

Approximating Human Reaching Volumes Using Inverse Kinematics

I. Rodríguez^a, M. Peinado^a, R. Boulic^b, D. Meziat^a

inma@aut.uah.es, manupg@aut.uah.es, ronan.boulic@epfl.ch, meziat@aut.uah.es

^aComputer Engineering Department, University of Alcalá
Campus Universitario. Escuela Politécnica. Alcalá de Henares, Madrid, Spain

^bVirtual Reality Laboratory. Swiss Federal Institute of Technology.
Lausanne, Switzerland

Abstract

This paper presents a system to analyse the reaching capabilities of the human body. Our research is motivated by the need for a system which helps the user to manage data relative to the most frequent tasks involved in human activity: reaching tasks. Depending on the type of reaching that we are dealing with we establish a set of constraints that characterise the reaching task. We have designed a data structure that let us approximate reachable volumes. A set of reaching strategies is defined, and reachable volumes are generated and stored for each of these strategies. In addition, an automatic strategy-selection process decides, in real time, which strategy is the most suitable for a certain reaching task. To demonstrate the usefulness of our system, we show how it can be applied to a study of seated reaching.

1. Introduction

Virtual Humans are a valuable medium for gaining knowledge and understanding about the human body. In several fields such as biomechanics or medicine, computer animation systems are an important help to the research on topics related with humans. Handicapped persons, ancient or injured people may have limited movements that don't let them to perform usual tasks in a normal way. Computer systems would allow to modify human body parameters and would let researchers conduct simulations of different study cases [Eui S. Jung, 92]. For example, the reachable volume for a person with some physical impairment should be different than the one corresponding to a person without it.

This paper presents a reachability analysis system for the human body. Our research is motivated due to the necessity of systems which help to manage data relative to the most frequent tasks involved in human activity, reaching tasks. Our main objective is to model near reach, in which the person doesn't have to step toward the target. We consider two varieties of reach: *standing reach*, which is useful in computer animation where virtual humans have to interact with virtual environments, and *seated reach*, which is important in ergonomics studies where the workplace evaluation assesses the reachability of all the workspace.

Our approach is directed toward hand reaching, but it could easily be adapted for foot, knee or head reaching. Think about the situation where a person is carrying a load with both hands and he is not able to close a door, and thus he needs his feet, knee or even elbow to close the door.

The system has been conceived in two steps. In the first one, the Inverse Kinematics technique and an efficient data structure are utilised for the generation and storage of the reachable space. After having stored these data, we exploit them in a method that selects the most adequate strategy among the available ones. The strategy depends on the position of the target to reach and on the body part to be used (i.e., the left hand).

This paper is organised in 6 sections. Section 2 is a brief introduction to the state of the art and to the related work in the field of computer reaching analysis. In section 3 we describe some of the techniques we have employed to study reach. These techniques yield some results that are used in the strategy selection algorithm presented in section 4. Finally, section 5 shows how our system is applied to the research of seated reach, and section 6 contains some conclusions and guidelines for future work.

2. Related Work

Some authors have tried to shed light on the brain foundations of goal-directed movements. Jennerod, for instance, presents a detailed investigation of the subject, integrating psychology and physiology [Jennerod, 88]. Following a different approach, the work in [Hestenes, 94] introduces an algebraic method for formulating and analysing the kinematics of reaching. A general solution is provided that parametrizes the joint variables in terms of the wrist position. Only the motion of the arm is contemplated, but the authors claim that their technique should generalise easily to model the entire skeleto-muscular system.

Seated reaching has been an area of relatively active research. For instance, a recent study undertakes in-vehicle dynamic posture prediction, proposing a method based on differential inverse kinematics [Zhang et Chaffin, 00]. Its results lead to implications with regard to the underlying control strategies of human reaching movements.

Other approaches to reaching use interpolation synthesis to produce movement from a mixture of prerecorded data motions. Wiley and Hahn specifically followed this approach. They solved the problem of obtaining natural reach postures by interpolating motion captured data [Wiley et Hahn, 97].

