Mirrored Motion: Augmenting Reality and Implementing Whole Body Gestural Control using Pervasive Body Motion Capture based on Wireless Sensors

Philip Smit, Peter Barrie, Andreas Komninos and Oleksii Mandrychenko

School of Engineering and Computing, Glasgow Caledonian University, 70 Cowcaddens Road, Glasgow, United Kingdom

{philip.smit, peter.barrie, andreas.komninos, oleksii.mandrychenko}@gcal.ac.uk

Abstract. There has been a lot of discussion in recent years around the disappearing computer concept and most of the results of that discussion have been realized in the form of mobile devices and applications. What has got lost a little in this discussion is the moves that have seen the miniaturization of sensors that can be wirelessly attached to places and to humans in order to provide a new type of free flowing interaction. In order to investigate what these new sensors could achieve and at what cost, we implemented a configurable, wearable motion-capture system based on wireless sensor nodes, requiring no special environment to operate in. We discuss the system architecture and discuss the implications and opportunities afforded by it for innovative HCI design. As a practical application of the technology, we describe a prototype implementation of a pervasive, wearable augmented reality (AR) system based on the motion-capture system. The AR application uses body motion to visualize and interact with virtual objects populating AR settings. Body motion is used to implement a whole body gesture-driven interface to manipulate the virtual objects. Gestures are mapped to correspondent behaviours for virtual objects, such as controlling the playback and volume of virtual audio players or displaying a virtual object's metadata.

Pre-print final approved draft - LNCS Whole Body Interaction - Springer
October 2010

Introduction To Motion Capture

Humans have always had difficulties interacting effortlessly with computers. The difference in language is perhaps too great to ensure natural and graceful communication; therefore it could be supposed that interaction may be improved in some ways by taking away some of the physical barriers between the machine and the user. Today many artificial intelligent technologies like speech and image recognition systems are commercially available to make people feel that the device is reacting to them in a more intuitive way. We took this concept a step further and investigated how a wireless sensor-based system could be implemented to allow the capture of human body movement and gestures in real time.

Motion-capture is not limited to man-machine interfacing only, but also has applications in a diverse set of disciplines, for example in movie and computer game production, sports science, bioengineering and other sciences to which the analysis of human body movement is a major focal point. Motion capture systems have tended to be complex, expensive, purpose-built setups in dedicated and strictly controlled environments that maximize their efficiency. However, in the context of pervasive computing, the design of a system to capture motion at any time and any place, is constrained by several parameters that are not considered in traditional systems. Such constraints are the durability, wearability (and discreetness of the system when worn), independence from specially configured environments, power consumption and management and connectivity with other pervasive systems. We aimed to address these problems in our study, and as such began to think about how to develop a low-cost, real-time motion-capture system.

The approach we took was to use sourceless sensors to establish the orientation of the human anatomical segments, from which posture is then determined. Sourceless sensors do not require artificial sources (e.g. IR illumination or artificial magnetic fields) to be excited. Instead, they rely on “natural phenomena”, e.g. the earth’s magnetic field and gravity, to act as stimulus [1]. Such sensors need to report their readings so these can be processed and translated into body movement. To achieve this, we thought it would be appropriate that wireless technology was used to connect the sensors, thus forming a Wireless Body Area Network (WBAN). Wireless sensors make the system unobtrusive, increase its wearability and compared to a wired solution, allow for a much wider range of applications. In the following sections we present our investigation into the development of a low-cost, low-power WBAN of sensors, as an enabler for HCI applications. We also present an outline of applications where this has been successfully used and discuss future opportunities for this system.

Background To Motion Capture

Wearable sensor systems have been used in the past with success in several contexts of which particular focus seems to have been placed within the domains of Pervasive Healthcare [2,3,4] and Interaction with Mixed or Virtual Reality systems [5,6], and Mobile Systems [7]. Wearable sensors have also been used to investigate Interaction in such domains as Computer Gaming [8] and the acquisition of varying levels of Context Awareness [9,10]. In such respects, while much progress has been made, this progress only partially fulfills the objective of capturing of full body motion in pervasive computing landscapes. There are only a few systems we are aware of which meets this objective; one is in [11], although this system relies on a set of wired sensors and a heavy backpack to power it, limiting its wearability and configurability, as sensors have to be used as a complete set. Two commercial systems work on a similar principle with [11], using sets of sensors wired to a hub, which transmits aggregated data wirelessly using Bluetooth or 802.15.4 (XSens[1], EoBodyHF[2]). Wired sensors limit the wearability of these systems.

