Newcastle University ePrints

Ushaw G, Davison R, Eyre J, Morgan G. [Benchmarking Motion Sensing Devices for Rehabilitative Gaming.](#)

In: The 30th ACM/SIGAPP Symposium On Applied Computing. 2015, Salamanca, Spain.

Copyright: © 2015 ACM. This is the author's accepted manuscript of an article, presented at The 30th ACM/SIGAPP Symposium On Applied Computing. 2015, Salamanca, Spain available from <http://www.acm.org/conferences/sac/sac2015/>

Always use the definitive version when citing.

Further information on publisher website: <http://www.acm.org/>

Date deposited: 3rd February 2015

Version of file: Author

ePrints – Newcastle University ePrints
<http://eprint.ncl.ac.uk>

Benchmarking Motion Sensing Devices for Rehabilitative Gaming

Gary Ushaw¹, Richard Davison¹, Janet Eyre² and Graham Morgan¹

¹School of Computing Science ²School of Neuroscience

Newcastle University, Newcastle upon Tyne, UK

{ gary.ushaw, richard.davison4, janet.eyre,
graham.morgan}@newcastle.ac.uk

Abstract

The diversity of commercially available motion detection devices has yielded much interest in utilising action video games in the field of rehabilitation of impaired dexterity. Identifying the most suitable input device for a rehabilitative game is of paramount importance, but no comparative study currently exists. For the first time, a benchmark is introduced for the quantitative assessment of motion capture devices for gaming. A range of devices are considered, in the context of rehabilitative gaming projects, presenting results on the fidelity of motion capture that each provides. Recommendations are then made on the viability of domestic motion capture devices in terms of both patient involvement and fidelity of data capture.

1 Introduction

Video game technology has been utilised in numerous works to promote rehabilitative exercise [3], [8], [9]. A vital decision which must be made early in developing such projects is identifying the most appropriate input device. This paper presents the first comparative benchmarking study of motion sensing devices for gaming and their suitability to rehabilitative exercise.

Rehabilitative gaming has a number of benefits. Firstly, the game should provide a more involving context, so that the patient enjoys the rehabilitative process and is more likely to be engaged in it. Secondly, the exercises can easily be carried out at home as gaming systems which include motion-capture devices are relatively commonplace and affordable. Finally, data on the patient's progress and recovery can be collected as the software monitors the patient's actions. The suitability of domestic gaming systems to meet the first two requirements of engagement and accessibility seem intuitively feasible. However there is currently no comparative measure of the devices' ability to meet the fidelity requirements of data collection for medical investigation and monitoring.

Many motion capture devices are available for domestic use, including Microsoft Kinect, Nintendo Wiimote, Sony PS Move and Sixense Hydra. In this paper we introduce a benchmark for domestic motion capture devices, and use it to measure and compare the fidelity of the data and the reliability of the systems. Our intention is that the results should serve as a recommendation for future projects utilising domestic gaming devices for rehabilitation.

2 Background and Related Work

Gaming devices have been successfully employed in rehabilitative studies since motion capture devices became commercially viable. Relevant studies are discussed here, with emphasis on the fidelity of results. A brief overview of each device considered in this study is first presented.

2.1 Motion Capture Devices

2.1.1 Wii MotionPlus

Nintendo's Wii system consists of a sensor bar, which has an infra-red LED cluster at each side, and the Wiimote motion controllers which contain an infra-red camera for tracking their position relative to the sensor bar. The controllers also contain three-axis accelerometers and three-axis gyroscopes which are used for dead-reckoning calculation of position, velocity and orientation.

