University of Pennsylvania
ScholarlyCommons

Center for Human Modeling and Simulation

Department of Computer & Information Science

June 1995

3D Human Body Model Acquisition from Multiple Views

Ioannis A. Kakadiaris
University of Pennsylvania

Dimitris Metaxas
University of Pennsylvania

Follow this and additional works at: <https://repository.upenn.edu/hms>

Recommended Citation

Kakadiaris, I. A., & Metaxas, D. (1995). 3D Human Body Model Acquisition from Multiple Views. Retrieved from <https://repository.upenn.edu/hms/94>

Postprint version. Published in *Proceedings of the Fifth International Conference on Computer Vision*, June 1995, pages 618-623.

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/hms/94>
For more information, please contact repository@pobox.upenn.edu.

3D Human Body Model Acquisition from Multiple Views

Abstract

We present a novel motion-based approach for the part determination and shape estimation of a human's body parts. The novelty of the technique is that neither a prior model of the human body is employed nor prior body part segmentation is assumed. We present a Human Body Part Identification Strategy (HBPIS) that recovers all the body parts of a moving human based on the spatiotemporal analysis of its deforming silhouette. We formalize the process of simultaneous part determination, and 2D shape estimation by employing the Supervisory Control Theory of Discrete Event Systems. In addition, in order to acquire the 3D shape of the body parts, we present a new algorithm which selectively integrates the (segmented by the HBPIS) apparent contours, from three mutually orthogonal views. The effectiveness of the approach is demonstrated through a series of experiments, where a subject performs a set of movements according to a protocol that reveals the structure of the human body.

Comments

Postprint version. Published in *Proceedings of the Fifth International Conference on Computer Vision*, June 1995, pages 618-623.

3D Human Body Model Acquisition from Multiple Views

Ioannis A. Kakadiaris and Dimitri Metaxas*

General Robotics and Active Sensory Perception (GRASP) Laboratory

Department of Computer and Information Science

University of Pennsylvania, Philadelphia, PA, 19104

email: ioannisk@grip.cis.upenn.edu, dnm@central.cis.upenn.edu

Abstract

We present a novel motion-based approach for the part determination and shape estimation of a human's body parts. The novelty of the technique is that neither a prior model of the human body is employed nor prior body part segmentation is assumed. We present a Human Body Part Identification Strategy (HBPIS) that recovers all the body parts of a moving human based on the spatiotemporal analysis of its deforming silhouette. We formalize the process of simultaneous part determination, and 2D shape estimation by employing the Supervisory Control Theory of Discrete Event Systems. In addition, in order to acquire the 3D shape of the body parts, we present a new algorithm which selectively integrates the (segmented by the HBPIS) apparent contours, from three mutually orthogonal views. The effectiveness of the approach is demonstrated through a series of experiments, where a subject performs a set of movements according to a protocol that reveals the structure of the human body.

1 Introduction

Computer vision has begun to play an increasingly important role in applications like anthropometry, human factors design, ergonomics, teleconferencing, virtual reality and performance measurement of both athletes and patients with psychomotor disabilities. All of these areas require identification of the parts of a human body and estimation of their shape and motion parameters. The main difficulties in developing algorithms for human shape and motion analysis stem from the complex 3D non-rigid motions of humans and the occlusion among body parts.

To recover the degrees of freedom associated with the shape and motion of a moving human body, most of the existing approaches introduce simplifications by either using a model-based approach [10, 4, 1, 15, 14] or employing assumptions on the kinds of motions they can analyze [16, 2, 17, 3, 7, 8]. Most of these techniques use non-deformable models that can only approximate the human body (e.g., generalized cylinders) and cannot adapt to different body sizes. To overcome this limitation other

researchers assume prior segmentation of the image data into parts [12] and then fit deformable models that can adapt to data from humans with different anthropometric dimensions [11, 9]. Thus, the process of segmentation and the process of shape and motion estimation are decoupled, leading to possible inaccuracies and lack of robustness. Moreover, no technique exists to automatically acquire a concise 2D model of the human body and its parts using vision sensors. Finally, the very difficult problem of 3D shape and motion estimation has not been addressed so far due to the difficulty in integrating multiple viewpoints and dealing with occlusions between the body parts.

