Fusion of Physically-Based Registration and Deformation Modeling for Nonrigid Motion Analysis

Leonid V. Tsap, Dmitry B. Goldgof, *Senior Member, IEEE*, and Sudeep Sarkar, *Member, IEEE*

Abstract—In our previous work, we used finite element models to determine nonrigid motion parameters and recover unknown local properties of objects given correspondence data recovered with snakes or other tracking models. In this paper, we present a novel multiscale approach to recovery of nonrigid motion from sequences of registered intensity and range images. The main idea of our approach is that a finite element (FEM) model incorporating material properties of the object can naturally handle both registration and deformation modeling using a single model-driving strategy. The method includes a multiscale iterative algorithm based on analysis of the undirected Hausdorff distance to recover correspondences. The method is evaluated with respect to speed and accuracy. Noise sensitivity issues are addressed. Advantages of the proposed approach are demonstrated using man-made elastic materials and human skin motion. Experiments with regular grid features are used for performance comparison with a conventional approach (separate snakes and FEM models). It is shown, however, that the new method does not require a sampling/correspondence template and can adapt the model to available object features. Usefulness of the method is presented not only in the context of tracking and motion analysis, but also for a burn scar detection application.

Index Terms—Corresponding points, deformable models, Hausdorff distance, multiscale methods, nonrigid motion analysis, physically-based vision.

I. INTRODUCTION

NONRIGID motion analysis is comprised of a large body of research directions and approaches. This is justified considering the variety and complexity of motion. Successful motion recovery is motivated by important applications such as stereo processing, gesture and face recognition for human–computer interaction, and medical imaging. Nonrigid motion analysis includes establishing point correspondences necessary for tracking, estimating motion and, finally, understanding the reasons why motion occurred in the observed way and not in any other possible way. Two of the major classes of techniques for nonrigid motion analysis include snakes and finite element models.

Snakes, or active contours (first introduced by Kass *et al.* [1]), are energy-minimizing splines which can find and reliably track salient image contours. Snakes are guided by both internal and external energies. The internal energy is the part that depends on intrinsic properties of the snake, such as its length or curvature, and on the image forces that drive the snake toward the desired image contours. The external energy depends on factors such as image structure, and particular constraints the user has imposed. Snakes have specific properties used for very precise tracking, yet not related to material properties or the internal structure of the object. Snakes continue to be the focus of many recent models [2]–[6]. McEachen, II and Duncan [2] tracked feature points over an entire cardiac cycle. Androutsos *et al.* [4] added image gradient direction to the energy functional and applied it to flow trace images. Yezzi *et al.* [7] unified the curve evolution approaches for active contours and the established energy formulation. Gunn and Nixon [5] used two contours to search the image space from both inside and outside of the target feature. Tagare [6] proposed a formulation that achieves reduction in the search space by precomputing orthogonal curves to deform the template. Peterfreund [8] modified the active contour model by applying velocity control for real-time tracking. Chandran and Potty [9] developed a strategy to avoid local minimas as a dynamic programming solution for snake energy minimization. Amini *et al.* [10] applied coupled B-spline snake grids to magnetic resonance images and validated results with a 3-D cardiac motion model. Recent developments in deformable model techniques are summarized in [11].

Unlike snakes, finite element models usually include material properties of the object and precise understanding of its structure, but have no inherent way to track objects' features. Finite element models include knowledge of forces applied to the object and can calculate response to these forces.

A variety of finite element models were proposed in the context of vision research. Sclaroff and Pentland [12] used FEM to obtain a parametric description of nonrigid motion in terms of its similarity to known extremal views. Young and Axel [13] built a finite element model of the left ventricle to fit material points tracked in biplanar views. Metaxas and Koh [14] used local adaptive finite elements to represent three-dimensional (3-D) shapes efficiently. Cohen and Cohen [15] presented a finite element method to track a series of two-dimensional (2-D) slices of heart ventricles and to make a 3-D reconstruction of the inside surface of the ventricles. A finite element model that learns the correct physical model of human lips by training from real data

L. V. Tsap is with the Center for Lawrence Livermore National Laboratory, Applied Scientific Computing, University of California, Livermore, CA 94551 USA (e-mail: tsap@llnl.gov; http://marathon.csee.usf.edu/~tsap; http://www.llnl.gov/CASC/people/tsap).

D. B. Goldgof and S. Sarkar are with the Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620 USA (e-mail: goldgof@csee.usf.edu; sarkar@csee.usf.edu; http://marathon.csee.usf.edu/~goldgof/; http://marathon.csee.usf.edu/~sarkar/sarkar.html.).

was proposed by Basu and Pentland [16]. Martin *et al.* [17] employed finite element computation of analytic modes describing shape variation of structures within the human brain. Tsap *et al.* [18], [19] used nonlinear finite element models to recover motion, material properties of nonrigid objects and resulting strain distributions, and applied it to burn scar assessment and repetitive stress injury analysis. Feature points used in finite element analysis were tracked with snakes.

Therefore, to accomplish defined motion analysis goals, it is necessary to achieve tracking similar or even better in quality than can be done with snakes, and also examine additional aspects not readily obvious from images. These features can add to the knowledge of the object (material properties, applied forces and detailed structure of the object) when using a more natural (with respect to objects' properties) physically-based model, such as a finite element model that can explain the deformation process.

Most closely related works such as [20], [18], and [21] accomplished these goals using two separate models: snakes to find tag positions in images and finite element models [18], [20] (or similar physically-based models [21]) to compute deformation parameters (such as displacements) and strain distributions. Although FEM models produced precise solutions in terms of both displacements and strains, they utilized only information at points where tag lines intersected.

Other related approaches included following works. Bajcsy and Kovacic [22] proposed a multiresolution elastic matching algorithm for medical applications. Contours of the brain atlas and the CT (Computerized Tomography) image of the brain from three different scales were matched in a coarse-to-fine manner. The finest level solution was then used to incrementally deform the model using elastic constraint equations. Christensen *et al.* [23] estimated probabilistic viscous fluid transformations to templates for neuroanatomy registration with individual studies. More recently, Papademetris *et al.* [24] estimated soft tissue deformations from sequences of left ventricular MR images using a linear elastic model. The main difference in our approach is having a true-physics-based model, integrating or allowing us to recover true material properties of an object.

Recently, a number of hybrid approaches were developed. A framework for combining complementary techniques (registration and deformable models) was proposed by Montagnat and Delingette [25]. Another hybrid solution (2-D) based on modal analysis, employed by Tao and Huang [26], blended finite-element-computed modes with template matching. Deformable models with parameter functions capable of adequately addressing local shape variations were proposed by Park *et al.* [27] and O'Donnell *et al.* [28]. A shape modeling approach that used multiresilution transformations from local to global models was introduced by Vemuri and Radisavljevic [29]. Again, the originality of the proposed combination includes integration of true elastic properties in the model.