2.1.2 PlayStation Move

The PlayStation Move from Sony consists of a RGB camera and a motion controller "wand". The wand contains an illuminated sphere, which the camera tracks in three dimensions (the distance from the camera is measured from the size of the sphere in the image), as well as a three-axis accelerometer, three-axis gyroscope and a geomagnetic sensor. Consequently the device can track both the position and orientation of the wand. Furthermore the geomagnetic sensor is used to calibrate the measurements against the Earth's magnetic field, thereby addressing cumulative errors from the sensors. The combination of camera tracking and dead-reckoning from the wand's sensors provide motion tracking whether the wand is visible to the camera or not. As the image processing software is searching for a circular shape of a known colour, the latency is much less than in other systems where more complex calculations are required. A potential drawback of applying the PlayStation Move to domestic rehabilitation is that the PlayStation3 is used as a server, in addition to the PC that is running the application, leading to an unwieldy set-up for home use.

2.1.3 Kinect and Kinect 2

Microsoft's Kinect motion sensing system consists of a RGB camera, an infra-red laser pattern projector and an infra-red camera. The device directly monitors

the user's body, so no hand-held devices are required. Three dimensional points in space are triangulated using the stereo pair of the infra-red projector and camera. Three outputs are made available by the system: an infra-red image, a colour image, and an interpolated inverse depth image. The Kinect SDK also provides libraries for tracking the skeleton of the user. This consists of positional information, and orientation can be inferred from the joint hierarchy. Depending on the lighting conditions, and the complexity of the image, the system delivers motion tracking information for up to 20 skeletal joints per user at up to 30 frames per second [10].

An improved Kinect was released for Microsoft's Xbox One in 2013 incorporating a wider field of vision camera, capable of calculating the position and orientation of 25 joints per user, as well as their heart rate and facial expression. The colour map resolution has been improved to 1080 x 1920, and the depth map has been augmented by time-of-flight technology contained in the new camera hardware affording a unique depth value per pixel. The combination of this technology and resolution should result in more accurate tracking as compared to the original Kinect device.

2.1.4 Sixense

The Sixense motion control system from Sixense Entertainment uses magnetic motion tracking to provide continuous position and orientation information. The use of electromagnetic fields is well established technology for reliably measuring three-dimensional space [4]. The other console-based devices use cameras making them prone to interruption due to line-of-sight issues; such problems do not occur with magnetic tracking. The Sixense comprises a base unit, which is connected to the PC via USB, and wireless controllers which the player holds while performing the movements. The base unit contains three orthogonally orientated magnetic coils which emit an electromagnetic field, providing the reference for the position and orientation of the controllers. Similarly each controller contains three smaller orthogonal coils whose position and orientation are measured relative to the emitter's coils. Consequently each controller broadcasts three-dimensional position and orientation data to the base unit. The sampling rate is 60Hz. We elect to use three controllers; one in each hand, and one tucked into the belt or pocket at waist level to provide further data on the patient's base position and motion.

2.2 Rehabilitative Gaming

Work to date on implementing rehabilitative regimes through video game technology has taken two distinct tacks – either existing commercial games have been used as the basis of the study [5], [6], or bespoke games have been developed as part of the research project [2], [12].

Utilising an existing commercial video game in this context has limitations as the game itself has not been developed with rehabilitation or monitoring in mind. The movements which are practised are those chosen by the game

designers to most suit their game-play, so are not necessarily the movements the medical professionals would choose to promote rehabilitation. Furthermore extracting data from a commercial game is impossible on a console, or without access to the source code and tool-chain on a PC. However, commercial games tend to be engrossing and immersive; factors which encourage the patient to persevere with the exercise.

Conversely games developed specifically for the purpose of rehabilitation are designed to recognise and react to prescribed exercise movements. Also, as the software is bespoke, it can measure and collate data pertaining to the patient's progress. A game developed for rehabilitation offers the opportunity to include adaptive game-play, whereby the game detects when the patient is performing less well, and dynamically changes parameters of the game to suit the patient's abilities [7], [1]. The downside of developing a bespoke application is that the quality of the game itself tends to be significantly lower than a commercial title for obvious budgetary reasons, so may not be as enjoyable and involving as a commercially developed game.