---

[1] XSens http://www.xsens.com

[2] EOWave Systems http://www.eowave.com

Our work's fundamental aim is to investigate the use of a low-cost distributed computing infrastructure with sensors to provide a means of capturing environmental and human activity as part of our research group's current interest areas (pervasive healthcare, mobile spatial interaction and mobile audio-video interaction). For HCI researchers there are exciting opportunities due to the standardization, miniaturization, modularization and economies of scale presented by the new technologies available for the creation of wireless sensor networks. Of special interest is wireless body area network (WBAN) technology. Using modern silicon Micro-Electro-Mechanical Systems (MEMS) manufacturing techniques, sensors (such as gyros, magnetometers and accelerometers) have become inexpensive, small and can now be worn on the body or integrated into clothing [12]. Such sensors, coupled with low power processors that may integrated the necessary wireless componentry, (such as the 32-bit Freescale MC1322x platform), provide the basic fabric for increasingly powerful wireless sensor networks.

System Design

From reviewing the existing literature, we identified a set of heuristics against which a pervasive motion capture system must perform well. Our criteria were as follows:

- **Connectivity:** Pervasive systems do not work in isolation. Any sensor-based system must allow its components to communicate with each other and coordinate its behavior. It must, however, also be able to communicate its components' status to external systems in the environment.
- **Power:** A pervasive system must not rely on external sources of power, as these are not omnipresent. It should have its own power source and appropriate power management features that allows it to operate for lengthy periods of time.
- **Performance:** The performance and responsiveness of a pervasive motion capture system must be such that it affords the real-time capture of bodily motion and its transmission to external systems with minimal lag and delay.

- **Wearability:** Systems must be light, easily wearable and discreet. Discretion can be achieved by embedding sensors in everyday objects or garments, or by designing them so that they can be easily concealed.

In designing the Mirrored Motion demonstrator, we considered these heuristics as appropriate to informing our system characteristics.

Connectivity & power

Our system is comprised of sensor "nodes" that can be attached to key locations on a user's body, monitoring the movement of major body parts (limbs, torso, head). One of the off-body nodes acts as a "coordinator", gathering data from all nodes and relaying to external systems for further processing. To coordinate the communication between the peripheral and the coordinator nodes, the Bluetooth and IEEE 802.15.4 standards were considered suitable candidates. We also considered 802.11x (Wi-Fi) but this was quickly rejected, as its power consumption is too high for continuous use. A shortcoming of Bluetooth is that it is limited to eight nodes per network, which would be insufficient for covering even just the basic major parts of a human body. In contrast, IEEE 802.15.4 can have 65536 nodes in a network (star or mesh topologies) and can work over similar node-to-node distances as Bluetooth. It can operate with a smaller network stack size, reducing the embedded memory footprint. For the flexible and extensible HCI applications to be considered, the larger node count is useful to create networks that integrate on and off-body nodes and have potentially multiple interacting users. IEEE 802.15.4 data rates are in the range of 20 – 250 Kbps, although in actual use the higher rate cannot be attained due to protocol overheads. Although lower than Bluetooth, this data rate has been shown in our experiments to be sufficient for body-motion frame rates. Because of its characteristics in allowing multiple node connectivity and very low power consumption, we selected 802.15.4 as the preferred communication protocol. The wireless module used in the system is a Panasonic PAN4555.

Wearability & performance

Sensors used in each node for the first prototype were a 3-axis accelerometer and a magnetometer per node. A magnetometer-accelerometer sensor can produce accurate orientation information when the only force experienced by the sensor is gravity. However, any additional forces will result in the reference vector produced by the accelerometer to be inaccurate. In a revision to our original design, miniature MEMS gyros were added to the sensor pack. Gyros measure angular velocity and this helps to reduce the effects of non-gravity forces.

