IMMERSIVE PRODUCTION AND DELIVERY OF INTERACTIVE 3D CONTENT

DELIVERABLE D6.3:
SECOND DEMONSTRATOR & EVALUATION

Contractual Date of Delivery:	November 30, 2014
Actual Date of Delivery:	November 30, 2014
Work Package:	WP6
Dissemination Level:	Public
Nature of Deliverable:	Report
Lead Contractor:	BBC
Coordinating Person:	Graham Thomas Graham.Thomas@bbc.co.uk
Project Leader:	BBC
Work Package Leader:	ART
Authors:	Graham Thomas, Florian Schweiger (BBC) Marco Volino, Peng Huang (Surrey), Vincent Jantet (Artefacto)

Classification and Approval

Classification: Public

This document has the status 'Public' and may be published as a whole including this page. Publication of parts of the document must be in unanimous agreement within the RE@CT Steering Board and subsequent EC approval/agreement.

Disclaimer

Neither the RE@CT Consortium nor any of its officers, employees or agents shall be responsible or liable in negligence or otherwise howsoever in respect of any inaccuracy or omission herein.

Without derogating from the generality of the foregoing neither the RE@CT Consortium nor any of its officers, employees or agents shall be liable for any direct or indirect or consequential loss or damage, personal injury or death, caused by or arising from any information, advice or inaccuracy or omission herein.

Acknowledgements

All partners of the RE@CT project contributed to this report during the last months.

RE@CT Consortium Overview

Participant no. *	Participant organisation name	Short name	Country
1 (Coordinator)	British Broadcasting Corporation	BBC	UK
2	Fraunhofer HHI	HHI	Germany
3	INRIA	INRIA	France
4	University of Surrey	Surrey	UK
5	Artefacto	ART	France
6	OMG	OMG	UK

Abstract

The RE@CT project aimed to revolutionise the production of realistic characters and significantly reduce costs by developing an automated process to extract and represent animated characters from actor performance capture in a multiple camera studio. The key innovation introduced by RE@CT is the development of methods for analysis and representation of 3D video to allow reuse for real-time interactive animation. This will enable efficient authoring of interactive characters with video quality appearance and motion.

RE@CT demonstrated its results in two application scenarios: an augmented reality application to demonstrate usage for serious gaming in education and entertainment, and a TV-like production to create interactive learning material for an interactive website about ballet.

This deliverable describes the second of the project's two demonstrators, including how the ballet material was captured, a description of the interactive web application, and a summary of the public demonstration of the work at CVMP2014 in London. An evaluation of the project's results is also included, which looks at the reduction in cost (in terms of production time) that the RE@CT technology offers over conventional modelling and animation approaches, and compares the quality of the rendered models to an equivalent video version. Comments from end users are also included, which give some insights into other benefits of the approach developed within the project.

Table of Contents

1 Introduction

In order to provide a set of tangible use cases to focus the project's work, and to provide material to support dissemination and user testing, the RE@CT project included two public demonstration events, each preceded by the preparation of a demonstrator application. The first of these was based around a cultural heritage augmented reality game. It was publicly demonstrated at the Mirage 2013 conference and described in Deliverable D6.2, published in June 2013. The second demonstrator was based around ballet, with a solo dancer being captured as a re-animateable 3D model, to provide an interactive on-line application that teaches the user about ballet and allows them to choreograph their own dance sequence. It was presented publicly for the first time at CVMP 2014, in London.

This deliverable describes the preparation of the final demonstrator and its presentation at CVMP 2014 to attendees from the VFX, media production and academic communities in Month 36 of the project (November 2014). It also presents an evaluation of the technologies developed in the project, in terms of the potential cost reduction over conventional approaches, the visual quality of the rendered characters, and feedback from users.

2 Use case for the final demonstrator

Interactive, or games-like, applications can be particularly valuable in educational entertainment, where they can help to keep a viewer engaged; the first project demonstrator (described in deliverable D6.2) was an example of such an application. For the project's final demonstration, we decided to explore this application area further, using RE@CT technology both to allow a viewer to explore 3D content and to create their own animation (based on the well-known idea of 'learning by doing').

The idea for the demonstrator was proposed by the BBC's Knowledge and Learning department, who have recently developed a set of web-based material under the 'iWonder' banner[1,2]. The 'iWonder' pages include text, images and video, and the idea was to explore how more interactive material might be used.

The demonstrator is based around an iWonder-style page that seeks to explain the rudiments of ballet. Working with a choreographer (Jack Thorpe-Baker) and a ballet dancer (Caroline Crawley), a 1.5-minute dance piece was recorded that presented a number of different moves. The first part of the demonstrator allows the viewer to watch this dance, freezing the action at various points and exploring the dancer's pose by freely moving the viewpoint around. Explanatory text can be overlaid at relevant 3D locations, to explain the details of particular moves. Once the viewer has learned the rudiments from this, they move on to the second part, where they are given the option of choreographing their own dance, using a set of five individual dance moves. These moves can be put together in any order, producing a seamless dance that can be viewed in 3D from any location around the dancer.

[1] http://www.bbc.co.uk/iwonder

[2] http://www.bbc.co.uk/blogs/internet/posts/BBC-iWonder-home-page-launched

3 Capture of content

Content for the final demonstrator was captured in the BBC's multi-camera studio on 23rd September 2014, following a rehearsal day the previous week. Full details of the capture process were given in Deliverable 2.3. As D2.3 was a non-public deliverable and this deliverable is public, a short summary of the key points are included below.

3.1 Full-body capture

The general capture setup was similar to the one used in previous RE@CT productions, with the exception that the studio had been relocated and now had different, slightly smaller dimensions. The lower ceiling and the less spacious truss particularly made the installation of cameras and studio lights more challenging.