To date little quantitative comparison of the devices has been published, with no benchmark known to the authors. [11] provides a comprehensive functional comparison of the devices, but contains no quantitative analysis. In [2] a study of the Kinect is presented in comparison to the much more expensive Optitrack Optical Motion Capture System. The results are promising for wrist and elbow joint movements, although no quantitative analysis of the accuracy of the captured data is presented.

3 Experiment

A benchmark test is required to compare the accuracy and suitability of the motion capture devices. Such a benchmark does not exist in the current literature. We decided to utilise three easily defined and easily repeated movements as the benchmarks of our study. The first benchmark movement is to move the hand in a vertical circle, at full arms length, in front of the subject (i.e. a circle in the coronal plane). The second is to hold the arm at full length in front of the subject and turn 360 degrees on the spot, so that the hand moves in a horizontal circle at shoulder height around the subject's position (i.e. a circle in the transverse plane). The third involves the subject swinging the arm in a vertical circle, keeping the arm at full length (i.e. a circle in the sagittal plane). A number of different subjects were used in the testing of each device. The devices must be able to cope consistently with patients of various shapes and sizes if they are to be used successfully in rehabilitation.

Markers were placed at 1m intervals up to 4m, with the device placed at a height of 120cm. Users stood with their heels inline to each of these markers in turn. At each distance, the user stretched their arm fully forward, and rotated it fully 360° in each axis - for the coronal and sagittal planes, this took the form of solely shoulder rotation, while the transverse plane necessitated the user to rotate their entire body using their legs and feet, while keeping their

back straight. For each plane, the rotation was performed three times, with the second rotation used to calculate results, in order to reduce side effects from the user speeding up, slowing down, or otherwise moving the controllers differently as the trial started and came to an end.

3.1 Metrics

Each device calculates 3D position and orientation in some way, so it is desirable to measure the quality of both. For a measurement of positional accuracy, only the relevant two axes for each plane from each trial-run were considered, creating a sequence of 2D coordinates that represent the vertices of a 2D polygon. Each trial's data was then centred around the 2D origin, according to the centroid of the polygon, calculated from all n positions p in the sequence s as:

$$Centroid(s) = \frac{p_1 + p_2 + \dots + p_n}{n}$$

Once the captured data has been transformed in this way, we determine two metrics by which to rate the quality of the data: circularity, and orientation drift. The circularity of a data set s is defined using the following common shape factor (where a is the area of the shape, and p its perimeter):

$$Circularity(s) = \frac{4\pi a}{p^2}$$

As we are dealing with 2D coordinates, the area and perimeter of each polygon is calculated using the Shoelace formula:

$$Area(s) = \frac{1}{2} \left| \sum_{i=1}^{n-1} p_i x p_{i+1} y + p_n x + p_1 y - \sum_{i=1}^{n-1} p_i y p_{i+1} x + p_n y + p_1 x \right|$$

$$Perimeter(s) = \sum_{i=1}^{n-1} \|\overrightarrow{p_{i+1} - p_i}\| + \|\overrightarrow{p_n - p_1}\|$$

The accuracy of orientation is calculated in a similar way. By transforming the relevant forward basis vector for the input device by an orientation from the recorded sequence, a vector pointing in the direction of the device is created. By projecting this vector onto the plane being tested, a 2D position is formed. For directions that lie exactly on the plane, the point lies at a distance of 1.0. Deviations in the direction away from the plane result in projected points closer to the origin. The projected direction vectors over the course of one movement form a 2D polygon that are tested using the circularity measurement described previously. This method does not take into account deviations in rotation *around* the axis, but still serves as a useful metric of orientation quality.

As the circularity calculation provides values in the range $[0,1]$, the circularity of position and orientation can be used to calculate a measure of data quality

as a simple arithmetic mean. Finally, a base of expected noise is measured, by recording a 30s sample of the device at rest. The standard deviation in position on all three axes, and cosine of angle difference from the forward basis vector are then calculated.