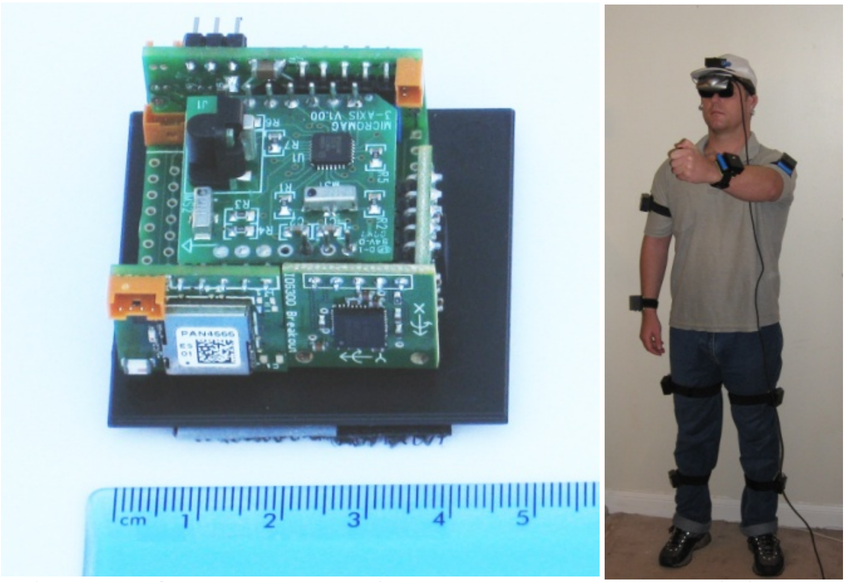

Figure 1. Our custom-designed sensor pack containing 3-axis accelerometer, gyro, magnetometer and 802.15.4 comms (left). On the right, a user demonstrating the small size and wearability of the packs, Velcro-strapped on his body. The cable attaches his VR headset to the host PC.

These sensors were originally packaged in a rather large form, roughly the size of an average mobile phone, as pre-configured development kits were used to prove the concept. Once satisfied with the performance of the system, we re-designed the hardware and created custom sensor packs that were optimized for size. Each pack is relatively small (less than 4 x 3”). They are attached to the user’s body with Velcro straps, making them easy to wear and remove. Their small size makes them easy concealable under normal clothing. Because this is an experimental platform we created a modular construction allowing the removal and addition of the sensor and wireless components. The necessary connectors and modules take up extra space. A custom version could be created with a smaller footprint, with all parts integrated onto a single PCB. Sensor nodes are placed on each of the tracked human limbs (upper and lower arm, head, torso, upper and lower leg) to track the orientation of each. The raw data acquired by the sensor WBAN is transmitted wirelessly to an external system (in our experiments, a typical PC). We set a data acquisition target for our system to achieve real-time performance at a sampling rate of 30Hz, as this would, in theory, allow us the re-creation of a user’s skeletal model on an external system with a refresh rate that would yield about 25-30 FPS, which is adequate for real-time video. The posture of the skeleton is calculated in real-time through forward kinematics. Kinematics simplifies computations by decomposing any geometric calculations into rotation and translation transforms. Orientation is obtained by combining (or fusing) these information sources into a rotation matrix – an algebraic format that can be directly applied to find the posture of the user. The result is a simple skeleton model defined as a coarse representation of the user.

The sensor network

The sensor nodes were successfully tested at a 30 Hz sample rate but this appeared to be the upper limit. Our empirical results show that the coordinator could handle up to 360 packets per seconds (i.e. up to ~12 nodes) with latency between 5 and 25 ms for the coordinator (using a simple 8-bit processor) to collect and forward any given frame to the external systems (PC). We would like to point out however that in our current system the packet rates are dependent on the constraints of the simple processing hardware and the application running on it. A lightweight application or better processor will probably handle much higher packet rates. In order to provide a synthesis of human movement and position within the system, a skeletal model was developed on the PC receiving the motion data. Similar models have been used successfully in the past [13,14,15]. Our model uses the lower torso as the root link and tracks the position of each limb as a set of links connected to each other starting at the root. The skeleton model we produced is easily extensible and can be augmented to incorporate many more nodes, such as to track palm, finger or foot movement. Because the receiver (coordinator) node on the PC is connected using a serial USB connection, it is possible to have multiple WBANs on the user's body, each with up to 12 sensor packs (in order to maintain very low latency levels). Our system is, in this respect, very highly configurable, as not all of the nodes need to be attached to the body or activated in order for the system to work. It is possible to arrange the system in such manner as to detect only arm movement, torso movement, leg movement or any combination of these, simply by strapping on the appropriate sensor packs and indicating to the capture interface which sensors are being work by checking the relevant boxes (see Figure 2).

