

Dynamic Shape and Appearance Modeling via Moving and Deforming Layers

Jeremy D. Jackson

Anthony J. Yezzi

Stefano Soatto

Abstract – We propose a model of the shape, motion and appearance of a scene, seen through a sequence of images, that captures occlusions, scene deformations, arbitrary viewpoint variations and changes in its radiance. This model is based on a collection of overlapping layers that can move and deform, each supporting an intensity function that can change over time. We discuss the generality and limitations of this model in relation to existing ones such as traditional optical flow or motion segmentation, layers, deformable templates and deformation. We then illustrate how this model can be used for inference of shape, motion, deformation and appearance of the scene from a collection of images. The layering structure allows for automatic inpainting of partially occluded regions. We illustrate the model on synthetic and real sequences where existing schemes fail, and show how suitable choices of constants in the model yield existing schemes, from optical flow to motion segmentation and inpainting.

1 Introduction

We are interested in modeling video sequences where changes occur over time due to viewer motion, motion or deformation of objects in the scene – including occlusions – and appearance variations due to the motion of objects relative to the light sources. A suitable model will trade off generality, by allowing variations of shape, motion and appearance, with tractability, by being amenable to inference and analysis. The goal of modeling is to support inference, and depending on the application one may be more interested in recovering shape (e.g. in shape analysis, classification, recognition, registration), or recovering motion (e.g. tracking, optical flow), or appearance variations (e.g. segmentation) including restoration (inpainting). Traditionally, the modeling task has been approached by making strict assumptions on some of the unknowns in order to recover the others, for instance the brightness-constancy assumption in optical flow, or the affine warping in shape analysis and registration. This is partly justified because in any image-formation model there is ambiguity between the three factors – shape, motion and appearance – and therefore the most general inference problem is ill-posed. In some applications, for instance video compression, the ambiguity is moot since all that matters is for the model to capture the sequence as faithfully and parsimoniously as possible. Nevertheless, since all three factors affect the generation of the image, a more germane approach would call for modeling all three jointly, then letting complexity dictate the responsibility of each factor, and the application dictate the choice of suitable regularizers to make the inference algorithms well posed. We therefore concentrate our attention on modeling, not on any particular application. So, this is not yet another paper on tracking, nor on motion segmentation, nor on optical flow, nor on shape registration. It is a little bit of all.

We propose a model of image formation that is general enough to capture shape, motion and appearance variations (Section 2), and simple enough to allow inference (Section 4). We want to be able to capture *occlusion phenomena*, hence our model will entail a notion of *hierarchy* or *layering*; we want to capture image variability due to arbitrary *changes in viewpoint* for non-planar objects, hence our model will entail *infinite-dimensional deformations* of the image domain. Such deformations can be due to changes in viewpoint for a rigid scene, or changes of shape of the scene seen from a static viewpoint, or any combination thereof. Our model will not attempt to resolve this ambiguity, since that requires higher-level knowledge. Furthermore, we want to capture large-scale motion of objects in the scene, as opposed to deformations, hence we will allow for a choice of a finite-dimensional group, e.g. Euclidean or affine, separate from infinite-dimensional deformations. Finally, we want to capture changes in appearance, hence scene radiance will be one of the unknowns in our model. Changes in radiance can come from changes in reflectance or changes in illumination, including changes in the mutual position between the light sources and the scene; again we do not attempt to resolve this ambiguity, since that requires higher-level knowledge. The image-formation model we propose is not the most general that one can conceive; far from it. Indeed, it is far less general than the simplest models considered acceptable in Computer Graphics, and we illustrate the lack of generality in Section 3. Nevertheless, it is more general than any other model used so far for motion analysis in Computer Vision, as we discuss also in Section 3, and is complex enough to be barely tractable with the analytical and computational tools at our disposal today. We pose the inference problem within a variational framework, involving partial differential equations, integrated numerically in the level set framework [15], although any other

computational scheme of choice would do, including stochastic gradients or Markov-chain Monte Carlo. The point of this paper is to propose a model and show that it can be inferred with at least one particular computational scheme, not to advocate a particular optimization technique.

1.1 Relation to existing work

This work relates to a wide body of literature in scene modeling, motion estimation, shape analysis, segmentation, and registration which cannot be properly reviewed in the limited space available. In Section 3 we illustrate the specific relationship between the model we propose and existing models. These include Layers [21, 12], which only model affine deformations of the domain and can therefore only capture planar scenes under small viewer motion or small aperture, and where there is no explicit spatial consistency within each layer and the appearance of each layer is fixed. As we will illustrate, our model allows deformations that can model arbitrary viewpoint variation, model layer deformation and enforce spatial coherence within each layer. One could think of our work as a generalization of existing work on Layers to arbitrary viewpoint changes, or arbitrary scene shape, and to changes in radiance (texture), all cast within a principled variational framework.

