

Using Multiple Cues for Hand Tracking and Model Refinement

Shan Lu, Dimitris Metaxas
CS Department
Rutgers University
Piscataway, NJ 08854
{shanlu, dnm}@cs.rutgers.edu

Dimitris Samaras
CS Department
S.U.N.Y. Stony Brook
Stony Brook, NY 11794
samaras@cs.sunysb.edu

John Oliensis
CS Department
Stevens Institute of Technology
Hoboken, NJ 07030
oliensis@cs.stevens-tech.edu

Abstract

We present a model based approach to the integration of multiple cues for tracking high degree of freedom articulated motions and model refinement. We then apply it to the problem of hand tracking using a single camera sequence. Hand tracking is particularly challenging because of occlusions, shading variations, and the high dimensionality of the motion. The novelty of our approach is in the combination of multiple sources of information which come from edges, optical flow and shading information in order to refine the model during tracking. We first use a previously formulated generalized version of the gradient-based optical flow constraint, that includes shading flow i.e., the variation of the shading of the object as it rotates with respect to the light source. Using this model we track its complex articulated motion in the presence of shading changes. We use a forward recursive dynamic model to track the motion in response to data derived 3D forces applied to the model. However, due to inaccurate initial shape the generalized optical flow constraint is violated. In this paper we use the error in the generalized optical flow equation to compute generalized forces that correct the model shape at each step. The effectiveness of our approach is demonstrated with experiments on a number of different hand motions with shading changes, rotations and occlusions of significant parts of the hand.

1. Introduction

In this paper we present a model based approach to high degree of freedom articulated motion tracking and model refinement, based on the integration of visual cues and apply it to the problem of hand tracking using a single camera sequence. Hand tracking has received significant attention in the last few years, because of its crucial role in the design of new human computer interaction methods, gesture analysis and sign language understanding. Glove based devices

capture human hand motion directly, but are expensive and hard to use. Vision-based hand tracking is a cost-effective, non invasive alternative. Serious challenges lie in the high number of degrees of freedom and the problem of occlusions.

Two general approaches have been suggested for this problem. Model based approaches try to estimate the position of a hand by projecting a 3-D hand model to image space and comparing it with image features (fingertips [25, 24, 28], line segments [25]). Spline- and quadrics-based hand shape models were used in [23, 27] to minimize differences between the silhouette of the projected model and the data. Others [31, 25] have used stereo to avoid occlusions. Appearance based approaches estimate hand postures directly from the images after learning the mapping from image feature space to hand configuration space [30, 29]. Such systems are more useful for recognizing discrete hand states than for general purpose hand tracking.

The study of motion and shading together has been formalized [20, 22] recently and extended to multiple views [21]. Our approach is model-based and hence can work with a single view. Our first contribution is in the combination of cue forces from edges, optical flow and shading. In particular we introduce in deformable model theory a generalized version of the gradient-based optical flow constraint, that includes shading flow i.e., the variation of the shading of the object as it rotates with respect to the light source. This constraint unifies the shading and the optical flow constraints and degenerates to each one of them when the other is not present. Although optical flow and edges in deformable models have been used in the past [18], as well as shading [17], these two methods were applied to different problem domains (moving and static objects respectively). In this paper we combine them to correct for the errors due to the brightness constancy assumption. We use cue information from the entirety of the hand and we are able to track its complex articulated motion in the presence of shading changes. Given the model-based formulation we augment the optical flow constraint with shading information.

The hand can have as many as 26 degrees of freedom when we model it as a multiple open chain structure. The dynamic/kinematic problem of such a large system, which contains not only open chains but also closed chains, can be modeled as a sub-problem of robotic mechanisms. There are many forward and inverse dynamics simulation techniques for human and robotic motion [14, 15, 16, 10, 13]. Using such a formulation we limit the allowable motion of the fingers with the use of recursive dynamics constraints. The model's driving forces are computed from image cues such as edges, optical flow and shading.

In our formulation we compute 2D data-based forces from edge, optical flow and shading cue constraints. The perspective camera model is used to convert these 2D forces into 3D forces that drive the hand model. These 3D forces are then used to calculate the acceleration of our dynamic hand, its new velocity and new position. Since this is a second order dynamic hand model we use it to predict finger motion from one frame to the next so that we are closer to the data in the next frame. To avoid unnecessary calculations of the shading constraint we monitor the intensity changes in several hand areas during tracking and use it only if these changes are significant.

Since we are using a deformable model framework we can take advantage of the error in the combined generalized flow and edge constraints to improve the shape of the hand model we use. The second contribution is that we further generalize our method by introducing at each integration step a model shape refinement process based on the error from the cue constraints. Since the shape of the hand and fingers is made of articulated piecewise rigid parts we employ this shape correction step in the first few frames to improve the hand shape. The use of cue error for shape correction has already been employed [19, 7, 26] for model shape correction. However, the previous methods used a limited number of cues.