Our long term goal is to provide an integrated framework for tracking the body parts of a human moving in 3D which copes with occlusion between the body parts. For some tasks (e.g., determining if a person is moving towards or away from you) information about the movement of the centroid of the silhouette is adequate. For other tasks though (e.g., virtual reality), detailed motion information for the body parts is required. Since humans have widely varied anthropometric dimensions the first part of the tracking process should be the acquisition of the subject's model.

This paper makes three contributions to the topic above. The first contribution is the development of an algorithm that first fits only one model to the image data, but then it segments it to the different body parts as the person moves. The second contribution is a new approach to modeling and controlling the processes in HBPIS. The Supervisory Control Theory of Discrete Event Systems [13] allows us to encapsulate both the discrete and the continuous aspects of the fitting and segmentation processes. Estimating the 3D shape of the body parts, for any human, is the third contribution of this paper.

2 2D Human body model acquisition

As a human moves and attains new postures its apparent contour changes dynamically and large protrusions emerge as the result of the motion of the limbs (figure II(1-3)). To represent large protrusions or concavities and their shape evolution in a compact and intuitive way, we introduce a new shape representation based on the parametric composition of primitives. Using this representation, we can describe compactly the shape of the union (in the case of protrusions) or intersection (in the case of concavities) of two primitives. For simplicity, we formulate the theory of

*This work has been supported in part by the following grants: ARO Grant DAAL03-89-C-0031PRI, ARPA Grants N00014-92-J-1647 and DAAH-049510067, NSF Grants CISE/CDA-88-22719, MIP94-20397, IR193-09917 and MIP94-20393.

composition in 2D, to represent the shape of the boundary of the union of primitives. Let $\mathbf{x}_0$ and $\mathbf{x}_1$ be two 2D parametric primitives, positioned in space so that $\mathbf{x}_1$ (*intersecting primitive*) intersects $\mathbf{x}_0$ (*root primitive*) at points A and B (figure I(a-c)). The shape $\mathbf{x}$ of the composed primitive ($\mathcal{C}: [v_b, v_e] \rightarrow \mathbb{R}^2$) can be defined in terms of the parameters of the defining primitives $\mathbf{x}_0$ and $\mathbf{x}_1$ as follows (see [5] for more details):

$$\mathbf{x}(v; v_0^A, v_0^B, c) = (1 - \delta(v; h_0^{-1}(v_0^A), h_0^{-1}(v_0^B), c)) \mathbf{x}_0(h_0(v)) + \delta(v; h_0^{-1}(v_0^A), h_0^{-1}(v_0^B), c) \mathbf{x}_1(h_1(h_0(v))).$$

Since humans have widely varied anthropometric dimensions the first part of the tracking process should be the acquisition of the subject's body model. Our goal is to automatically segment body apparent contours of moving humans and estimate their shape without assuming a prior model of the human body or body part segmentation. To achieve this goal we perform a controlled experiment. We request that the individual under observation performs a set of movements (protocol movA in [5]) that reveals the structure of the human body and ensures that all the major parts of the human body become visible. The result of processing the image data from this set of movements will be a 2D model of the subject's body including its parts and their shape.

Human Body Part Identification Strategy (HBPIIS)

Step 1: Initially, assume that the subject's body consists of a single part. Create a list of deformable models L (with one entry initially) that will be used to model the subject's body parts. In addition, create a graph G with one node. The nodes of the graph G denote the body parts recovered by the algorithm. The edges of the graph denote which parts are connected by joints.

Step 2: If not all the frames of the motion sequence have been processed fit the models of the list L to the image data using the physics-based shape and motion estimation framework [9] and execute steps 3 and 4. Otherwise, output L and G .

Step 3: For each model in L , determine:

- a: if the *Parametric Composition Invocation Criterion* is satisfied (e.g., this criterion is satisfied when the subject lifts her arm towards the horizontal position, as the apparent contour of the arm protrudes from the one of the torso).
- b: if the *Part Decomposition Criterion B* is satisfied (e.g., this criterion is satisfied during the movement of the legs).
- c: if the *Part Decomposition Criterion C* is satisfied (e.g., this criterion is satisfied when the subject bends her elbow).

For each composed model in L , determine:

- d: if the *Part Decomposition Criterion A* is satisfied (e.g., this criterion is satisfied during the later stages of the movement of the arm with respect to the torso).

