N^2$. \square

ACKNOWLEDGMENTS

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

REFERENCES

- [1] D. Nguyen, R. Fedkiw, and H. Jensen, "Physically Based Modeling and Animation of Fire," *Proc. ACM SIGGRAPH*, pp. 721-728, 2002.
- [2] J. Stam and E. Fiume, "Depicting Fire and Other Gaseous Phenomena Using Diffusion Processes," *Proc. ACM SIGGRAPH*, pp. 129-136, 1995.
- [3] R. Fedkiw, J. Stam, and H.W. Jensen, "Visual Simulation of Smoke," *Proc. ACM SIGGRAPH*, pp. 15-22, 2001.
- [4] T. Hawkins, P. Einarsson, and P. Debevec, "Acquisition of Time-Varying Participating Media," *Proc. ACM SIGGRAPH*, pp. 812-815, 2005.
- [5] N. Foster and D.N. Metaxas, "Modeling the Motion of a Hot, Turbulent Gas," *Proc. ACM SIGGRAPH*, pp. 181-188, 1997.
- [6] A. Lamorlette and N. Foster, "Structural Modeling of Flames for a Production Environment," *Proc. ACM SIGGRAPH*, pp. 729-735, 2002.
- [7] N. Chiba, K. Muraoka, H. Takahashi, and M. Miura, "Two-Dimensional Visual Simulation of Flames, Smoke and the Spread of Fire," *J. Visualization and Computer Animation*, vol. 5, no. 1, pp. 37-54, Jan.-Mar. 1994.
- [8] Y. Takai, K. Ecchu, and N.K. Takai, "A Cellular Automaton Model of Particle Motions and Its Applications," *The Visual Computer*, vol. 11, no. 5, pp. 240-252, 1995.
- [9] R.T. Baum, K.B. McGrattan, 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.
- [10] D.P. Correia, P. Ferrão, and A. Caldeira-Pires, "Advanced 3D Emission Tomography Flame Temperature Sensor," *Combustion Science and Technology*, vol. 163, pp. 1-24, 2001.
- [11] A. Schwarz, "Multi-Tomographic Flame Analysis with a Schlieren Apparatus," *Measurement Science & Technology*, vol. 7, pp. 406-413, 1996.
- [12] 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.
- [13] G.W. Faris and R.L. Byer, "Beam-Deflection Optical Tomography," *Optics Letters*, vol. 12, no. 3, pp. 72-74, 1987.
- [14] 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.
- [15] B.W. Albers and A.K. Agrawal, "Schlieren Analysis of an Oscillating Gas-Jet Diffusion Flame," *Combustion and Flame*, vol. 199, pp. 84-94, 1999.
- [16] T. Otsuka and P. Wolanski, "Particle Image Velocimetry (PIV) Analysis of Flame Structure," *J. Loss Prevention in the Process Industries*, vol. 14, no. 6, pp. 503-507, 2001.
- [17] A. Schödl, R. Szeliski, D.H. Salesin, and I. Essa, "Video Textures," *Proc. ACM SIGGRAPH*, pp. 489-498, 2000.
- [18] D. Drysdale, *An Introduction to Fire Dynamics*, second ed. John Wiley and Sons, 1998.
- [19] P.J. Narayanan, P.W. Rander, and T. Kanade, "Constructing Virtual Worlds Using Dense Stereo," *Proc. Int'l Conf. Computer Vision*, pp. 3-10, 1998.
- [20] A.C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [21] T. Frese, C.A. Bouman, and K. Sauer, "Adaptive Wavelet Graph Model for Bayesian Tomographic Reconstruction," *IEEE Trans. Image Processing*, vol. 11, no. 7, pp. 756-770, July 2002.
- [22] R. Szeliski and P. Golland, "Stereo Matching with Transparency and Matting," *Proc. Int'l Conf. Computer Vision*, pp. 517-524, 1998.
- [23] I. Ihrke and M. Magnor, "Image-Based Tomographic Reconstruction of Flames," *Proc. ACM SIGGRAPH/Eurographics Symp. Computer Animation*, pp. 367-375, 2004.
- [24] G. Wolberg, *Digital Image Warping*. IEEE CS Press, 1990.
- [25] A. Mohammad-Djafari, "Image Reconstruction of a Compact Object from a Few Number of Projections," *Proc. Int'l Conf. Signal and Image Processing*, pp. 325-329, 1996.
- [26] V. Elangovan and R.T. Whittaker, "From Sinograms to Surfaces: A Direct Approach to the Segmentation of Tomographic Data," *Proc. Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 213-223, 2001.
- [27] G.T. Herman and A. Kuba, Eds., *Discrete Tomography: Foundations, Algorithms, and Applications*. Birkhäuser, 1999.
- [28] F. Dacheux, K. Mueller, and A. Kaufman, "Volumetric Back-projection," *Proc. IEEE Symp. Volume Visualization*, pp. 109-117, 2000.
- [29] J.D. Bonet and P. Viola, "Roxels: Responsibility Weighted 3D Volume Reconstruction," *Proc. Int'l Conf. Computer Vision*, pp. 418-425, 1999.
- [30] R. Bhotika, D.J. Fleet, and K.N. Kutulakos, "A Probabilistic Theory of Occupancy and Emptiness," *Proc. European Conf. Computer Vision*, vol. 3, pp. 112-132, 2002.
- [31] A. Reche, I. Martin, and G. Drettakis, "Volumetric Reconstruction and Interactive Rendering of Trees from Photographs," *Proc. ACM SIGGRAPH*, pp. 720-727, 2004.
- [32] G.W. Faris and R.L. Byer, "Beam-Deflection Optical Tomography of a Flame," *Optics Letters*, vol. 12, no. 3, pp. 155-157, 1987.
- [33] V. Kwatra, A. Schödl, I. Essa, G. Turk, and A. Bobick, "Graphcut Textures: Image and Video Synthesis Using Graph Cuts," *Proc. ACM SIGGRAPH*, pp. 277-286, 2003.
- [34] S. Soatto, G. Doretto, and Y. Wu, "Dynamic Textures," *Proc. Int'l Conf. Computer Vision*, pp. 439-446, 2001.
- [35] G. Doretto and S. Soatto, "Editable Dynamic Textures," *Proc. Computer Vision and Pattern Recognition*, vol. 2, pp. 137-142, 2003.