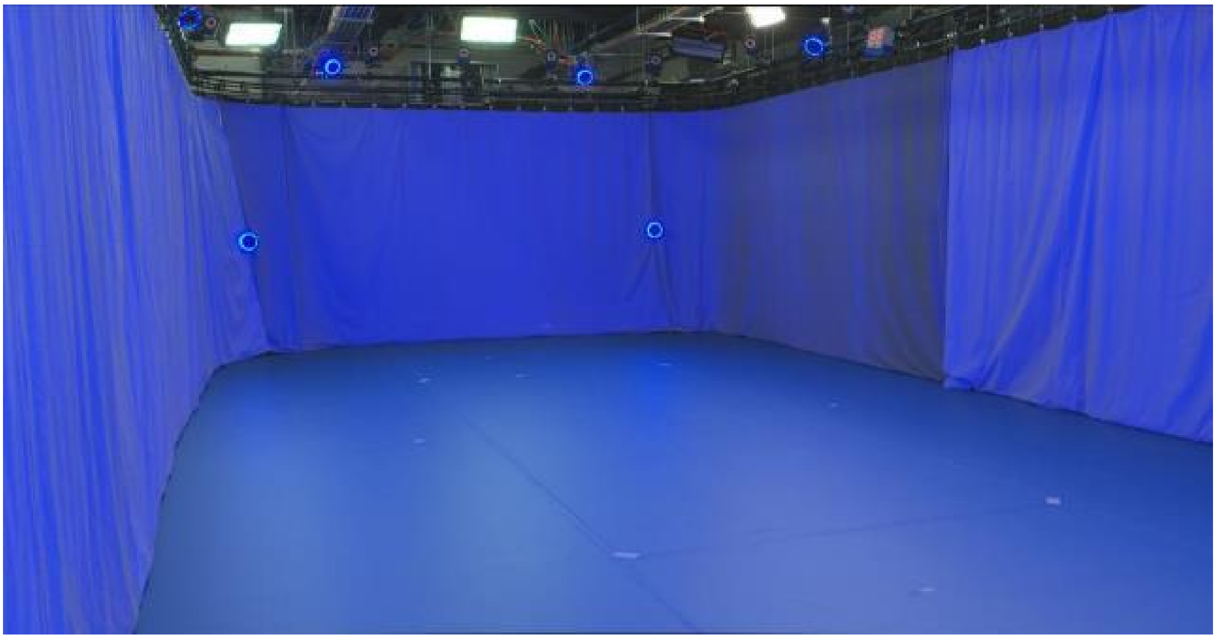

Illustration 1 - the capture studio

Illustration 1 shows the empty studio immediately before a capture run, the space effectively available measuring 11 by 5.7 metres, with the curtain rails at a height of 2.8 metres. As previously, the space was enclosed with retro-reflective curtains to facilitate keying. Due to the requirements of this capture session, a professional dance floor with vinyl flooring was used instead of using retro-reflective floor tiles.

The camera setup comprised nine full-HD cameras (1920x1080/25p), evenly distributed around the space just above the curtain rail, and four additional UHD cameras (4096x2160/50p). The latter were placed at hip height on tripods in the corners of the studio and replaced the manually operated pan-tilt-zoom cameras that were used in previous productions to obtain higher resolution textures. The advantage of this static configuration was that the curtains could be fully closed and the cameras hidden (see Illustration 1) so as to avoid disturbing background motion. Due to the limited ceiling height, a central overhead camera pointing down was not used this time. All cameras were equipped with blue LED rings to illuminate the retro-reflective curtains. Illustration 2 and Illustration 3 show examples of images from the full HD and UHD cameras.