Step 4: For the protocol of movements movA, for each model in L at most one of the criteria specified in step 3 will be satisfied. Depending on the outcome, determine which of the following algorithms to invoke.

For each model in L

- If 3a is satisfied, invoke the *Parametric Composition* algorithm.
 - If 3b is satisfied, invoke the *Part Decomposition B* algorithm.
 - If 3c is satisfied, invoke the *Part Decomposition C* algorithm.
- For each composed model in L

- If 3d is satisfied, invoke the *Part Decomposition A* algorithm.

The Part Decomposition Criterion C and the related algorithm have been described in [6]. In the following, we present the other criteria and the related algorithms in more detail.

HBPIIS - Step 3a: As a subject moves and attains new postures the apparent contour changes dynamically and large protrusions are emerging as the result of the motion of the limbs (figure II(1-3)). If there is no hole (by hole we denote the existence of a closed contour within the apparent contour) present within the apparent contour and there is a significant deformation of the apparent contour, we represent the protrusions as the result of composition of two primitives.

Parametric Composition Invocation Criterion: Signal the need for parametric composition of primitives if no hole is evolving within the apparent contour and $\|\mathbf{p}_0(v, t) - \mathbf{p}_0(v, t_{\text{init}})\| > k_I$, where k_I is an a priori defined constant, and $\mathbf{p}_0(v, t)$, $\mathbf{p}_0(v, t_{\text{init}})$ represent the current and the initial shapes (w.r.t. the model-centered reference frame) of the model $\mathbf{m}_0$ of the apparent contour.

Algorithm: Parametric Composition

Step 1: Determine the interval $\mathcal{I}_0$ of the material coordinate v of the model $\mathbf{m}_0$ in which maximum variation of the shape over time is detected.

Step 2: Perform an eigenvector analysis on the data points that correspond to the area $\mathcal{I}_0$ to approximate the parameters of a new deformable model $\mathbf{m}_1$ which can fit these data points.

Step 3: Construct a composed parametric primitive $\mathbf{m}$ with $\mathbf{m}_0$ and $\mathbf{m}_1$ as defining primitives.

Step 4: Update the list L by replacing the model $\mathbf{m}_0$ with the composed model $\mathbf{m}$.

HBPIIS - Step 3d: The models of the updated list L are continuously fitted to the time-varying data. If the parameters of a composed deformable model indicate that its defining primitives are moving with respect to one another, then this signals the presence of two distinct parts.

Part Decomposition Criterion A: If the generalized coordinates a composed model $\mathbf{m}$ (whose defining primitives are the models $\mathbf{m}_0$ and $\mathbf{m}_1$) satisfy the relation $\|\Delta \mathbf{q}_c(t) - \Delta \mathbf{q}_c(t_{\text{init}})\| > k_{A0} \vee \|\Delta \mathbf{q}_\theta(t) - \Delta \mathbf{q}_\theta(t_{\text{init}})\| > k_{A1}$, where k_{A0} and k_{A1} are two a priori defined constants, $\Delta \mathbf{q}_c(t) = \mathbf{q}_{c1}(t) - \mathbf{q}_{c0}(t)$ and $\Delta \mathbf{q}_\theta(t) = \mathbf{q}_{\theta1}(t) - \mathbf{q}_{\theta0}(t)$ represent the relative translation and orientation of the model-centered coordinate systems of $\mathbf{m}_0$ and $\mathbf{m}_1$, then decompose $\mathbf{m}$.

Algorithm: Part Decomposition A

Step 1: Construct two new models $\mathbf{n}_0$ and $\mathbf{n}_1$ using the parameters of the defining models of the composed model $\mathbf{m}$.

Step 2: Update G and L , by replacing $\mathbf{m}$ with $\mathbf{n}_0$ and $\mathbf{n}_1$.

HBPIIS - Step 3b: The visual event of a hole evolving within the apparent contour which is being tracked indicates that parts which were initially occluded are gradually becoming visible. During this process though, the regions of the initially occluded part that become visible are not contiguous, thus the appearance of the hole. Therefore,

when a hole is evolving within the apparent contour whose shape has changed considerably, we do not invoke the parametric composition algorithm. We monitor the evolution of the hole, and when the hole ceases to exist (and the shape of the model exhibits large deformation, the Part Decomposition B algorithm is invoked.

