

# The Use of Reality Models in Augmented Reality Applications

Gudrun Klinker, Didier Stricker, and Dirk Reiners

Fraunhofer Project Group for Augmented Reality at the ZGDV  
Arabellastraße 17 (ECRC)  
D-81925 Munich, Germany  
{klinker, stricker, reiners}@igd.fhg.de  
<http://www.igd.fhg.de/www/igd-a4/>

**Abstract.** Augmented reality (AR) is a technology by which a user's view of the real world is augmented with additional information from a computer model. AR applications require a very accurate model of the environment (a *reality model*) to augment the current view seamlessly with synthetic information (the *virtual model*). In this paper, we report on the problems we encountered with image data from real exterior construction sites. We discuss quality requirements for reality models in order to be useful in AR applications, and we outline potential further needs for reality models.

## 1 Introduction

With AR technologie, users can work with and examine real 3D objects while receiving additional information about those objects or the task at hand [3, 5, 6, 25]. The virtual objects need to coexist in physically plausible manners with the real world: they occlude or are occluded by real objects, they are not able to move through real objects, and they cast shadows on such objects.

The automatic construction of reality models is a long-standing issue in computer vision research. In the context of the European CICC-project [9], we explore very applied, pragmatic approaches which are closely related to the requirements of rather realistic application pilots in the exterior construction. Other approaches towards semi-automatically generating architectural models from images have been reported by Debevec [11] and by Faugeras [12].

### 1.1 Reality models for exterior construction applications


Exterior construction applications impose very demanding challenges on the robustness and usability of evolving AR technologies. Real construction sites are huge. Information from many views, at close range and from long distances, has to be integrated. Furthermore, construction environments are not well structured. Natural objects such as rivers, hills, trees, and also heaps of earth or construction supplies are scattered around the site. Typically, no exact detailed 3D

information of such objects exists making it difficult to generate a precise model of the site. Even worse, construction sites are in a permanent state of change. Buildings and landscapes are demolished, new ones are constructed. People and construction equipment move about, and the overall conditions depend on the weather and seasons. *Reality modelling* and *reality tracking* are thus very complex and demanding tasks. AR applications thus need to identify suitable simplified approaches for generating and dynamically maintaining appropriate models of the real environment.

## 1.2 AR applications in exterior construction projects

The first approach augments video sequences of large outdoor sceneries with detailed models of prestigious new architectures, such as TV towers and bridges that will be built to ring in the new millennium (see Figure 1a). Since such video sequences are very complex, we currently pre-record the sequences and employ off-line, interactive calibration techniques to determine camera positions. Given all calibrations, the augmentation of the images with the virtual object is performed live, i.e., the virtual model can be altered and transformed while it is being seen in the video sequence.

The second approach operates on live video streams, calibrating and augmenting images as they come in. To achieve robust real-time performance, we need to use simplified, "engineered" scenes.



(a) Virtual bridge across a real river.

(b) Virtual wall and grid in a real room.

**Fig. 1.** Interactive vs. automatic video augmentation.

In particular, we place highly visible markers at precisely measured locations to aid the tracking process (Figure 1b). Such use of special markers is becoming a common practice [21, 24, 28, 31].

## 2 Reality models

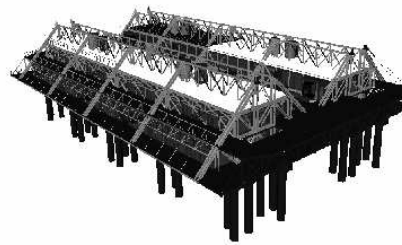
For optical camera calibration and to enforce physical interaction constraints between real and virtual objects, augmented reality systems need to have a precise description of the physical scene. Reality models don't need to be as complex as, for example, Virtual Reality (VR) models. VR models are expected to synthetically provide a realistic immersive impression of reality. Thus, the description of photometric reflection properties and material textures is crucial. AR, on the other hand, can rely on live optical input to provide a very high sense of realism. However, AR reality models have to be much more precise than VR models since users have an immediate quantitative appreciation of the quality of the integration between reality and augmentations.

### 2.1 Use of existing models

The most straightforward approach to acquiring 3D scene descriptions is to use existing geometric models, such as CAD data, output from GIS systems, and maps.