Our paper is organized as follows. The dynamic hand model is described in Sec. 2. Sec. 3 presents model initialization and generation of image forces. Sec. 4 introduces illumination information on the optical flow constraint. Sec. 5 presents the recursive dynamics formulation of the hand model and the constraints on the allowable motion. Sec. 6 formulates the use of residual error during tracking for model shape correction. Tracking experiments are shown in Sec. 7, including complex palm-finger tracking with significant rotation and model-shape correction.

2. Hand Model

In our forward dynamics formulation, the hand model (Fig. 1(a)) consists of a base link (palm), and five link-chains (fingers) connected to the base link through five two-degree-of-freedom revolute joints. Each finger is three links

connected by two one-degree-of-freedom revolute joints. The finger parts are modeled as cylinders and the palm is modeled as a six-rectangle-side-solid.

A two-degree-of-freedom revolute joint can be simplified as two one-degree-of-freedom revolute joints connected by a zero length and zero mass link, (dummy link)[4]. In the hand model there are 21 links including 5 dummy links and 20 one-degree-of-freedom revolute joints. We number the palm (base link) as link 0. For each finger there are 4 links including one dummy link and 4 joints. The joint connecting the finger to the palm is joint 1, and link 1 connects joint 1 and joint 2 (Link 1 is the dummy link). Joint i connects link $(i - 1)$ and link i ; link i links joint i and joint $(i + 1)$. Each link has a local coordinate frame fixed to its starting end.

The above geometric model is based on the measurements of an average male. The user specifies approximately the joint locations in the image to initialize the model. Using our proposed method we then correct this basic shape during tracking to fit the data.

3. Image Based Cues

3.1. Fitting the 3-D Model to 2-D images

This approach needs a geometric 3-D model to transform 2-D forces into 3-D ones which will be applied on the dynamic model. Initially the model is fitted to a known pose of the hand, as can be seen in Figure 1(b). At this stage of the work, we assume knowledge of the camera parameters.

At each frame visibility checking is performed in order to match correctly image and model points. The computation of the relative motion to the palm of occluded fingers, is based on the rigid motion of the hand. When the relative motion is not too large, we pick up the finger edges when they reappear. This method will fail when the fingers undergo significant relative motions when occluded. In order to track them successfully in that case, other methods should be integrated in the framework, such as appearance based methods, which is outside the scope of this paper.

3.2. Force Calculation for Dynamic Model

The 3-D finger motion is recovered by fitting the model to image-derived data. The external forces are applied on the dynamic model, then the rotation and translation of finger joints are calculated. Figure 1(c) shows two kinds of typical finger motion. We obtain the forces by calculating displacements using the following procedure.

- Extract the finger edges using the Canny edge operator.
- A curvature-finding operator [6] is used to find the base points such as B_i, B_j shown in Figure 1(d) for the fingers between the end fingers (little finger and thumb).

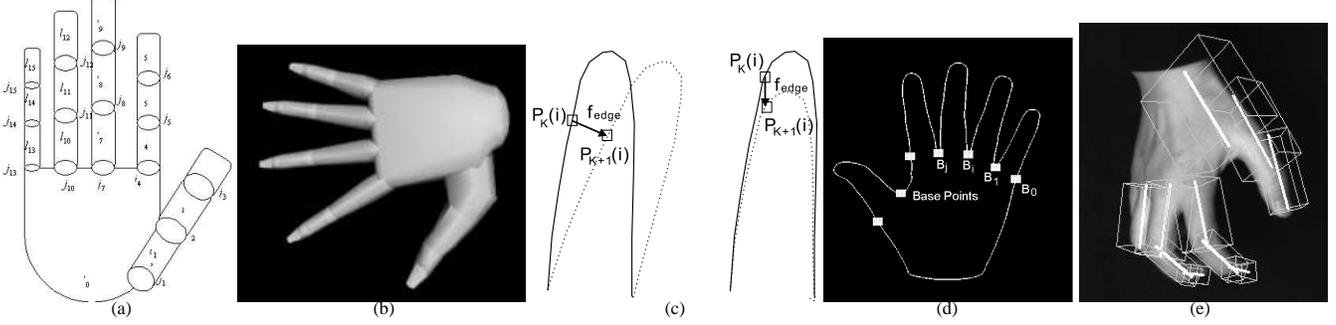


Figure 1. (a) Dynamic Model of Hand. (b) Initial posture of hand model. (c) Finger motion and force from edge displacement. (d) finger segmentation and base points. (e) Representing the projection of the model's articulated segments by their medial axis (thick white line)

For the end fingers we use symmetry to find the beginning of the finger on the outer side (B_0), where curvature can not be used. The beginning is the same as on the inner side of finger where the base point (B_1) is defined based on curvature.

- The edges between B_i and B_j correspond to the finger segment. The edge points of sub-segments can be derived from the corresponding 3-D points in the 3-D model during tracking.
- Because the hand motion will result to the change of base-point position between the current- and after-frame, a normalization process is necessary to match the base-points in current- and after- frame according to the distance of two base-points and the length of finger segment.
- Let $p_k(i)$ and $p_{k+1}(i)$ corresponding edge points in k -th frame and $k+1$ -th frame. The 2-D force f_{edge} from edge displacement can be calculated by the equation.

$$f_{edge}(i) = p_{k+1}(i) - p_k(i) \quad (1)$$

Another force f_{opt} can be directly derived from the optical flow of the image. In the optical flow equation:

$$I_x u + I_y v + f_t = 0, \quad (2)$$

the temporal differential $\mathbf{e} = (u, v)$ at position (x, y) will be considered as the external force. The optical flow of hand motion is computed by the Lucas-Kanade method[9].