Few researches are dedicated to the real-time computer selection of suitable reaching strategies. We can cite a study that distinguished three reach areas depending on the distance from the hand to the target [Mas et al, 97]. The author defined different strategies for each area taking into account factors like the need to control the center of mass or the need for additional supports.

3. Mechanisms to work on reachable volumes

3.1. Volume Generation and storing: Inverse Kinematics and Octal Trees

We have designed a data structure that let us approximate a certain volume. It is called Volume Approximation Tree (VATree). A simple and fast algorithm tests whether a given point lies inside the approximated volume. The proposed data structure is an octal tree (nodes have eight children). Octal trees (usually referred to as octrees) are commonly used for representing volumes or surfaces [Libes, 91]. A more detailed description of this data structure is out of the scope of this paper [Rodriguez et al, 03].

A set of constraints has to be defined in order to generate the reachable volume corresponding to each strategy. Table 1 shows the set of constraints that define three different strategies. The performance of a standing reaching needs some forward bending of the trunk, thus requiring postural stability provided by a controlled center of mass. Other constraints are a positional constraint applied to the body part which does the reach, and a look-at constraint that makes the subject gaze at the target. All these constraints define what we call a *direct near reach*. If the goal is so low that the subject needs to crouch, in addition to those constraints previously mentioned, flexing legs and a change on the root of motion are also required. Finally, for seated reach, in which the person has to reach an object starting from a seated posture, the center of mass needs not be controlled, but it is necessary to change the location of the motion flow root.

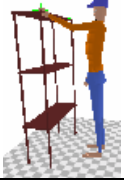


SET OF CONSTRAINTS	STRATEGY
Center of mass control Positional constraint (i.e hand) Look-at constraint Motion flow root (pelvis)	STANDING REACH: DIRECT 
Center of mass control Position constraint (i.e hand) Looking constraint Flexing legs Change root of motion (i.e foot)	STANDING REACH: CROUCH 
Position constraint (i.e hand) Looking constraint Change root of motion (i.e foot)	SEATED REACH 

Table 1. Constraints that define reaching strategies

As shown in Table 1, depending on the type of reaching that we want to generate, we establish a set of constraints that characterise the task. The distance between the virtual human and the volume to approximate determines a near or far reaching. The height of the volume with respect to the virtual human also determines differences in the sort of reaching. A very low volume situated at feet level will allow studying crouch reach. Note that during the generation of reachable spaces, we don't need visual feedback; we obtain data about reachability that are stored and will be exploited in a posterior phase.

Figure 1 describes the process of labelling and storing a reachable volume. Given the set of constraints that define a strategy, a mechanism queries the Inverse Kinematics engine, asking for reachability. Inverse Kinematics (IK) is a technique to create animations of articulated structures [Baerlocher, 98]. An initial voxel is specified and the reachability query is applied to its eight vertices. The IK engine replies telling whether the voxel is reachable or unreachable. We say that a voxel is reachable when the strategy is adequate to reach all the points in its interior. An unreachable voxel, on the other hand, is entirely made of unreachable points. When a voxel has a mixture of reachable and unreachable parts, it is divided into eight new, smaller child voxels and the same process is applied to each of them.

