

Characterizing full-body reach duration across task and viewpoint modalities

Damien Maupu, Ronan Boulic, and Daniel Thalmann

VRLab

Faculté Informatique & Communications

Ecole Polytechnique Fédérale de Lausanne (EPFL)

EPFL IC ISIM VRLAB, Station 14, 1015 Lausanne, Switzerland

email: {damien.maupu, ronan.boulic, daniel.thalmann}@epfl.ch

www: <http://vrlab.epfl.ch>

Abstract

The full-body control of virtual characters is a promising technique for application fields such as Virtual Prototyping. However it is important to assess to what extent the user full-body behavior is modified when immersed in a virtual environment. In the present study we have measured reach durations for two types of task (controlling a simple rigid shape vs. a virtual character) and two types of viewpoint (1st person vs. 3rd person). The paper first describes the architecture of the motion capture approach retained for the on-line full-body reach experiment.

We then present reach measurement results performed in a non-virtual environment. They show that the target height parameter leads to reach duration variation of $\pm 25\%$ around the average duration for the highest and lowest targets. This characteristic is highly accentuated in the virtual world as analyzed in the discussion section. In particular, the discrepancy observed for the first person viewpoint modality sug-

gests to adopt a third person viewpoint when controlling the posture of a virtual character in a virtual environment.

Keywords: motion capture, real-time interaction, virtual prototyping, reaching tasks

1 Introduction

Interactive control of a virtual character, or avatar, has applications in fields like virtual prototyping, workspace design, and computer animated puppetry or for interacting with general purpose virtual environments. Traditional interaction devices, such as a mouse, are of limited usability in such a context because they provide too few degrees of freedom at a time. While they are perfect for defining high level parameters to control an avatar (e.g. walking or running direction), they are inadequate to interactively define simultaneously the numerous parameters of an arbitrary posture. An enactive way to control avatars would be to use one's ability to execute full body movement. Full-body postural input has been so far limited to off-line motion capture (mocap) or to virtual puppetry where interactions with the environment are limited. However recent advances in technology make motion capture mature enough for on-line precise-tasks full-body control of virtual character (tasks such as reaching).

We propose an approach to reconstruct on-line full-body (no hand or facial animation) movements using prioritized inverse kinematics (Section 3). Our solution is easy to calibrate, use a reduced set of markers

Digital Peer Publishing Licence

Any party may pass on this Work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via Internet at <http://www.dipp.nrw.de/>.

First presented at the 4th International Conference on Enactive Interfaces 2007, ENACTIVE07 extended and revised for JVRB

and is able to recover a large panel of movements from reach to locomotion.

We present then two studies that we have conducted to study human full-body reach in the real world (Section 4) and in the virtual world under different contexts (various viewpoints and various graphical representations in the virtual world) (Section 5). We want to assess the influence of controlling an avatar on reaching tasks. Similarities and discrepancies are discussed in Section 6.

2 Related work

Motion capture systems provide the information of position and/or orientation of sensors that can be further used to recover the posture of a human performer. In the framework of skeletal based animation, a virtual human is composed of a skeleton (i.e. a hierarchy of joints starting at a root) and a deformable skin (i.e. a mesh) that follows the underlying skeleton. Therefore, an avatar posture is defined by a vector of joint state where all joints store a local orientation and the root joint stores also the position of the skeleton root to control its global position. One key problem faced by motion capture is to determine the minimum number of sensors that allow mapping transparently the measured data into joints orientation.

Most often, the set of sensors is chosen so that combining them allows to unambiguously recover the local orientation of all skeleton rigid segments [Men99] (i.e. three position sensors per rigid segment). Such an approach is still in use in most motion capture studios working with optical systems, either using passive or active markers technology delivering 3D position data only. The main drawback is that the markers placement takes a long time. Indeed, about forty optical markers have to be placed very precisely as explained in [Uni08]. Badler et al. [BHG93] have explored ways to reduce the number of sensors. They present a way to drive an avatar torso using inverse kinematics and a set of four magnetic sensors delivering position and orientation data. Molet et al. [MBT99] proposed an analytic method to robustly distribute magnetic sensors orientation data over multiple joints. Chai et al. [CH05] recently introduced a motion capture technique employing video cameras and a small set of markers that interrogate a database of prerecorded movements and output the most suitable one. The main limitation is the restriction to such a database.