Our work relates to a plethora of variational algorithms for optical flow computation, for instance [18, 1, 8] and references therein, except that we partition the domain and allow arbitrary smooth deformations as well as changes in appearance (that would violate the brightness constancy constraints that most work on optical flow is based on, with a few exceptions, e.g. [10]). It also relates to various approaches to motion segmentation, where the domain is also partitioned and

allowed to move with a simple motion, e.g. Euclidean or affine, see for instance [7] and references therein. Such approaches do not allow deformations of the region boundaries, or changes in the intensity within each region. Furthermore, they realize a partition, rather than a hierarchy, of domain deformations, so our model can be thought of as motion segmentation with moving and deforming layers with changes in intensity and inpainting [3]. In this, our work relates to [19], except that we allow layers to overlap. So, our work can be thought of as a layered version of Deformation with changes in region intensities. Also relevant to our work is [17] where one distance function is registered to another using rigid and non-rigid transformations. Our work relates to deformable templates [9, 14], in the sense that each of our layers will be a deformable template. However, we do not know the shape and intensity profile of the template, so we estimate that along with the layering structure. A one layer version of our work is similar to [20] where the author describes energies on the manifold $G \times M$ where $g \in G$ is a group action (possibly a C^∞ diffeomorphism or an affine transformation) and M is a manifold consisting of a collection of landmark points or images). For the example of G being the set of C^∞ diffeomorphisms and M being the set of images, the geodesic between two points $(g_1, m_1), (g_2, m_2) \in G \times M$ describes metamorphoses from one (group, image) pair to another. Our work is also related to active appearance models [6, 2], in that we seek the same goal, although rather than imposing regularization of shape and appearance by projection onto suitably inferred linear subspaces we employ generic regularizers. One can therefore think of our work as a generalization of active appearance models to smooth shape and intensity deformations, cast in a variational framework. Of course this work relates more generically to active contours, e.g. [4, 13, 5, 16] and references therein. In the next section we

introduce our model, and in Section 4 we illustrate our approach to infer its (infinite-dimensional) constitutive elements.

2 Modeling

We represent a scene as a collection of L overlapping *layers*. Each layer, labeled by an index $k = 1, \dots, L$, is a function that has associated with it a domain, or *shape* $\Omega^k \subset \mathbb{R}^2$, and a range, or *radiance* $\rho^k : \Omega^k \rightarrow \mathbb{R}^+$. Layer boundaries model the occlusion process, and each layer k undergoes a *motion*, described by a (finite-dimensional) group action g^k , for instance $g^k \in \mathbb{S}E(2)$ or $\mathbb{A}(2)$, and a *deformation*, or *warping*, described by a diffeomorphism $w^k : \Omega^k \rightarrow \mathbb{R}^2$, in order to generate an image I at a given time t . The warping models changes of viewpoint for non-planar scenes, or actual changes in the shape of objects in the scene. Since each image is obtained from the given scene after a different motion and deformation, we index each of **the image’s corresponding variables** by t : g_t^k , w_t^k , and I_t . Finally, since layers occlude each other, there is a natural ordering in k which, without loss of generality, we will assume to coincide with the integers: Layer $k = 1$ is occluded by layer $k = 2$ and so on. But since this occlusion model could change, say layer $k = 2$ goes behind layer $k = 3$ and then later layer $k = 2$ is in front of layer $k = 3$, there is a function $l = \max\{k \mid x \in \Omega^k\}$. This function l indicates the layer that will contribute to the intensity at a pixel in a given image which is the frontmost **layer** that intersects the warped domain. For simplicity we assume that $\Omega^0 = \mathbb{R}^2$ (the backmost layer, or “the background”). With this notation, the model of how the value of the generic image $I_t : \Omega^0 \rightarrow \mathbb{R}^+$ at the location $x \in \Omega^0 \subset \mathbb{R}^2$ is generated can be summarized as $I_t(g_t^l \circ w_t^l(x)) = \rho^l(x)$, with $x \in \Omega^l$, $l = \max\{k \mid x \in \Omega^k\}$. To

simplify the notation, we call $x_t^l \doteq g_t^l \circ w_t^l(x)$, which sometimes we indicate, for simplicity, as x_t ,

so that

$$\boxed{\begin{cases} I_t(x_t^l) = \rho^l(x), & x \in \Omega^l \\ x_t^l = g_t^l \circ w_t^l(x), & l = \max\{k \mid x \in \Omega^k\}. \end{cases}} \quad (1)$$

Our goal in this work is to infer the radiance family $\{\rho^k\}_{k=1,\dots,L}$, the shape family $\{\Omega^k\}_{k=1,\dots,L}$, the motions $\{g_t^k\}_{k=1,\dots,L;t=1,\dots,N}$ and the deformations $\{w_t^k\}_{k=1,\dots,L;t=1,\dots,N}$ that minimize the discrepancy of the measured images from the ideal model (1), subject to generic regularity constraints.