The approach proposed in this paper encompasses advantages of both techniques in a single model-driving strategy. Both detection/tracking and accurate object model estimation are merged to provide a more comprehensive basis for nonrigid motion analysis. We propose that FEM model can naturally handle registration and modeling. Our model combines two types of deformable models. Control points used for tracking are also a part of a finite element model (that includes an available knowledge of material properties). Therefore, not only snake intersection points (as it often occurred), but also additional tracked points are included in the model. Furthermore, the method includes a multiscale scheme based on evaluation of the undirected Hausdorff distance to speed up the process of matching features between two frames since large deformations are considered. This criteria is similar to work by Huttenlocher [30] and Rucklidge [31]. However, in their approach, a set of image pixels in next frame formed a new model. We perform actual model transformation that simulates nonrigid motion of the object.

A. Overview

Major contributions of this work can be described along the following directions:

- combination of physically-based registration and deformation modeling;
- integration of material, geometry and appearance properties into a single FEM model;
- multiscale approach to correspondence recovery.

The approach assumes that a sequence of registered intensity and range images of a deforming elastic object with visible surface features (such as a grid in Figs. 4 and 6 or irregular and natural features in Section III-D), and a physically-based model [Fig. 2(b)] are available (Section II-A). The main idea of our approach is that a finite element model can naturally handle both registration and modeling using a single model-driving strategy. Previously, snakes were often used to track intensity features (grid lines in Figs. 4 and 6); recovered correspondences (for instance, intersection points) were then incorporated into finite element models that computed deformation parameters. Both processes were iterative: snakes by definition, and a nonlinear finite element solution as a number of linear approximating steps (FEM iterations). This work combines both stages in a single iterative process when the model tracks object features and recomputes physical parameters at the same time (some initial results have previously appeared in [32]).

The second important aspect of the strategy is an efficient data utilization. As much available data as necessary is used. The model consists of a number of control points. The goal is defined as a correct matching of control points with grid points in the next frame. Matching occurs at different resolution levels—using 9, 49, and 217 control points (Fig. 1). The multilevel (multiscale) strategy is described in detail in Section II-C. The matching task is similar for all scales: given the coordinates of control points, find the Hausdorff distance (defined in Section II-B) between the model and the image (next frame), and use it to structure possible correspondences between control points and feature points in images (as discussed in Section II-B). The selected set of correspondences drives deformation of the model.

Section III describes application of the proposed method to motion analysis of man-made elastic materials, human skin, and burn scar detection application. Objects with a grid are used for comparison with a traditional technique utilizing active contours

Fig. 1. Configuration of control points using three different scales.

(a) (b) (c)

Fig. 2. (a) Region of interest. (b) Finite element model. (c) Threshold (T_2) calculation.

and FEM separately. It is also shown (Section III-D) that the proposed method does not need a grid and can take advantage of the available irregular object features or even natural features (in skin experiments). The last section summarizes the results of this research.

II. DESCRIPTION OF THE METHOD

A. Data, Modeling, and Assumptions

Data acquisition, general modeling principles and necessary assumptions are discussed first. Data sequences are acquired using a K2T structured light range scanner. During acquisition, registered intensity and range images of stretching elastic objects are taken (Figs. 4 and 6). Only part of the object (elastic material or human arm) with the region of interest is considered. In first sets of experiments this region includes the grid which is produced with a simple stamp and aids in producing trackable features. Other sets of images contain irregular or even natural features. Let us assume that grid is separated on the intensity image (for instance, using thresholding) and the model is aligned with it. Since the data was collected initially for a different project which employed snakes, the stamp produced overlapping lines (which allows us to compare results of both approaches). For the purpose of this method they are not considered (a conventional method with snakes used original images for the performance comparison). Of course, a fully automatic method would require a different stamp. Hence, only the area bound by outside grid lines is considered [Fig. 2(a)]. Therefore, the finite element model used to describe it is local. It consists of 3-D elastic shells with assigned corresponding material properties (properties of elastic materials are obtained by their mechanical testing; average skin properties are found in the recent literature). Since the geometry of the stamp is known in advance, it allows for the advance model construction necessary for the success of the method [Fig. 2(b)]. Since in this research the emphasis is placed on multiscale use of control points rather than

on the finite element model itself, a current set of control points (Fig. 1) from now on is referred to as our model. The method assumes alignment with the first frame in the sequence and consistency in point inter-relationships so that points do not overlap (occlude) each other.

B. Undirected Hausdorff Distance as Dissimilarity Measure

Control points (which are also FEM nodes or keypoints) provide a natural way to locate corresponding points in the next frame (model registration) and apply distances between them as displacements (model deformation). Control points are guided by the Hausdorff distance [30] between the model M (fitted to the current frame) and the next frame in the sequence F_{n+1}

$$H(M, F_{n+1}) = \max(h(f(M, F_{n+1})), h(f(F_{n+1}, M))) \tag{1}$$

where $h(f(M, F_{n+1})$ is the forward distance (the distance from the model to the image) and $h(f(F_{n+1}, M))$ is the reverse distance.

To compute the forward distance, differences are identified between each control point m_i in the finite element model M and the nearest point a_i in the next frame F_{n+1}, and then the largest distance is selected

$$h(f(M, F_{n+1})) = \max_{m_i \in M} \min_{a_i \in F_{n+1}} \|m_i - a_i\| \tag{2}$$

where f denotes some transformation that occurred as a result of the motion or deformation, and $\|.\|$ is the Euclidean distance. The resulting control point m_i is, therefore, the furthest control point from any range object point in F_{n+1}. The reverse distance $h(f(F_{n+1}, M))$ is defined similarly.

The goal is to use the Hausdorff distance as a measure of mismatch between the model and the object, and then to reduce such differences by applying displacements to the model. This approach belongs to the class of reverse problems when the results (displacements) are given instead of the cause (body loads). Every time displacements are applied, a new configuration of the model is recomputed. As opposed to tracking with snakes (which is a separate physically-based model), the undirected Hausdorff distance can be easily combined with a finite element model. No separation into a motion-detection-oriented model and an object-properties-oriented model is necessary. A single physics-based model is used; it includes material information and is driven by the multiscale analysis of possible correspondences using the Hausdorff distance. Correspondence recovery at each step is followed by displacement calculations and their application to the model. This represents a single iteration of the method.

C. Multiscale Approach to Correspondence Analysis and Model Deformation

Expected range of motion is addressed by the multiscale approach to correspondence analysis and model deformation. Larger motion necessitates the use of coarser alignment models before finer aspects of object deformation are addressed. Multiscale strategy discussed in this section is applicable to a large object and motion domain; however, the number of scales

is based on the magnitude of size or motion and, obviously, may change for different objects.