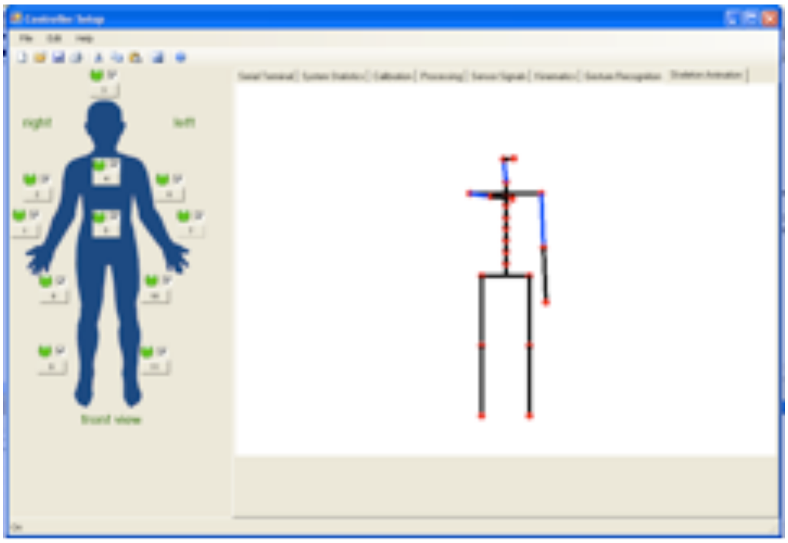

Figure 2. The motion capture interface (PC). A user can indicate which sensor packs are being worn by checking the relevant boxes on the human outline shape. The green status dots turn red when the sensor pack is not transmitting. The skeleton model on the right is constructed in real-time.

A calibration procedure has to be enacted at the start of a motion capture session by the user. Posture calibration is performed with the user assuming a predefined reference posture (standing up straight, arms down), as in [15]. The calibration takes approximately 2-3 seconds to complete, which can be considered to be a low overhead for the human actor. The captured data is sent from the coordinator to the PC and is then processed through a configurable low-pass filter before going through the skeletal transformation. At this stage, the PC can then display a stick-figure animation as shown in Figure 2. The calibration interface and sensor placement guide on the human is also shown in Figure 2.

Whole Body HCI

Achieving motion capture solves only one part of the problem in creating novel human-computer interfaces. We developed a demonstration application based on our sensor system, in which the movement is captured from the user and then the skeleton is covered with a digital skin, using DirectX and integrated into a synthetic 3D environment as shown in Figure 3. In this demonstrator, the user is equipped with a VR headset as well as the motion capture system. The 3D world is the start of the experimentation with interaction. This experiment provides smooth motion tracking from first (with 2/3D head mounted display) and third person perspectives that immerse the users in a synthetic experience using real movements and synthetic visual feedback, so, for example, when the user holds up their hands in front of them in the real world they see their hands in a 3D virtual world (videos of this can be viewed on our website at http://www.mucom.mobi). In another application, we augmented our nodes with an optical proximity sensor, to allow a sensor node to be mounted within a shoe to undertake a field investigation of foot motion. Further to this, we began investigating how our equipment can be used to accurately detect gait and foot clearance for elderly persons, helping solve and investigate issues in fall prevention. This is particularly important as until now, people could only be monitored in specialized labs (with expensive video equipment); now it is feasible to monitor an elderly person in their own environment and for extended periods of time, at relatively low-cost. A recent laboratory-based trial compared an existing video-tracking system with our foot-mounted sensor system. The results show a high degree of correspondence between the two data sets.

Continuing in the domain of pervasive healthcare, we also produced a prototype of a Marble Maze game that was used with a wobble board. The user stands on the board and makes small movements in order to guide the marble through the virtual maze, helping improve body balance and posture for rehabilitation patients. We used one sensing node to detect the movement of the wobble board with a high level of success.

Introduction to an Augmented Reality Application

There is significant interest in the development of more "natural" methods for Human Computer Interaction. Keyboards, joysticks, mice, displays and other devices are hardware interfaces in widespread use. Many intelligent technologies like speech and image recognition systems are commercially available to facilitate interaction through the use of naturalistic human-computer interface modalities. One interaction modality that has been the focus of considerable research lately is that of Gestural Interaction, where commands from mouse and keyboard might be replaced with a user's gestures [17].

Figure 3: Real-time mapping of user body movement to a 3D virtual avatar in an immersive world.

Virtual reality (VR) has been a focus of research and commercial practice for a number of years, not only for entertainment purposes but also for industrially relevant applications such as 3D product design and visualization. The approach of Augmented Reality, where virtual worlds and objects, or worlds and metadata are mapped on to views of the real world, mixing the real with the artificial, has emerged in the computer science world in addition to VR. However, both types of visualization suffer from problems of control – how can a user manipulate virtual objects as naturally as they would manipulate real physical ones? We aimed to examine the concept of naturalistic interaction with virtual objects in an AR setting by investigating how our wireless-sensor-based system could be used to recognize gestures made by the user's body and help create a wearable AR system that could be deployed and used without the need for fixed infrastructure.

The approach we took was to develop a system based on the Mirrored Motion system, a VR display headset and a web camera attached to the user’s head. The sensors provide raw data subsequently used for the recognition of the user's gestures, whilst the camera gives a live video feed on which virtual objects are superimposed. The web camera works with the sensor on the user’s head to obtain the camera’s orientation and as such, synchronize the panning and rotation of the virtual world to match the web camera movements.