Such a discrepancy is measured by a cost functional $\phi(\Omega^k, \rho^k, w_t^k, g_t^k)$ to be minimized

$$\begin{aligned} \phi \doteq & \sum_{t=1}^N \int_{\Omega^0} \left(I_t(x_t) - \rho^l(w_t^{l-1} \circ g_t^{l-1}(x_t)) \right)^2 dx_t + \zeta \sum_{k=1}^L \int_{\partial\Omega^k} ds \\ & + \lambda \sum_{k=1}^L \int_{\Omega^k} \|\nabla \rho^k(x)\|^2 dx + \mu \sum_{k,t=1}^{L,N} \int_{\Omega^l} r(w_t^k(x)) dx \end{aligned} \quad (2)$$

subject to $l = \max\{k \mid x \in \Omega^k\}$.

Here r is a regularizing functional, for instance $r(w) \doteq |\dot{w}| + \frac{1}{|\dot{w}|}$ where $|\dot{w}|$ is the determinant of the jacobian of w . Since it is desirable to keep w to be a one-to-one function this regularizer r keeps $|\dot{w}|$ close to one. If $|\dot{w}|$ deviates from 1 then either of terms $|\dot{w}|$ and $\frac{1}{|\dot{w}|}$ gets bigger. λ , μ , and ζ are positive constants. Note that l is a function, specifically $l : \Omega^0 \rightarrow \mathbb{Z}^+$. We have chosen the two-norm for the data-dependent term and the regularizer for simplicity, but other choices would of course do as well.

3 Generality of the model

It can be easily shown that eq. (1) models images of 3-D scenes with piecewise smooth geometry exhibiting Lambertian reflection with piecewise smooth albedo¹ viewed under diffuse illumination from an arbitrarily changing viewpoint. It does not capture global or indirect illumination effects, such as cast shadows or inter-reflections, complex reflectance, such as specularities, anisotropies or sub-surface scattering. These are treated as modeling errors and are responsible for the discrepancy between the model and the images, which is measured by ϕ in eq. (2). We lump these discrepancies together with sensor errors and improperly call them “noise.” Although far from general, (1) is nevertheless a more ambitious model than has ever been used in the context of motion estimation and tracking. In fact, many existing models are special cases of (1).

We start by showing how the model includes traditional **optical flow** as a special case. In particular, if we assume a single layer to represent the whole image domain (i.e. $L = 0$), a trivial group action (i.e. $g = Id$) and no regularity in the modeled radiance $\rho = \rho^0$ (i.e. $\lambda = 0$) then the resulting minimization problem includes only the radiance ρ and the warps $w_1 = w_1^0$ and $w_2 = w_2^0$ as unknowns (we consider the case of just two images I_1 and I_2 for now). We are therefore left with the much simpler energy

$$\phi(\rho, w_1, w_2) = \sum_{t=1}^2 \int_{\Omega^0} (I_t(x_t) - \rho(w_t^{-1}(x_t)))^2 dx_t + \mu \sum_{t=1}^2 \int_{\Omega^0} r(w_t(x)) dx. \quad (3)$$

If our goal is just to find the warp $w = w_2 \circ w_1^{-1}$ that registers I_1 to I_2 (through the common

¹The model can be further generalized by allowing ρ^l to be vector-valued to capture a set of radiance statistics such as the coefficients of a filter bank or other texture descriptors, but this is beyond the scope of this paper.

radiance model ρ), then we may further simplify things by setting $w_1 = Id$ and $w_2 = w$, thereby eliminating yet another unknown and yielding (approximately, as we omit the Jacobian of ω)

$$\phi(\rho, w) = \int_{\Omega^0} (I_1(x) - \rho(x))^2 + (I_2(w(x)) - \rho(x))^2 dx + \mu \int_{\Omega^0} r(w(x)) dx. \quad (4)$$

Since we have omitted the smoothness penalty on ρ , it is straightforward to show for a given choice of w that (4) is minimized by setting $\rho(x)$ to the mean of $I_1(x)$ and $I_2(w(x))$. Thus, in this special case (no smoothness on ρ) we may replace the joint optimization in (4) with a direct optimization of w through this substitution of ρ . The resulting energy

$$\phi(w) = \frac{1}{2} \int_{\Omega^0} (I_1(x) - I_2(w(x)))^2 dx + \mu \int_{\Omega^0} r(w(x)) dx, \quad (5)$$

depending upon the choice of the regularizer r (note that r typically depends on the derivatives of w rather than its direct values), corresponds to either the classical optical flow in [11] or to one of its many variants.