Optical flow near finger edges is not as reliable due to possible mismatches of edge points, so we will only consider the optical flow of the inside area of the finger segment (obtained from the projection of the 3-D model in the image plane). For optical flow computation, we select points with significant gradient magnitude only. In Fig. 2 we see the edge forces and the optical flow forces, applied to different regions of the image.

3.3. Force transformation from 2-D to 3-D

We assume a perspective projection model. Therefore, the point $\mathbf{x} = (x, y, z)$ in the world coordinate system and the point $\mathbf{x}_c = (x_c, y_c, z_c)^T$ in the camera coordinate system ensure the following equation.

$$\mathbf{x} = \mathbf{R}_c \mathbf{x}_c + \mathbf{T}_c, \quad (3)$$

where, \mathbf{T}_c and \mathbf{R}_c are translation and rotation matrices.

In the deformable model formulation presented in [8], by taking the time derivatives of the perspective projection equation, with an image point \mathbf{x}_p they get $\dot{\mathbf{x}}_p = \mathbf{H} \dot{\mathbf{x}}_c = \mathbf{H}(\mathbf{R}_c^{-1} \dot{\mathbf{x}})$, with

$$\mathbf{H} = \begin{bmatrix} f/z_c & 0 & -x_c/z_c^2 f \\ 0 & f/z_c & -y_c/z_c^2 f \end{bmatrix} \quad (4)$$

The focal length f is obtained by pre-calibration of the camera. According to deformable model theory these 3D forces are converted to generalized forces $\mathbf{f}_q = \mathbf{J}^T \mathbf{f}_{3d}$ on the model parameters \mathbf{q} , with $\mathbf{J} = \partial \mathbf{x}(x, y, z)/\partial \mathbf{q}$ the Jacobian of the model points, by $\dot{\mathbf{q}} = \mathbf{f}_q$. Consequently, the generalized forces calculated from 2-D images will be $\mathbf{f}_q = (\mathbf{J}_p \mathbf{J})^T \mathbf{f}_{2d}$ with $\mathbf{J}_p = \mathbf{H} \mathbf{R}_c^{-1}$ the Jacobian of the model points under perspective projection.

To apply the external forces on the dynamic model, we transform the individual forces obtained from edges and the optical flow within every hand segment into one total force and torque to be used in the recursive dynamic framework. The total force and torque for each hand segment are $\mathbf{F} = \sum_{i=1}^n f_i, \sum_{i=1}^n \mathbf{r}_i \times \mathbf{f}_i$, respectively. \mathbf{f}_i and \mathbf{r}_i are the individual force vectors and force position vectors, respectively.

4. Extending the Optical Flow Constraint

In [17] a methodology was developed for the incorporation of illumination constraints (any type that is differentiable w.r.t. the model parameters) in a deformable model

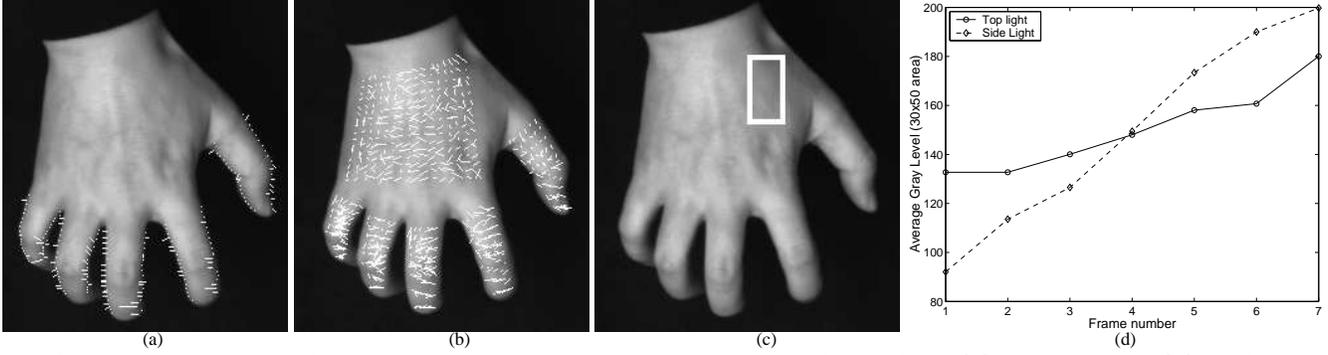


Figure 2. Forces applied to the hand model, and the effects of shading. (a) Edge forces (b) optical flow forces in the interior of the model. (d) is the change in average intensity in a small smooth area of the hand (depicted in (c), when the illumination comes from the top (blue line) and from the side (green dashed line) respectively.

formulation. In that work, the fitting of the model was done based on a static image, i.e. that data did not change during the fitting process. Hence, any partial derivatives with respect to time in the illumination constraint C were zero. In this paper we generalize this constraint formulation to include image motion. Instead of one image, the fitting process will be guided by a sequence of moving images.