Background to Augmented Reality

AR technology enhances a user’s perception of and interaction with the real world. The key feature of the augmented reality technology is to present auxiliary information (visual, audio, haptic etc) in the sensory space of an individual, though in our work we concentrate on augmenting the environment with visible virtual objects. The virtual objects display information that the user cannot directly detect with his or her own senses. The information conveyed by the virtual objects helps the user to perform real-world tasks. This new form of human-computer interaction can be applied to various industry areas [16]. AR is an emerging technology area and as such, applications that could benefit from AR technology are still not fully explored. Typical examples are seen in engineering, military and educational domains. AR technology for digital composition of animation with real scenes is being explored to deliver adverting content or bring new digital entertainment experience to users.

Our system represents an exciting opportunity to engage in interaction design research. For the purposes of AR, orientation sensors coupled with a web camera provides evident opportunity for orientation in a virtual world accordingly to the direction that camera faces. The skeletal model built from our sensors’ data supplies the receiver with rotation matrices and linear angles that can be used to recognize human gestures [18]. We aimed to extend the surrounding spatial environment with supplementary information through AR. We wanted to use the system not only to help visualize virtual objects for AR, but also interact with the objects through gestures.

System design

AR technology is not a new concept. Apart from studies in AR visualization, many applications already exist in advertising, industry, robotics, medicine, sport, military and many other spheres. Additionally, several researchers have proposed to use gesture recognition in conjunction with AR [12,20]. However, we are not yet aware of any AR systems based on full body motion capture and that utilize gesture interaction, which do not require extensive infrastructure support and which can be used in pervasive computing settings. From reviewing the existing literature, we identified defined two goals [8] [10] to be implemented:

- **Gesture recognition.** The proposed system must recognize user's gesture in 3D space independently on the user's location. Gestures can be trained before recognition.
- **Extending reality.** The system must provide means for presenting auxiliary information in the field of view of a user's VR headset. Providing the particular gesture is recognized the system is to change the state of correspondent virtual object.

In designing the AR demonstration, we considered these goals as appropriate to inform our system characteristics.

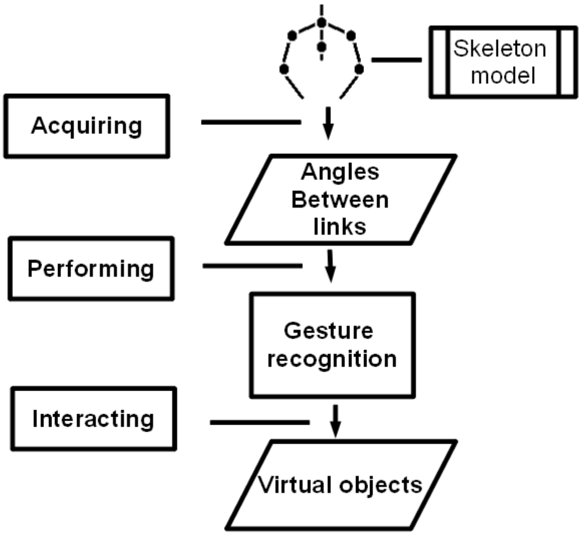

Figure 4. Gesture recognition architecture.

Gesture recognition

As described earlier, our system is comprised of sensor "nodes" that can be attached to key locations on a user's body, monitoring the movement of major body parts (limbs, torso, and head). One of the off-body nodes acts as a "coordinator", gathering data from all nodes and relaying to external systems (e.g. a PDA, a server or a desktop computer) for further processing. The approximate frequency of streaming data is 20 Hertz. While our system is capable of full body motion capture, in this application we used an upper body set of sensors, as we were more interested in torso and hands gesture recognition. An internal processing system provides us with an updatable skeleton model of the user which is a method also used by other researchers, e.g. [7]. In general terms, gesture recognition consists of several stages, like feature extraction, preprocessing, analyzing and making a decision. Our experimental method consists of using linear angles between any two links in the skeletal model as a dataset that is fed into the gesture recognition algorithms described below (See Figure 4).

At the preprocessing stage we perform work to filter the noise caused by rapid movements and inaccuracy of the measurements (around 3-5 degrees). A magnetometer-accelerometer-gyro sensor can produce accurate orientation information when the forces experienced by the sensor are gravity or low accelerated movements. Any additional forces will result in the reference vector produced by the accelerometer to be inaccurate.