Our model has the advantage of not enforcing global regularization (regularization is imposed within layers, but not across layers), of not comparing images to each other, but to an underlying model (this carries significant advantages when it comes to robustness to noise, as we illustrate with experiments), and of having an explicit model of the appearance of the scene, which allows “inpainting” individual layers while preserving their motion boundaries.

Choosing $L = 1$, $w = Id$, $\lambda = 0$, $\mu = 0$ yields **motion segmentation**, that has also been addressed by many, see for instance [7] and references therein for the case of affine motion $g \in \mathbb{A}(2)$.

In motion segmentation one partitions the domain into a number of individually moving segments, each of which is assumed to move with a constant (finite-dimensional) motion. Like in optical flow, there is no model of appearance, and the data-dependent term consists of the brightness constancy constraint which forces direct image-to-image comparison:

$$\phi(g_t, \Omega_t) = \sum_t \int_{\Omega_t} (I_t(x) - I_{t+1}(g_t x))^2 dx + \nu \mathcal{L}(\Omega) \quad (6)$$

where \mathcal{L} denotes the length of the boundary of the region Ω_t . Note that, in this case, we have allowed Ω to be one of the unknowns since w_t is no longer part of the inference, although one could easily define $\Omega_t \doteq w_t(\Omega_0)$, as we have discussed in the previous section.

Choosing $L = 1$, $\rho = \text{const}$, $r(w) = \|w\|$ yields a model called **Deformation** in [19], and has also been extended to grayscale images $L = 1$, $r(w) = \|w\|$. Our work is the natural extension of Deformation to layers.

Choosing $L > 1$, $w = Id$, Ω^k unconstrained and $g \in \mathbb{A}(2)$ would yield a variational version of the **Layers** model [21], that to the best of our knowledge has never been attempted. Note that this is different than simpler variational multi-phase motion segmentation, since in that case the motion of a phase affects the shape of neighboring phases, whereas in the model (1) layers can overlap without distorting underlying domains. One can think of the Layer model as a multi-phase motion segmentation with inpainting [3] of occluded layers and shape constraints.

The model also relates to **deformable templates**, where $\rho = \text{const}$ in the traditional model [9] and $\rho = \text{smooth}$ in the more general version [14]. Another relevant approach is **Active Appearance Models** where the regions, warping and radiances are modeled as points in a linear space.

$$w_t^k(x) = w_0^k(x) + W^k(x)s_t \quad (7)$$

where $w_0 : \Omega^k \rightarrow \mathbb{R}^2$ and $W^k : \Omega^k \rightarrow \mathbb{R}^n$ denotes a set of basis functions or principal components, and $s_t \in \mathbb{R}^n$, $t = 1, \dots, N$ is a vector of shape coefficients. Similarly,

$$\rho_t^k(x) = \rho_0^k(x) + P^k(x)\alpha_t \quad (8)$$

where $\rho_0^k : \Omega^k \rightarrow \mathbb{R}$ and $P^k(x) : \Omega^k \rightarrow \mathbb{R}^n$ is a vector of principal components, and $\alpha_t \in \mathbb{R}^n$ a vector of appearance parameters. Note that the functions P^k and W^k have to satisfy orthogonality constraints, and these have to be enforced during the inference of the bases. The model (1) does not impose such restrictions, and render the problem well-posed by generic regularization instead.

Finally, by virtue of the regularization imposed on ρ , our scheme relates to **image inpainting**, except that we perform inpainting *both* by layer transfer from multiple images and by regularization. The advantage of our method is that it can exploit whatever information is there: If multiple views are available, their contribution is weighted relative to the harmonic interpolation term. If only one image is available, then intensity regularization dictates the filling process.