Prioritized Inverse Kinematics (PIK) has been used

in the past to do (desktop) interactive posture edition [BB04]. This approach relies on a linearization of the system that may introduce some artifact as discussed in [BPC05]. Nevertheless Peinado et al. [PHW⁺04] have successfully recovered a clarinet musician performance from a very partial set of six markers with the help of additional constraints on the center of mass, the hands (linked to the instrument), the head (by defining an angle with the clarinet) and the feet. The fact that the musician was not moving the feet was exploited to define the highest priority constraint to ensure the permanent contact with the floor. A qualitative analysis of this output has revealed how to improve the balance constraint to produce a more believable motion by allowing a limited swaying of the center of mass both in the sagittal and the lateral planes [MPB⁺06]. In this paper, the motion is reconstructed from less a-priori knowledge on the user motions. Therefore, the space of possible movements increases.

Besides, Boulic et al. have explored a video-based approach for on-line motion capture of the upper body posture [BVU⁺06]. However the nature of the input data still prevents its use for studying full body movements in particular reaching tasks where the hands may get occluded.

As a consequence we propose in this paper a methodology to synthesize transparently and on-line a large panel of full-body motion relying on a reduced set of active optical markers using PIK. The proposed sensor set-up allows free movements including steps.

3 On-line virtual mannequin control

3.1 Prioritized Inverse Kinematics

Prioritized Inverse Kinematics (PIK) allows constraints (i.e. effectors) to be associated with a priority level so that important properties are enforced first (e.g. feet stay on the ground) while less important adjustments are made in the remaining solution space (for more details, please refer to [BB04]). The overall algorithm works as follow (Figure 1). First effectors are defined: type (i.e. constraints in position or in orientation), parent joint, set of recruited joints, priority and goal. Then in a real-time loop, motion capture data is acquired, in order to feed the effectors goals. The PIK solver finds an optimal posture variation to progress toward minimizing effector errors according to their relative priority. Finally the virtual human is rendered using the graphic engine Mvisio [PVT08].

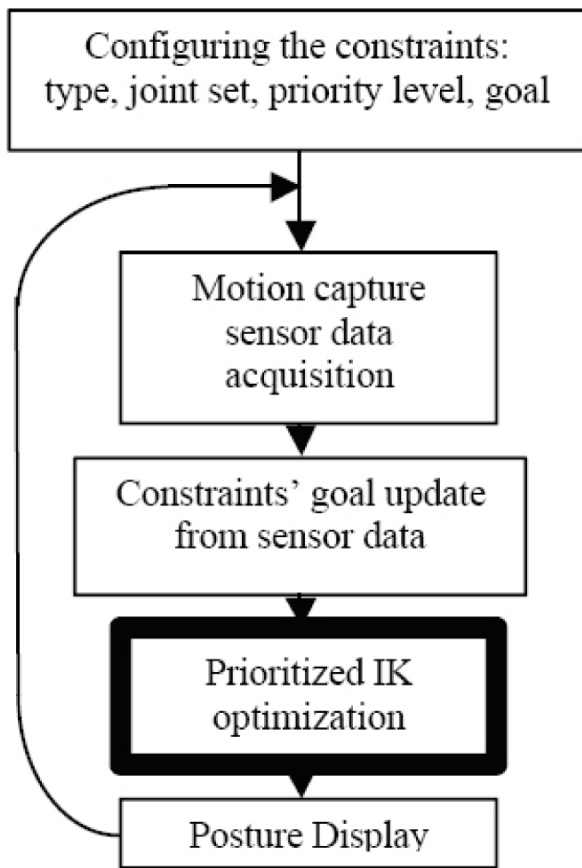


Figure 1: Exploiting Prioritized Inverse Kinematics within an on-line motion capture loop

Nearly each marker is used to feed a corresponding “position” effector that is located at the same place on the virtual skeleton. In some cases a few markers are grouped to control the orientation of a body part.