- τ_i is the probability of remaining in an emitting state, $T(E_i, E_i)$
- η_i is the probability of going to the next emitting state $T(E_i, E_i + 1)$
- ς_i is the probability of skipping at least one emitting state $T(E_i, N_i+1)$
- k_i is the probability of skipping an additional emitting state $T(N_i, N_i+1)$
- p_i is the probability of ending a skip sequence $T(N, E_i+1)$

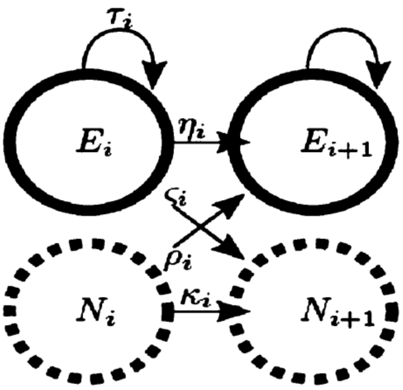

Figure 5. Transition probability parameters for the HMM model.

Analyzing sequences of linear angles and performing the gesture recognition itself was implemented with the help of AMELIA general pattern recognition library [21], which we used as a basis to implement our own customized Hidden Markov Model. Hidden Markov models (HMMs) are the basis of most gesture recognition algorithms used today. However, traditional HMM-based gesture recognition systems require a large number of parameters to be trained in order to give satisfying recognition results. In particular, an **n**-state HMM requires n^2 parameters to be trained for the transition probability matrix, which limits its usability in environments where training data is limited [14,19]. The reduced model that was used in our system uses a constant number of parameters for each state to determine transition probabilities between all states. As there are many different notation conventions in use for Hidden Markov Models, here we utilize a convention we believe makes our model easy to understand. We thereby define our augmented hidden Markov model ($S = \{E,N\}$, S_b, S_e, T,O) by a set of states S, a designated beginning state S_b, a designated ending state S_e, a state transition probability function T, and an observation probability function O. The augmented HMM behaves essentially the same as a regular HMM, with only a few points of departure. The set of states S is divided into disjoint sets of emitting states E and non-emitting states N. The difference between the two is that when entered, emitting states emit an observation belonging to the observation set θ according to the observation probability function $O : E \times \theta \rightarrow [0,1)$. The model always begins in the beginning state $S_b \in S$, and until it ends up in the ending state $S_e \in S$ it makes transitions according to the transition probability function $T : (S - S_e) \times S \rightarrow [0,1)$. T must also satisfy that the sum of transition probabilities out of any state is 1. In the reduced parameter model, we use the following parameters, depicted also in Figure 5.

Our system allows users to record their own gestures for predefined actions that control the behaviour of virtual objects (e.g. selecting/deselecting an object, turning on and off a virtual appliance such as a virtual radio, controlling the behaviour of a virtual object such as start/stop playback of music), some of which are depicted in Figure 6. As such, different actors may use diverse gestures for the same action. Typically, to record one gesture an actor repeats it for 3-4 times, as in [15,19]. Once a few “recordings” of a gesture have been made, the system is then trained on the captured motion data set in order to be able to recognize the gestures. A general gesture tends to be 2 to 3 seconds in time. After training, the user can perform gestures in different sequences as well as performing actions that are not gestures. Our system recognizes gestures with the probability of 80-90% (determined experimentally). Examples of our gesture recognition systems are available to view online in video form from our website (http://www.mucom.mobi).

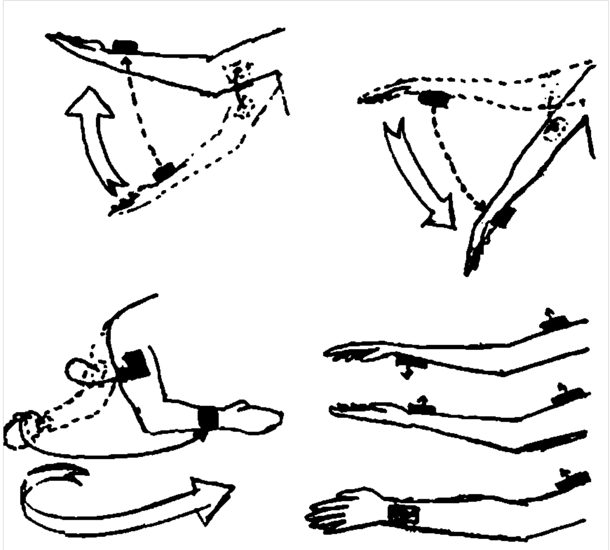

Figure 6. Examples of naturalistic gestures designed to control a virtual radio in the AR system. The top gestures show raising the (left) and lowering (right) the volume. The bottom left shows skipping a station. By modifying the position of just one node (carpal), we can achieve a large number of distinct gestures (bottom right).