Table 1: Device Steadiness Comparison

Device	Sum Dist	Sum Angle	Dist Trav	Angle Trav
KinectV1	0.0842	0.693	0.00369	0.00961
KinectV2	0.954	149.878	0.000908	0.719
MoveMe	2101.93	8.835	8.944	0.00169
Sixense	227.936	1.219	4.271	0.00825
MotionPlus	4281.13	2.254	4272.98	0.202

3.2 Data Capture

A sequence of 3D positions p and orientations q are recorded, with timestamps of each sample. Positions are converted from the internal metric of the device to metres, while orientations are stored as quaternions. Orientations are transformed such that the local-space 'forward' reference frame of the device points along the user's outstretched arm - this is necessary as the Wii MotionPlus uses a coordinate system where the z axis is 'up', rather than 'forward'.

The Kinects are entirely camera based. This results in a lack of information when the camera cannot see the user's hand. The MoveMe uses camera tracking for positional data, so it too suffers from a drop-out in positional information when the Move 'wand' is occluded. Although orientation can still be tracked from the sensors contained within the Move device, for the purposes of this experiment they were ignored when the position was not determinable, to preserve the 1:1 mapping of positions to orientations. This is expected to be most evident when recording movements in the transverse plane, as the user's body obscures the camera's view of the hand / wand, with some additional occlusion expected with the sagittal plane recordings. While any loss of data is undesirable, the ability for the devices to redetect and track the desired movement after obscuration is a useful metric when considering the efficacy of motion devices.

4 Results and Evaluation

4.1 Steadiness

Steadiness measurements were made at a distance of 2m. Table 1 shows the sum of the changes between the samples in a 30s recording of the device at rest, and the change between the first and last samples. Immediately obvious is the poor result in movement from the Wii RemotePlus - the sum distance is several orders of magnitude larger than any other, due to accumulated error in double

integration of position from acceleration. Both the MoveMe and Sixense have large sum distances recorded, but, due to the low change in final position, it can be inferred that this was due to a high frequency jitter, rather than a large drift over time. The Wii RemotePlus reports an acceleration value without any processing for gravity; this was accounted for by taking a short sample of the accelerometer vector at rest, and subtracting this value, transformed by the orientation of the device. Evidently this was not sufficient to reliably remove gravity from the accelerometer reading, and further processing would be required to create a stable result.

The Kinect2 shows a large accumulated sum angle, despite a low distance travelled. As the orientation is determined from the skeleton (calculated by the Kinect software) this is a curious result, as it suggests that the 'parent' joint of the skeletal hierarchy has had a large drift or jitter in its position. The low distances for the Kinect devices are at odds with the other tested devices, suggesting additional filtering of position. The low changes in angle over time as compared to position for all devices is a good indicator of the relative difficulties in determining these values.

4.2 Data Sample Rate

The devices communicate with the host PC in a variety of ways: Kinect via direct USB connection, Sixense wirelessly to a USB connected base unit, and PS Move via Bluetooth to a PlayStation3 running the MoveMe server, which then contacts the host computer via TCP/IP. Due to this, and the differing ways in which position and orientation are derived, it is useful to compare update rates as reported from the interface software. The mean sample counts for each device across all recorded trial runs are collated in Table 2.

Table 2: All devices Sample Count Chart

Device	Coronal	Transverse	Sagittal
KinectV1	70.125	117.0	56.625
KinectV2	36.381	75.864	34.9
MoveMe	140.583	220.231	144.333
Sixense	137.667	188.267	142.933

The Kinect devices update at a significantly lower rate than the MoveMe and Sixense devices, which both report at a similar rate. The effects of this are discussed later.

4.3 Data Circularity

The circularity metric calculations for each device are shown in Figure 1.