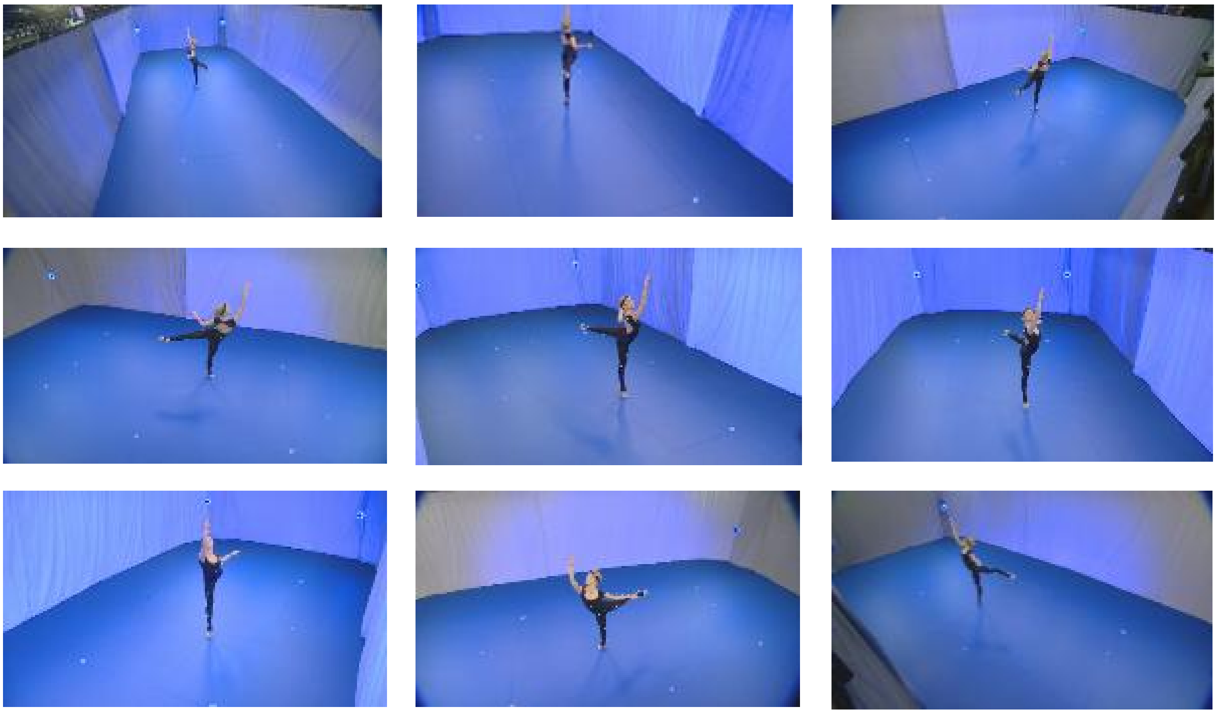

Illustration 2 - Stills from HD cameras

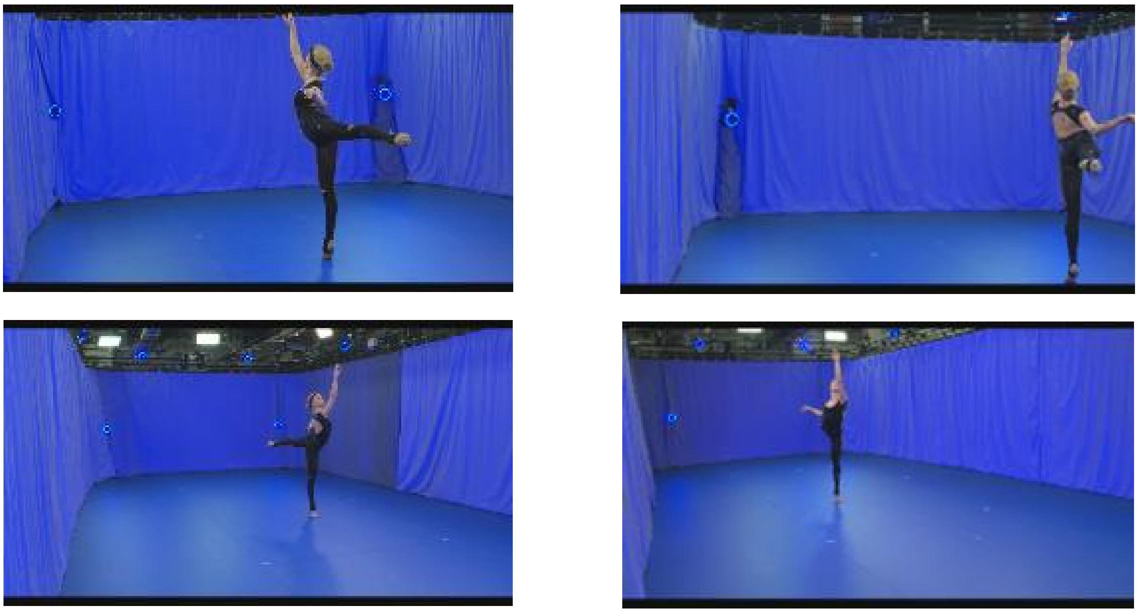

Illustration 3 - Stills from UHD cameras

To act as ground-truth for evaluation of surface capture methods, separate recordings were made with the dancer wearing motion capture markers, to obtain high-resolution marker-based full body kinematics using a Vicon motion capture system. This data was not needed for the demonstrator itself, but will make the overall data set much more valuable for future research. An example of this process is shown in Illustration 4.

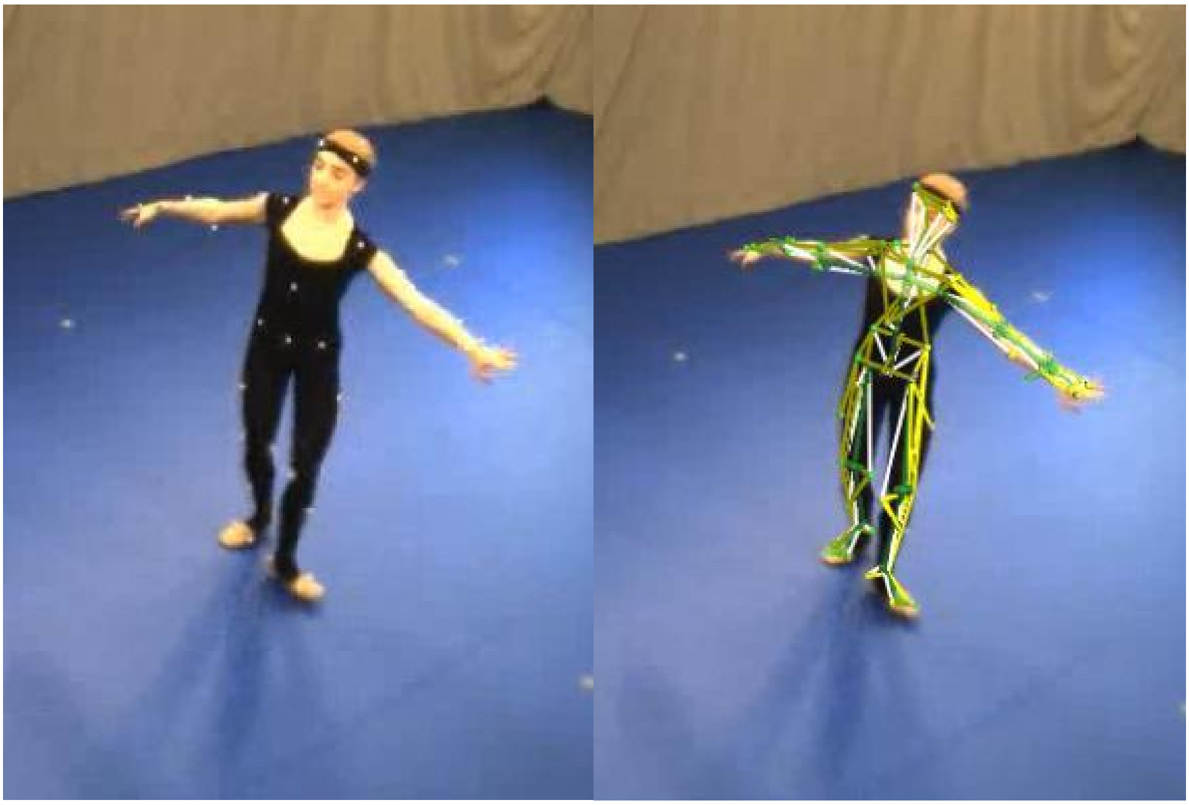

Illustration 4 - A still from the Vicon video reference footage. On the right a calibrated kinematic model, tracked from markers, has been overlaid.

3.2 Face capture

To provide content from which a detailed face/head model could be created, a set of stereo image pairs was captured from 7 DSLR-pairs (Canon EOS 550D) placed around the performer to cover roughly 270° of frontal and profile view. We also captured the performer from behind, providing the means to create a 360° model.

Illustration 5 shows the capture process and an example of an image set.

Illustration 5 - Capturing stereo pairs for head modelling (left) and a set of captured images (right)

In addition to the stills capture, we also recorded video streams from the four UHD cameras of the performer displaying different emotional states likely to be needed in a re-animation session. This will allow us to provide more natural texture sequences for these emotions. Example frames from one of the captured sequences are shown in Illustration 6.

Illustration 6 - Images from the UHD video capture of facial expressions

4 Demonstration application

From the captured HD and UHD ballet footage, several full 4D datasets have been produced using the algorithms described in previous deliverables. The key steps included foreground-background segmentation, silhouette-based voxel carving, tracking of rigid body parts, and the stereo-based reconstruction of face details.

In order to facilitate the visualisation of these datasets a WebGL-based Character Animation Engine (CAE) has been developed which allows real-time, client-side rendering and interactive control of virtual RE@CT characters. WebGL was chosen as the ideal technology for this due to its ever increasing cross platform and device support which now spans computers, mobile devices and set-top boxes. Illustration 7 shows a system overview of the WebGL CAE.