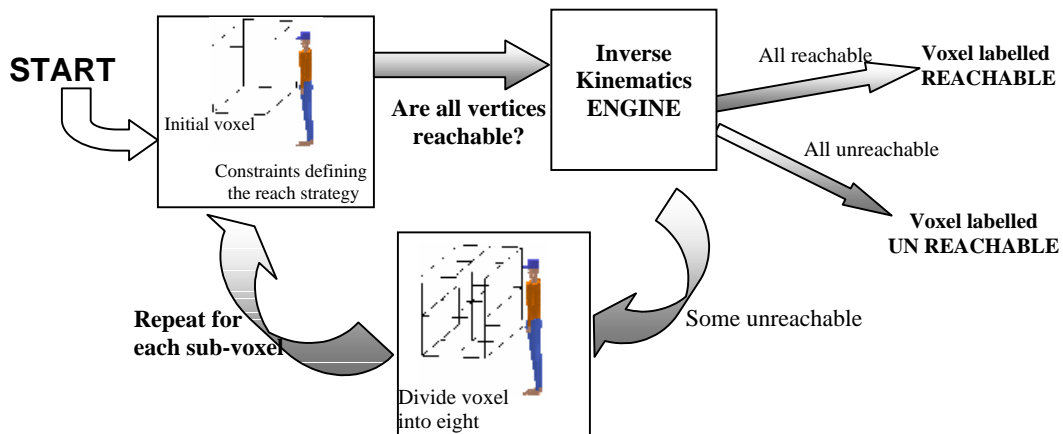


Figure 1. Computation of a reachable volume

3.2. Exploitation: High Level Inverse Kinematics

Once reachable volumes for the different implemented strategies have been computed and stored, a high level Inverse Kinematics module is provided that processes that reachability data in order to create reaching tasks. Figure 2 depicts the interaction between the Inverse Kinematics module used to create reaching animations and the reachability module that determines which strategy is the most appropriate for a particular reaching task. A 3D point, i.e, the target, and a body member involved in the reaching task are enough to obtain a response from the reachability module saying if there is a strategy adequate to reach the target. A simple mechanism of priorities is established so that if the goal is unreachable with one strategy, another one can be tried.

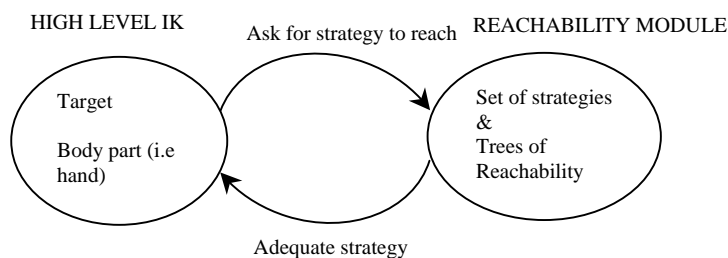


Figure 2. Interaction between the Inverse Kinematics and Reachability modules

4. Strategy Selection

As shown earlier, our system consists of two stages: in the first several reaching strategies are created, and a reachable volume is generated and stored for each of them. This volume represents the portion of space that is reachable using the strategy.

In a second stage, our system can decide which from the set of available strategies is most appropriate for a certain reaching task. This decision is made based on the information generated and stored in the first stage (i.e., the reachable volumes of each strategy). However, this information is not sufficient, since the reachable volumes of different strategies might have regions of intersection. Goals located in such regions would be reachable using different strategies, so a mechanism is needed to select one of those “candidate” strategies.

We solve this problem by classifying the strategies according to a priority scheme: each strategy is given a unique priority, and when two or more strategies are valid for a reaching task, the one with the highest priority is chosen. We assign higher priorities to those strategies that have a lower “cost” for the subject. The concept of cost is difficult to define precisely in this context. In general, we assign priorities by considering that a task has a greater cost when it attaches more constraints to the inverse kinematics engine. In addition, other factors that increase the cost assigned to a strategy, thus decreasing its priority, are the fraction of body mass it requires moving and whether it needs control of the center of mass in order to maintain balance.

Another issue we have taken into account when assigning priorities is *naturalness*. When two strategies has the same cost, according to the definition of cost stated above, we assign a higher priority to the strategy that appears more natural, i.e., to the one that would be adopted by a human in normal circumstances. For instance, both the strategies *direct reaching* and *tip-toe reaching* have similar costs, for both of them affect the whole body, need control of the center of mass and use the same amount of IK constraints. Nevertheless, in those goals for which both strategies make the reaching possible, it is apparent that the adoption of a tip-toe strategy would seem less natural, and therefore a lower priority is assigned to this strategy than to the direct one.