At this point in time, our system has two limitations: Firstly, saving of the recorded gestures training data is not yet implemented (due to development-time constraints) but we consider it as a simple goal. Secondly, our current recognition model does not allow a gesture to stop in the actor's relaxed position. For example, if a user stands still and tries to record a gesture, finishing it at the relaxed posture, the recognition system will not determine when the gesture ends. However, this limitation will be removed in the near future.

Extending reality using Whole Body Interaction

There are differing approaches to augmenting reality and presenting synthetic visual information overlayed on real world views. Magic Lens applications rely on the use of a camera enabled device that acts as a viewer, through which additional information pertaining to real objects or completely virtual 3D objects can be viewed. Another mode is the use of special glasses, on which simple graphics or text is rendered. For our approach, we used a set of VR goggles connected to a webcam. Live video that comes from a web camera is constantly placed in front of a viewer in a 3D world. In order to ensure that the 3D world's game camera corresponds with some fidelity to the live video feed from the webcam, the system must be calibrated by starting at a pre-determined real-world location whose coordinates are mapped to a pre-determined point in the virtual world. The user sees a combined image from real video and virtual objects. Virtual objects are placed in front of the dynamic web camera feed. The coordinates of the video are not updated, therefore live view always stays on the same place - in front of the viewer, whereas coordinates of the virtual objects are updated. We combined a web camera with the head sensor, which helps map the camera orientation (and hence the user's view of the real world) in 3D space. As a user moves his or her head, the virtual world moves accordingly. The virtual objects that are in front of the human actor will come in and out of the user's field of view, when the viewer looks to the left or to the right (up or down). In order to provide a synthesis of live video feed and virtual objects to the user, so that an augmentation of reality can be implemented, Microsoft's XNA gaming development tools were used. In our AR application, a user sees a mixture of real and virtual objects. In order to interact with virtual objects or metadata pertaining to real objects, these must somehow be selected. To select a virtual object, we used data that comes from the sensor on the right hand. We transform pitch and rotation to the Y and X movements of a cursor in the virtual world. To select a virtual object, the user thus uses their right arm to "point" a crosshair cursor to the virtual object they want to select. Every virtual object has its own bounding form. For simplicity, we used bounding spheres only. We took advantages of the XNA ray technique to understand whether a ray (game camera – cursor – infinity) intersects with the bounding spheres of virtual objects. When the cursor line of sight intersects and hovers over an object, it becomes selected (Figure 7). We found this method of selection in preliminary tests easy to understand and one that is well received as it affords more precise and flexible control than using head direction for selecting (one can look and point in different directions).

Figure 7. The user points an arm-controlled cursor (visualized to the user as an X) at the virtual object (marked by the arrow), which is then highlighted. Metadata for that object is subsequently displayed.

Conclusions & Further Work

We have described how we defined a set of criteria for a pervasive body motion capture system and created a system informed by these, which was then used to investigate whole and partial body interaction in a series of demonstrators. Throughout our development we aimed to make use of easily available, low-cost components, keeping the cost per node to approximately £150. Given the many different environments (e.g. healthcare, gaming, VR, AR etc) in which we wished people to interact with and benefit from our work we needed to ensure that the system was additionally highly configurable, to allow a wide range of interaction opportunities to be investigated. Overall we were successful in delivering a high-performance, truly pervasive, extensible and highly wearable system that fulfils the criteria for such systems.

We have described how we implemented gesture recognition with the pervasive body motion capture system and created augmented reality, which might be used in different fields such as entertainment, sports, military etc. Throughout our development we aimed to make use of our existing low-cost nodes. Overall we were successful in delivering a high-performance, truly pervasive, extensible and highly wearable system that fulfils the criteria for augmented reality systems. In Figure 4, the user's only restriction to mobility is the headset connection, in this picture connected to a desktop PC, but equally able to be connected to a portable laptop/tablet. However, our system at the moment does not support the motion of the user's body between locations in the real/virtual world. We assume that the user remains fixed and as such we have only used the upper body sensor set as a means to trap gestures. In the near future, we plan to take advantage of our ability to capture motion from the entire body, in order to allow the user to move through the AR world. We would be particularly interested in examining how our MEMS-based system performs in inferring user location (e.g. while walking) and how the accuracy of our system might be enhanced through the fusion of GPS data, where available. Additionally, a hybrid positioning system as described would be of great interest to examine in scenarios where indoor-outdoor transitions occur for the user.