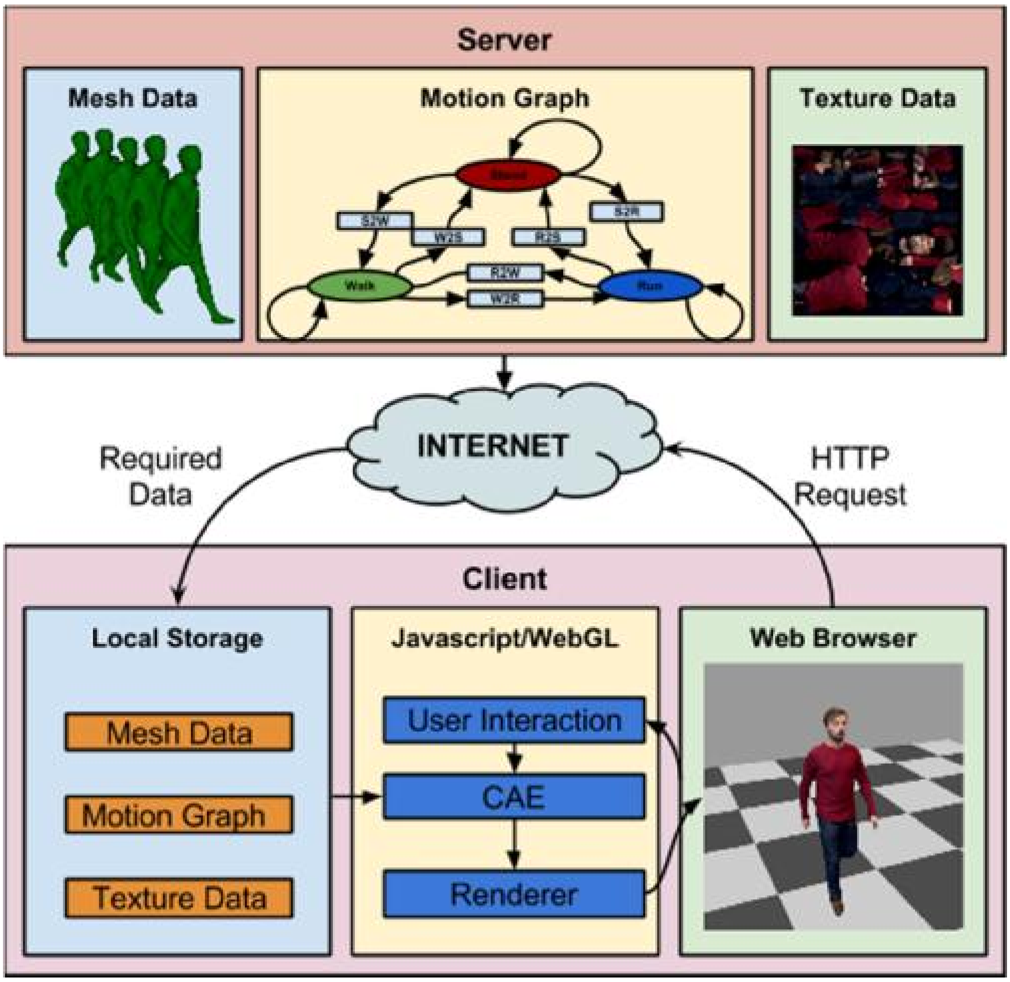

Illustration 7 – Overview of the WebGL character animation engine

Geometry, texture and motion graph data are stored on a web server. When a client requests the webpage the data is transferred via the HTTP protocol and stored locally. The motion graph-based CAE maintains and updates the current state of the model based on user interaction. Frames are sent to the renderer to be displayed in the browser window. Parametric control of the motion is achieved by blending selected pairs of related frames using hardware-based geometry blending. Suitable frame pairs and transitions are determined offline using a shape similarity measure and stored within the motion-graph itself. This allows real-time parametric control of virtual characters and forms the foundation of online game-like environments using 4D content from the RE@CT project.

Based on the WebGL character animation engine and the reconstructed 4D ballet data, a two-part interactive demonstrator has been implemented. Following the aims outlined in Section 2, both parts have been embedded in a website in the style of the BBC's educational iWonder product.

Illustration 8 – Annotation of a dance routine in the first part of the demonstrator

The first part of the demonstrator allows the exploration of a ballet performance in 3D. The user is presented with a video of the dance in a standard web video player. At key moments, he or she is offered the possibility to pause the video and enter a 3D exploration mode which causes a cross-fade from the video pane to an underlying WebGL canvas. The latter is controlled by the character animation engine and hence reacts to mouse and keyboard input allowing the user to move the virtual camera around the stationary 3D model of the dancer. The 3D world is augmented with annotations, prepared with authoring tools described in deliverable D5.4, providing additional information about the ballet routine that is being examined.

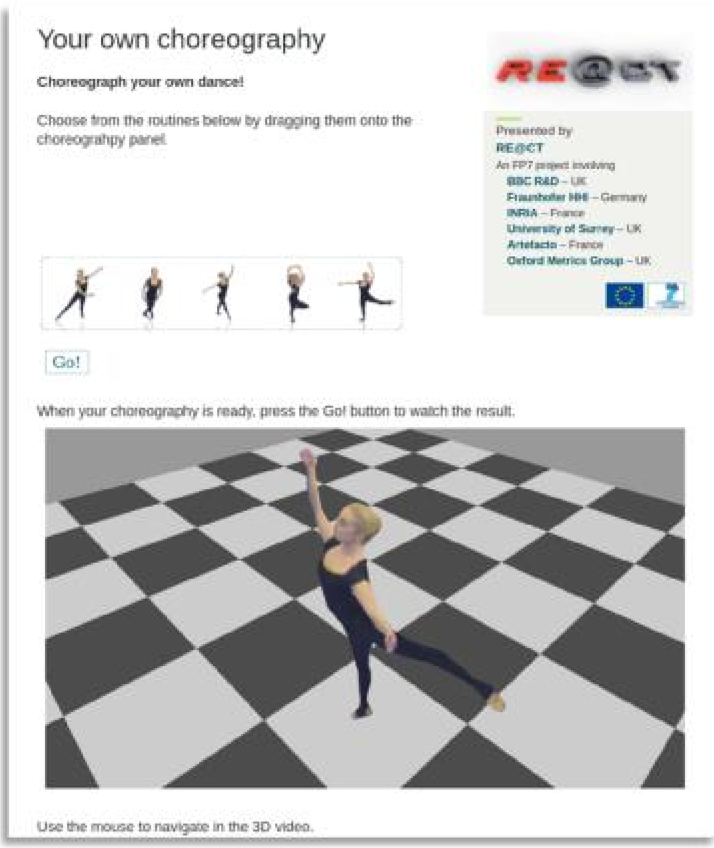

Illustration 9 – Drag-and-drop interface of the second part of the demonstrator

The second part of the demonstrator consists of a drag-and-drop interface with icons that represent several basic ballet routines. These can be freely rearranged to create an individual choreography which is then rendered with the WebGL character animation engine. Thanks to the transition information stored in the CAE's motion graph, blending from one dance routine to the next is seamless. Using mouse and keyboard commands, the user can again move the virtual camera around the 4D model of the dancer while the choreography is executed.

5 Final dissemination event

It was decided to hold the final dissemination event in conjunction with the European Conference on Visual Media Production (CVMP[3]), held on 13-14th November 2014. This event is in its 11th year, and is held annually in central London. It brings together a cross-section of academic and industrial researchers with practitioners in the film, TV and games production communities. It has a particularly strong attendance from those focused on visual effects, helped by its proximity to the post-production community in Soho. The industrial attendees therefore cover a market segment highly relevant to the commercial application of RE@CT technology.