Once we have developed a way to hierarchically classify the different strategies, the process of selecting the most suitable strategy for a given reaching task is relatively straightforward: starting from the strategy with the highest priority, a query is sent to the strategy to find whether it is valid to reach the goal. If the answer is affirmative, the selection process ends and this strategy is chosen. Otherwise, the strategy with the next priority is considered, and the process starts again. If the end of the list is hit, and no strategy has been found that achieves the goal, the process concludes and the goal is considered unreachable.

It is worth mentioning that, since the strategies are arranged in order of a priority scheme, the selected strategy for a given reaching task will always be the most natural, or in other words, the one most likely to be adopted by a person under the same conditions.

5. Results

The main emphasis of our work has been on two specific varieties of reach: *standing reach* and *seated reach*. In both of them the goal is reached with one of the hands. The difference lies in the initial posture: in standing reaching the motion starts with the subject standing upright, while in seated reaching the subject is seated on a chair. Many strategies exist that can be applied to both kinds of reach. In this paper we will focus on seated reach. For a more detailed analysis of standing reach, see [Rodriguez et al, 03].

The strategies we have included in the study of seated reaching are the following:

Naive sit. In this strategy motion is limited to the upper body, not including the pelvis (which remains fixed). This is the only strategy that can be implemented in IK systems where the motion flow root can lay nowhere but at the pelvis.

Normal sit. This strategy represents a major improvement over the *naive* strategy, since the pelvis is also involved in the reaching motion (i.e. motion exploits the pelvic joint to lean forward).

Sit & rise. The two strategies above may not be sufficient for distant goals. For such cases we have devised a strategy, inspired by experimental observation, by means of which the subject can lift slightly off the chair when needed.

Next we will show how the strategies introduced above behave in practice. Figure 3 shows the resulting postures of performing a reaching task using the *naive* (left) and *normal* (right) strategies. In the first one a certain rigidness can be observed, which leads to an unnatural stretch of the arm. With the *normal* strategy, on the other hand, the pelvis takes part in the overall body movement, thus permitting a reach posture that is closer to what an individual would adopt.

Let us refer now to the two snapshots in Figure 3.a. They show the result of attempting the reaching of a goal situated in a relatively high and distant position. In the left image the *normal* strategy is being used, and the subject is unable to reach the goal. Faced this situation, most individuals would attempt to rise their buttocks slightly off the chair. This is the exact behaviour of the *sit & rise* strategy, as shown in the left image of Figure 3.b. It can be observed that this subtle lift is enough to make the reaching possible in this particular case.




Figure 3. (a) Naive sit strategy (left) vs Normal sit strategy
 (b) Normal sit strategy (left) vs Sit & rise strategy

A convenient way to evaluate the feasibility of a strategy is to visualize its reachable volume. Besides, this technique allows us to compare several strategies in order to see which one will be the most appropriate to reach a certain region of space. The reachable volumes of the three strategies we have seen so far are shown in Figure 4 (all results shown are for left hand reaching). With regard to the sit & rise strategy, this figure shows that this strategy permits to reach farther goals, especially in those regions located over the subject's head.

As can also be seen in this Figure 4, the *naive* strategy not only yields less natural postures, as stated before, but has also a more limited reach when compared with the other two. Therefore, in our final system only two strategies are integrated: *normal sit* for goals situated relatively near, and *sit & rise* for more distant goals. Which strategy to use, given a certain goal, is decided through a process like the one shown in Section 4.




Figure 4. Reachable volumes for the different seated reaching strategies.

6. Conclusions and Future Work

We have introduced a framework that serves as a valuable tool for modelling and analysing reaching in virtual humans. In addition, it allows for the real time selection of the most suitable reaching strategy from a set of predefined ones. The selection is made in order to a priority scheme that seeks naturalness in the resulting posture. This framework has been employed successfully to design and test more realistic reaching strategies, as well as to produce realistic reaching animations in a simple way.

Nevertheless, there is still room for improvement in several areas. For instance, an issue that deserves to be looked into in more depth is the research of new strategies and the refinement of the existing ones. This way we plan to obtain postures that resemble more closely the behaviour of persons. A strategy we are currently working on consists of a torsion of the upper body in order to reach goals situated on the sides of the subject. We plan to apply this strategy both to standing and seated reach.