(a) Building under construction



(b) Virtual model

**Fig. 2.** Building under construction at the Expo 98 site in Lisbon.

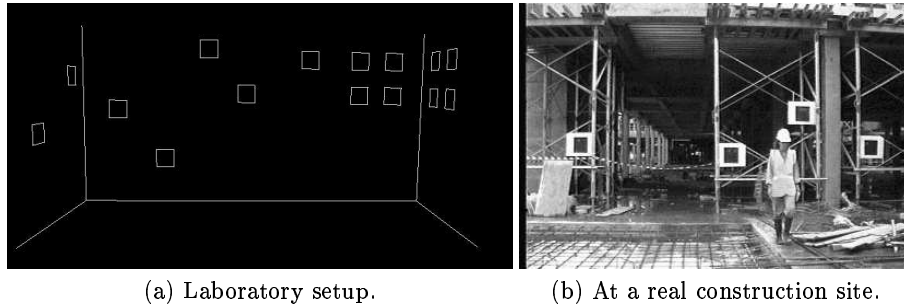
Once a building has entered the construction phase, the virtual model of the building itself can begin serving as a reality model, as shown in Figure 2.

Yet, they cannot always be used for AR. The data in CAD models doesn't necessarily coincide with discernible features in images. Furthermore, the models often don't describe the evolving reality of a construction site in enough detail.

### 2.2 Manual approach

The manual approach involves obtaining 3D measurements within the real world, using data bases and physical instruments. The 3D data points are entered into a small model. The approach works well when only very sparse reality models are needed but it cannot be used to generate elaborate descriptions of the real world.

Figure 3a shows a thus generated reality model of our "tracking laboratory", a room with several carefully measured targets on its walls.



(a) Laboratory setup.

(b) At a real construction site.

**Fig. 3.** Example and use of manually created reality models.

Figure 3b shows use of a similar simplistic reality model of a small area at the Bluewater construction site in Kent, UK. The model data was collected with a 3D laser pointer that was attached to a differential GPS system. The location of the upper right corner of each black square was determined by orienting the 3D laser pointer at the square, thus yielding the orientation and distance (time-of-flight) between the square and the pointer.


Yet, the approach is prohibitively time-consuming, if thousands of points are to be measured this way to generate suitable surface approximations for occlusion handling. Furthermore, the approach depends upon the availability of professionals and special equipment. Thus, models cannot be expected to be obtainable on short notice.

### 2.3 Interactive approach

In the European CICC project, we build a very sparse initial model of a landscape from external information such as maps and geodesic measurements. Our system helps us to interactively extend this model by superimposing it on a calibrated image. Models of new objects can then be entered into the model, using their two-dimensional position in a map and estimating their height from their alignment in several images.

From this model, we generate an initial camera calibration for a few site photos, interactively indicating how features in the image relate to the model (see Figure 4a).

Once an image has been successfully calibrated, the model is overlaid on the image, showing good alignment of the image features with the model features. Figure 4b illustrates the calibrated insertion of a new house into the model, using an initial much too large guess for the height of the building.



**Fig. 4.** Interactively created and extended reality model of the city of London.

## 2.4 Towards automatically generated models

Computer vision techniques are designed to automatically acquire three-dimensional scene descriptions from image data. Much research is currently under way, exploring various schemes to optically reconstruct a scene from multiple images, such as structure from motion [2, 12, 29, 32], (extended) stereo vision [11, 18, 20], and photogrammetric techniques [17].

In the context of the European project Cumuli [10], we explore to what extent automatically generated scene models can support AR and VR applications. In collaboration with INRIA and Lund University, we are developing and testing tools which exploit epipolar relationships between features in several images, geometric constraints on architectural structures, as well as city maps, to determine a set of progressively more precise projective, affine and finally euclidean properties of points in the three-dimensional scene.