Figure 1: Device Circularity Comparison Chart

4.3.1 Wii MotionPlus

The position of the Wii MotionPlus device can be determined in two ways: Either via double integration of its accelerometer data over time, or by processing the position of the IR light sources from the tracker bar, as seen by the Wii's onboard camera (broadly, the closer together the points, the further away the device is from the tracker bar). Unfortunately, the experiment as designed is unsuitable for position determination via the tracker bar, as the bar will never be in the line of sight when rotating in the coronal plane, and only briefly when rotating in the transverse or sagittal planes. Accurate derivation of position from accelerometer data is notoriously difficult, as any noise or bias in the accelerometer data will quickly accumulate into large errors in position. As gravity will be detected by the accelerometers, it must be accounted for, and any inaccuracy in this calculation will again result in large errors in position. As positional data could not be reliably calculated, no further tests were performed using the Wii MotionPlus device.

4.3.2 Kinect

Trials of the Kinect devices reveal that both have practically identical performance at 2m and 3m, however the Kinect 2 benefits from updated optics and processing that allow it to track users at 1m and 4m. Figure 2 illustrates this, with blue bars representing Kinect V2, and orange bars Kinect V1.

The Kinect scores at 2m are noticeably lower than those of the other tested devices (0.65 for Kinect 1, 0.52 for Kinect 2, compared to 0.81 for MoveMe, and 0.85 for Sixense). Splitting up the quality metric on a per-axis basis in Table 3, clarifies where the weaknesses in the Kinect sensor lie.

Table 3: Kinect Quality Per Axis

Kinect	Coronal	Transverse	Sagittal
Kinect 1	0.839	0.204	0.807
Kinect 2	0.834	0.154	0.784

Figure 2: Kinect Quality Distance Comparison

Both Kinect devices display poor mean performance in the transverse plane (Kinect 1 quality 0.204 Stdev 0.23, Kinect 2 quality 0.154 Stdev 0.143). This is unsurprising as rotations in this plane lead to the tracked hand becoming obscured. However, comparison with the MoveMe in the transverse plane (mean quality 0.68, Stdev 0.288), a similarly camera-based solution, reveals particular weakness in this plane.

Figure 3: Plot of Kinect 2 at 2 metres in transverse plane

The method whereby each device calculates its data must be considered. Both Kinects calculate an entire hierarchical skeleton of joint data, with the position of the wrist bone relying on the rest of the arm being correctly tracked, thus requiring more of the users arm to be visible than with the MoveMe, where the bright, uniquely coloured ball of the PS Move is found via simple computer vision techniques.

The quality scores of both Kinect devices are limited by the low sample rate of the devices. Kinect 2 in the transverse plane at 2m is shown in Figure 3. When compared to Figure 4, the lower sample density becomes apparent. The data also shows a phenomenon unique to the Kinect devices: as the user rotates

Table 4: Kinect Sample Count Per Axis

Kinect	Axis	1 M	2 M	3 M	4 M
1	Coronal	-	80.80	52.33	-
1	Trans.	-	133.40	89.67	-
1	Sagittal	-	62.40	47.00	-
2	Coronal	34.80	35.56	40.17	29.00
2	Trans.	69.83	79.00	77.00	78.00
2	Sagittal	36.00	31.57	34.50	38.50

such that their hand becomes obscured, the Kinect devices attempt to seek out the hand elsewhere in the image, frequently picking up the left hand or other part of the body, thus the cluster of data points towards the centre of the plot. This is in contrast to the MoveMe, which simply stops tracking until the ball is detected.

4.3.3 PlayStation Move

During the MoveMe tests, the 1m point was not successfully recorded, as the field of view of the camera was insufficient. Recorded orientation information was excellent, with orientation circularity of 0.95 in all 3 planes at 2m and 3m, dropping to 0.76 at 4m. Unlike the other devices, orientation is derived from sensors on board the handheld device, rather than a camera image or magnetic field, and so is unaffected by distance. However a reduction in image position tracking leads to fewer orientation samples, impacting the circularity. Position circularity at 2m is good (circularity 0.646, stdev 0.178), but displays a rapid drop-off at distance.