Another issue that needs to be dealt with is self-interference detection. At the moment we have a method that lets us detect which postures are invalid for containing collisions between the upper limbs and the torso of the subject. In most cases these collisions involve the opposite arm to the reaching one, so we believe we can solve

them, once detected, by adding an additional positional constraint to the IK engine, as described in [Boulic et al, 97].

Although we have focused on the generation of the final reach posture, a question that needs to be answered is how to produce the intermediate frames in a reaching animation, i.e., how to obtain the sequence of frames that smoothly takes the subject from the initial posture to the final one. Our current approach is simply to generate the animation from a new IK simulation. This approach has a couple of important drawbacks: first, its cost can be excessive, since IK computations tend to be very expensive. The second disadvantage is that it can yield unrealistic results, since no control can be exercised over the resulting reaching motion. Such control would be needed, for instance, to create a curvilinear reach trajectory such as described in [Faraway, 01], or a bell-shaped velocity profile stated by Fitt's Law.

We think we can get this kind of effect by using a motion interpolation scheme to generate intermediate frames. This will be much faster than performing an IK simulation, and its only drawback will be the need to store all the final postures in the reachability tree of each strategy. Note that once this new method is put into practice, our framework will work at two different levels: first, IK will be used to compute the reachability trees and the reach postures for each vertex of the tree; then, at the animation level, motion interpolation will provide the intermediate frames of the definite reaching motion.

References

[Eui S. Jung, 92] Eui S. Jung, Dohyung Kee and Min K. Chung, *Reach Posture Prediction of Upper Limb for Ergonomic Workspace Evaluation*, Proc. of the Human Factors Soc.-36th Ann.M., 702-706, 1992.

[Baerlocher et Boulic, 98] Baerlocher P., Boulic R., *Task-priority Formulations for the Kinematic Control of Highly Redundant Articulated Structures*, Proc. of IROS'98, Victoria, Canada, Oct. 1998 .

[Rodriguez et al, 03] Rodriguez I, Peinado M., Boulic R., Meziat D., *Reaching Volumes Generated by Means of Octal Trees and Cartesian Constraints*. Computer Graphics International. Japan. To appear.

[Mas et al, 97] Mas, R., Boulic, R., Thalmann, D., *Extended Grasping Behaviour for Autonomous Humans Agents*. First ACM Conference on Autonomous Agents. Los Angeles, 1997.

[Boulic et al, 97] Boulic, R., Huang, Z., Thalmann, D., *A Comparison of Design Strategies for 3D Human Motions*, included in Human Comfort and Security of Information Systems; Advanced Interface for the Information Society, K. Varghese and S. Pflieger (Eds), Research Report ESPRIT, ISBN 3-540-62067-2, Springer Verlag, 1997.

[Zhang et Chaffin, 00] Zhang, X. and Chaffin, D.B., *A Three-Dimensional Dynamic Posture Prediction Model for Simulating In-Vehicle Seated Reaching Movement: Development and Validation*, Ergonomics, 43:9, 1314-1330, 2000.

[Hestenes, 94] Hestenes, D. *Invariant Body Kinematics: II. Reaching and Neurogeometry*, Neural Networks, Vol. 7 No. 1, 79-88.

[Jennerod, 88] Jennerod, M. *The Neural and Behavioural Organization of Goal-directed Movements*, ASIN 0198521960, Oxford University Press., 1991.

[Wiley et Hahn, 97] Wiley, D.J. and Hahn, J.K. *Interpolation Synthesis of Articulated Figure Motion*, IEEE Computer Graphics and Applications, Vol. 17, No. 6, Nov./Dev. 1997.

[Libes, 91] Libes, D. *Modeling Dynamic Surfaces with Octrees*, Computer & Graphics Magazine, Vol. 15, No. 3, 1991.

[Faraway, 01] Faraway, J.J., *Modeling Hand Trajectories During Reaching Motion*. TR-383. Department of Statistics. University of Michigan, 2001.