Three scales (defined in terms of control points) are adopted for the experiments:

1) ICA (initial coarse alignment, nine control points);
2) GGD (general global deformation, 49 control points);
3) CLD (complex local deformation, 217 control points).

Although model deformations start at the ICA scale (using only nine control points), the initial distance estimation is done at the GGD scale. This allows for more precise computation of the undirected Hausdorff distance (or the partial distance [31] for noisy sequences) used as a first threshold (an upper threshold T_1) employed by the method. The meaning of this threshold is an estimate of the largest allowed motion in a given experiment (later applied to control points).

The initial analysis used to determine possible correspondences is performed at the ICA scale. Euclidean distances are calculated (Fig. 3, box 4) between the closest model and image points (if the forward Hausdorff distance was larger), or between the closest image and model points (if the reverse Hausdorff distance was larger). These distances are sorted in decreasing order (Fig. 3, box 5). T_1 is then applied (box 6) to weed out erroneous matches which are possible at any scale. However, at coarse scales, displacements greater than T_1 are simply infeasible (by definition of the Hausdorff distance).

Another threshold (a lower threshold T_2) is then introduced to deal with erroneous matches resulting in small displacements (results of noise and incorrect matches). At each point we find the slope of the tangent to the curve $Displ = f(i)$, where i is the index of the correspondences sorted in decreasing order with respect to resulting 3-D displacements ($\sqrt{dx_i^2 + dy_i^2 + dz_i^2}$). When an absolute value of this slope (or function derivative at a point) at least triples [see Fig. 2(c)], the corresponding 3-D displacement value is chosen automatically as threshold T_2 (usually it increases 4–4.5 times). This threshold can also be computed by using a second derivative or analysis of consecutive differences between (sorted) displacements. T_2 allows for separation of displacements representing another group of erroneous correspondences, namely, those with small displacements, also addressing noise sensitivity issues. This group contains a number of wrong matches, especially during the first few method iterations [Fig. 2(c)].

Therefore, a sparse model is employed to select only 1 : 1 correspondences and discard the rest. Of course, a number of potentially useful correspondences are rejected during this step. This does not matter since at this stage alignment that accounts mostly for translation is more important. Finer model scales process more data that explains nonrigid deformation of the object. If threshold T_2 cannot be found, then the forward Hausdorff distance (at the coarse scale only) is used to align the model uniformly with the next frame data. It is a good approximation of translation between the frames for the considered subset of nonrigid motion.

A change in scale occurs when a current scale no longer improves the alignment. This means that all control points have been assigned correspondences and there is no mismatch between them and the area of interest in the next frame (Fig. 3,
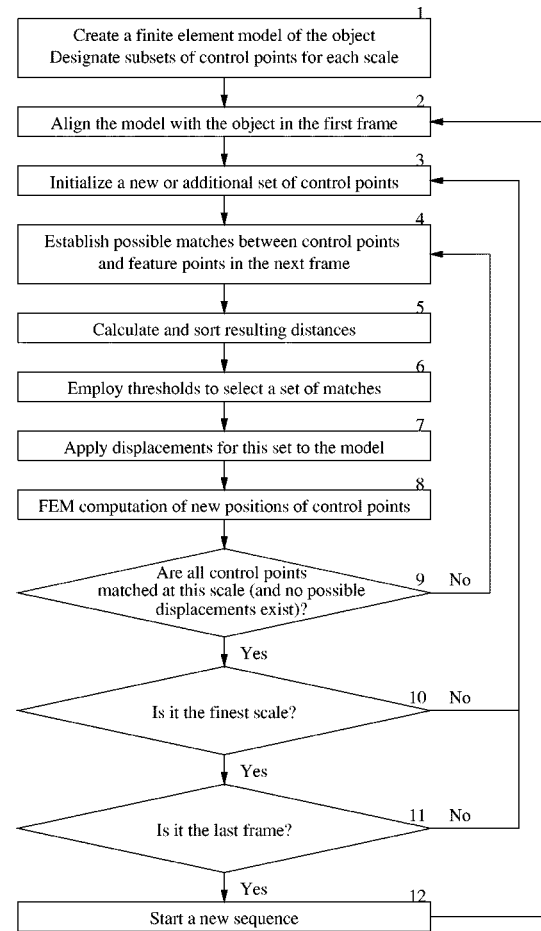

Fig. 3. Algorithm of the multiscale approach.

box 9). Therefore, increase in model scale at this point produces possibility for improving tracking quality.

GGD (49 control points) and CLD (217 control points) scale processing are similar to the coarse scale iterations, except that T_1 is not needed. Since finer aspects of object motion are analyzed, concern for filtering out abnormally large displacements is not justified.

The GGD scale addresses effects of elastic motion (stretching) of the object. It results in a better alignment and accounts for most nonrigid deformations. The CLD scale does not improve tracking significantly if the force is distributed along some real or imaginary line/surface rather than represented by a concentrated loading. In the latter case GCD processing alone cannot account for more complex deformations of grid lines. The steps applicable to all scales are shown in Fig. 3.

The model is displacement-driven; when correspondences are established, displacements are calculated and applied to control points of the model (box 7). The process is repeated during each iteration. Again, in these experiments, the motion of the object is elastic deformation. At any scale, the process can be summarized as follows:

- Hausdorff distance is computed;
- for each control point, possible displacements are found and applied;
- model is incremented accordingly;

- process iterates until the difference between the model and the object is minimal (for each frame).

III. EXPERIMENTAL RESULTS

This section presents an application of the proposed method to motion analysis of man-made elastic materials and human skin. Usefulness of the method is evaluated not only for tracking and motion analysis, but also for a specific application to strain analysis in the burn scar detection procedure.

The model used in experiments described below is local; it covers only a part of the object with the region of interest (the grid in Sections III-A–III-C and irregular features in Section III-D). Second, the model is nonlinear, although an iterative nature of the approach (multiple method iterations for every deformation) allowed us to reach solutions in a single FEM iteration every time. Thus, registration and deformation analysis iterations are combined in a single process. The model consists of thin elastic shell elements defined in a 3-D space (a total of 324 elements and 361 nodes). The same mesh and solution are used in all experiments described in this section. Nonlinearity of materials is considered using multiple model iterations (steps) for every deformation. For more details on model-building, finite element calculation, implementation using ANSYS [33] package and skin parameter selection (material properties and thickness), please see [18]. Having correct material properties in the model is a necessary condition for this method to succeed; however, if precise values of certain materials or tissues are not immediately available, they can also be recovered with the same model as described in [19]. For example, for the scar assessment application (Section III-C), it is not currently feasible to obtain person-specific skin elasticity and thickness precisely, and average properties are sufficient for the strain map computation. Additional experiments show that, varying Young's modulus in the range specified by the literature (see [18] for a literature review), namely from 10 to 100 kPa, does not change the relationship between the scar elasticity and the elasticity of surrounding areas, which is crucial for the strain distribution analysis. To test this, we started from different values of Young's modulus within this range (if multiple areas were involved, then we multiplied their respective Young's modulus by a constant) and still ended up with exactly the same strain maps [Fig. 9 was tested for both approaches in (a) and (b)]. Therefore, we compute elastic properties of burn scars relative to the surrounding areas. This is a sufficient indication of the success of treatment required by the medical community. Poisson's ratio is set equal to 0.49 since most papers agree about it.