Our mocap approach does not request markers to be carefully placed. A calibration phase similar to the one proposed by Molet et al. [MBT99] is performed once at initialization time. During that stage the performer has to adopt the H-anim calibration posture (Figure 2, right: with the arms along the body) where all joint states are the identity [HA08]. In that state each effector and its associated marker coincide. The calibration phase allows measuring the offset distance between the effector and its parent joint in the virtual skeleton. All offsets remain constant during the on-line interaction.

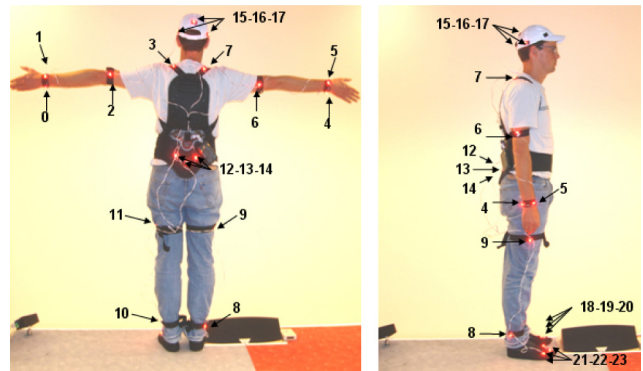


Figure 2: Setup of active markers

3.2 Markers & effectors set up

Although our solution could be adapted to work with any particular mocap technology, in our current implementation we use an active optical motion capture system (8 cameras) from Phasespace [Pha08]: markers are infra-red LEDs with constant identification number.

Our mocap set up (resulting from several tests) consists of 24 markers that define the goal for 19 effectors: thirteen in position control and six in orientation control.

Eight markers drive the arms. The two on the wrist help to recover the forearm orientation but does not allow recovering the wrist state (the hand is rigidly linked to the forearm as can be seen in Figure 3 top line). Only the two markers closest to the clavicles are exploited to recover shoulder shrugging movements. The spine is controlled by LEDs on the clavicles and on the spine base. Controlling the spine base orientation with a group of three markers prevents producing unrealistic curvature of the back and helps the avatar to adopt flexed knees postures. The subset of 4 lumbar and thoracic vertebrae is coupled to enforce the spine anatomic behavior [RB07]. Legs are controlled with ten markers. The constraints on the feet have the most important priority whenever they are in contact with the floor. Otherwise, their priority is decreased on the fly during the on-line performance. This is detected when the foot markers are above a predefined height threshold. Such an approach reduces foot sliding. The head is controlled only in orientation. Joints limits inequality constraints are activated. In case a marker is occluded, the last known position is kept.

Figure 2 shows the user equipped with the markers. Table 1 sums up the effectors’ attributes.

3.3 Results

Figure 3 presents some results showing the large panel of possible postures our system is able to reconstruct on-line. Each iteration of inverse kinematics cost about 6 to 9 milliseconds on a 3.2 GHz Intel Xeon processor with 1 GB RAM.

3.4 Exploring convergence speedup possibilities

In the overall, the system offers a sufficiently fast convergence of the postural control so that the user feels the avatar's posture consistently reflects his or her posture over the on-line interaction.

However, there is a limitation in the frequency bandwidth of the movements. For example, if the user swings his arms fast, the avatar will not be able to track those movements accurately. This is mainly due to the local nature of the PIK solution which incrementally converges towards an optimal solution. If the position error is too large; the solver does not have sufficient iterations per displayed posture to converge towards the optimal solution. This phenomenon can be seen as a kind of low-pass filtering but is mainly not perceptible for most actions in particular for reaching tasks.

To be able to do more iterations, we have ex-




Figure 3: A large panel of reach postures can be reconstructed on-line by the system

Marker n°	Constraint type	Controlled body part / recruited joint set	Priority rank
18-19-20 21-22-23	Position & Orientation	both toes / recruit joints until Root	1
12	Position	spine base / only Root	2
12-13-14	Orientation	spine base / only Root	3
0-4	Position	both wrists / elbow, shoulder & clavicle	3
0-1 4-5	Orientation	both wrists / only elbow twist DOF is recruited	4
2-6	Position	both shoulders / only clavicles	5
3-7	Position	both clavicles / recruit joints until Root	5
9-11	Position	both knees / recruit joints until Root	6
8-10	Position	both ankles / recruit joints until Root	7
15-16-17	Orientation	head / recruit only cervical joints	8