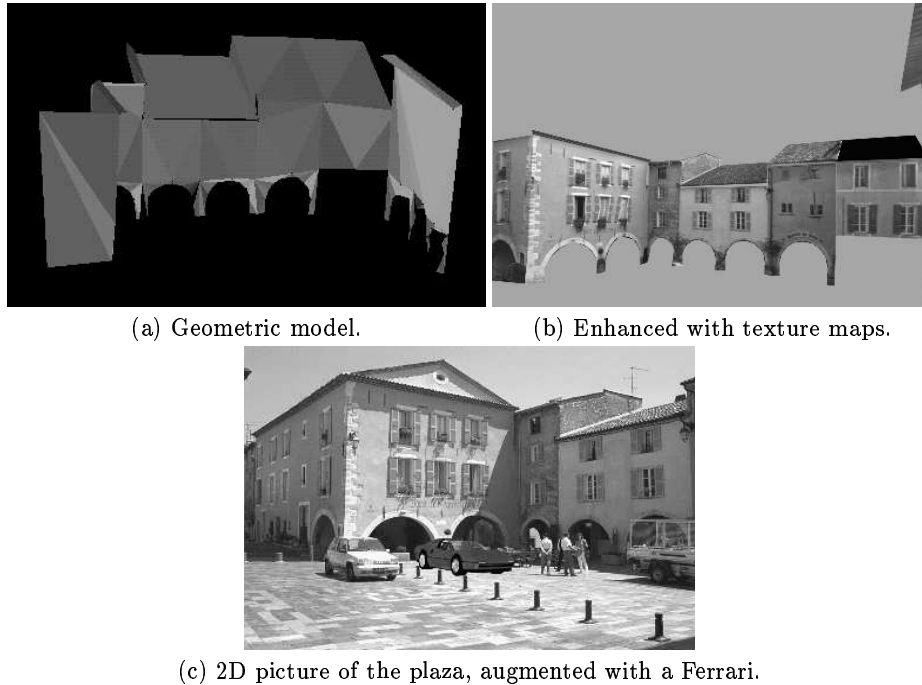
Figure 5 shows a reconstructed model of the Arcades of Valbonne. Figure 5a shows the reconstructed geometric model. In Figure 5b, the model has been enhanced by mapping textures from the original image data onto the surfaces. Figure 5c illustrates how photo of the area can be augmented with synthetic objects, such as a Ferrari, once the images have been analyzed and calibrated.

## 2.5 Range data models

Alternatively to motion-based scene recognition, the RESOLV project uses three-dimensional range sensors to conduct a 3D survey of a building. The environment is scanned from a number of capture positions and reconstructed into a model, unifying measurements from all viewing positions (Figure 6). Surfaces are recognized by processing the range data and are textured from camera images.

## 2.6 Use of reality models for camera calibration


Precise camera calibration is a key issue in AR. Calibration algorithms are inherently sensitive to noise and to specific alignments of features in the reality



**Fig. 5.** Automatically generated model of the arcades of Valbonne.

model. For example, houses in cityscapes tend to be aligned along a road or river. Many target features are thus approximately coplanar – considering a distance of several hundred meters between the camera and the houses – and cannot supply good 3-dimensional depth cues for camera calibration. In our work, we are emphasizing pragmatic concepts to cope with such real problems.


- Reality models should use targets in a nicely spread three-dimensional volume. For example, the inclusion of distant high rises and power poles in a model can greatly stabilize calibration results. The targets also need to be easily detectable and precisely locatable in image data.
- To help users correctly position image features, our system determines automatically which image feature currently has the largest influence on a calibration misalignment. By moving that feature by one pixel up, down, left, or right, a new calibration generates a much smaller mismatch between image features and projected scene features.
- We use as much externally available information as available, such as known internal camera parameters. Our algorithm can be further constrained when the approximate camera location and orientation is known from other tracking devices [19].



**Fig. 6.** a) RESOLV trolley b) Automatically generated model of part of the interior of the Royal Institute of Chartered Surveyors, London.


### 3 Augmenting reality

With AR, such virtual geometric objects can be integrated into the real environment during all phases of the life cycle of a building. Before the construction project is started, AR can support marketing and design activities to help the customer visualize the new object in the environment (Figure 7). During construction, AR can help evaluate whether the building is constructed according to its design (Figure 8).




**Fig. 7.** Side view of a new footbridge, planned to be built across the river Wear in Sunderland, UK.

After the construction is completed, maintenance and repair tasks benefit from seeing hidden structures in or behind walls (Figure 9).



**Fig. 8.** A virtual wall at a real construction site



(a) Original image.