A. Application of the Method to Skin Motion Analysis: A Closer Look at Scales and Iterations

This section presents application of the method to skin motion analysis. The experiment presented in this section addresses our current burn scar assessment research described in [18]. The proposed method substitutes the two previously used separate models such as snakes and FEM. In this section we use images with the grid so that the new method can be compared with an old approach. However, later in Section III-D it is shown that the


Fig. 4.   Range and intensity images of skin motion.

TABLE I
SUMMARY OF ITERATIONS (BETWEEN TWO FRAMES IN FIG. 4) AT DIFFERENT SCALES

Scale	ICA scale			GGD scale		CLD scale	
Iteration number	1	2	3	4	5	6	7
Average distance, millimeters	7.191	3.388	0.812	0.523	0.192	0.110	0.097
Average error, %	70.17	33.04	7.92	5.10	1.87	1.07	0.95

proposed method does not need a grid and can take advantage of the natural features (such as birthmarks).

One range and two intensity images of a region of skin being stretched containing a burn scar are shown in Fig. 4. The range image in Fig. 4(a) corresponds to the intensity image before stretching [Fig. 4(b)]. The accuracy of range images allows estimating true 3-D deformations. The presence of a burn scar contributes to nonuniformity of elastic motion.

Control points of the generic grid model are manually aligned with the first frame using both intensity and range data. Then the method proceeds automatically using the available data, model and the strategy described in Section II-C. All three defined scales are used. Thresholds are determined using the undirected Hausdorff distance (T_1) and the jump in distances computed between corresponding points (T_2) similarly to Fig. 2(c). The change in scales occurs when all such distances are equal to zero. The solution required a total of seven method iterations. Control points of the model are determined automatically and moved as follows (although the grid is slightly rotated clockwise, for the simplicity of explanation we will refer to control points as leftmost and rightmost as if grid lines were vertical):

- iteration 1—three leftmost control points moving toward the left side of the grid (ICA scale);
- iteration 2—three rightmost control points moving toward the right side of the grid (ICA scale);
- iteration 3—correspondences and motion for the remaining model points (ICA scale);
- iteration 4—motion of new topmost control points (GGD scale);
- iteration 5—remaining correspondences responsible for general deformation aspects (GGD scale);
- iteration 6—better approximation of the leftmost line (CLD scale);
- iteration 7—other local deformation aspects (CLD scale).

Results are shown in Table I. Both iterations and scales are included. The average real motion between feature points in two frames is 10.253 mm. Of course, the motion is not uniform, it is greater for the areas closer to the place where the force is applied. The average distance between control points in the model and corresponding points of the grid is calculated for validation purposes. It is used to compute the average error (a ratio of recovered and real motion of feature points). Performance of the new method was compared to the traditional approach that

Fig. 5.   Results of the skin motion experiment. (a) Magnitude based vectors representing motion of control points. (b) Results in terms of grid motion between two frames.


Fig. 6.   Range and intensity images of the elastic material during stretching.

Fig. 7.   Magnitude based vectors representing motion of control points.

Fig. 8.   Results in terms of grid motion between consecutive pairs of frames.

uses a separate snake model to recover a sparse set of correspondences (grid intersections) and a dense set using FEM model. The new approach produced not only a lower final average error (0.95% vs. 1.54%) with a smaller standard deviation (0.05% versus 0.11%), but also a better execution time on a SunUltra-SPARC 300 MHz/512K cache/128 MB RAM (24 s versus 1 min 7 s). Increase in efficiency is achieved by replacing two iterative processes (snakes and nonlinear FEM solution both requiring multiple iterations in two separate loops) with a merged iterative process when the model tracks object features and recomputes physical parameters (every step is a single FEM iteration for all experiments described) at the same time.

The error is reduced because CLD scale better accounts for a nonuniform grid line curvature near the stretching force. The total motion of control points is shown in Fig. 5(a) using magnitude based vectors (CLD scale, last iteration). We can visualize grid motion between frames by connecting control points at the finest (CLD) scale [Fig. 5(b)].

B. Results of Motion and Structure Recovery of Elastic Objects: Performance Analysis for Longer Sequences

Sequences of intensity and range images depicting the stretching of an elastic material are utilized for experimental performance assessment of the method. A sequence containing seven intensity images and range images represents input to the algorithm (only one range image is shown in Fig. 6) along with a generic grid model fitted to the initial frame (Fig. 1).

The deformation is produced by incrementally increasing the force causing it. The force is introduced in the second frame.

It changes from 1 newton (N) in the second frame to 3.5 N in the last frame in 0.5 N increments. This sequence allows us to investigate elastic motion in the intervals of material behavior where it can be explained by an elastic model with some geometric nonlinearities (solved in single FEM iterations per model iterations in most cases since model iterations approximate a deformation step incrementally). Material properties are computed experimentally (using a conventional mechanical engineering technique [19]) and included in the model. Magnitude based vectors representing motion of control points are shown in Fig. 7. Recovered grid motion for corresponding iterations is displayed in Fig. 8.

Summary of results for all frames are displayed in Table II. Results are shown per frame, for final method iterations only. The proposed approach performs better than the conventional (snake computation as a part of the traditional technique for this

TABLE II
PERFORMANCE COMPARISON (PER FRAME, FOR FINAL ITERATIONS ONLY).
METHOD (1) IS A CONVENTIONAL APPROACH (SEPARATE SNAKE AND FINITE
ELEMENT MODELS) AND (2) DENOTES THE PROPOSED MULTISCALE
SINGLE-MODEL METHOD

Frame number	2	3	4	5	6	7
Applied force, N	1.00	1.50	2.00	2.50	3.00	3.50
Real average distance, millimeters	7.553	8.440	5.199	7.646	7.465	4.803
Average error, % - method (1)	0.78	0.72	0.73	0.75	0.84	0.87
Average error, % - method (2)	0.77	0.78	0.74	0.70	0.79	0.84