Table 1: Sum up of the set of constraints

plored a way to reduce the number of controlled variables in position control. Usually, a position control is performed in 3D to ensure correct convergence, especially when close to the target. However, when the distance from the effector to the goal is higher than a threshold, we have experimented to reduce the controlled dimension to 1D along the effector→goal direction. The computing cost being linear with the number of controlled dimensions, the gain can be significant if this happens simultaneously to multiple position effectors, e.g. during fast user movements. Although we did observe some better convergence

results, it was not as significant as expected over a large range of users' activities ; it appears that the effector subset that benefit from the dimension reduction remain small in proportion of all other effectors. Also, switching the controlled dimension brings some perceptible movement discontinuities. For this reason we prefer to use this improvement only for off-line movement reconstruction rather than on-line interaction for which we favor intuitive and transparent postural control.

Another frequent improvement of the Prioritized IK convergence is to attract the current posture towards a preferred posture through the lowest priority task [BB04][PHW⁺04][BPC05]. This is motivated by the fact that the IK solution is local and the controlled system is redundant, hence after a few minutes of interaction an unnatural postural drift may be observed [RB07]. Attracting the posture towards a preferred posture solves this issue if the remaining solution space is large enough. This approach makes sense whenever there is a small number of effectors (i.e. less than 5) because the solution may converge towards mathematically optimal but unnatural postures [PHW⁺04]. However in the present context the set of effectors is large (nineteen), hence the dimension of the remaining solution space for the posture attraction is really low. No differences were observed between an avatar attracted to a reference posture and one that is not. On the positive side, the large number of effectors ensures that no postural drift is observed either; therefore such lowest level improvement is not necessary and the related computation cost can be spared.

4 Reach behavior in the real world

The first experience we conducted aim to study human full-body reach in the real world.

4.1 Method

4.1.1 Participants

Seven male students aged from 25 to 30, measuring 1.68m to 1.91m participated to the study. None had counter-indication for standing-up over the duration of the study.

4.1.2 Apparatus

Let H be the height of the subject. During the experiment, each subject was standing at a distance $H/3$ to a 2.4m large screen on which a 0.1m diameter white target was displayed (Figure 4). The height of the target was varying from $0.2H$ to $1.1H$ by increment of $0.1H$. The subject was holding a joystick device with both hands in order to ensure a clear starting and intermediate posture (both hands on the abdomen), to ensure a symmetry in the reach posture and to allow the subject to signal the end of the task through a button action.



Figure 4: Experimental setup for the study in the real world. The subject is at a distance $H/3$ from the screen and must reach targets at height going from $0.2H$ to $1.1H$. H being the subject height.

4.1.3 Procedure

Each subject has to go through series of 20 reaches (each target's height was presented twice in a random order). He was free to flex the legs or to reach the target on their toes but feet must stayed on the ground. Each reach task required to bring the joystick device four times on the target and back to the starting posture while counting each contact with the screen. This procedure was intended to reduce the performance variability of each subject by offering a sufficiently long activity which duration was regulated by the verbal activity (i.e. counting aloud).

4.2 Results

Figure 5 presents the average task duration normalized by the average task series duration as a function of the normalized height for one subject (as well as a standard deviation interval). A clear augmentation of about 25% is observable on both ends of the nor-

normalized height interval. Figure 6 shows results for all subjects and confirms that augmentation across them.




Figure 5: Single subject with one standard deviation




Figure 6: All subjects

5 Reach behavior in the virtual world

The second experiment studies human ability to perform reach in the virtual world while controlling either a simple rigid shape (i.e. proxy) or a virtual avatar.

5.1 Method

5.1.1 Participants

Eleven subjects, male and female, from 25 to 30 measuring 1.63m to 1.91m participated to the study. None

had counter-indication for standing-up over the duration of the study.