Illustration 10 - The RE@CT demo are at CVMP 2014

The project was given a demonstration area containing two large tables, a large monitor and poster boards, shown in Illustration 10. The following demonstrations were shown:

- The WebGL-based 'choreograph your own dance' demonstrator (BBC/all)
- A PC application for rendering animated sequences of captured actors using motion graphs and 4D Video Textures (Surrey)
- The authoring tool for creating sequences of animations, showing how these can be uploaded to an augmented reality iPad application to display the animated character standing on the game board from the first demonstrator (Artefacto)
- Video sequences showing rendered body and head models (INRIA and HHI)
- The 'Cara' head-mounted camera system and associated software (OMG)
- A video playing on the large screen showing behind-the-scenes shots from the final test shoot and examples of the captured models and animations (BBC/all)

Some 3D printed models created from data from the test shoots were also brought along; the level of detail in the head models attracted particular attention from a number of attendees.

In addition to the demonstration (which ran during morning, lunch and afternoon breaks on both days, as well as during the drinks reception on the first day), the project gave a 30-minute

[3] www.cvmp-conference.org

presentation during a special session on EU projects. The slides used in this presentation, and the ‘behind the scenes’ video, were subsequently made available via the project website.

The demonstration attracted considerable interest from attendees. Illustration 11 shows demonstrations being given to visitors including the CTO of a motion capture studio. Illustration 12 shows some 3D models that were printed from the capture sessions, and the Vicon ‘Cara’ head-mounted capture system.

Illustration 11 - CVMP attendees studying the demos

Illustration 12 - 3D printed models of performers (left) and the Vicon 'Cara' head camera system (right)

6 Evaluation

The project set out with the aim of reducing the cost for producing interactive characters whilst increasing the level of realism, aiming for a quality equivalent to that of a video recording of the actual actor's performance. This section reviews what the project has achieved in respect of these aims, and also looks at feedback from professional users of the content creation system.

6.1 Cost Reduction

It is difficult to make a comparison between an established tool pipeline and a prototype set of research tools - the direct comparison is likely to be highly misleading in terms of the time and resource required. Tools for the use of motion capture data in game character production have evolved over the past 20 years. Future generations of video-based production tools with improved interfaces as well as the processing steps are likely to considerably reduce the time and cost required vs. a prototype implementation.

Nevertheless, it is possible to provide some approximate figures for the time taken for various stages of the processing pipeline for both traditional animation and for the RE@CT pipeline. For the RE@CT pipeline, we looked at how long each process takes at present, and how long it might reasonably be expected to take if the processing and user interfaces were optimised. We based the figures on the experience gained when creating the animated characters for the first demonstrator (AR cultural heritage game), for which animated characters were produced using both traditional techniques and RE@CT techniques. The table below compares the time for each processing step:

Process	Traditional modelling & animation pipeline	Current RE@CT pipeline	Potential future optimised RE@CT toolset
Planning of capture process	5 days (same for all approaches as this is essentially a storyboarding / planning process). For the RE@CT approach this includes designing a motion graph for each character (including motion selection and motion path selection, e.g. walk->walk2stand->stand).		
Studio-based rehearsal and capture session	5 days (same for all approaches: most of the time is recording the actor's moves; the small amount of time spent on fixing markers to actors for mocap is roughly equivalent to time spent doing make-up for RE@CT approach). Processes differ slightly, e.g. need to calibrate bone lengths in conventional mocap using captured 'Range of Motion' moves.		
Review of captured material and selection of good "takes"	2 days (same for all approaches)		
Obtaining animation data	"Cleaning" mocap data: 2 days per set of sequences Adapting data to fit skeleton of animated character: 5 days for 10 animations per character	Segmentation, reconstruction, sequence alignment, tracking: 1 day per set of sequences Building motion graph for each character is done automatically (main processing time is	Segmentation, reconstruction, sequence alignment, tracking: 1-2 hours per set of sequences Building motion graph : about **1 hour** per character

Process	Traditional modelling & animation pipeline	Current RE@CT pipeline	Potential future optimised RE@CT toolset
		creating frame-to-frame similarity matrix) which takes 1-2 hours per character with 10 motions.	with 10 motions
Creating textured model	**10 days** for the character of "Richard Lion King"	Choosing how to map textured surface to 2D texture patches (per sequence): a few minutes using automated method (e.g. in Maya or Blender), or about an hour if manually cutting up the mesh. Generating texture maps: seconds to hours of computation time per frame depending on the approach. Assuming 2 minutes/frame, 4 CPUs in parallel and 10 2-second sequences at 25 fps would be around **4 hours**.	Assuming an average time of 15 secs per frame (1 minute per frame on a single CPU, running in parallel of 4 cores) would give a total of around **2 hours** for 10 animations[4]
TOTAL	**29 days**	**14-15 days**	**12-13 days**

Table 1 - Comparison of times taken for creating animated characters for the RE@CT first demonstrator

The table shows that the most significant time saving comes from creating the textured models: a manual process that took about 10 days using a traditional approach but only a few hours of automated processing time using the RE@CT pipeline.

6.2 Visual quality

One of the goals of the RE@CT project is to create realistic character animation. Illustration 13 shows an example of compositing a rendering of 4D video-based character animation into a real scene. The character successfully preserves the quality of captured video and seamlessly merges into the real scene. This demonstrates the potential use of the RE@CT production pipeline to deliver next-generation high quality character animation for gaming, film and TV broadcasting applications.

[4] 10 animations, each about 2 seconds long at 25 frames/sec, is a total of 500 frames.

Illustration 13 - Example of composition of 4D video-based character animation and real scene

Objective evaluation has been performed on four different multi-view video datasets shown in Illustration 14. Character1 and Dan were captured using eight cameras in a circle of 8 metre diameter giving full 360 degree coverage of the subject with reconstruction [CGA07] and temporal alignment [IJCV12]. Face and Cloth were captured using five cameras in a frontal configuration with reconstruction and alignment [ECCV12]. All video in datasets was captured at 1920x1080 HD-SDI at 25P and available for download for research purposes from cvssp.org/cvssp3d.