TABLE III
PERFORMANCE COMPARISON FOR A NUMBER OF DIFFERENT EXPERIMENTS
INVOLVING ELASTIC STRETCHING

Experiment	1	2	3	4	5
Average error, % - method (1)	1.54	0.87	2.17	1.53	0.64
Average error, % - method (2)	0.95	0.84	1.30	1.19	1.72
Std. dev., % - method (1)	0.11	0.06	0.15	0.10	0.09
Std. dev., % - method (2)	0.05	0.04	0.07	0.07	0.08

comparison used a fast greedy algorithm by Williams and Shah [34]) for frames with more complex local deformations where CLD scale or even finer scales are beneficial. Performance comparison for a number of different experiments (each uses a different sequence of intensity and range images) involving elastic stretching is shown in Table III (experiments #1 and #2 are ones presented in details in the previous and current subsections, respectively). Standard deviation is also reported. It seems that the proposed method has a lower standard deviation than a conventional approach even for an experiment where is underperformed (however, it performed better than conventional in the other four experiments). In some experiments, the new method slightly increased the number of local errors (due to minor mismatches), but it more than compensated for this by reductions in their magnitude.

The number of method iterations per frame differs; however, it is on average between 5 and 8. Each method iteration takes between 3 and 4 seconds on a UltraSPARC (300 MHz/512K cache/128 MB RAM). Therefore, solution requires less time (on average 23 seconds) than the old approach (more than 1 min per frame).

The reason why the number of method iterations varies, especially during the ICA step, can be explained with the following observations. The motion seems more or less uniform across all frames; however, analysis of displacements reveals that in the first two frames the grid predominantly translates (since the grid contains only a part of the stretching material), while the remainder of frames contain mostly elastic motion (stretching). Quantitatively it can be described as a ratio of displacements between opposite grid points along the force direction. Ratios close to 1 denote translation, ratios from 3 to 5 in our experiments indicate stretching.

That is why if the next frame is the frame where the bandage mostly translates (for instance, first frame), then the ICA part of the method proceeds faster. This reduction in the number of iterations facilitates the finding of almost all coarse scale correspondences during the first iteration as opposed to 3–4 iterations otherwise.
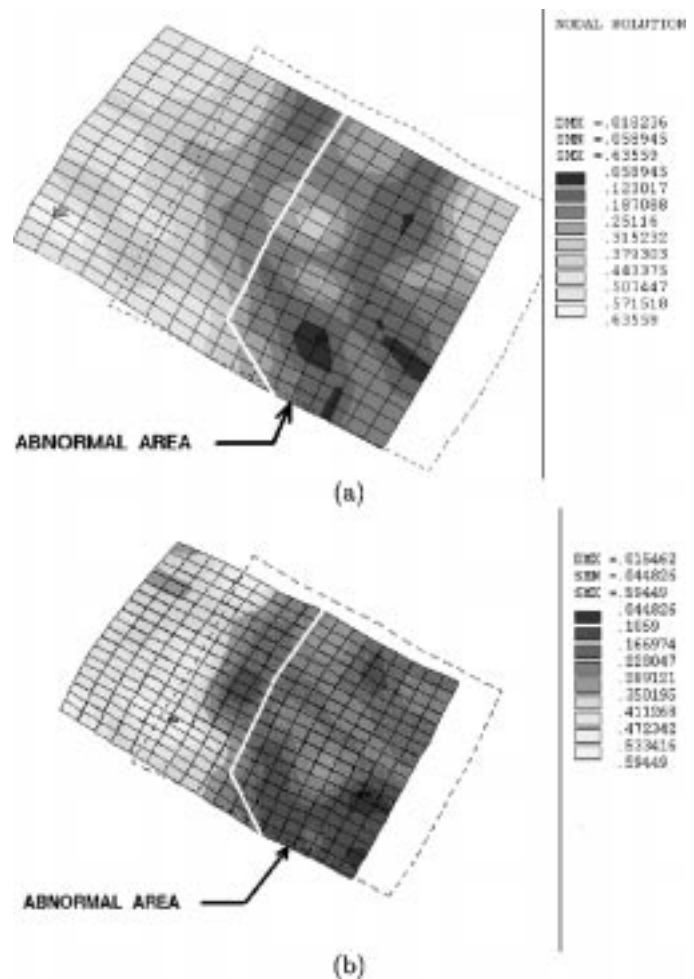

Fig. 9. Resulting strain distribution for the skin motion experiment computed using (a) conventional approach and (b) proposed method. Nine levels of strain are displayed in the grid area from the lowest (white) to the highest (black).

C. Applicational Effectiveness of the Method for Strain Analysis

This section demonstrates applicational value of the outlined method for strain analysis. The application addressed here is the computation of human skin response to applied load that reveals differences in underlying properties. For instance, it allows for the detection of burn scars and estimation of their relative properties [18] which are very useful to physicians for comparing and evaluating treatment options. Of course, accuracy of the structure and correspondence recovery is very important in such an application because it greatly influences resulting strain distributions that pinpoint differences in properties. Strain is recovered after the last method iteration since these differences are detected better using the entire range of motion. Resulting displacements are computed as the differences in positions of control points between the first and the last iterations: $\Delta x_{fin} = x_n - x_1$; $\Delta y_{fin} = y_n - y_1$; $\Delta z_{fin} = z_n - z_1$. Strain is then recovered throughout the surface of the model [18]. Scars restrict the motion, and, therefore, the method is identifying low strain areas (denoted as darker areas in Fig. 9). The legend column on the right of strain distributions shows (top to bottom): maximum displacement, minimum and maximum strain, and strain gradation from the lowest to the highest. These results

Fig. 10. (a) and (b) Intensity images of the elastic material before and after deformation. Feature points are marked as small black circles, validation points—as crosses; (c) range image (before motion); and (d) finite element model (in terms of elements) fitted to the data. Feature points are identified.

correspond to the skin motion experiment (Fig. 4). Fig. 9(a) is obtained with a conventional approach (using separate snake and finite element models, also applied to the available 3-D data [18]). Strain recovered using the proposed method is shown in Fig. 9(b). The resulting strain map is precise enough to identify abnormal areas such as scars (ground truth in the form of scar outlines was provided by physicians). A conventional approach identified correctly 93.83% of the burn scar area; the new method identified correctly 95.68% of the burn scar area. Methods were compared using five burn scar image sequences. The strain can also be used as an additional criteria to restrict impossible modes of motion.

D. Motion Recovery From Irregular and Natural Features

This section extends the use of the proposed method to deforming objects with irregular surface features. These features cannot be easily considered with a conventional approach. First, another piece of an elastic material (considered initially in Section III-B) is stretched (Fig. 10). Stretching is nonuniform affecting the upper side of the bandage much more than the lower.