4 Inference

Minimizing the cost functional in (2) is a tall order. It depends upon each domain Ω^k and its boundary (a closed planar contour), its deformation (a flow of planar diffeomorphisms) w_t^k , the radiance (a piecewise smooth function) ρ^k , all of which are infinite-dimensional unknowns. In

addition, it depends on a group action per layer per instant, g_t^k , and on the occlusion model, which is represented by the discrete-valued function $l(x) = \max\{k \mid x \in \Omega^k\}$, and all of this for each layer $k = 1, \dots, L$.

The first simplification is to notice that, as long as each layer is a compact region bounded by a simple smooth curve, there is no loss of generality in assuming that Ω^k are fixed. This is because each diffeomorphism w_t^k will act transitively on it. Therefore, we assume that each region Ω^k is a circle in some of the examples. While there is no loss of generality, there is a loss of energy, in that if we were allowed to also optimize with respect to the initial regions we would be able to reach each deforming layer faster and with less energy. This, however, does not enhance the generality of the model, hence we will forgo it for some examples (see Figure ?? for an illustration of this effect).

Apart from this simplification, we proceed by minimizing the functional (2) using simultaneous gradient flows with respect to the groups (motion), the radiances (appearance) and the diffeomorphisms (deformation). The detailed evolution equations are a bit complicated depending upon the number of layers and the occlusion structure between layers. To help avoid excessive subscripting and superscripting and multiple-case definitions according to occlusion relationships, we will outline some of the key properties of the various gradient terms for the case of a background layer Ω^0 , a single image I , and a single foreground layer Ω^1 . We will also, to help keep the illustration simple, assume that the group action g^0 and the warp w^0 for the background layer are simply the identity transforms. This is the simplest possible scenario that will allow us to still show the key properties of the gradient flows.

Let $\hat{x} = g^1(w^1(x))$ and $\hat{\Omega}^1 = g^1(w^1(\Omega^1))$. With this notation, we may write the image-dependent terms in our energy functional as follows.

$$E = \int_{\hat{\Omega}^1} (I(\hat{x}) - \rho^1(x))^2 d\hat{x} + \int_{\Omega^0 \setminus \hat{\Omega}^1} (I(x) - \rho^0(x))^2 dx \quad (9)$$

If g denotes any single parameter (e.g. horizontal translation) of the group g^1 , then differentiating yields

$$\begin{aligned} \frac{\partial E}{\partial g} = & \int_{\partial \hat{\Omega}^1} \left\langle \frac{\partial \hat{x}}{\partial g}, \hat{N} \right\rangle \left((I(\hat{x}) - \rho^1(x))^2 - (I(\hat{x}) - \rho^0(\hat{x}))^2 \right) d\hat{s} + \\ & + 2 \int_{\hat{\Omega}^1} (I(\hat{x}) - \rho^1(x)) \left\langle \nabla \rho(x), \text{inv} \left[(w^1)' \right] \frac{\partial}{\partial g} \text{inv}[g^1](\hat{x}) \right\rangle d\hat{x} \end{aligned} \quad (10)$$

where \hat{N} and $d\hat{s}$ denote the outward unit normal and the arclength element of $\partial \hat{\Omega}^1$ respectively. We are able to note two things. First, the update equations for the group involve measurements both along the boundary of its corresponding layer (first integral) as well as measurements within the layer's interior (second integral). Notice that this latter integral vanishes if a constant radiance ρ is utilized for the layer. We also see that it is not necessary to differentiate the image data I . Derivatives land on the estimated smooth radiance ρ instead, which is a significant computational perk of our model that results in considerable robustness to image noise.

A similar gradient structure arises for the case of the infinite dimensional warp w (boundary-based terms and region-based terms for each layer are similar to previous integrals). However, additional terms arise in the gradient flow equations for w depending upon the choice of regularization terms in the energy functional (smoothness penalties, magnitude penalties, etc.).

The curve evolution is also similar to the boundary-based term for the evolution of g :

$$\frac{\partial C}{\partial t} = -\left(\left(I(\hat{x}) - \rho^1(x)\right)^2 - \left(I(\hat{x}) - \rho^0(\hat{x})\right)^2\right) \hat{N} \quad (11)$$

Finally, the optimality conditions for the smooth radiance functions ρ^0 and ρ^1 are given by the following Poisson-type equations.

$$\Delta \rho^1(x) = \lambda(\rho^1(x) - I(\hat{x})), \quad x \in \Omega^1 \quad (12)$$

$$\Delta \rho^0(x) = \begin{cases} 0, & x \in \hat{\Omega}^1 \\ \lambda(\rho^0(x) - I(x)), & x \in \Omega^0 \setminus \hat{\Omega}^1 \end{cases} \quad (13)$$

Notice that the background radiance ρ^0 is “inpainted” in regions occluded by the foreground layer Ω^1 by harmonic interpolation from the boundary of $\hat{\Omega}^1$, since ρ satisfies Laplace’s equation $\Delta \rho^0 = 0$. Once all the terms are put together we can generate a gradient flow that simultaneously evolves all layer assignments, boundaries and intensities. In the next section we illustrate some of the features of the model and the resulting optimization, as it compares with existing schemes.