- [36] K. Bhat, S. Seitz, J. Hodgins, and P. Khosala, "Flow-Based Video Synthesis and Editing," *Proc. ACM SIGGRAPH*, pp. 360-363, 2004.
- [37] J. Duderstadt and W. Martin, *Transport Theory*. John Wiley and Sons, 1979.
- [38] C. Silva and J. Santos-Victor, "Intrinsic Images for Dense Stereo Matching with Occlusions," *Proc. European Conf. Computer Vision*, vol. 1, pp. 100-114, 2000.
- [39] J. Nocedal and S.J. Wright, *Numerical Optimization*. Springer, 1999.
- [40] J.-Y. Bouguet, "Camera Calibration Toolbox for Matlab," Sept. 2002, http://www.vision.caltech.edu/bouguetj/calib_doc/.
- [41] S.W. Hasinoff, "Three-Dimensional Reconstruction of Fire from Images," master's thesis, Dept. of Computer Science, Univ. of Toronto, 2002.
- [42] I. Ihrke and M. Magnor, "GrOVis-Fire: A Multi-Video Sequence for Volumetric Reconstruction and Rendering Research," http://www.mpi-sb.mpg.de/~ihrke/Projects/Fire/GrOVis_Fire/, 2005.
- [43] S. Seitz and C. Dyer, "View Morphing," *Proc. ACM SIGGRAPH*, pp. 21-30, 1996.
- [44] Digital Air Incorporated, <http://www.virtualcamera.com>, 2003.
- [45] G.W. Stewart, *Introduction to Matrix Computations*. Academic Press, 1973.



Samuel W. Hasinoff received the BS degree in computer science from the University of British Columbia in 2000, and the MS degree in computer science from the University of Toronto in 2002. He is currently a PhD candidate at the University of Toronto. His research interests include computer vision and computer graphics, with a focus on three-dimensional reconstruction, image-based rendering, and computational photography. In 2006, he received an honorable mention for the Longuet-Higgins Best Paper Award at the European Conference on Computer Vision. He is a student member of the IEEE.



Kiriakos N. Kutulakos received the BA degree in computer science from the University of Crete, Greece, in 1988, and the MS and PhD degrees in computer science from the University of Wisconsin, Madison, in 1990 and 1994, respectively. Following his dissertation work, he joined the University of Rochester where he was a US NSF Postdoctoral Fellow and later an assistant professor until 2001. He is currently an associate professor of computer science at the University of Toronto. He won the Best Student Paper Award at CVPR '94, the Marr Prize in 1999, a Marr Prize Honorable Mention in 2005, and a Best Paper Honorable Mention at ECCV '06. He is the recipient of a CAREER award from the US National Science Foundation, a Premier's Research Excellence Award from the government of Ontario, and an Alfred P. Sloan Research Fellowship. He served as program cochair of CVPR 2003 and is currently an associate editor of the *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.