Figure 4: Plot of MoveMe at 4m and 2m in transverse plane

Tracking of the MoveMe starts to degrade at 3m. Degradation of z axis tracking is most prominent, while movement in the coronal plane remains accurate. Figure 4 shows an example of the MoveMe in the transverse plane. Immediately obvious is the large gap in the top left quarter of the graph, with the beginning of the gap denoting the point at which the move was obscured, and the end the

point at which tracking was regained, with an additional tracking fault. Data is also more ovular, with an approximately 2:1 ratio between minimum and maximum point lengths in the plot axis (position circularity: 0.45). This can be compared to the right hand plot, in which a transverse plane recording at 2m shows an approximate 1:1 ratio (position circularity: 0.807). The position detection limit of the MoveMe software appears to be approximately 4m. It was noted that further movement in the z axis caused a cut off effect, where no further movement in the axis is seen.

4.3.4 Sixense

Figure 5: Sixense Quality Distance Comparison

The Sixense is unique among the controllers as it can track position without the use of a camera. Its performance degrades quickly with distance, however. This can be seen in a graph of its combined score metric vs. distance, in which a combined score of 0.96, drops by 50% at 3m (Figure 5). This performance degradation is the worst over distance, and beyond 2m the Sixense displays the worst quality metric of all devices. At a distance of 1m (or even less, as the device is not constrained by a camera's field of vision) however, the device is the best performer, generating almost perfect circularity for both orientation and position.

Figure 6: Plot of Sixense at 1m and 4m in sagittal plane

Table 5: Sixense Circularity Axis Comparison

Axis	1 Metre	2 Metre	3 Metre	4 Metre
Coronal	0.9592	0.9521	0.4361	0.6648
Trans.	0.9847	0.9512	0.7023	0.3588
Sagittal	0.9497	0.6622	0.1560	0.1740

Table 6: Sixense Quality Deviation Axis Comparison

Axis	1 Metre	2 Metre	3 Metre	4 Metre
Coronal	0.0257	0.0225	0.3634	0.3739
Trans.	0.0064	0.0243	0.1657	0.2119
Sagittal	0.0198	0.1837	0.2114	0.1075

The results collated in Table 5 reveal a weakness in the sagittal plane. Table 6 shows the standard deviation of collected samples for each axis and distance. The quality of data extracted from the Sixense at 1m has particularly low deviation in quality in the transverse plane. At 4m, the sagittal plane results in a consistently poor quality metric. The extent to which the Sixense degrades is best seen visually - Figure 6 shows recorded data in the sagittal plane at 1m and 4m, with a breakdown in accuracy, making detection of a circular gesture impossible at the greater distance.

5 Conclusions

Domestic motion capture devices have great potential for rehabilitative gaming. In this paper we have presented a benchmark for measuring the suitability of such devices. We have applied that benchmark to the commonly available systems and shown that each has strengths and weaknesses. We conclude by discussing those relative strengths and make recommendations as to how best to utilise them to the needs of a specific rehabilitative project.

Of note is the excellent quality of data obtained from the Sixense device when used at distances close to the base unit, showing high accuracy in both position and orientation across all three planes. This suggests the Sixense would make an excellent device to use in situations where a high degree of fidelity in movement detection is required, such as detection of specific limb movements and poses. The device begins to heavily degrade at distance, so careful considerations must be made as to the environment it is used in; applications deployed via a laptop or tablet would be ideal conditions for the Sixense device.

The primary drawback of the MoveMe system is its reliance on expensive external hardware and a network connection, but in cases where greater working distances are required, the device produces quality positional and orientation information at 2m, with a degradation in performance beyond that. At 2m, all trial participants could be fully detected with limbs outstretched, making the MoveMe a good choice for living room interaction on a larger screen. Further

work will be required to determine whether tighter control of the environment, in regards to any patches of light within the MoveMe’s field of vision, would result in improved performance at range.