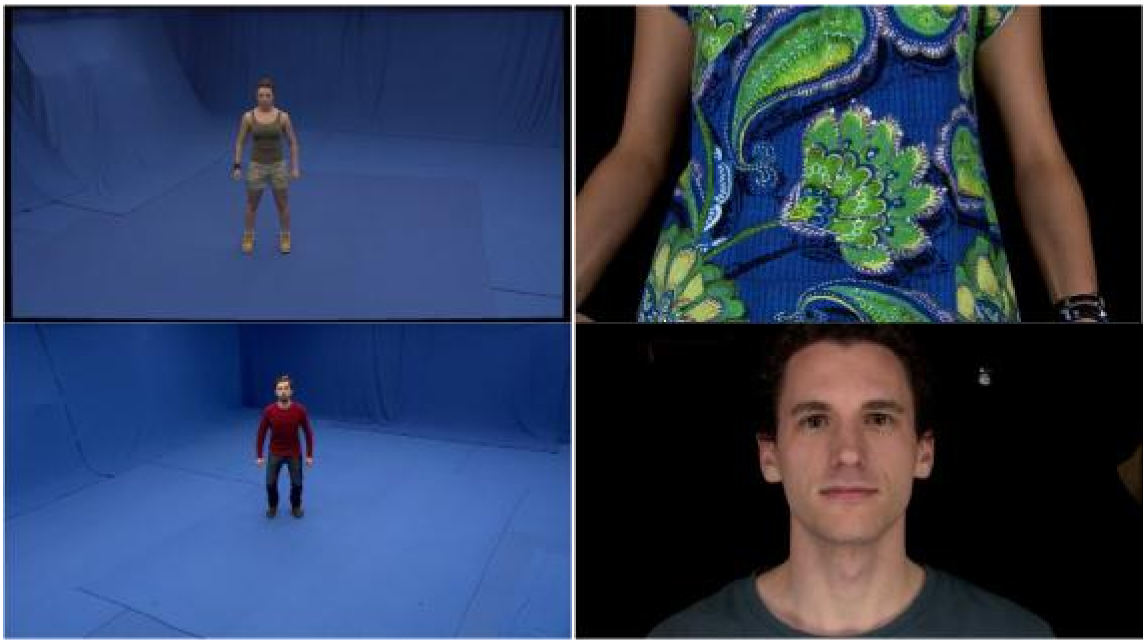

Illustration 14 - Single-camera image from evaluated datasets: Character 1; Cloth; Dan; Face.

Texture Alignment: Results of the surface-based optical flow alignment [BMVC14] are presented in Illustration 15 for the Dan, Face and Cloth examples: a conventional texture map is shown before (left) and after (right) alignment with a heat map (centre) highlighting the difference. Before alignment large misalignments exist between different views, visible in the close-ups, which produce ghosting and blurring artefacts during rendering. After alignment these errors are corrected resulting in a sharp texture as is visible in the close-ups for the cloth example.