5.1.2 Apparatus

Each subject was controlling either some simple rigid shapes (i.e. a proxy consisting of some capsules moving similarly to subject's hands by tracking some mocap markers as shown on figure 7) or a self-similar avatar (Figure 8). He/she was able to flex knees and go on his/her toes, but couldn't step forward or backward. When controlling the avatar, the motion was recovered using the motion capture setup described earlier (Section 3). The subject was successively immersed in the virtual world through two complementary setups. First, the subject was in front of a large screen (Figure 8) where he/she was presented a third person point of view consisting of two orthographic views: back and right side (Figure 9). We believe those two views offer the user the best feeling of depth. Indeed because our projecting device is not stereo we noticed that a perspective projection could confused the subjects. In addition, combining two views and rendering the avatar with some transparency effect (Figure 9) reduce the possibilities that the avatar occludes the target. Otherwise the subject wore an HMD in order to get a first person perspective viewpoint (Figure 7) and experiment the virtual world as if he were the avatar.

5.1.3 Procedure

The subject was asked to reach some target box in the virtual world (red or green box on figures 7, 8 and 9). Targets were virtually placed 0.5m from the subjects and at various heights of 0.5m, 0.95m, 1.4m and 1.85m. Each subject went through one series of reach (high and low targets height are repeated twice. Order is randomized) per viewpoint and per controlled entity (i.e. either proxy or an avatar) giving four different contexts. Before a series, the subject was getting a 30s period to get used to the current context. Between series the subject was given a 60s period to relax.

5.2 Results

The study consisted of 264 measurements of which we retain 254. We reject ten measurements that were above a time-out value of 9s (3.78%). Outliers were mostly due to incorrect posture reconstruction. In such cases, the user had to drive the avatar slowly so it adopt the correct posture. Figures 10, 11, 13 and 14 present




Figure 7: User driving simple rigid shapes at the hands in a first person viewpoint. One capsule has been enhanced for this document for visibility reasons. In this snapshot the first-person viewpoint is also displayed on the large screen for assessment purpose by an external operator.

some results for all subjects: reach durations normalized by the average time of each series is displayed as a function of the height normalized by the subject height. Data from figure 6 are repeated in figure 12 to match the scale of figures 10, 11, 13 and 14. Figure 15 displays the average duration for the low reach over all subject for each context.

6 Discussion

We now compare results from the two presented experiences in order to find out some similarities and discrepancies. First, it can be observed that data from figures 10, 11, 13 and 14 (i.e. 2nd experience) are much more scattered than data from figure 12 (i.e. 1st experiment). This results from the differences of the two reach protocols. In the 1st experiment, a reach task consists of four contacts with the target, while in the 2nd experiment a reach task consists of only one contact with the target. The first experiment measurements tend to reduce disparities. In the second experiment reaches are short and more sensitive to variation in the movement. However because durations and heights are normalized in similar ways for the two experiments, we can compare data at the same scale.




Figure 8: User driving the avatar in a third person viewpoint




Figure 9: Third person viewpoint. Targets' height are approximatively at the level of each part of the shelf rendered in grey.

The first experiment reveals that lower and higher reaches require more efforts than middle ones. This additional effort leads to extended durations of about 25% longer for extreme reaches (figure 12). Such a reach characteristic is also observable in the virtual world (figures 10, 11, 13 and 14). However in the virtual world, the difficulty of performing the lower and higher reaches is amplified.

For the first person viewpoint (figures 13 and 14), a positive bias is observable for lower reaches in comparison to third person viewpoint (figures 10 and 11). The low reaches durations are even more amplified than high ones. The duration increase is even better illustrated on figure 15: one can clearly observe that, for the low reach, durations are longer for the first person

viewpoint with or without controlling the virtual character. This bias can surely be explained by the limited field of view that HMDs offer. Indeed, the head (approximately between 1.60 and 1.80 depending on the subject) is closer to high reaches (1.85m) than to low reach (0.5m). Looking at figure 7 one can see that the low target is not in direct view because the subjects only sees two levels of the shelf fully represented on figure 8. As a consequence, low reaches require longer scene analysis before finding and reaching the target. Some other reasons could be an inadequate simulation of depth perception. However concerning depth simulation we fully trust the graphic engine [PVT08] which has proved its worth for various other scenarios.