Although the Kinects score less well, notably in the transverse plane, it should be pointed out that these devices provide data on the entire shape of the body. When used at a distance of 2m, with no hand occlusion, these devices produce quality data. At 2m, the two devices are very similar, both displaying a weakness in recalculating body tracking sufficiently to determine hand position during rotations in the transverse plane. As with the MoveMe, further testing is required to determine the sensitivity of the Kinect to differing lighting conditions; however, several trials with the Kinect 2 showed the device tracking the left hand as the user turned away from the camera, suggesting that processing of the camera image for facial detection to infer orientation could have a positive effect on reducing false readings.

It is hoped that our results will serve as a recommendation for future projects utilising domestic gaming devices for rehabilitation.

References

- [1] BORGHESE, N. A., PIROVANO, M., MAINETTI, R., AND LANZI, P. L. An integrated low-cost system for at-home rehabilitation. In *Virtual Systems and Multimedia (VSMM), 2012 18th International Conference on* (2012), IEEE, pp. 553–556.
- [2] CHANG, C.-Y., LANGE, B., ZHANG, M., KOENIG, S., REQUEJO, P., SOMBOON, N., SAWCHUK, A. A., AND RIZZO, A. A. Towards pervasive physical rehabilitation using microsoft kinect. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2012 6th International Conference on* (2012), IEEE, pp. 159–162.
- [3] GRAZIADIO, S., DAVISON, R., SHALABI, K., SAHOTA, K., USHAW, G., MORGAN, G., AND EYRE, J. Bespoke video games to provide early response markers to identify the optimal strategies for maximizing rehabilitation. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing* (2014), ACM, pp. 20–24.
- [4] HANSEN, P. K. Method and apparatus for position and orientation measurement using a magnetic field and retransmission, Feb. 10 1987. US Patent 4,642,786.
- [5] JOO, L. Y., YIN, T. S., XU, D., THIA, E., CHIA, P. F., KUAH, C. W. K., AND HE, K. K. A feasibility study using interactive commercial off-the-shelf computer gaming in upper limb rehabilitation in patients after stroke. *Journal of Rehabilitation Medicine* 42, 5 (2010), 437–441.

- [6] LANGE, B., FLYNN, S., AND RIZZO, A. Initial usability assessment of off-the-shelf video game consoles for clinical game-based motor rehabilitation. *Physical Therapy Reviews* 14, 5 (2009), 355–363.
- [7] PIROVANO, M., MAINETTI, R., BAUD-BOVY, G., LANZI, P. L., AND BORGHESE, N. A. Self-adaptive games for rehabilitation at home. In *Computational Intelligence and Games (CIG), 2012 IEEE Conference on* (2012), IEEE, pp. 179–186.
- [8] REGO, P., MOREIRA, P. M., AND REIS, L. P. Serious games for rehabilitation: A survey and a classification towards a taxonomy. In *Information Systems and Technologies (CISTI), 2010 5th Iberian Conference on* (2010), IEEE, pp. 1–6.
- [9] SAPOSNIK, G., LEVIN, M., ET AL. Virtual reality in stroke rehabilitation a meta-analysis and implications for clinicians. *Stroke* 42, 5 (2011), 1380–1386.
- [10] SMISEK, J., JANCOSSEK, M., AND PAJDLA, T. 3d with kinect. In *Consumer Depth Cameras for Computer Vision*. Springer, 2013, pp. 3–25.
- [11] TANAKA, K., PARKER, J., BARADOY, G., SHEEHAN, D., HOLASH, J. R., AND KATZ, L. A comparison of exergaming interfaces for use in rehabilitation programs and research. *Loading...* 6, 9 (2012).
- [12] USHAW, G., ZIOGAS, E., EYRE, J., AND MORGAN, G. An efficient application of gesture recognition from a 2d camera for rehabilitation of patients with impaired dexterity. In *8th International Conference on Health Informatics* (2013), Springer.