5 Experiments

In the first experiment we illustrate the capability of our model to track deforming layers. In Figure 1 we show three sequences of an image where a deflating balloon is undergoing a rather erratic motion while deforming from an initial waterdrop shape to a circular one, finally to a spermatozoidal shape. On the top row of Figure 1 we show the layer boundaries for a model that only

allows for rigid deformations of the initial contour (a circle) using a single scaling term. This is essentially a variational implementation of the model of [21]. As it can be seen, it captures the gross motion of the balloon, but it cannot capture the subtler shape variations. The second row shows the same three sample images with the boundary of the first layer superimposed, where the layer is allowed to deform according to the model we have introduced. The data fidelity term used is a Mumford-Shah term so the radiances representing each layer are smooth functions. As it can be seen, the layer changes shape to adapt to the deforming balloon, all while capturing its rather erratic motion. The average RMS error per image for the affine layer model is 30.87, whereas the residual for the case of the deforming layers is 5.51. More importantly, the phenomenology of the scene, visible in the figure, has been correctly captured.




Figure 1: *Tracking a Balloon: Three sample views are shown from a sequence of a deflating balloon moving with an erratic motion while changing its shape from a drop-like shape to a circle. In the top row we show the boundary of the first layer as estimated by a rigid layer model with a single scaling term that does not allow for layer deformation, akin to a variational implementation of traditional layer models. As it can be seen, the model tracks the motion of the layer, but it fails to capture its deformation. On the bottom row we show the same three images with the first layer superimposed, where the layer is allowed to both move (rigidly) and deform (diffeomorphically), yielding 82% lower RMS residual error, and capturing the subtler shape variations.*

In the next experiment we illustrate all the features of our model by showing how it recovers the background behind partially occluded layers while recovering their motion and deformation.

In Figure 2 we show a few samples from a dataset where the silhouette of a moving hand forms a victory sign while moving the relative position between the fingers. The background, which is partially occluded, is a spiral. Here we use the average shape as the initial shape of the foreground layer to find its affine motion, and then the diffeomorphic warp w_i . Again we assume smooth radiance within each layer, so when we recover the background layer we show a slightly smoothed version of the spiral (of course we could further segment the black spiral from the background and thus obtain sharp boundaries, but this is standard and would not help us illustrate the feature of the model, therefore we do not illustrate it here.)




Figure 2: *Victory sign, with deforming hand, moving in front of a partially occluded background portraying a spiral. The goal here is to recover the radiance of each layer (the spiral in the background and the constant black intensity of the hand), as well as the motion and deformation of the foreground layer. Note that current layer models based only on affine motion would fail to capture the phenomenology of this scene by over-segmenting the region into three regions, each moving with independent affine motion. Our model captures the overall motion of the layer with an affine group, and then the relative motion between the fingers as a deformation, as we illustrate in the next figure 3.*

In Figure 3 we illustrate the results of this experiments, arranged to summarize the modeling process. On the top row we show the recovered layers. Since we are assuming a smooth radiance within each layer, we can only recover a smoothed version of the spiral. These layers are deformed according to a diffeomorphism, one per layer, defined on the domain of the layer (second row) and

then moved according to an affine motion. The third row shows the image generated by the model, which can therefore be thought of as a generative (although deterministic) model since it performs comparison at the image level, not via some intermediate representation. The corresponding images are displayed in the last row, with the layers superimposed for comparison.

The first set of experiments, using standard sequences used for optical flow analysis, is designed to illustrate the difference between our model and standard optical flow. A representative set of results of the motion field estimated by optical flow (left) and our model (right) is reported in Figure 5. Our model does not rely on global regularization, but only regularization within each layer segmented in Figure 4. Therefore, the boundaries of the motion field are better resolved.

Naturally, our model is a superset of those commonly used for optical flow computation. We illustrate this point by reducing the weight of the smoothness term for ρ in Figure 6, which yields results closer to standard optical flow. In comparison to the ground truth vector field, the vector field given by optical flow has an average angular error of 8.12. Deformation with a smoothness weight λ of 200 gives an average angular error of 9.99. Reducing the smoothness weight λ to 20 gives an average angular error of 8.11 which is closer to the result of optical flow.