In this case model fitting procedure adapts the generic model so that it fits the range data and identifies feature points (shown as small circles). Validation points (shown as crosses) are not used during the computation, the difference in their positions before and after the motion is compared to model's estimates after the process completes. There is no single solution to model fitting. The model is acceptable as long as it contains feature points, uses adequate resolution to represent sensed data, and avoids abrupt changes in the element sizes from very fine to coarse (which can lead to ill-conditioning problems). Resolution scales and threshold selection techniques are the same as in previous experiments. Results (shown in Fig. 11) include magnitude based vectors representing motion of all model points and feature points only. The validation error is less than 3% for all validation points (Table IV).


Fig. 11. Magnitude based vectors representing motion of (a) all model points and (b) feature points only.

TABLE IV
MOTION ERROR FOR VALIDATION POINTS

| Validation point	1	2	3	4	5	6	7	Avg
Real distance, millimeters	3.40	3.72	7.21	7.67	8.03	12.91	15.85	8.40
Error, %	2.81	2.83	2.60	2.66	2.71	2.57	2.49	2.67

Fig. 12. (a) and (b) Intensity and (c) range images of skin stretching. (d) Motion tracking of feature points.

A similar experiments are conducted using natural features instead of marked points. Fig. 12(a)–(c) show intensity and range images of skin stretching. Note that there are no artificial markings on the skin. In this case birthmarks are chosen as features to include into the model (similarly to feature points in the previous experiment) and to use for motion analysis [resulting displacement fields for them are displayed in Fig. 10(d)]. Six points are used for validation (the average error is 4.72%). The error remains relatively small since physical model restricts impossible motions. The motion of internal nodes is constrained by the elasticity, and, therefore, predictions at these nodes reflect dynamic continuity. This shows that the method can be extended to other applications and domains, and simplify data acquisition and processing for many existing applications (such as a burn scar assessment application briefly described in Section III-C).

IV. DISCUSSION AND CONCLUSIONS

In this paper we presented a novel multiscale approach to recovery of nonrigid motion from sequences of registered intensity and range images. The main idea of our approach is that a finite element (FEM) model can naturally handle both registration

and deformation modeling using a single model-driving strategy. The method includes a multiscale iterative algorithm based on analysis of the undirected Hausdorff distance to recover correspondences. Our model can handle what previously was accomplished using two types of deformable models (snakes and finite element models). Increase in efficiency is achieved by replacing two separate iterative processes (snakes and nonlinear FEM solution both requiring multiple iterations in two independent loops) with a merged iterative process when the model tracks object features and recomputes physical parameters (every step is a single FEM iteration for all experiments described) at the same time. Efficiency of the method can be further improved by pre-computing the response of the model to all imposed displacements [35]. Control points used for tracking are also a part of the finite element model containing knowledge of an object's material properties that can lead to better analysis of the deformation process. Therefore, not only intersection points, but also additional tracked points are included in the model.

Such a model can explain observed motion effects (such as displacements) as well as nonobservable aspects (such as strains). Strain distributions reveal differences in material properties which can account for motion abnormalities. A good model incorporating available values of material properties is a necessary condition for this approach to succeed. However, it is possible to start with an approximate material information (see Section III for discussion of such cases) and to improve it during tracking (as shown in [19]).

The method includes a multiscale strategy based on evaluation of the undirected Hausdorff distance which represents a reliable error function. Wrong matches occur by mapping control points to nearby features, but they are corrected during subsequent iterations. This possible (yet insignificant) failure mode still results in acceptable (in terms of accuracy) motion field estimates (as shown during experiments). In a general case, a number of scales is object- and motion-dependent, similarly to other physically-based models. Choosing the number of control points at the finest scale is a trade-off between the efficiency of structure representation and effects of noise. (It has been found experimentally that for our setup the scanning error is between 0.5 mm and 1 mm.)

Even with today's range acquisition technology, many important applications (for instance, burn scar assessment) can benefit from the presented ideas. Recent availability of less expensive, faster range data makes it a feasible source of information for tracking. Triclops [36] color stereo vision system (manufactured by Point Grey Research, Vancouver, BC, Canada) is used to capture color and range sequences at a rate of several frames per second. If some prior information is available (for instance, if one is looking for human features or for material of certain colors), then "selective" range processing increases the rate to 16 frames/s [37].

The registration part is based on an interpolation technique: point correspondences are considered as imposed displacements on the finite element model. This approach can be affected by increased noise (which is not an issue in our experiments, but could be a problem in new faster range scanners). The proposed model is flexible so that a different correspondence finding approach (possibly based on an approximation scheme) can be in-

tegrated into the model. Also, such errors are local; they affect somewhat only the node of origin and even to a lesser extent nodes next to it (usually, nine out of 361 nodes), but have no effect on resulting strain distributions (Fig. 9). The difference in the burn scar area identification due to such local errors is less than 5% for all experiments.

The method was evaluated with respect to speed and accuracy. Noise sensitivity issues were addressed. Advantages of the proposed approach were demonstrated using man-made elastic materials and human skin motion. Experiments with regular grid features were used for performance comparison with a conventional approach (separate snakes and FEM models). It was shown, however, that the new method does not require a sampling/correspondence template and can adapt the model to available object features. Usefulness of the method was presented not only in the context of tracking and motion analysis, but also for specific applications such as burn scar detection. This work presents a significant step toward development of models that can inherently handle multiple processing functions, currently registration and deformation, and appearance in the near future.

APPENDIX
NONLINEAR FINITE ELEMENT COMPUTATION

The central concept of the FEM [38] is decomposition of a complex object into simpler components called finite elements. For each object, a mapping between the initial and deformed states can be computed. Availability of material properties aids such mapping resulting in a precise nonrigid motion tracking of this object. The *modulus of elasticity,* or Young's modulus, E, is defined as $\Delta\sigma/\Delta\epsilon$, where $\Delta\sigma$ is the stress change and $\Delta\epsilon$ is the strain change [39]. The stress, σ, can be viewed as force per unit area and the strain, ϵ, as changes of lengths per unit length. Strain measures how much a material has deformed. The *Poisson's Ratio,* μ, is defined as the ratio of the magnitude of the transverse strain to the magnitude of the axial strain.