(b) X-Ray view into the wall.

**Fig. 9.** Seeing the piping in the wall.

### 3.1 Occlusion handling using geometric reality models or depth maps

Occlusions between real and virtual objects can be computed by geometric rendering hardware by first drawing the reality model transparently and then rendering the virtual objects. Other mixing approaches initialize the Z-Buffer from depth maps obtained with a laser scanner or stereo computer vision. As a result, the user sees a picture on the monitor that blends virtual objects with live video, while respecting 3D occlusion relationships between real and virtual objects (Figure 10).

### 3.2 Simulation of illumination effects and physical constraints

Objects in the real world not only determine their own shading, they also have an influence on the appearance of other, distant objects by means of shadows and reflections [14].





**Fig. 10.** Virtual London bridge reflecting in the real water (partially occluded by the houses at the far side of the Thames).

With the availability of reality models, the geometry of shadows cast by virtual objects onto real ones can be computed [28]. Reflections are a more difficult topic that can be solved for many useful special cases, such as reflections of virtual objects in planar real mirrors (Figure 10). Difficult to handle are reflective virtual objects, as in general they would have to reflect things from the surrounding environment that are not visible in the image. Rendering the reflections from the real environment onto virtual objects requires the availability of a high-quality reality model.

For an augmented world to be realistic the virtual objects not only have to interact optically with the real world, but also physically. This applies to virtual objects when animated or manipulated by the user. For example, a virtual chair shouldn't go through walls when it is moved, and it should exhibit gravitational forces [7]. Given a reality model, this behavior can be achieved using collision detection and avoidance systems that are known from Virtual Reality systems [35].

These two laws make up the most important physical constraints. A full physical simulation including more aspects of the interaction between real and virtual objects, such as elastic behavior and friction, would be desirable. For off-line applications this is possible if enough information about the virtual objects and a complete enough reality model is available. For real-time applications most simulation systems are not fast enough. Yet, even simple implementations of the above rules will make the system much more realistic.

## 4 Future work

So far, we have used reality models for camera calibration, occlusion handling and a simple analysis of light reflections from virtual onto real objects. Yet, the concept of generating, using and updating reality models within AR applications goes far beyond this. We will now briefly allude to two other relevant aspects of AR that we expect to become important in the future.

### 4.1 Leaving reality behind

Augmented reality and virtual reality are not two discrete alternatives but rather part of a spectrum of mixed realities [25] with full virtual reality on one end and full physical reality on the other. Augmented Reality is in the middle, combining the best of both worlds. But sometimes it might be desirable to lean more in one direction or the other.

Since registered augmented reality by concept needs real images, its freedom of movement is limited to the places where an image recording device (possibly a human eye) can go. Virtual reality on the other hand allows complete freedom of movement, as computer generated images can be generated for every possible viewpoint. It may sometimes be desirable to leave the augmented reality behind and switch into the virtual reality to take a look from a point where it is physically impossible to go, e.g. from above.


When leaving reality behind, the view has to be constructed entirely from synthetic information, i.e., from the reality model plus new virtual objects. A very promising area of current computer graphics research in this direction is image based rendering [8, 15, 22, 23], which strives towards generating images from new viewpoints given some images from other viewpoints. A future system might employ a camera to record images while viewing the augmented scene and using them to incrementally refine the reality model. The ever-improving reality model allows the system to render increasingly realistic synthetic images from places that have not been visited by the user.

### 4.2 Diminishing reality

Many construction projects require that existing structures be removed before new ones are built. Thus, just as important as augmenting reality is technology to diminish it.

Figure 11a shows one of several pictures of TV-towers on Monte Pedroso near Santiago de Compostela, Spain. Prior to augmenting the image with a model of a new TV tower, the existing towers need to be removed (Figure 11b). To this end, a part of the sky has to be extrapolated into the area showing the TV-towers and the barracks. Then the new tower can be put into place (Figure 11c).

We currently use interactive 2D tools to erase old structures from images (Figure 14). This approach can only be used for static, individual photos, but not for video sequences from a live, dynamically moving camera.




**Fig. 11.** Monte Pedroso near Santiago de Compostela, Spain.

In principle, the problem of diminishing reality consists of two phases. First expiring buildings have to be identified in the image. When such structures are well represented in a reality model, they can be located by projecting the model into the image according to the current camera calibration.