We will start by taking the reflectance equation. Let us assume that we have a reflectance function of the general form $I_L = L(\mathbf{l}_p, \mathbf{q})$, where I_L is the observed image intensity, \mathbf{l}_p are the lighting model parameters (both light source parameters and shape reflectance properties such as the surface albedo of a Lambertian model), and \mathbf{q} are the hand model parameters. $L(\mathbf{l}_p, \mathbf{q})$ can be differentiated with respect to the model parameters \mathbf{q} . This means that the reflectance of the surface is locally computable and that there are no global illumination effects. We also assume that the illumination parameters do not change with time. The constraint equation is $C = I_L - L(\mathbf{l}_p, \mathbf{q})$, and we differentiate it w.r.t. time, and apply Baumgarte stabilization [3] in order to obtain

$$\dot{C}(\mathbf{q}, t) + \alpha C = \mathbf{C}_q \dot{\mathbf{q}} + \mathbf{C}_t + \alpha C = 0, \quad (5)$$

In this case we cannot ignore the partial derivatives \mathbf{C}_t w.r.t. time. Therefore, using the above formulas we expand Equation 5 to:

$$\frac{\partial \mathbf{I}_L}{\partial \mathbf{q}} \dot{\mathbf{q}} - \frac{\partial \mathbf{L}(\mathbf{l}_p, \mathbf{q})}{\partial \mathbf{q}} \dot{\mathbf{q}} + \frac{\partial \mathbf{I}_L}{\partial t} - \frac{\partial \mathbf{L}(\mathbf{l}_p, \mathbf{q})}{\partial t} + a(I_L - L(\mathbf{l}_p, \mathbf{q})) = 0 \quad (6)$$

We notice that if \mathbf{J} is the Jacobian of the model points, and \mathbf{J}_p is the Jacobian of the model points under perspective projection, as described in Sec. 3, then

$$\frac{\partial \mathbf{I}_L}{\partial \mathbf{q}} \dot{\mathbf{q}} + \frac{\partial \mathbf{I}_L}{\partial t} = \nabla \mathbf{I}_L \mathbf{J}_p \mathbf{J} \dot{\mathbf{q}} + \frac{\partial \mathbf{I}_L}{\partial t} \quad (7)$$

is the left hand side of the model based optical flow constraint equation that was presented in [18]. As explained in

this paper, in model based optical flow, motion field vectors are vectors of velocities of model points, and hence $\dot{\mathbf{x}} = \mathbf{J} \dot{\mathbf{q}}$ applies. Typically in the literature [11] this optical flow term is set to 0. This is correct in the case of ambient only illumination. For the case of light sources at infinity it is also correct for pure translational motion. For the simplest case of a Lambertian surface with a light source at infinity it can be shown [12] that if ω is the angular velocity of the rotational motion, \mathbf{n} is the normal of surface, \mathbf{l} is the light source direction, and ρ is the albedo of the surface, the magnitude of the error $|Dv|$ between the true motion field and the apparent (and computable) optical flow is

$$|Dv| = \rho \frac{|\mathbf{l}(\omega \times \mathbf{n})|}{\|\nabla E\|} \quad (8)$$

This error is small when the change of gradient ∇E is big, but in the case of smooth surfaces this effect becomes much more pronounced. Similarly $\frac{\partial \mathbf{L}(\mathbf{l}_p, \mathbf{q})}{\partial t} = 0$ since normals change based only on the model parameters \mathbf{q} .

This means that when there is no motion the constraint equation simplifies to the shading constraint. Therefore

$$\frac{\partial \mathbf{I}_L}{\partial \mathbf{q}} \dot{\mathbf{q}} - \frac{\partial \mathbf{L}(\mathbf{l}_p, \mathbf{q})}{\partial \mathbf{q}} \dot{\mathbf{q}} + \frac{\partial \mathbf{I}_L}{\partial t} - a(L(\mathbf{l}_p, \mathbf{q}) - I_L) = 0 \quad (9)$$

encompasses both constraints. In the case of a smooth moving object (9) allows to deal with errors due to directed illumination and offers the possibility of recovering the motion of relatively smoothly shaded surfaces. Fig. 2(c), (d) shows the change in average intensity in a small smooth area of the hand, when the illumination comes from the top and from the side respectively. In the second case, changes in the intensity of the points are dramatic.

5. Dynamic Tracking of Hand Motion

In our methodology we estimate the hand motion in response to the applied 3D forces on the hand as a Forward

Dynamics problem where given the external forces we want to compute the velocity and position of the palm.