General *nonlinear* deformation theory defines the displacement field as a combination of rigid-body motions and pure deformations. Former include translations and rotations. Their main property is that the distance between any pair of material points remains unchanged. Any quantity that measures the *change* in length between the neighboring points is a measure of pure deformation. Static nonlinear FEM problem can be expressed as

$$[K(\overline{u})]\overline{u} = F(\overline{u}) \tag{3}$$

where both a matrix of stiffness coefficients $[K(\overline{u})]$ and a force vector $\overline{F}(\overline{u})$ depend on the displacement vector $\overline{u}$. Geometric nonlinearities refer to the nonlinearities in the structure or component due to the changing geometry as it deflects. The stiffness changes because the shape changes and/or the material rotates. The material behaves elastically when the force exerted depends only on how much the material has been deformed. In this work we consider elastic stretching. We account for small geometric nonlinearities and finite strains. The applied loads on a body make it move (or deform) from the position $\overline{u_1}$ to the position $\overline{u_2}$. Hence, the displacement vector is

$$\overline{u} = \overline{u}_2 - \overline{u}_1. \tag{4}$$

The deformation gradient can be defined as

$$[G] = \frac{\partial \overline{u}_2}{\partial \overline{u}_1} \tag{5}$$

where deformation gradient includes the volume change, the rotation and the shape change. The volume change at a point is

$$\frac{dV_2}{dV_1} = \det[G] \tag{6}$$

where $\det$ denotes determinant of the matrix. The deformation gradient can be separated into a rotation and a shape change using the right polar decomposition theorem

$$[G] = [R][U_s] \tag{7}$$

where $[R]$ is the rotation matrix and $[U_s]$ is the right stretch (shape change) matrix. Once a stretch matrix is known, a logarithmic or Hencky strain measure is defined as

$$[\epsilon] = \ln[U_s] \tag{8}$$

or, equivalently, through the spectral decomposition of $[U_s]$

$$[\epsilon] = \sum_{i=1}^{3} \ln(\lambda_i) e_i e_i^{T} \tag{9}$$

where λ_i are eigenvalues of $[U_s]$ (principal stretches) and e_i are eigenvectors of $[U_s]$ (principal directions). Hence, from (7) the average rotation at a point can be calculated. Computationally, incremental approximation (defined in [40]) is used by the ANSYS [33] program utilized in this research for nonlinear FEM calculation. Computationally, incremental approximation is used [33] and increment of the deformation gradient at the current time step n is defined using the previous time step $n-1$

$$[\Delta G_n] = [G_n][G_{n-1}]^{-1}. \tag{10}$$

For more details about the FE solution and implementation, please refer to [18], [40].

ACKNOWLEDGMENT

REFERENCES

[1] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1988.

[2] J. C. McEachen, II and J. S. Duncan, "Shape-based tracking of left ventricular wall motion," *IEEE Trans. Med. Imag.*, vol. 16, pp. 270–283, June 1997.

[3] A. Amini, R. Curwen, and J. Gore, "Snakes and splines for tracking nonrigid heart motion," *Lecture Notes Comput. Sci.*, vol. 1065, pp. 252–261, 1996.

[4] D. Androutsos, P. E. Trahanias, and A. N. Venetsanopoulos, "Application of active contours for photochromic tracer flow extraction," *IEEE Trans. Med. Imag.*, vol. 16, pp. 284–293, June 1997.

[5] S. R. Gunn and M. S. Nixon, "A robust snake implementation; A dual active contour," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, pp. 63–68, Jan. 1997.

[6] H. D. Tagare, "Deformable 2-D template matching using orthogonal curves," *IEEE Trans. Med. Imag.*, vol. 16, pp. 108–117, Jan. 1997.

[7] A. Yezzi, Jr., S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum, "A geometric snake model for segmentation of medical imagery," *IEEE Trans. Med. Imag.*, vol. 16, pp. 199–209, Apr. 1997.

[8] N. Peterfreund, "The velocity snake," in *Proc. IEEE Nonrigid Articulated Motion Workshop*, San Juan, PR, June 1997, pp. 70–79.

[9] S. Chandran and A. K. Potty, "Energy minimization of contours using boundary conditions," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 546–549, May 1998.

[10] A. A. Amini, Y. Chen, R. W. Curwen, V. Mani, and J. Sun, "Coupled b-snake grids and constrained thin-plate splines for analysis of 2D tissue deformations from tagged MRI," *IEEE Trans. Med. Imag.*, vol. 17, pp. 344–356, June 1998.

[11] A. Singh, D. B. Goldgof, and D. Terzopoulos, Eds., *Deformable Models in Medical Image Analysis*. Los Alamitos, CA: IEEE Comput. Soc. Press, 1998.

[12] S. Sclaroff and A. P. Pentland, "Physically-based combinations of views: Representing rigid and nonrigid motion," in *Proc. IEEE Workshop Motion Non-Rigid Articulated Objects*, Austin, TX, Nov. 1994, pp. 158–164.

[13] A. Young and L. Axel, "Non-rigid wall motion using MR tagging," in *Proc. IEEE Computer Soc. Conf. Computer Vision Pattern Recognition*, Champaign, IL, June 1992, pp. 399–404.

[14] D. Metaxas and E. Koh, "Efficient shape representation uisng deformable models with locally adaptive finite elements," *Proc. SPIE*, vol. 2031, pp. 160–171, July 1993.

[15] L. D. Cohen and I. Cohen, "Finite-element methods for active contour models and balloons for 2-D and 3-D images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 1131–1147, Nov. 1993.

[16] S. Basu and A. Pentland, "A three-dimensional model of human lip motions trained from video," in *Proc. IEEE Nonrigid and Articulated Motion Workshop*, San Juan, PR, June 1997, pp. 46–53.

[17] J. Martin, A. Pentland, S. Sclaroff, and R. Kikinis, "Characterization of neuropathological shape deformations," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 97–112, Feb. 1998.

[18] L. V. Tsap, D. B. Goldgof, S. Sarkar, and P. Powers, "A vision-based technique for objective assessment of burn scars," *IEEE Trans. Med. Imag.*, vol. 17, pp. 620–633, Aug. 1998.

[19] L. V. Tsap, D. B. Goldgof, and S. Sarkar, "Nonrigid motion analysis based on dynamic refinement of finite element models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, pp. 526–543, May 2000.

[20] A. Young, D. Kraitchman, L. Dougherty, and L. Axel, "Tracking and finite element analysis of stripe deformations in magnetic resonance tagging," *IEEE Trans. Med. Imag.*, vol. 14, pp. 413–421, Sept. 1995.

[21] S. Kumar and D. B. Goldgof, "Automatic tracking of SPAMM grid and the estimation parameters from Cardiac MR images," *IEEE Trans. Med. Imag.*, vol. 13, pp. 122–132, Mar. 1994.

[22] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," in *Comput. Vis. Graph. Image Process.*, vol. 46, 1989, pp. 1–21.

[23] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Adeformable neuroanatomy textbook based on viscous fluid mechanics," *Proc. Conf. Information Sciences Systems*, pp. 211–216, Mar. 1993.

[24] X. Papademetris, P. Shi, P. Dione, A. J. Sinusas, and J. S. Duncan, "Recovery of soft tissue object deformation using biomechanical models," in *Proc. Information Processing Medical Imaging*, Visegrad, Hungary, June 1999, pp. 352–357.