A beneficial side-effect of having an explicit model of the scene, simple as it is (a regular irradiance pattern, with smoothness controlled by λ), is the possibility of comparing individual images to a (noiseless) model, rather than comparing noisy images to each other. The effects are visible in Figure 8, where the flow field obtained with our model on artificially corrupted sequences (Figure 7) is far closer to the cleaner version of the sequence than using standard optical flow.

The comparison with optical flow illustrates the necessity for partitioning the domain into in-

dependently moving objects. This is a motion segmentation task. Therefore, here we compare our model with more standard ones that partition the flow into affine segments, while still relying on the brightness constancy constraint and without an explicit model of the appearance of the scene. Such models can be obtained simply by increasing the regularization of the layer deformation (i.e. the entire layer moves with the same finite-dimensional motion: translational, Euclidean or affine). Figure 9 illustrates this effect.

Note that our model, by virtue of having an explicit representation of the appearance of each layer, can automatically fill in the appearance of underlying layers, as we illustrate in Figure 10.

In Figure 11 we illustrate inpainting using our model. In this example there is some camera jitter, which makes it so the whiteboards in the two images are not quite lined up. Also there has been some corruption of the images which is modeled as the foreground layer that is moved around via an affine group. The whiteboard (background layer) is recovered with its own affine registration and the inpainted whiteboard is shown.

The conclusion we would like to draw from these experiments is that our model, being a superset of existing schemes (optical flow, motion segmentation, deformation, inpainting), allows the user to apply existing algorithms simply by proper choice of constants. Naturally the price to pay for such flexibility and for the added power stemming from a richer model is computational complexity. However, all the experiments we have shown have been run on a pentium M 2GHz PC and takes five minutes per 1000 iterations.

6 Discussion

We have presented a generative model of the appearance (piecewise smooth albedo), motion (affine transformation) and deformation (diffeomorphism) of a sequence of images that exhibit occlusions. We have used this model as a basis for a variational optimization algorithm that simultaneously tracks the motion of a number of overlapping layers, estimates their deformation, and estimates the albedo of each layer, including portions that were partially occluded. Where no information is available, the layers are implicitly inpainted by their regularizers.

This model generalizes existing layer models to the case of deforming layers. Alternatively, one can think of our algorithm as a layered version of deformable tracking algorithms, or as a generalized version of optical flow or motion segmentation where multiple layers are allowed to occlude each other without disturbing the estimate of adjacent and occluded ones.

Our numerical implementation of the flow-based algorithm uses level set methods, and is realized without taking derivatives of the image, a feature that yields significant robustness when compared with boundary-based approach to estimating optical flow. We have illustrated our approach on simple but representative sequences where existing methods fail to capture the phenomenology of the scene by either over-segmenting it, or by failing to capture its deformation while only matching its affine motion.

References

- [1] L. Alvarez, J. Weickert, and J. Sanchez. A scale-space approach to nonlocal optical flow calculations. In *In ScaleSpace '99*, pages 235–246, 1999.
- [2] S. Baker, I. Matthews, and J. Schneider. Image coding with active appearance models. Technical report, Carnegie Mellon University, The Robotics Institute, 2003.
- [3] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 417–424. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
- [4] A. Blake and M. Isard. *Active contours*. Springer Verlag, 1998.
- [5] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int. J. of Computer Vision*, 22(1):61–79, 1997.
- [6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *Proc. of the Eur. Conf. on Comp. Vis.*, pages 484–496, 1998.
- [7] D. Cremers. Multiphase levelset framework for variational motion segmentation. In *Intl. Conf. on Scale-Space Theories in Computer Vision*, pages 599–614, June 2003.
- [8] R. Deriche, P. Kornprobst, and G. Aubert. Optical flow estimation while preserving its discontinuities: a variational approach. In *Proc. of ACCV*, volume 2, pages 290–295, 1995.
- [9] U. Grenander. *General Pattern Theory*. Oxford University Press, 1993.

- [10] H. W. Haussecker and D. J. Fleet. Computing optical flow with physical models of brightness variation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(6):661–673, 2001.
- [11] B. K. P. Horn and B. G. Schunk. Determining optical flow. *Artificial Intell.*, 17:185–203, 1981.
- [12] S. Hsu, P. Anandan, and S. Peleg. Accurate computation of optical flow by using layered motion representations. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recogn.*, pages 1621–1626, 1992.
- [13] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision*, pages 810–815, Washington, DC, USA, 1995. IEEE Computer Society.
- [14] M. I. Miller and L. Younes. Group action, diffeomorphism and matching: a general framework. In *Proc. of SCTV*, 1999.
- [15] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi equations. *J. of Comp. Physics*, 79:12–49, 1988.
- [16] N. Paragios and R. Deriche. Geodesic active contours and level sets for the detection and tracking of moving objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3):266–280, 2000.
- [17] Nikos Paragios, Mikael Rousson, and Visvanathan Ramesh. Non-rigid registration using distance functions. *Comput. Vis. Image Underst.*, 89(2-3):142–165, 2003.