The outdated image pixels then need to be replaced with new pixels. There is no general solution to this problem since we cannot know what a dynamically changing world looks like behind an object at any specific instant in time – unless another camera can see the occluded area. Yet, some heuristics can be used to solve the problem for various realistic scenarios. We can use morphological operators to extrapolate properties of surrounding “intact” areas (e.g: a cloudy sky) into outdated areas. Furthermore, when a building is to be removed from a densely populated area in a city, particular static snapshots of the buildings behind it could be taken and integrated into the reality model to be mapped as textures into the appropriate spaces of the current image. First results of such “X-ray vision” capabilities are shown in Figure 9b.

For video loops of a dynamically changing world, computer vision techniques can be used to suitably merge older image data with the new image. Faugeras et al. have shown that soccer players can be erased from video footage when they occlude advertisement banners: For a static camera, changes of individual pixels can be analyzed over time, determining their statistical dependence on camera noise. When significant changes (due to a mobile person occluding the static background) are detected, “historic” pixel data can replace the current values [36].

We use geometric constraints to compute pixelwise correspondences between regions in several images that outline a particular object (Figure 15). From such correspondances, we can trace specific points on the object across all images and we can decide in which images it is visible or occluded. Accordingly, occluded pixels can be replaced by visible ones, effectively removing the occluding object from the image. In more general schemes using mobile cameras, such techniques can lead towards incremental techniques to diminish reality. While moving about in the scene, users and cameras see parts of the background objects. When properly remembered and integrated into a three-dimensional model of the scene, such “old” image data can be reused to diminish newer images, thus increasingly effacing outdated objects from the scene as the user moves about.



**Fig. 12.** Automatic removal of a lego block.

## 5 Conclusions

Reality models are a crucial aspect of Augmented Reality work. Implicitly or explicitly, every AR application relies on knowledge about its surroundings in order to augment the reality appropriately. As reported in this paper, manual and semi-automatic techniques are commonly used to set up the models – a time consuming and complex task. Tools to automatically generate such models, e.g., via computer vision algorithms will greatly improve the quality and flexibility of AR applications. As a first step, the off-line generation of static scene descriptions will suffice. Yet, the long-term goal must be to automatically update and improve the models in real time while the application takes its course. At that point, AR applications will be able to deal with moving people and objects in the scene, such as objects that are being disassembled as part of the task. Users will also be able to temporarily leave reality behind to explore aspects otherwise unreachable. Furthermore, they will not only be able to augment reality but also to diminish it, virtually removing objects by showing what is behind them – according to the current 3D information in the reality model.

## Acknowledgments

The research is financially support by the CICC project (ACTS-017) in the framework of the European ACTS programme and by the Cumuli project (LTR-21914) in the ESPRIT programme. The laboratory space and the equipment is provided by the European Computer-industry Reseach Center (ECRC).

We are grateful to our current and former colleagues at Fraunhofer IGD, ZGDV and ECRC for many useful comments and insights which helped us develope and refine our work. Particular thanks go to Dieter Koller and Eric Rose

who both participated in the research group in the past. Both the CICC and the Cumuli consortium have deeply influenced our approach.

## Sources of graphical material

**Sunderland Newcastle, UK** (with Ove Arup and Partners): Figures 1a and 7a,b show a picture of the river Wear, Sunderland Newcastle, UK. The bridge model was provided by Sir Norman Foster and Partners.

**Thames river, London, UK** (with Ove Arup and Partners): Figures 4a,b, and 10 show pictures of the river Thames, London near St. Paul's Cathedral. The 3D model of a designed milleneum footbridge was provided Sir Norman Foster and Partners. The model was acquired for the CICC project by Ove Arup and Partners.

**Bluewater Kent, UK** (with Bovis and Trimble Navigation Limited): Figures 3b and 8b show a picture from a video sequence.

**Santiago de Compostela, Spain** (with Ove Arup and Partners): Figures 11a,b,c show a picture of Monte Pedroso near Santiago de Compostela, Spain. The model of the TV-tower was provided by Sir Norman Foster and Partners.