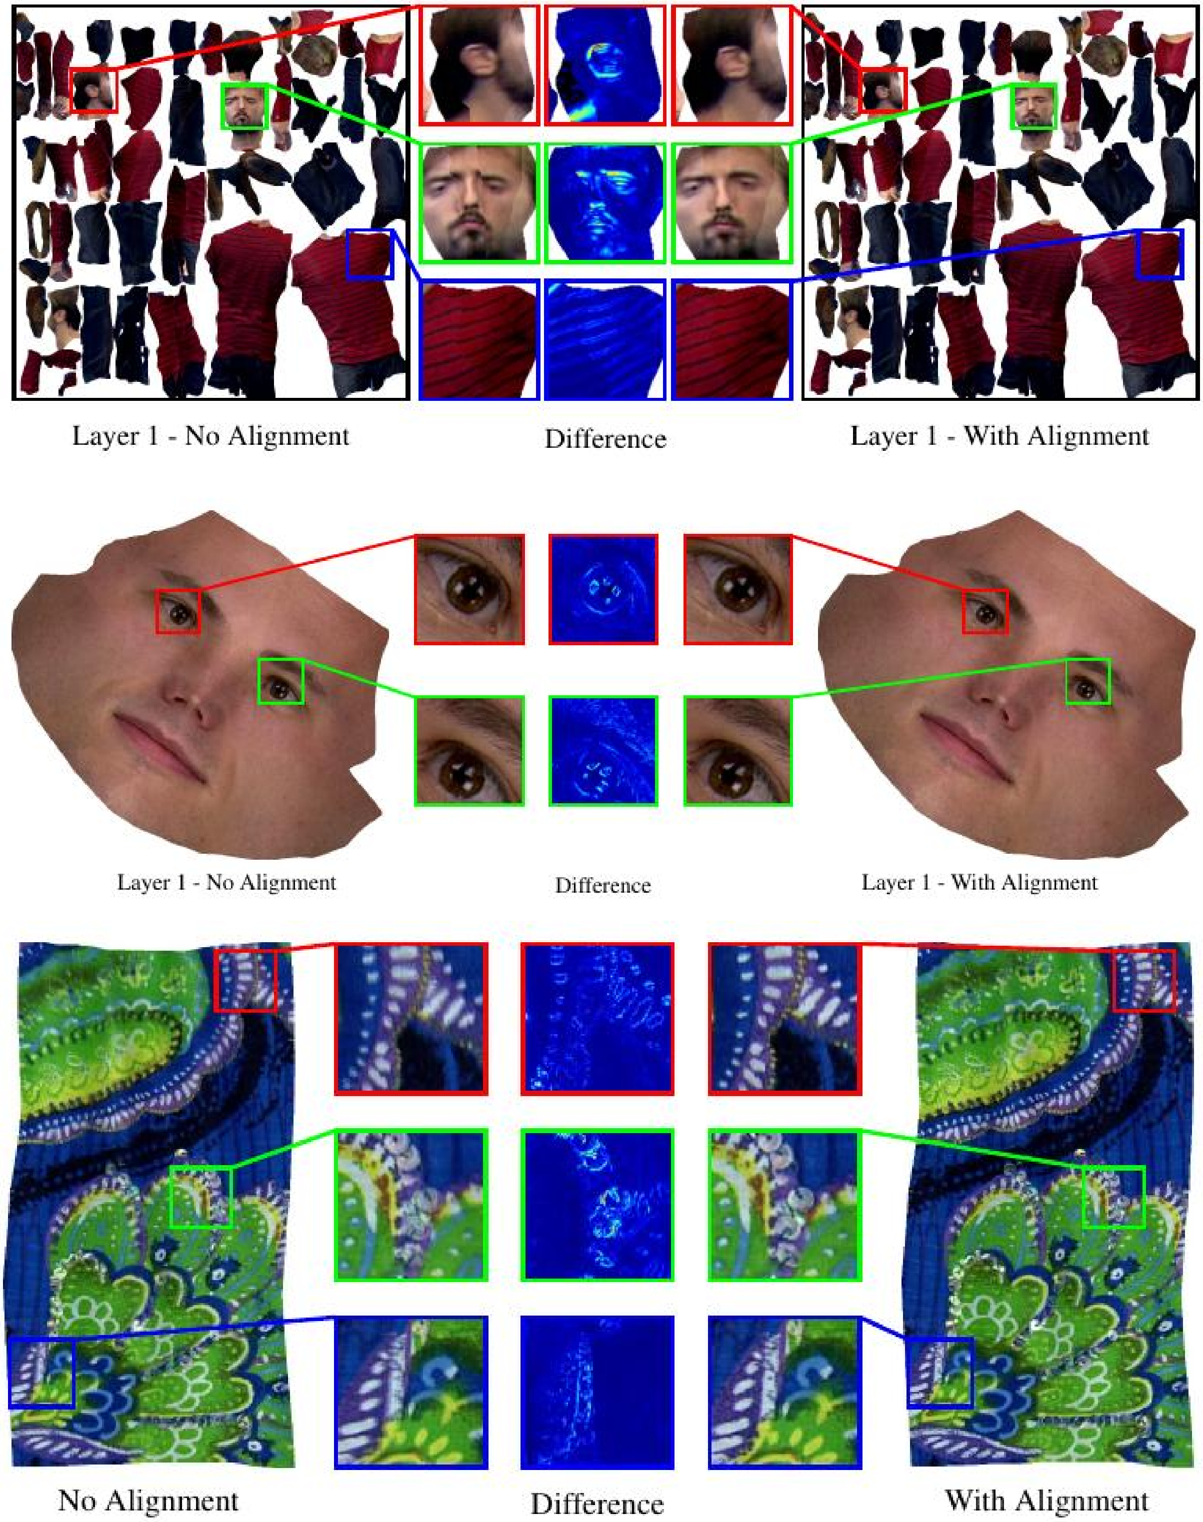

Illustration 15 - Example frames for free-viewpoint video texture with and without optical flow guided texture alignment

Rendering Quality: Free-viewpoint Video Rendering (FVR) quality with the multi-layer texture [BMVC14] vs. direct resampling of the captured multi-view video was evaluated using an open-source render FVVR [JG09] as a benchmark. The Structural Similarity Index Measure (SSIM) [TIP04], which has been shown to correlate with perceived image quality, was used to evaluate the rendering quality. Evaluation was performed for views mid-way between the capture cameras to test the hardest FVR case. Illustration 16 presents two evaluations of rendering quality for Dan/Cloth datasets: (a) with respect to resampling optimisation approach (no-optimisation (NO), spatial optimisation only (SO) and spatio-temporal optimisation (TO)); and (b) with respect to texture image resolution using TO. This demonstrates that the optimisation method has no effect on the rendering quality for number of layers (N_L), $N_L > 2$, as it is essentially the same texture information just assigned to different layers. Secondly evaluation of texture resolution shows that rendering quality remains the same for sizes > 1024x1024. Importantly for the Dan character dataset captured, with surrounding cameras, rendering quality remains constant for $N_L \geq 3$ indicating only 3 layers are required.

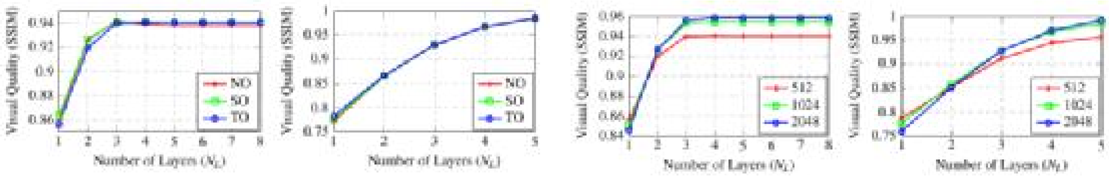

(a) Evaluation of optimisation: Dan at 512 resolution (left); Cloth at 1024 resolution (right)

(b) Evaluation of resolution: Dan using TO (left); Cloth using TO (right)

Illustration 16 - Evaluation of optimisation and resolution for Dan and Cloth datasets

User Aesthetic Preference: 51 non-expert participants undertook a web-based survey to evaluate aesthetic preference for videos and stills generated by the 4D video texture (4DVT) framework [EG14]. Participants were asked to compare 30 randomized separate rendering pairs, each pair produced under two different algorithm configurations, and rate them from 1 (preference for the left) to 5 (preference for the right), 3 for no preference. Identical pairs were also inserted as a control, and we observed preferences of 3.15±0.67 indicating natural variations in response. Perceptual impact of any deterioration in rendering quality induced by 4DVT layered representation is evaluated in Illustration 17 (left). Overall slight preference for FVVR (3.41±0.94) was observed in still imagery. This may be explained by a slight drop in visual fidelity caused by 4DVT in still imagery. However this difference becomes less perceptible in video (3.25±0.75), see supplementary video. We next evaluated 15 pairs of real 4D captured character (rendered using FVVR) alongside a synthesized 4DVT character, Illustration 17 (right). Overall the output is judged to be approximately equal (3.30±1.10) indicating the plausibility of the synthetic character. Interestingly, the Knight dataset, which was captured in a studio with lower resolution cameras than the others, presents greater preference for real imagery (3.61±1.03). This is caused by the increased noise in the optical flow due to low-res textures, resulting in poorer 4DVT alignment. Nevertheless, the overall mean is approximately neutral (3.30) with a high standard deviation (1.10), indicating little perceivable difference between synthetic output and real imagery. Finally, a synthetic 4DVT animation was shown to the participants, who were asked to rate it from 1 (very artificial) to 5 (very realistic). The overall score of 3.97±0.69, indicates good level of acceptance.

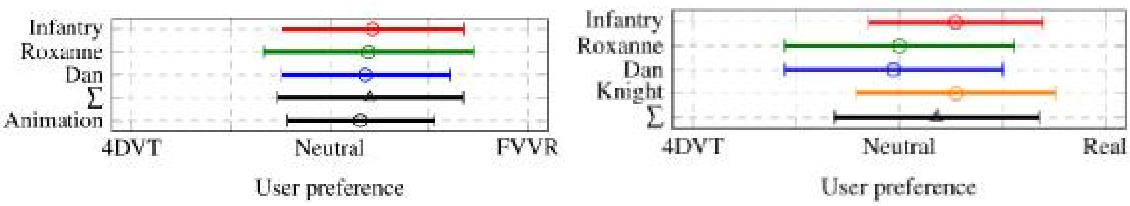

Illustration 17 - User preference for 4DVT vs. FVVR (left) and 4DVT vs. Real Capture (right)

6.3 Feedback from users

The following comments were made by members of the BBC production team involved in the final test production, when asking them for their feedback:

How did the recording compare to making conventional TV, and were there any aspects that could be improved?

- The planning process was similar to a conventional production in terms of booking and working with on-screen talent. The process of blocking through a performance before recording was not at all unusual, despite the fact that there were tens of cameras to cater to, and a very fixed capture volume.
- Although the small capture volume could have posed difficulties for a dance sequence, in practice it was extremely helpful that the performer could pause, reset their position within the volume and then resume the motion, in the knowledge that the full sequence could be later reconstructed.
- Working in a conventional TV studio, and being able to use existing skills such as wardrobe, hair and make-up artists, was very useful with being able to fit in with traditional TV workflows (as distinct from having to employ CG designers for clothing and make-up).
- Having the action recorded from multiple viewpoints was actually incredibly useful to both the dancer and the choreographer, as it was possible to play back the action from any desired angle, to troubleshoot any aspects of the performance which required it.
- It would have been useful to have some sort of time-stamped 'out of volume!' flag recorded with the video, so that it was straightforward to find, review and adjust any problem areas of the performance.