- [18] C. Schnörr. Computation of discontinuous optical flow by domain decomposition and shape optimization. *Int. J. of Computer Vision*, 8(2):153–165, 1992.
- [19] S. Soatto and A. Yezzi. Deformotion: deforming motion, shape average and the joint segmentation and registration of images. In *Proc. of the Eur. Conf. on Computer Vision (ECCV)*, volume 3, pages 32–47, 2002.
- [20] A. Trounev and L. Younes. Metamorphoses through lie group action. *Foundations of Computational Mathematics*, 5(2):173–198, 2005.
- [21] J. Wang and E. Adelson. Representing moving images with layers. In *IEEE Trans. on Image Processing*, volume 3(5), pages 625–638, 1994.




Figure 3: *Multiple Layers Mapping onto Multiple Images: The inference process returns an estimate of the albedo in each layer (top). Since we are assuming smooth albedo, the spiral is smoothed. The deformation of each layer is estimated (second row) together with its affine motion, to yield an approximation of the image (third row). This is used for comparison with the measured images (bottom row) that drives the optimization scheme.*




Figure 4: Rotating Sphere and segmentation obtained using deformation




Figure 5: Optical Flow vs. Ground Truth vs. Deformation: Standard optical flow (left) imposes global regularization, which results in errors at the boundary (the vector field is more spread out than the model proposed, on the right). The ground truth is in the middle. The average angular errors for optical flow and deformation are 11.49° and 6.31° respectively. The standard deviation for the angular errors are 1.37° and 1.44° . The parameters and regularization constants used were $dt= 0.2$, $iterations= 10000$, $\alpha = 10$ (data fidelity), $\mu = 0.5$, (smoothness of w), $\lambda = 5$ (smoothness of ρ).




Figure 6: Optical flow (left) can be obtained from the general model (right) by allowing $\lambda \rightarrow 0$. Compare the results with parameters $dt= 0.028$, $iterations= 71000$, $\alpha = 20$, $\mu = 0.55$, $\lambda = 20$ on the bottom row with 200 on the middle row. Note that the two models (left and right) are closer on the bottom row. In comparison to the ground truth vector field, the vector field given by optical flow has an average angular error of 8.12. Deformation with a smoothness weight of 200 gives an average angular error of 9.99. Reducing the smoothness weight to 20 gives an average angular error of 8.11 which is closer to the result of optical flow.




Figure 7: Yosemite Sequence with added Gaussian noise of zero mean and variance of 0.05 and ground truth in the middle.




Figure 8: Flow Fields for Optical Flow, Ground Truth, and Deformation: One of the beneficial effects of having a model of the appearance of the scene is better robustness to noise. In standard optical flow images are compared to each other directly, resulting in fragility to independent noise sources. We show results with added Gaussian noise with zero mean and variance of 0.05. The benefit or robustness to noise is offset by the added complexity of having to estimate a smooth function to approximate scene radiance. (Parameters used for deformation: $dt= 0.2$, iterations= 10000, $\alpha = 2$, $\mu = 1.0$, $\lambda = 4$.) (Parameters used for optical flow: $dt= 0.083$, iterations= 24000, $\alpha = 2$, $\mu = 2$.) The parameters used provided the best solution for each method. The average angular errors for optical flow and deformation are 26.24 and 20.24 respectively.




Figure 9: The model proposed can be used to perform motion segmentation by increasing the regularization μ of the domain deformation for each layer (parameters used: $dt = \frac{0.2}{4000}$, iterations = 20000, $\alpha = 20$, $\mu = 4000$, $\lambda = 2000$).



Figure 10: Our model yields “inpainted” layers. The top row shows the boundaries of layers, the middle row the reconstructed appearance of the layers (ρ) and the bottom row the warpings (w). Parameters used: $dt= 0.2$, $iterations= 2000$, $\alpha = 20$, $\mu = 1.0$, $\lambda = 0.8$, $\zeta = 3.0$ (arclength weight) .




Figure 11: Image Inpainting with our model. First two images: corrupted images of a teacher's whiteboard with some camera jitter, Last image: Image inpainting result.