**Gmunder Straße, Munich, Germany** (with Philipp Holzmann AG, Germany): Figures 1b and 9a,b show indoor pictures of a bathroom under construction. Figure 8a shows an outdoor snapshot.

**Valbonne, France** (with INRIA Sophia-Antipolis): Figures 5a,b,c show a picture and a reconstructed model of the Arcades in Valbonne, France.

**Royal Institute of Chartered Surveyors, London, UK** (work by U. Leeds, JRC and BICC): Figure 6 shows pictures of the RESOLV trolley and the reconstructed model of the Royal Institue of Chartered Surveyors. Courtesy of the RESOLV project.

## References

1. K.H. Ahlers, A. Kramer, D.E. Breen, P.-Y. Chevalier, C. Crampton, E. Rose, M. Tuceryan, R.T. Whitaker, and D. Greer. Distributed augmented reality for collaborative design applications. In *Proceedings of Eurographics '95*, 1995. Also available as technical report ECRC-95-03 (<http://www.ecrc.de>).
2. A. Azarbajejani and A.P. Pentland. Recursive estimation of motion, structure, and focal length. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 17(6):562–575, June 1995.
3. R.T. Azuma. A survey of augmented reality. *Presence, Special Issue on Augmented Reality*, 6(4):355–385, August 1997.
4. M. Bajura and U. Neumann. Dynamic registration correction in video-based augmented reality systems. *IEEE Computer Graphics and Applications*, 15(5):52–60, 1995.
5. W. Barfield and T. Caudell. *Augmented Reality and Wearable Computers*. Lawrence Erlbaum Press, 1998.
6. J. Bowskill and J. Downie. Extending the capabilities of the human visual system: An introduction to enhanced reality. *Computer Graphics*, 29(2):61–65, 1995.

7. D.E. Breen, E. Rose, and R.T. Whitaker. Interactive occlusion and collision of real and virtual objects in augmented reality. Technical Report ECRC-95-02, ECRC, Arabellastr. 17, D-81925 Munich, <http://www.ecrc.de>, 1995.
8. S.E. Chen and L. Williams. View interpolation for image synthesis. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 279–288, August 1993.
9. CICC. Collaborative integrated communications for construction. ACTS AC-0017, 1995-1998, <http://www.hhdc.bicc.com/cicc/>, 1995.
10. CUMULI. Computational Understanding of Multiple Images. Esprit LTR-21914, 1996-1999, <http://www.inrialpes.fr/CUMULI/>, 1996.
11. P.E. Debevec, C.J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *Proc. SIGGRAPH*, pages 11–20, New Orleans, August 4-9 1996. ACM.
12. O. Faugeras, S. Laveau, L. Robert, G. Csurka, and C. Zeller. 3D reconstruction of urban scenes from sequences of images. In A. Gruen, O. Kuebler, and P. Agouris, editors, *Automatic Extraction of Man-Made Objects from Aerial and Space Images*. Birkhauser, 1995.
13. S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality. *Communications of the ACM (CACM)*, special issue on augmented reality, 30(7):53–62, July 1993.
14. A. Fournier. Illumination problems in Computer Augmented Reality. *Journee Analyse/Synthese d'Images (JASI)*, pages 1–21, January 1994.
15. S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen. The lumigraph. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 43–54. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
16. A. Hildebrand, S. Mueller, and R. Ziegler. REALISE - Computer Vision basierte Modellierung fuer Virtual Reality. In *Internationaler Workshop MVD'95: Modeling - Virtual Worlds - Distributed Graphics*, pages 159–168, Sankt Augustin, Germany, November 1995. infix.
17. B.K.P. Horn and M.J. Brooks. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
18. T. Kanade, A. Yoshida, K. Oda, H. Kano, and M. Tanaka. A stereo machine for video-rate dense depth mapping and its new applications. In *Proc. 15th Computer Vision and Pattern Recognition Conference (CVPR)*, San Francisco, June 18-20 1996. IEEE.
19. G. Klinker, D. Stricker, and D. Reiners. Augmented reality for exterior construction applications. In W. Barfield and T. Caudell, editors, *Augmented Reality and Wearable Computers*. Lawrence Erlbaum Press, 1998.
20. G.J. Klinker, K.H. Ahlers, D.E. Breen, P.-Y. Chevalier, C. Crampton, D.S. Greer, D. Koller, A. Kramer, E. Rose, M. Tuceryan, and R.T. Whitaker. Confluence of computer vision and interactive graphics for augmented reality. *Presence: Teleoperators and Virtual Environments*, special issue on Augmented Reality, 6(4):433–451, August 1997.
21. D. Koller, G. Klinker, E. Rose, D. Breen, R. Whitaker, and M. Tuceryan. Automated camera calibration and 3d egomotion estimation for augmented reality applications. In *7th Int'l Conf. on Computer Analysis of Images and Patterns (CAIP-97)*, Kiel, Germany, September 10–12, 1997, G. Sommer, K. Daniilidis, and J. Pauli (eds.), Lecture Notes in Computer Science **1296**, Springer-Verlag, Berlin, Heidelberg, New York, 1997.