What are your thoughts about the realism of the results?

- The dance sequence we have reconstructed is very believable in terms of the realism of its motion. There is a fluidity to it which is very strongly reminiscent of the original performance. Although the 3D characters we are producing at present are not highly detailed, they come across as being low-resolution reconstructions of real people rather than computer generated characters. There is some novelty in viewing such content in free-viewpoint, as it is rare at present, and I believe the reconstructions themselves are the more engaging for it being apparent that the reconstruction is of a live, human performance.
- My biggest concern at the moment is that the models and textures are not very detailed, which will theoretically be addressed by camera technology becoming better/cheaper. It is also a concern that occlusion (e.g. by some outfits) of parts of the actor, results in imperfect reconstructions, and the dance reconstructions I have seen seem to show some limbs bending unnaturally. At present the nature of the technology requires careful planning and adjustment in some cases, of the performance. I would like to see the technology improve to the extent that the performer can move naturally with the entire range of motions available to them in life.

What other application areas might there be for RE@CT technology, or free-viewpoint video in general?

- Instructional films - view something from every angle
- Sport - for the same reason
- Constructing stories which can be experienced through the viewpoint (eye-position) of any given character
- If multiple people can be captured simultaneously, then in theory, it would be possible to recreate the viewpoint of someone in a crowd, leading the viewer to feel that they were present during the action

- Impossible shots – e.g. sports footage where the camera flies around the participants of a race, single-take, extreme, moving close-ups of a fight sequence where, at present a camera would simply get in the way.

Comments were gathered by attendees at CVMP who used the authoring tool developed by Artefacto. Visitors tended first to comment on the lack of detail in the oldest reconstructed dataset. They were more impressed by the automated reconstruction pipeline, and became enthusiastic when they started to configure the motion, and blending the sequences. They made encouraging remarks about the latest reconstructed model: "mixing the body reconstruction from wide shots with face reconstruction from close-ups drastically improves the realism". They enjoyed playing with the real-time functionality of the software, observing the smooth transitions between the sequences: "The resulting trajectories are customizable but always keep a natural and realistic motion". They tended not to explore the timeline capability of the Authoring Tool in much detail.

7 Conclusion

This Deliverable has presented a summary of the final demonstrator from the RE@CT project and the presentation of project results at CVMP in London in November 2014. The project demos attracted interest from a range of visitors working in the post-production and visual effects industries.

An analysis of the production time for an animated character using RE@CT techniques compared to a traditional modelling and motion capture techniques showed a likely drop in production time from 29 days to 12-13 days – a saving of over 50%.

The results of an objective and subjective evaluation of the quality of the rendered characters were also presented. They showed that the 4D Video Texture approach with just 3 texture layers gave visual quality very close to that from a free-viewpoint video render using all available cameras, and there was little or no statistically-significant difference in the perceived video quality.

Comments from production staff at BBC, and industry attendees at CVMP, showed that the content capture process fitted in well with conventional TV workflows, and that the system opened up many new possibilities for innovative forms of interactive content. The main area for improvement that was identified was the texture resolution on the animated characters, when they are viewed close-up. This should improve over time as camera technology improves (indeed, these judgements were based on models that did not use any UHD cameras for texturing; if these were used, the texture would be expected to be sharper).

These evaluations show that the project succeeded in its aim of reducing the cost for producing interactive characters whilst increasing the level of realism, achieving a quality equivalent to that of a video recording of the actual actor's performance.

8 References

[IJCV12] C. Budd, P. Huang, M. Klaudiny, and A. Hilton. Global Non-rigid Alignment of Surface Sequences. International Journal of Computer Vision, 102(1-3):256–270, 2012.

[ECCV12] M. Klaudiny, C. Budd, and A. Hilton. Towards optimal non-rigid surface tracking. In ECCV, pages 743–756, 2012.

[CGA07] J. Starck and A. Hilton. Surface capture for performance-based animation. IEEE Computer Graphics and Applications, pages 21–31, 2007.

[TIP04] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Processing, 13(4):600—612, 2004.

[EG14] D. Casas, M. Volino, J. Collomosse, and A. Hilton. 4d video textures for interactive character appearance. Computer Graphics Forum (Proc. EUROGRAPHICS 2014), 33 (2):371–380, 2014.

[BMVC14] M. Volino, D. Casas, J. Collomosse and A. Hilton. "Optimal Representation of Multi-View Video", BMVC 2014.

[JG09] J. Starck, J. Kilner, and A. Hilton. A Free-Viewpoint Video Renderer. Journal of Graphics, GPU, and Game Tools, 14(3):57–72, 2009.

Appendix – CVMP presentation material

Illustration 18 - Posters used at CVMP

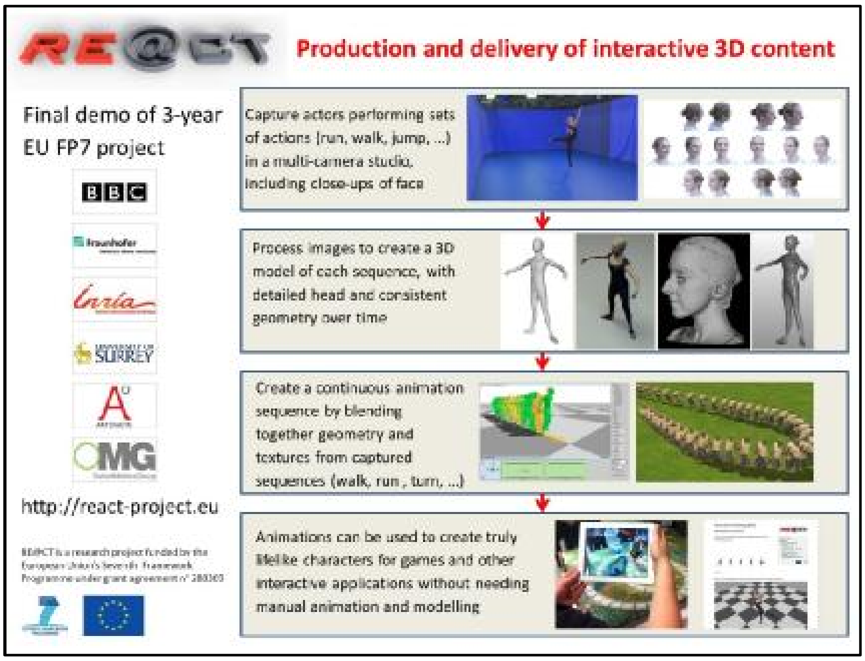

Illustration 19 - Slide used to introduce demo