The longer durations observed for the first-person viewpoint, in addition to the well-known limitations of HMD (e.g. comfort) make us recommend not to use an HMD when one wants to evaluate reach durations in a complex virtual environment. In addition, in first-person viewpoint the user of the system doesn't see the whole body and cannot adapt in case the avatar adopt a wrong posture (for example because of an obstacle in the scene). To summarize, we consider that the third-person viewpoint offers a better visualisation when one wishes to control the posture of a virtual character that is intended to interact with its environment.

7 Conclusion

In this paper, we have presented a method to do on-line full-body motion capture with a reduced set of markers using a PIK solver. It relies on hard constraints (joint limits, coupled spine) and soft constraints (position and orientation effectors). Our approach is easy to calibrate and allows performing a wide range of full-body movements. The set up we introduced can be seen as an enactive interface in the sense that it allows controlling transparently a virtual human. The reconstructed motion reflects correctly the one of the user. However, there is a limitation in the frequency bandwidth of the movements. Such an interface will allow us to conduct other researches in the framework of virtual prototyping and explore PIK-based on-line collision avoidance as in [PMM⁺07].

As a future work we plan to first investigate better strategies to resolve temporary markers occlusion. In addition, Kalman filter could be a good way to smooth the orientation retrieved from groups of markers. We also plan to compare the motion we reconstruct to

the one produced with an analytic IK solver or with some commercial products, such as Autodesk Motion-Builder. Finally, adding a few more markers on the fingers should allow simple manual interactions such as grasping and releasing a virtual object (full-body mocap setups are usually limited for hands movements, i.e. only wrist rotation).

We also presented two studies about human ability to reach in the real and in the virtual worlds. For the virtual world, we have explored different setups: subjects controlling either simple shapes at the hands level or a self-similar avatar, subjects immersed via a third or first person viewpoint. The reach duration is clearly dependent of the target height: low and high targets require longer durations (about +25% in the real world). This characteristic is highly accentuated in the virtual world and depends also on the viewpoint modality (i.e. in first person viewpoint reach duration is even more increased). In future works, we plan to extend the study with other modalities: for example it could be interesting to pursue the study by adding 3D stereo to our large screen or to use a CAVE (with/without stereo).

With current virtual prototyping software, it is possible to include virtual actors in virtual scene. However, such a task can be quickly cumbersome because each degree of freedom's trajectory must be specify one by one. On-line motion capture put the user in the loop and let him use his body posture to direct and pose the avatar. From our experiments we tend to conclude that a third-person viewpoint is preferable over a first-person viewpoint when posing a virtual mannequin interacting with its virtual environment. Such a third-person helps to have a global vision of the scene. The user can adopt his/her body configuration to pose correctly the avatar taking into account the whole environment and more specifically obstacles. However third-person viewpoint breaks the end-user experience which is one of the advantage of first-person viewpoint. One can be interested to know what the avatar can see and therefore first-person is clearly a better choice in such a context. Ideally, we could imagine a interface that lets the user switch between first-person and third-person viewpoints. A 3D stereo CAVE seems to be a good candidate for such an application because one can quickly alternate between first-person viewpoint with large field of view and a more global scene view (i.e. third-person viewpoint). To evaluate a virtual environment the user could first navigate in first-person viewpoint with a joystick (only

head will be tracked to adjust stereo). Then the user could adopt a third-person viewpoint to correctly pose the character using on-line motion capture. Finally, the user could switch back to first-person viewpoint to evaluate what the avatar see.

8 Acknowledgments

The authors would like to thank Daniel Raunhardt for his support of the inverse kinematics library (Swiss National Science Foundation under the grant 200020-109989). This research was partially supported by the European network of Excellence ENACTIVE.