Part Decomposition Criterion B: If an evolving contour within the apparent contour ceases to exist and $\|\mathbf{p}_0(v, t) - \mathbf{p}_0(v, t_{\text{init}})\| > k_B$, where k_B is an *a priori* defined constant, and $\mathbf{p}_0(v, t)$, $\mathbf{p}_0(v, t_{\text{init}})$ represent the current and the initial shapes (w.r.t. the model-centered reference frame) of the model $\mathbf{m}_0$ of the apparent contour, then invoke the Part Decomposition B algorithm.

Algorithm: Part Decomposition B

Step 1: Determine the interval $\mathcal{I}_0$ of the material coordinate v of the model $\mathbf{m}_0$, in which maximum variation of the shape over time is detected.

Step 2: Determine the interval $\mathcal{I}_1$, $\mathcal{I}_1 \subseteq \mathcal{I}_0$, for which consistent change of the orientation of the finite elements within the area $\mathcal{I}_0$ of the deformable model is detected.

Step 3: Perform an eigenvector analysis on the data points that correspond to the area $\mathcal{I}_1$, to approximate the parameters of a new deformable model $\mathbf{m}_1$ which can fit these data points.

Step 4: Update the list L by replacing the model $\mathbf{m}_0(t)$ with the models $\mathbf{m}_0(t_{\text{init}})$ and $\mathbf{m}_1$.

For step 2 of the HBPIS, we fit the models included in list L to the image data, using our weighted-force assignment algorithm [6] which allows the overlap between models and therefore, the instant rotation center of a part can be determined (see [5]).

2.1 Implementation of the HBPIS

In this section we describe the modeling and the control of the processes involved in the system implementing the HBPIS, using the Discrete Event Systems theory. A discrete event system is a dynamical system in which changes in the state occur at discrete instances of time. The Supervisory Control Theory of Discrete Event Systems [13] allows us to encapsulate both the discrete and the continuous aspects of the fitting and the segmentation processes. In addition, it allows us to investigate the behavior of the participating processes which operate concurrently, to formulate in a compact and elegant way the basic concepts and algorithms for our problem, and to decompose the overall design into components modeled by Finite State Machines (FSM).

In terms of the DES theory, there are three components to the system that implements HBPIS: the *plant*, the *supervisor* and the *observer(s)* (Fig. I(d)). The set of all deformable models that are being fitted to the data at each time instant constitutes the *plant*. The observers are processes that monitor the quantitative changes to the shape of the modeling primitives and send messages to the *supervisor*. The supervisor is the process that controls the behavior of the system, handles the complex interactions

between the fitting processes and invokes the appropriate part composition or decomposition algorithms depending on the input from the observers. Therefore, the observers provide feedback to the supervisor regarding the state of the plant. Each participating process of the system is modeled as a nondeterministic finite state machine. The states of this finite state machine correspond to a particular discretization of the evolution of the task over time. Transitions between these states are caused by discrete events. Some of these events can be expressed in the form (*guard* $\rightarrow$ *operation*), where *guard* is a boolean valued expression. If the value of the *guard* is true, the operation (e.g., send/receive message) is executed. The rest of the events are expressed in the form (*guard*), where the value of *guard* controls the transition between states. By $p_i ? q_i$ a *communication* operation is denoted, in which a process receives a message or command q_i from channel p_i . $p_i ! q_i$ also denotes a communication operation in which a process sends a message or command q_i through channel p_i . In the implementation of the HBPIS, for each of the deformable models used at every time instant to fit the data, there is a corresponding fitting process whose design is shown in figures I(e,f). These processes run in parallel and are synchronized by a global clock. In addition, they receive commands from the supervisor and send messages back. The states of each fitting process correspond to a discretization of the fitting task based on which degrees of freedom of the model are used (e.g., global and/or local parameters). For each fitting process there is an observer process (figures I(g,j)) that monitors the parameters of the deformable model in order to infer the state of the deformable model. Our goal is to be able to drive a deformable model from any state (the error state and the events associated with the error state have been omitted for clarity) in the set $\{0, 1, 2, 3, 4, 5\}$, back to state 0. This is the desirable state since if all the fitted deformable models are in this state, then the structure of the data has been captured. This is because in that state there is no significant change from the initial shape and no evolving holes. To drive each deformable model to the desirable state 0, we design a supervisor. The function of the supervisor (figures I(h,i)) is to select the appropriate action depending on the state of the observer processes and enable, disable or force controllable events in order to ensure correct behavior of the plant.