Since we use a recursive dynamic formulation we will use Featherstone's[2] spatial notation to model our kinematic and dynamic variables. We integrate the constraint of Eq. 9 in the above formulation to determine the vector \mathbf{q} of the model's degrees of freedom which includes the joint variables, global rotation and translation based on the method of Lagrange multipliers. In addition we also use generalized edge forces computed from image-based edge forces as outlined in Sections 3.2 and 3.3. In summary we solve for

$$\dot{\mathbf{q}} = \mathbf{f}_q \quad (10)$$

subject to:

$$\text{Equation (9) and } \mathbf{f}_{edges} \quad (11)$$

where \mathbf{f}_q are generalized forces computed from the two constraints above.

Furthermore, human fingers are not ideal dynamic links, their joints have upper and lower bounds. Therefore, we need to solve the above dynamic equations under joint limit constraints. These joint limits which constrain the relative motion of fingers together with our dynamic formulation which does not allow the inter-penetration of fingers make hand tracking significantly more robust. Our method has the following steps:

1. At time t , mark the joints that reach their joint limits.
2. Solve the dynamic equations of the hand (i.e., solve for \mathbf{q}) at time $t + dt$ recursively using (10).
3. For each finger, starting at joint 1 (the joint that connects the palm and the finger), mark the first joint that keeps at its joint limit during the time period from t to $t + dt$. If there is no such joint, go to step 6.
4. Fix the joints marked at step 3, and merge two links connected by a fixed joint to one link. Update the dynamic hand model.
5. Go back to step 2.
6. Output the status of the dynamic model of the hand at time $t + dt$. Increase time $t = t + dt$, and go to step 1.

6. Model Shape Correction from Cue Residual Error

Based on the above methodology we estimate the model's rotation and translation parameters \mathbf{q} at each time step. In our presentation of the approach we have so far assumed that the model shape is known. We now relax this assumption and extend our method to allow the estimation of the model shape. This extension will allow us to start with a rough approximation of the model shape which we will refine over time. The approach is as follows. Assuming that the initial hand model is not accurate, we will compute

the change of its shape based on the residual error from the cues, i.e., the generalized optical flow and the edges. Our assumption here is that the residual error is primarily caused by errors in the initial hand shape. This assumption is justified since we do not allow any abrupt changes in illumination during tracking. Based on the above our algorithm is as follows at a time step i :

1. Compute the first component of the residual error at time step i from the generalized optical flow constraint equation (9). In this equation the error is computed based on the estimated motion of the model at time step i and the difference between the projected estimated model's intensities and the image data at time step i .
2. Compute the second component of the residual error at time step i from the edge difference between the projected estimated model location at step i and the image data at time step i . This computation is described in Section 3.2.
3. Add the two components of the residual error and use them as a constraint to modify the shape of the model in a fashion similar to the one described in the previous section. The difference is that instead of modifying the motion of the hand model we modify its shape. For the purposes of this paper we allowed non-isotropic scaling deformations and we found them sufficient as is shown in our results section.

The shape correction step is done at every step i right after the estimation of the model's motion parameters. However, in our experiments, approximately 10 – 20 frames after the initial step we stop correcting the model shape. Since the hand model is comprised of articulated rigid parts 10 – 20 frames are usually sufficient to correct for the shape of the hand.

7. Experiments

We have performed experiments to test our method with a variety of hand motions. All our experiments run on a P4 1.0GHz processor at approximately 4 frames per second. The experimental results are shown in Fig. (3) and Fig.(4), respectively.

The first experiment (Fig.3) demonstrates the process of fitting the model to the image by simultaneous motion tracking and model shape change. From the rotational motion of whole hand, the width and the length of palm and the fingers are modified to fit the image. (the first and second rows in Fig. 3.) In the third and fourth rows in Fig. 3, individual fingers bend towards the direction of the thumb. The thickness of the fingers also changes to fit the images. At the last frames in the sequence (the fifth row in Fig. 3), the shape of