References

- [BB04] Paolo Baerlocher and Ronan Boulic, *An inverse kinematics architecture enforcing an arbitrary number of strict priority levels*, *Visual Computer* **20** (2004), no. 6, 402–17, ISSN 0178-2789.
- [BHG93] Norman I. Badler, Michael J. Hollick, , and John P. Granieri, *Real-Time Control of a Virtual Human Using Minimal Sensors*, *Presence Teleoperators and Virtual Environments* **2** (1993), no. 1, 82–86, ISSN 1054-7460.
- [BPC05] Ronan Boulic, Manuel Peinado, and Benoit Le Callennec, *Challenges in exploiting prioritized inverse kinematics for motion capture and postural control*, *Gesture in Human-Computer Interaction and Simulation*, *Lecture Notes in Computer Sciences*, 2005, ISBN 978-3-540-32624-3.
- [BVU⁺06] Ronan Boulic, Javier Varona, Luis Unzueta, Manuel Peinado, Angel Suescun, and Francisco Perales, *Evaluation of on-line analytic and numeric inverse kinematics approaches driven by partial vision input*, *Virtual Reality* **10** (2006), no. 1, 48–61, ISSN 1359-4338.
- [CH05] Jinxiang Chai and Jessica K. Hodgins, *Performance animation from low-dimensional control signals*, *ACM Trans. Graph.* **24** (2005), no. 3, 686–696, ISSN 0730-0301.
- [HA08] H-Anim, *Humanoid Animation Working Group*, 2008, www.h-anim.org, Last visited July, 2008.
- [MBT99] Tom Molet, Ronan Boulic, and Daniel Thalmann, *Human motion capture driven by orientation measurements*, *Presence* **8** (1999), no. 2, 187–203, ISSN 1054-7460.
- [Men99] Alberto Menache, *Understanding Motion Capture for Computer Animation and Video Games*, 1 ed., Morgan Kaufmann, 10 1999, ISBN 0124906303.
- [MPB⁺06] Barbara Mazzarino, Manuel Peinado, Ronan Boulic, Marcelo Wanderley, and Antonio Camurri, *Improving human movement recovery using qualitative analysis*, *Proceedings of the 3rd International Conference on Enactive Interfaces*, ENACTIVE06, November 2006, pp. 177–178.
- [Pha08] Phasespace, *Impulse Mocap system*, 2008, www.phasespace.com, Last visited July, 2008.
- [PHW⁺04] Manuel Peinado, Bruno Herbelin, Marcelo Wanderley, Benoit Le Callennec, Ronan Boulic, Daniel Thalmann, and Daniel Méziat, *Towards Configurable Motion Capture with Prioritized Inverse Kinematics*, *Third International Workshop on Virtual Rehabilitation (IWVR2004)*, 2004, pp. 85–96.
- [PMM⁺07] Manuel Peinado, Daniel Méziat, Damien Maupu, Daniel Raunhardt, Daniel Thalmann, and Ronan Boulic, *Accurate on-line avatar control with collision anticipation*, *VRST '07: Proceedings of the 2007 ACM symposium on Virtual reality software and technology* (New York, NY, USA), ACM, 2007, ISBN 978-1-59593-863-3, pp. 89–97.
- [PVT08] Achille Peternier, Frederic Vexo, and Daniel Thalmann, *The Mental Vision Framework - A Platform for Teaching, Practicing and Researching with Computer Graphics and Virtual Reality*, pp. 242–260, Springer-Verlag, Heidelberg, 2008, ISBN 978-3-540-69744-2.

- [RB07] Daniel Raunhardt and Ronan Boulic, *Exploiting coupled joints - anatomic control of the spine with ik through linearly coupled joints*, GRAPP 2007, Proceedings of the Second International Conference on Computer Graphics Theory and Applications, Barcelona, Spain, March 8-11, 2007, vol. AS/IE, 2007, ISBN 978-972-8865-72-6, pp. 13–20.
- [Uni08] Carnegie Mellon University, *CMU Graphics Lab Motion Capture Database*, 2008, mocap.cs.cmu.edu/-markerPlacementGuide.pdf, Last visited July, 2008.

Citation

Damien Maupu, Ronan Boulic, and Daniel Thalmann <i>Characterizing full-body reach duration across task and viewpoint modalities</i> , Journal of Virtual Reality and Broadcasting, 5(2008), no. 15, November 2008, urn:nbn:de:0009-6-16205, ISSN 1860-2037.




Figure 10: 3rd person viewpoint in the virtual world when subjects control some simple shapes.