We believe that our system will prove an extremely useful tool for a range of interaction opportunities; aside from our previous projects we are working on applying our system in several areas. We are particularly interested in its potential in mixed reality situations for gaming, We also wish to investigate issues in human-human interaction through embodied agents, controlled through the motion capture system. We are looking into the control of VR agents, as well as robotic agents for which the metaphor of "transferring one's soul" will be used to investigate response and interaction with other humans. Finally, we are interested in pursuing applications in tangible interfaces and semi-virtual artifacts, as well as gesture-based whole-body interaction with large situated displays. We hope to be able to create new types of human-computer interfaces for manipulating program windows, arranging or opening files using ad-hoc large projected or semi-transparent situated displays.

REFERENCES

1. Bachmann E, (2000) Inertial and Magnetic Angle Tracking of Limb Segments for Inserting Humans into Synthetic Environments. Ph.D thesis, Naval Postgraduate School, Monterey.
2. Bonato P, (2003) Wearable sensors/systems and their impact on biomedical engineering. Engineering in Medicine and Biology Magazine, IEEE, 22(3), p.18-20.
3. Ouchi K, Suzuki T, Doi M, (2002) LifeMinder: a wearable healthcare support system using user's context. In Proceedings of the 22nd International Conference on Distributed Computing Systems, p. 791-792.
4. Jovanov E (2005) Wireless Technology and System Integration in Body Area Networks for m Health Applications. Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Shanghai, China.
5. Martins T, Sommerer C, Mignonneau L, Correia N (2008) Gauntlet: a wearable interface for ubiquitous gaming. In Proceedings of the 10th international conference on Human computer interaction with mobile devices and services. Amsterdam, The Netherlands: ACM, pp. 367-370.
6. Svensson A, Björk S, Åkesson KP (2008) Tangible Handimation: Real-time Animation with a Sequencer-based Tangible Interface. In Proceedings of the 5th Nordic Conference on Human Computer Interaction, Lund, Sweden.
7. Crossan A, Williamson J, Brewster S, Murray-Smith R, (2008) Wrist rotation for interaction in mobile contexts. In Proceedings of the 10th international conference on Human computer interaction with mobile devices and services. Amsterdam, The Netherlands: ACM, pp. 435-438.
8. Antifakos S, Schiele B (2002) Bridging the gap between virtual and physical games using wearable sensors. In Proceedings of the Sixth International Symposium on Wearable Computers, 2002. (ISWC 2002), p. 139-140.
9. Clarkson BP (2002) Life patterns: structure from wearable sensors. Ph.D Thesis, Massachusets Institute of Technology
10. Seon-Woo L, Mase K (2002) Activity and location recognition using wearable sensors. Pervasive Computing, IEEE, 1(3), p.24-32.
11. Vlasic D, Adelsberger R, Vanucci G, Barnwell J, Gross M, Matusik W, Popovic J (2007) Practical motion capture in everyday surroundings. ACM Trans. Graph., 26(3), 35.
12. Tognetti A, Lorussi F, Tesconi M, Bartalesi R, Zupone G, De Rossi D (2005) Wearable kinesthetic systems for capturing and classifying body posture and gesture. Conf Proc IEEE Eng Med Biol Soc. 2005; 1:1012-5.
13. O'Brien JF, Bodenheimer RE, Brostow GJ, Hodgins JK (2000) Automatic Joint Parameter Estimation from Magnetic Motion Capture Data. Graphics Interface 2000 Montréal, Canada. pp. 53-60.
14. Molet T, Boulic R, Thalmann D (1999) Human Motion Capture Driven by Orientation Measurements. Presence: Teleoperators and Virtual Environments Vol.8(2), pp.187-203.
15. Bodenheimer B, Rose C, Pella J, Rosenthal S (1997) The process of motion capture: Dealing with the data. Computer Animation and Simulation. pp. 3-18
16. Lyu MR, King I, Wong TT, Yau E, Chan PW, (2005) ARCADE: Augmented Reality Computing Arena for Digital Entertainment. 5th IEEE Aerospace Conference, Big Sky, MT, USA

17. Azuma R (1997) A survey of augmented reality, Presence: Teleoperators and Virtual Environments, Vol. 6, No. 4, pp. 355-385, August 1997.
18. RajkoS,Oian G, (2008) HMM Parameter Reduction for Practical Gesture Recognition. IEEE International Conference on Face and Gesture Recognition, 2008.
19. Rajko S,Oian G, Ingalls T, James J (2007) Real-time Gesture Recognition with Minimal Training Requirements and On-line Learning. IEEE Conference on Computer Vision and Pattern Recognition.
20. BuchmannV, Violich S, Billinghurst M, Cockburn A (2004) FingARtips: gesture based direct manipulation in Augmented Reality. In: GRAPHITE 2004: Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Australasia and South East Asia.
21. AMELIA: A generic library for pattern recognition and generation: http://ame4.hc.asu.edu/amelia/ (link valid 08/2010)