3 3D Shape Estimation by Integrating Multiple Views

In this section we present how to combine 2D information from three mutually orthogonal views (figure II(30)) in order to estimate the 3D shape of a subject's body parts. Due to the complexity of the shape of the human body and occlusion among the body parts, the subject is requested to perform a different set of movements (protocol movB

in [5]) that incrementally reveals the structure of the body.

Our approach is to first build a single 3D model of the human standing, and then to incrementally refine it by extracting the 3D models of the different parts as they become visible to the different views. At each stage of the algorithm the data from a specific movement are processed. The stages of the algorithm are summarized in figure I(k).

For each movement, the apparent contours from each active view are fitted using the techniques described in the previous section. Therefore, for each active view there is a set of 2D deformable models that fit the corresponding data. Due to occlusion, the number of deformable models in each of the views may not be the same. Consequently, depending on the type of movement a new part or parts may be detected in some of the 2D apparent contours. The new part is either a previously unseen body part or a subpart of a part whose 3D model has already been recovered. In the first case, the 2D models of the corresponding apparent contours of a part are integrated using the algorithm described below (Fig. II(42)). In the second case, the model of the apparent contour of the subpart is intersected with the model for the 3D shape of the part to obtain two new 3D shapes, the 3D shape of the sub-part and the 3D shape of the rest of the part. Due to the fact that more than one parts may be simultaneously determined in 2D (e.g., left and right arms), the extracted spatial location of the part in each of the two views is employed to determine the apparent contour correspondence. It should be noted also that the 3D shape of the new part is obtained at the end of the appropriate set of movements. This is a desirable property, since the shape fitting is done in 2D and the 3D shape estimation is done only once.

In the following, the algorithm for the integration of the 2D models of two apparent contours is described. The input to the algorithm is the 2D models of the apparent contours of the part as observed from two mutually orthogonal views, and the spatial relation between the views. A 3D deformable model is initialized and the nodes that lie on its meridians whose plane is parallel to the planes of the apparent contours are fitted to the nodes of the 2D models. Local deformations are employed to capture the exact shape of the 2D models, while the rest of the 3D shape is interpolated. Due to the fact that the local deformations may be large in magnitude and to avoid shape discontinuities between the fitted nodes and the rest of the shape, a thin-plate deformation energy is imposed during the fitting process. Therefore, the deformable model fits accurately the 2D models, while its shape in between the apparent contours (in the absence of any data) varies smoothly.

4 Experimental Results

This section presents selected results from a series of experiments (a detailed description is offered in [5]) that

were carried out on real image sequences in order to evaluate the effectiveness of the proposed approach in general and of the specific algorithms in particular.

The first experiment was designed to evaluate the performances of the Part Composition Invocation Criterion, and of the Part Decomposition A algorithm. In this experiment, the subject lifted her left arm, until the arm reached the horizontal position (figures II(1-3)). In figure II(5) notice the deformation of the initial shape of the model (figure II(4)), which is quantified in figure II(6). The parametric composition invocation criterion was satisfied and therefore the Parametric Composition algorithm was invoked. The defining primitives of the composed deformable model are shown in figure II(7). In a later frame, the Part Decomposition Criterion A was satisfied and the Part Decomposition A algorithm was invoked in order to recover the underlying parts. Figure II(8) shows the recovered models for the parts, and their fitting to the data in figure II(3). The second experiment was designed to demonstrate the Part Decomposition criterion B and Part Decomposition B algorithm. In this experiment, the subject extended her left leg the front until she reached a comfortable position (figures II(9-12)). While fitting the data in figure II(11,15) the Part Decomposition Criterion B was satisfied and therefore the Part Decomposition B algorithm was invoked. The two models recovered are depicted in figure II(16), and their fitting to the data is shown in figure II(12). The purpose of the third experiment was to demonstrate the results of the Human Body Part Identification Strategy when the subject moves according to the protocol MovA. Figures II(18-21,24-29) depict some of the models extracted during the motion. Finally, the fourth experiment was designed to demonstrate the feasibility of estimating the 3D shape of body parts by selectively integrating 2D outlines from multiple views (figure II(30)). Figure II(31) shows the model recovered for the left arm, while figures II(32-35) show the models for the hand, lower arm and upper arm and their layout. Figure II(36) shows the recovered model of the torso and figures II(37-41) the corresponding parts of the left leg.