22. M. Levoy and P. Hanrahan. Light field rendering. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 31–42. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
23. L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 39–46. ACM SIGGRAPH, Addison Wesley, August 1995. held in Los Angeles, California, 06-11 August 1995.
24. J.P. Mellor. Realtime camera calibration for enhanced reality visualization. In *Proc. of Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, pages 471–475, Nice, France, April 1995. IEEE.
25. P. Milgram and F. Kishino. A taxonomy of mixed reality visual displays. *IEICE Transactions on Information Systems*, E77-D(12), December 1994.
26. RESOLV. REconstruction using Scanned Optical Laser and Video. ACTS AC-0021, 1995-1998, <http://www.hhdc.bicc.com/resolv/>, 1995.
27. E. Rose, D. Breen, K.H. Ahlers, C. Crampton, M. Tuceryan, R. Whitaker, and D. Greer. Annotating real-world objects using augmented reality. In *Proc. Computer Graphics: Developments in Virtual Environments*. Academic Press Ltd, 1995.
28. A. State, G. Hirota, D.T. Chen, W.F. Garrett, and M.A. Livingston. Superior augmented reality registration by integrating landmark tracking and magnetic tracking. In *Proc. SIGGRAPH*, pages 429–438, New Orleans, Aug 4-9 1996. ACM.
29. R. Szeliski and S.B. Kang. Recovering 3d shape and motion from image streams using non-linear least squares. Technical Report CRL 93/3, Cambridge Research Lab, Digital Equipment Corporation, One Kendall Square, Bldg. 700, March 1993.
30. M. Tuceryan, D. Greer, R. Whitaker, D. Breen, C. Crampton, E. Rose, and K. Ahlers. Calibration requirements and procedures for a monitor-based augmented reality system. *IEEE Transactions on Visualization and Computer Graphics*, 1, September 1995.
31. M. Uenohara and T. Kanade. Vision-based object registration for real-time image overlay. In *Proc. of Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, pages 13–22, Nice, France, April 1995. IEEE.
32. VANGUARD. Visualisation Across Networks using Graphics and Uncalibrated Acquisition of Real Data. ACTS AC-0074, 1995-1998, <http://www.esat.kuleuven.ac.be/konijn/vanguard.html>, 1995.
33. A. Webster, S. Feiner, B. MacIntyre, W. Massie, and T. Krueger. Augmented reality in architectural construction, inspection, and renovation. In *Proc. ASCE Third Congress on Computing in Civil Engineering*, pages 913–919, Anaheim, CA, June 17-19 1996.
34. M. Wloka and B. Anderson. Resolving occlusion in augmented reality. In *Proc. of the ACM Symposium on Interactive 3D Graphics*, pages 5–12. ACM, 1995.
35. G. Zachmann. Real-time and exact collision detection for interactive virtual prototyping. In *Proc. of the 1997 ASME Design Engineering Technical Conferences*, Sacramento, CA, Sept 14–17 1997. ASME. CIE-4306.
36. I. Zoghlami, O. Faugeras, and R. Deriche. Traitement des occlusions pour la modification d'objet plan dans une sequence d'image. Private communication; see also [http://www.inria.fr/robotvis/ personnel/zimad/Orasis6/Orasis6/html](http://www.inria.fr/robotvis/personnel/zimad/Orasis6/Orasis6/html), 1996.