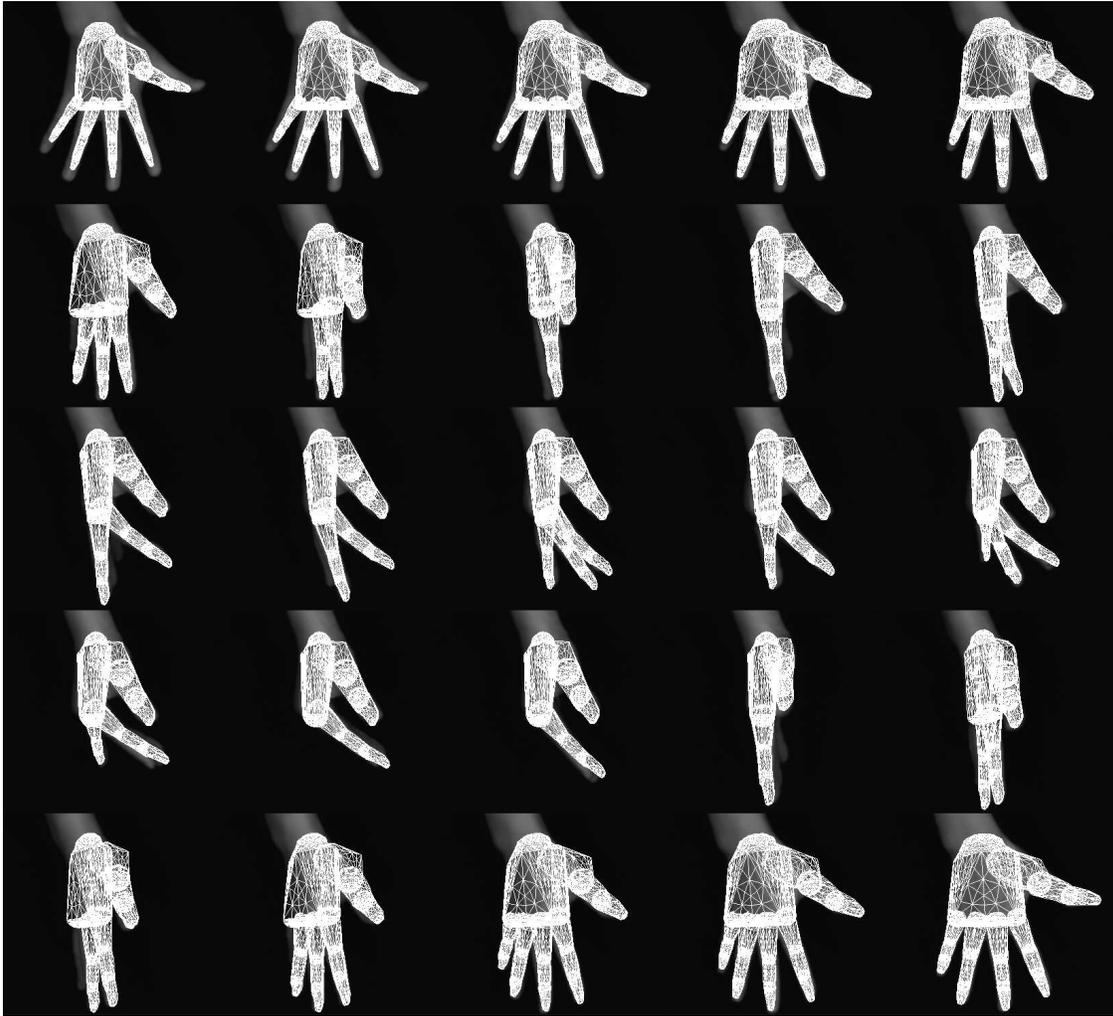


Figure 3. Projected wire frame images from a sequence tracking hand rotation and finger movements. The shape of the palm and the fingers are being dynamically modified from the initial shape during tracking. The wire frames of the model projected on the original images show that the modified shape fits very well the image data.

hand model does not change anymore, and the whole hand rotates back to the original position.

The second experiment depicted in Fig.4 used the fitted model shown in Fig.3. The sequence has been taken with the same camera position as in Fig.3. From the first to the third row, the ringer finger and little finger bended largely away from camera. Finally, the little finger has been occluded completely. Then, the whole hand rotates at about 90 degrees. To show the accuracy of the tracking we project the wire frames of the hand model back onto the image.

The full sequences and the tracking results are available as movie files and can be found at <http://www.dcis.rutgers.edu/~shanlu/hand/> for the above two experiments.

The above experiments demonstrate the successful tracking of complex hand and finger motions involving large rotations and large relative finger motions. The dynamic estimation of the hand shape model significantly improved the tracking accuracy and robustness.

8. Conclusions

In this paper we have augmented traditional optical flow and replaced it with a more general equation that includes shading information. We then further extended this formulation to use the cue residual error to correct for the model's shape. We have used this formulation within a deformable model framework and we were able to track difficult hand

motions under a variety of illumination conditions. Our dynamic hand model formulation allows the integration of multiple cues and for robustness we also use edges in our tracking. We have shown tracking results for simple and complex palm and finger motions. Future work includes better occlusion recovery handling using Kalman Filtering and the incorporation of other sources of visual information such as color, in order to work on cluttered backgrounds.

References

- [1] G. Engeln-Mullges, F. Uhlig. Numerical Algorithms with C. Springer, 1996
- [2] R. Featherstone. Robot Dynamics Algorithm. Kluwer Academic, Boston, 1987
- [3] J. Baumgarte. Stabilization of constraints and integrals of motion in dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 1:1–16, 1972.
- [4] G. Huang, D. Metaxas, and J. Lo. Human Motion Planning Based on Recursive Dynamics and Optimal Control Techniques. *Computer Graphics International 2000*, pp. 19-28.
- [5] J. Lo and D. Metaxas. Recursive Dynamics and Optimal Control Techniques for Human Motion Planning. CA'99. Geneva, Switzerland, May 26-29, 1999
- [6] S.B. Kang and K. Ikeuchi, "Toward Automatic Instruction from Perception – Recognizing a Grasp from Observation", *IEEE Trans. of Robotics and Automation*, pp.432-443, Aug. 1993.
- [7] R. Koch. Dynamic 3-D Analysis through Synthesis Feedback Control. IEEE PAMI, 15(6), pp. 556-568, June 1993.
- [8] D.N. Metaxas, *Physics-Based Deformable Models – Applications to Computer Vision, Graphics and Medical Imaging*, Kluwer Academic Publishers, 1997.
- [9] B. Lucas and T. Kanade, "An Iterative Technique of Image Registration and Its Application to Stereo", *Proc. 7th IJCAI*, pp.674-679, August, 1981.
- [10] J. Angeles and O. Ma. Dynamic Simulation of n -Axis Serial Robotic Manipulators Using a Natural Orthogonal Complement. *The International Journal of Robotics Research*, 7(5):32-47, October 1988.
- [11] B.K.P. Horn. *Robot Vision*. 1986.
- [12] A. Verri and T.A. Poggio. Motion field and optical flow: Qualitative properties. *PAMI*, 11(5):490–498, May 1989.
- [13] H. Brandl, R.Johanni, and M.Otter. An Algorithm for the Simulation of Multibody Systems with Kinematic Loops. *IFTToMM Seventh World Congress on the Theory of Machines and Mechanisms*, Sep. 1987.
- [14] J.K. Hodgins, W.L. Wooten, D.C. Brogan, and J.F. O'Brien. Animation of Human Athletics. *SIGGRAPH 95*.
- [15] A.J. Stewart and J.F. Cremer. Beyond keyframing: An algorithmic approach to animation. *Graphics Interface*, 1992
- [16] J. Wilhelms and B. Barsky. Using Dynamic Analysis to Animate Articulated Bodies such as Humans and Robots. In *Graphics Interface*, 1985
- [17] D. Samaras and D. Metaxas. Incorporating Illumination Constraints in Deformable Models. CVPR 1998, pp. 322-329.
- [18] D. DeCarlo and D. Metaxas. Optical Flow Constraints on Deformable Models with Applications to Face Tracking. IJCV, July 2000, 38:2, pp. 99-127.
- [19] D. DeCarlo and D. Metaxas. Adjusting Shape Parameters Using Model-Based Optical Flow Residuals. IEEE PAMI, 24(6), pp. 814-823, June 2002.
- [20] Negahdaripour, S., Revised Definition of Optical Flow: Integration of Radiometric and Geometric Cues for Dynamic Scene Analysis, PAMI(20), No. 9, Sep. 1998, pp. 961-979
- [21] Carceroni, R.L., Kutulakos, K.N, Multi-View Scene Capture by Surfel Sampling: From Video Streams to Non-Rigid 3D Motion, Shape and Reflectance, ICCV01(II: 60-67)
- [22] Haussecker, H., and D. J. Fleet, 2000. Computing optical flow with physical models of brightness variation, PAMI(23), No. 6, pp. 661-673, 2001.
- [23] JJ Kuch and T. S. Huang. Vision-based hand modeling and tracking for virtual teleconferencing and telecollaboration. In ICCV95, pg.666-671,1995.
- [24] J. Lee and T. Kunii. Model-based analysis of hand posture. IEEE CGA, 15:77-86, Sept. 1995.
- [25] J. Rehg and T. Kanade. Model-based tracking of self-occluding articulated objects. In ICCV '95, pg 612-617.
- [26] Shimada, N., Shirai, Y., Kuno, Y. and Miura, J., Hand Gesture Estimation and Model Refinement Using Monocular Camera: Ambiguity Limitation by Inequality Constraints. Procs. Automatic Face and Gesture Recognition Workshop, pp. 268-273, 1998.
- [27] B. Stenger, P.R. Mendonca and R. Cipolla. Model-based Hand Tracking Using an Uncented Kalman Filter. Procs. British Machine Vision Conference, Vol. 1, pp. 63-72, Sept 2001.
- [28] Ying Wu and T. S. Huang. Capturing articulated human hand motion: A divide-and-conquer approach. In ICCV '99, pg 606-611, Corfu, Greece, Sept. 1999.
- [29] Ying Wu; Lin, J.Y.; Huang, T.S. Capturing natural hand articulation. ICCV '01 (II:426-432)
- [30] Rosales, R., Athitsos, V., Sigal, L., and Sclaroff, S., 3D Hand Pose Reconstruction Using Specialized Mappings, ICCV01
- [31] Q. Delamarre and O. Faugeras, "Finding pose of hand in video images: a stereo-based approach", AFGR 98.