5 Conclusion

This paper presented a novel integrated approach to the 3D shape model acquisition of a subject's body parts from multiple views. The extraction of a subject's model can be considered as the first part of a tracking process. Our long term goal is to provide an integrated framework for tracking the body parts of a human moving in 3D which copes with occlusion between body parts. The experiments that were carried out have proven that it is possible to determine the parts of a human body from an image sequence and estimate their shape in 3D.

Figure I: An example of composition of two superellipsoids: (a) depicts $\mathbf{x}_1(v_1)$ which intersects $\mathbf{x}_0(v_0)$ at points A and B , (b) depicts their composition $\mathbf{x}(v)$, (c) depicts the composition function $\delta(v; 0.65\pi, 0.77\pi, 10)$. (d-j) explain the FSM of the components of the implementation of HBPIS. (k) 3D shape estimation (IG stands for the Integration algorithm and IN for the Intersection algorithm).

Figure II: (1-29) refer to 2D Human body model acquisition, (30) depicts the placement of the cameras w.r.t the subject, (31-41) refer to the 3D shape estimation of body parts, and (42) depicts the integration algorithm.

References

- [1] K. Akita. Image sequence analysis of real world human motion. *Pattern Recognition*, 17:73–83, 1984.
- [2] C. I. Attwood, G. D. Sullivan, and K. D. Baker. Model based recognition of human posture using single synthetic images. In *Proceedings of the Fifth Alvey Vision Conference*, pages 25–30, Reading, UK, 1989. University of Reading.
- [3] Z. Chen and H. J. Lee. Knowledge-guided visual perception of 3d human gait from single image sequence. *IEEE Transactions on Systems, Man and Cybernetics*, 22(2):336–342, 1992.
- [4] D. Hogg. Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.
- [5] I. A. Kakadiaris and D. Metaxas. 3D Human body model acquisition from multiple views. Technical Report MS-CIS-95-17, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, April 1995.
- [6] I. A. Kakadiaris, D. Metaxas, and R. Bajcsy. Active part-decomposition, shape and motion estimation of articulated objects: A physics-based approach. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 980–984, Seattle, WA, June 21–23 1994.
- [7] S. Kurakake and R. Nevatia. Description and tracking of moving articulated objects. In *Proceedings of the IEEE International Conference on Pattern Recognition*, pages 491–495, 1992.
- [8] M. K. Leung and Y. H. Yang. An empirical approach to human body analysis. Technical Report 94-1, University of Saskatchewan, Saskatchewan, Canada, 1994.
- [9] D. Metaxas and D. Terzopoulos. Shape and nonrigid motion estimation through physics-based synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):580–591, June 1993.
- [10] J. O’Rourke and N. I. Badler. Model-based image analysis of human motion using constraint propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(6):522–536, 1980.
- [11] A. Pentland. Automatic extraction of deformable part models. *International Journal on Computer Vision*, 4:107–126, 1990.
- [12] R. J. Quian and T. S. Huang. Motion analysis of articulated objects. In *Image Understanding Workshop*, pages 549–553, San Diego, CA, January 1992.
- [13] P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE*, 77(1):81–97, January 1989.
- [14] J. M. Rehg and T. Kanade. Visual tracking of high DOF articulated structures: An application to human hand tracking. In Jan-Olof Eklundh, editor, *Third European Conference on Computer Vision*, pages 35–46, Stockholm, Sweden, 2–6 May 1994. Springer-Verlag.
- [15] K. Rohr. Towards model-based recognition of human movements in image sequences. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 59(1):94–115, January 1994.
- [16] J. A. Webb and J. K. Aggarwal. Structure from motion of rigid and joined objects. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 686–691, 1981.
- [17] M. Yamamoto and K. Koshikawa. Human motion analysis based on a robot arm model. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 664–665, Hawaii, June 3–6 1991.