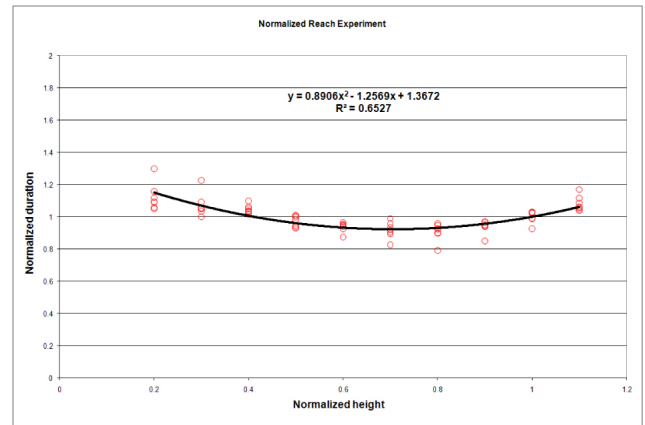


Figure 12: Normalized full-body reach duration as a function of the normalized target height (real setup)

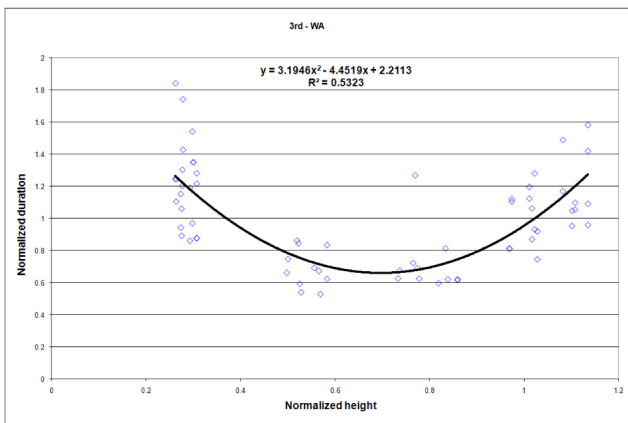


Figure 11: 3rd person viewpoint in the virtual world when subjects control an avatar.

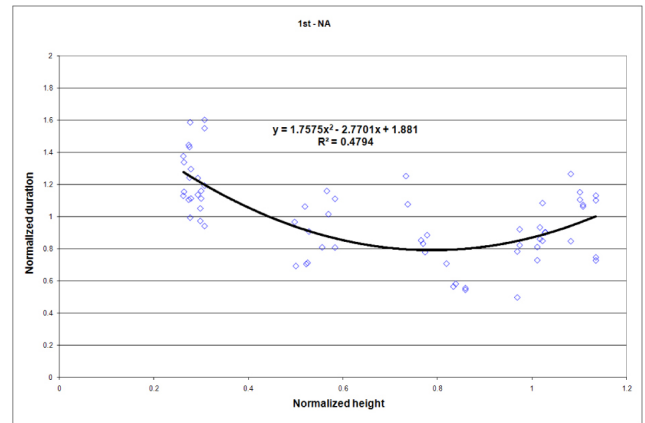


Figure 13: 1st person viewpoint in the virtual world when subjects control some simple shapes.




Figure 14: 1st person viewpoint in the virtual world when subjects control an avatar.

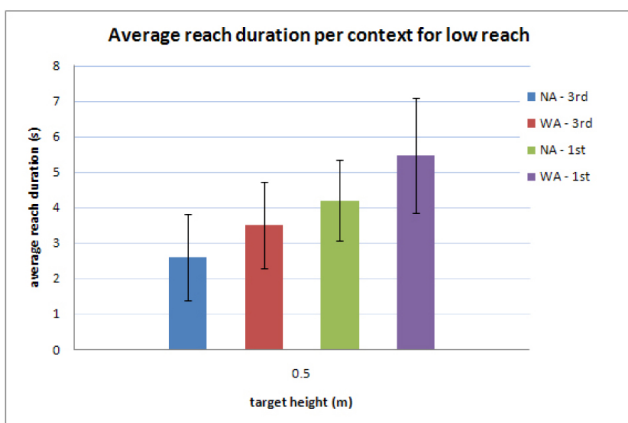


Figure 15: Absolute duration per context (viewpoint - control entity) for low reach in the virtual world (NA stands for “no avatar”, i.e. simple rigid shapes, WA stands for “with avatar”, 1st and 3rd indicate the viewpoint).