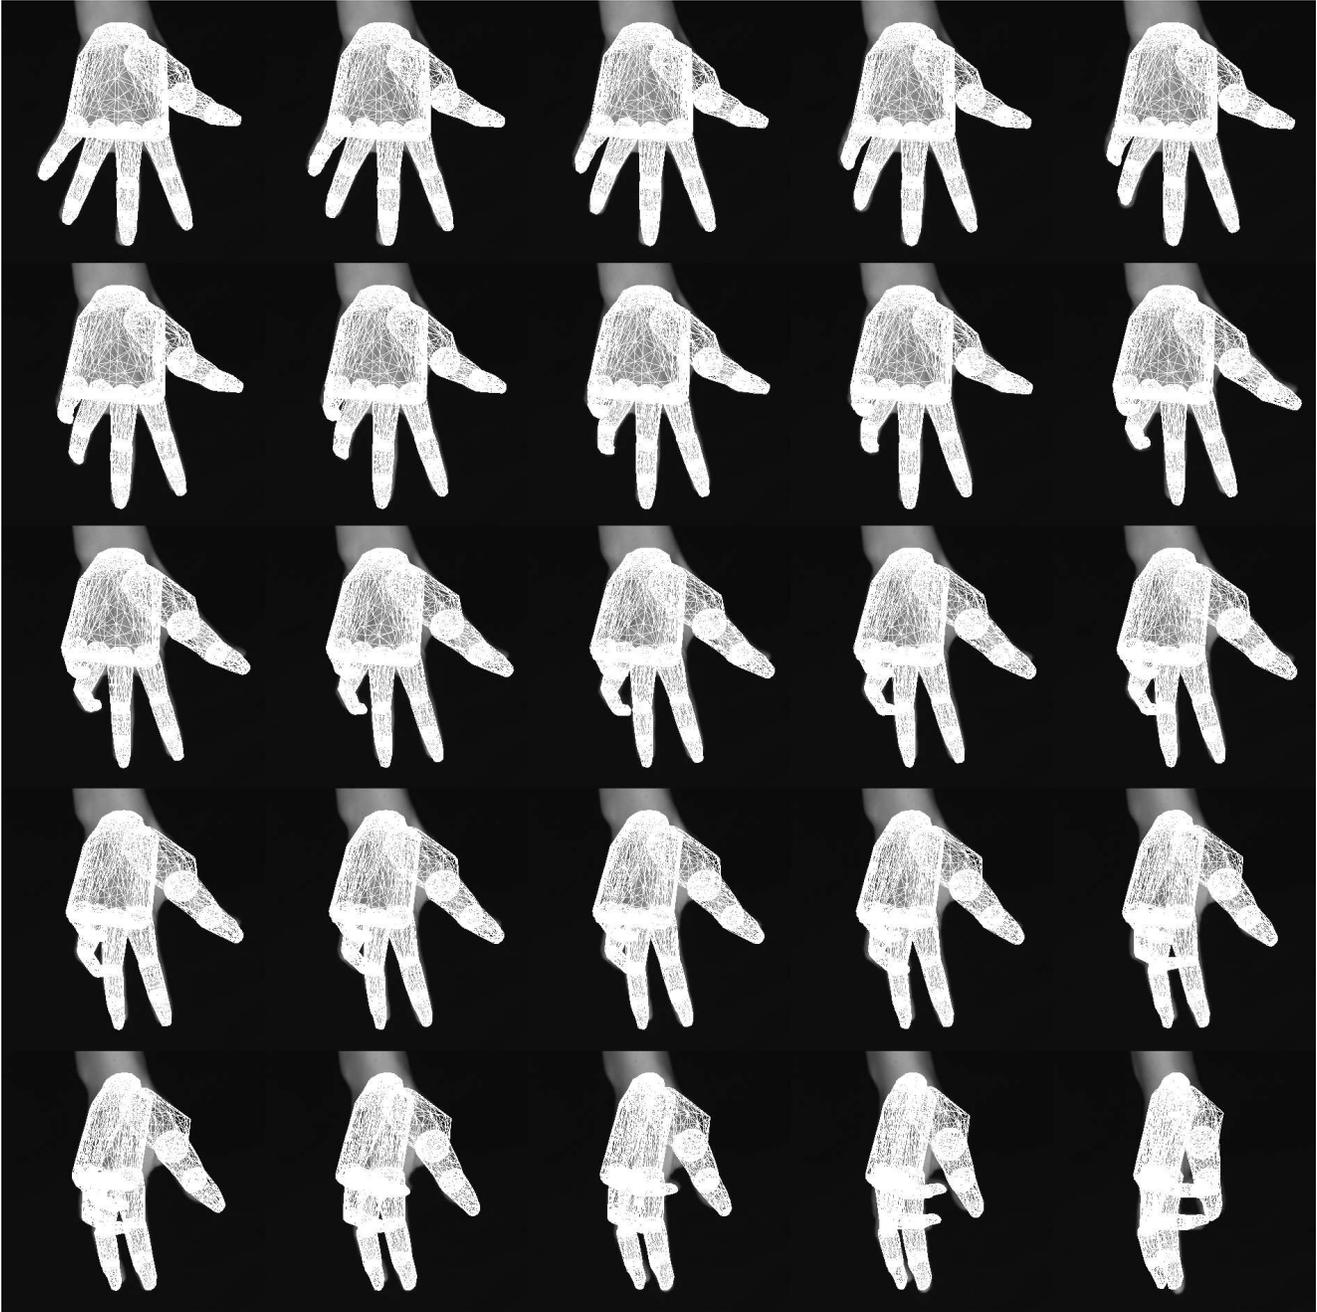


Figure 4. 25 Frames from a sequence, tracking flexing of fingers and hand rotation. After the initial shape estimation as shown in Figure 3, the movements of the hand and the fingers have been tracked accurately. The tracking results are shown by projecting the wire frame of the hand model on the original images.