[25] J. Montagnat and H. Delingette, "A hybrid framework for surface registration and deformable models," in *Proc. IEEE Comput. Soc. Conf. Computer Vision Pattern Recognition*, San Juan, PR, June 1997, pp. 1041–1046.

[26] H. Tao and T. S. Huang, "Connected vibrations: A modal analysis approach to nonrigid motion tracking," in *Proc. IEEE Comput. Soc. Conf. Computer Vision Pattern Recognition*, Santa Barbara, CA, June 1998, pp. 735–740.

[27] J. Park, D. Metaxas, A. Young, and L. Axel, "Deformable models with parameter functions for cardiac motion analysis from tagged MRI data," *IEEE Trans. Med. Imag.*, vol. 15, pp. 278–289, June 1996.

[28] T. O'Donnell, T. Boult, and A. Gupta, "Global models with parametric offsets as applied to cardiac motion recovery," in *Proc. IEEE Comput. Soc. Conf. Computer Vision Pattern Recognition*, San Francisco, CA, June 1996, pp. 293–299.

[29] B. C. Vemuri and A. Radisavljevic, "From global to local, a continuum of shape models with fractal priors," in *Proc. IEEE Comput. Soc. Conf. Computer Vision Pattern Recognition*, June 1993, pp. 307–313.

[30] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the Hausdorff distance," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 850–863, Sept. 1993.

[31] W. J. Rucklidge, "Locating objects using the Hausdorff distance," in *Proc. Int. Conf. Computer Vision*, Boston, MA, June 1995, pp. 457–464.

[32] L. V. Tsap, D. B. Goldgof, and S. Sarkar, "Multiscale combination of physically-based registration and deformation modeling," in *Proc. IEEE Comput. Soc. Conf. Computer Vision Pattern Recognition*, vol. 2, Hilton Head Island, SC, June 2000, pp. 422–429.

[33] Swanson Analysis System Inc., "ANSYS User's Manual Rev. 5.3," Swanson Analysis System, Inc., Houston, TX, 1996.

[34] D. J. Williams and M. Shah, "A fast algorithm for active contours and curvature estimation," *CVGIP: Image Understand.*, vol. 55, no. 1, pp. 14–26, 1992.

[35] S. Cotin, H. Delingette, and N. Ayache, "Real-time elastic deformations of soft tissues for surgery simulation," *IEEE Trans. Visual. Comput. Graph.*, vol. 5, pp. 62–73, Jan.–Mar. 1999.

[36] Point Grey Research Inc., "Triclops Stereo Vision System Version 2.1, User's guide and command reference," Point Grey Research, Vancouver, Canada, 1999.

[37] L. V. Tsap, "Gesture-tracking in real time with dynamic regional range computation," Real-Time Imag., to be published.

[38] D. S. Burnett, *Finite Element Analysis*. Reading, MA: Addison-Wesley, 1988.

[39] P. C. Chou and N. J. Pagano, *Elasticity: Tensor, Dyadic, and Engineering Approaches*. New York: Dover, 1992.

[40] T. J. R. Hughes, "Numerical implementation of constitutive models," in *Theoretical Foundation for Large-Scale Computations for Nonlinear Material Behavior*. Dordrecht, The Netherlands: Martinus Nijhoff, 1984.

Dmitry B. Goldgof (S'84–M'87–SM'93) received the M.S. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1985 and the Ph.D. degree in electrical engineering from the University of Illinois, Urbana-Champaign, in 1989.

He is currently a Professor and Graduate Program Director in the Department of Computer Science and Engineering at the University of South Florida, Tampa, and a member of H. Lee Moffitt Cancer Center.

Dr. Goldgof's research interests include motion analysis of rigid and nonrigid objects, computer vision, image processing and its biomedical applications, and pattern recognition. He has graduated nine Ph.D. and 19 M.S. students, and has published more than 40 journal and 80 conference publications, 14 books chapters, and one book. He is North American Editor for *Image and Vision Computing Journal*.

Dr. Goldgof is a winner of the 19th Annual Pattern Recognition Society Award. He is currently Electronic Associate Editor for IEEE TRANSACTIONS ON SYSTEMS, MAN AND CYBERNETICS, PART B Electronic Section, and a member of International Association of Pattern Recognition Education Committee.

Leonid V. Tsap received the B.S. degree in computer science from the Kiev Civil Engineering Institute, Ukraine, in 1991, and the M.S. and Ph.D. degrees in computer science from the University of South Florida (USF), Tampa, in 1995 and 1999, respectively.

He also held full-time professional positions ranging from Computer Programmer to Senior Analyst, and taught a number of classes at USF. He is currently with the Center for Applied Scientific Computing, University of California Lawrence Livermore National Laboratory. His current research interests include image analysis/computer vision, image synthesis and visualization, pattern recognition, perceptual user interfaces, physically-based modeling, and biomechanics. More specifically, he is interested in nonrigid motion analysis and its applications to medical diagnostics, intuitive data exploration and molecular dynamics. His recent research resulted in 21 refereed publications. He is a member of the editorial board of *Pattern Recognition*.

Dr. Tsap is a member of the IEEE Computer Society and ACM. He is a three-time winner of the annual University of South Florida USPS Scholarship Award, and a recipient of the Provost's Commendation for Outstanding Teaching by a Graduate Student. He also received University of South Florida Graduate Council's Outstanding Dissertation Prize.

Sudeep Sarkar (S'90–M'92) received the B.Tech degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1988, where he was judged the best graduating electrical engineer. He received the M.S. and Ph.D. degrees in electrical engineering, on a University Presidential Fellowship, from The Ohio State University, Columbus, in 1990 and 1993, respectively.

Since 1993, he has been with the Computer Science and Engineering Department, University of South Florida, Tampa, where he is currently an Associate Professor. His research interests include perceptual organization in single images and multiple image sequences, probabilistic reasoning, Bayesian Networks, low-level image segmentation, color-texture analysis of burn scars, nonrigid modeling of impact of burn scars, and performance evaluation of vision systems. He is the co-author of the book *Computing Perceptual Organization in Computer Vision* (Singapore: World Scientific). He also the co-editor of the book *Perceptual Organization for Artificial Vision Systems* (Norwell, MA: Kluwer). He was the guest co-editor of the *Computer Vision and Image Understanding (CVIU) Journal Special Issue on Perceptual Organization in Computer Vision* in October 1999. He is on the editorial boards of the *Journal of Pattern Recognition* and *Pattern Analysis and Applications Journal*.

He is the recipient of the 1994 National Science Foundation CAREER Award, the 1997 USF Teaching Incentive Program Award for Undergraduate Teaching Excellence, and the 1998 Outstanding Undergraduate Teaching Award. He is presently serving on the editorial board of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE.