

Hybrid Rendering – A New Integration of Photogrammetry and Laser Scanning for Image Based Rendering.

Bruce Lamond and Gordon Watson.
EdVEC (Edinburgh Virtual Environment Centre),
University of Edinburgh.
<http://www.edvec.ed.ac.uk>
{Bruce.Lamond, [Gordon.C.Watson](mailto:Gordon.C.Watson@ed.ac.uk)}@ed.ac.uk

Abstract.

Constructing accurate architectural models of interiors is a non-trivial task. Photogrammetry can be used to recover the large-scale features but has difficulty in modelling sufficient detail. Laser scanning on the other hand produces highly detailed depth data but is difficult to turn into a globally consistent model. We present a hybrid system for the creation of photorealistic and geometrically efficient and accurate models through the integration of photogrammetric modelling techniques with laser acquired range data. Our Hybrid Rendering software builds on the foundation of the Facade photogrammetric modelling system, but improves the degree of complexity of this representation by incorporating range data for those parts of the scene that are unmodelled. The salient planar features in the scan are extracted and used to align the models via a two-stage quaternion-based and ICP-based registration. The process allows for straightforward integration of multiple range scans. Rendering of novel views is achieved using view-dependent texture mapping techniques combined with point-based rendering.

1. Introduction.

Recent advances in the availability of low cost memory, processors and graphics hardware for the PC market means that users expect ever-increasing realism in graphics in computer games, cinema, TV and medical applications. The resources required in displaying highly accurate polygon-based surface models at sub-pixel scales still pose problems in this area in computer graphics. Automatic and accurate reconstruction of scenes from sparse-view multiple-baseline stereo is also likewise problematic for the computer vision community. In contrast, the constant resource cost and high degree of photorealism seen in

image-based rendering (IBR) techniques have shown this approach to be extremely effective within certain context. The central goal in IBR is to reconstruct a continuous representation of the plenoptic function from a number of samples of the function [1, 10]. Samples are usually in the form of digitised stills and reconstruction is achieved through interpolation between views. A major characteristic of all IBR algorithms is that they should use an accurate geometric precursor or proxy [4]. In practice, this usually proves difficult to achieve from a sparse set of input images. Inevitably this leads to an inaccurate or primitive proxy. In this work we attempt to subsume problems associated with simple geometric proxies by registering data from a laser scanner with the proxy for a higher degree of accuracy in architectural scenes. We use a proxy derived from the Facade photogrammetric modelling system [5].

Detail is then added to the proxy in the form of laser range data. The data sets are registered approximately from a small number of point correspondences via computation of a quaternion-based single transformation. We then extract planar regions from the range image of the laser data and use these planes to compute an accurate transformation of the laser data on to the Facade proxy using the Iterative Closest Point (ICP) algorithm. By using the façade model as a global framework the operation of constructing a model from multiple range images is much simplified in comparison to constructing the model from range images alone [15]. Using a global proxy also enables multiple scans to be registered. In addition holes inevitably present in the laser range data due to occlusions or dark surfaces are automatically filled by the model to guarantee a consistent model.

We believe our approach to alignment using both photogrammetric and range scanning approaches is more robust than an approach based on range data alone. This is due to salient features in images such as edges and corners being poorly sampled by the laser disc and hence inaccurately represented in the range

data. Our alignment strategy utilises image edges only in the photogrammetric stage to form the proxy model.

Once the alignment is done we are able to discard the majority of the laser data where that data forms those planes used in registration, typically this removes 80% of the imaged points. The remaining range samples are rendered using point-based rendering techniques and the original model is rendered as polygons.

The model is rendered using view-dependent texture mapping of the original photographic images and occlusions are resolved with ray-casting.

For this reason we call our system Hybrid Rendering. The intention is to create geometrically accurate and photorealistic models which can be built easily on any typical PC. Following a brief examination of contemporary work in this area, we describe in detail how we have implemented our system and present some of the results we have obtained from it. We finish with a discussion on how we intend to improve the system in the future.

1.1. Background & Related Work.

This project contains aspects of photogrammetric modelling, IBR and 3D photography. The Facade photogrammetric modelling system [5] has received notable attention in this area. It recreates photorealistic architectural scenes from a small number of representative photographs of the scene, which are rendered using a view-dependent texture-mapping scheme. The user is directed to define the scene in terms of related geometric solids, referred to as blocks, in much the same manner as a 3D modeller constructs a scene. The edges of these blocks are then matched to corresponding edges in the photographs and a model of the scene with the generative camera parameters is reconstructed via the minimisation of an objective function which sums the disparity between projected model edges and those marked in the scene. The Facade system reduces the complexity of calculating the scene by explicating the relative positioning of blocks. The recovered cameras are converted to slide projectors which reproject the original images back onto the model. To improve the texture mapping, a weighting function is used to calculate the contribution image pixels make based on their angle to the novel view angle. In this paper Debevec notes further detail can be added via displacement maps derived from stereopsis of the base images, although in practice this method has proven un-reliable and requires high accuracy in the recovered positions of cameras. Provided there exists a sufficient number of available block shapes to use, *façade* has proven to be robust and works very well in practice.

The best-known and wide-ranging work in laser

scanning in the last decade has been the Digital Michelangelo Project [9]. A number of pieces of *in situ* statuary have been scanned at high resolution to produce detailed model records, using a modified laser triangulation scanner and gantry to capture detail at the sub-millimetre scale. The scale of the subjects means that they have had to align and merge multiple scans and then perform hole filling on unmodelled portions of the mesh. Alignment is done approximately by hand in the initial phase before accurate alignment using the Iterative Closest Point algorithm on overlapping pairs of scans. Surface reconstruction involved creating a dense grid of voxels. Each voxel near a surface is visited in turn and viewed on a line of sight towards the range camera. Voxels in front of the surface are marked as empty. In this way an isosurface is extracted from unseen voxels and even extended to include boundaries between empty and unseen regions for a completely filled surface. Colour reflectance values are mapped onto the mesh by correcting the colour images geometrically and radiometrically. We avoid the problems associated with hole-filling and low-fidelity acquisition of colour data associated with this type of laser scan by backing the laser data with the closed polygonal photogrammetric model and by using digital photographs respectively.

More recent similar work has been done in [12] on model reconstruction from range and image data. Multiple range scans are segmented into planar regions, and line segments extracted from these regions are used to register overlapping data sets. Segmented and registered scans are swept into a solid volume and intersected to produce a complete model of the scene. Edge detection is used to extract linear features in the image data and these features are matched to the line segments extracted from the range scans allowing the models to be textured. The registration process in our system is different in that it is less complicated to match up the data sets and registration is done relative to the texture mapped geometric proxy.

The work presented here is an adaptation of the main idea presented in [14] who postulated the idea of a system to extend the Facade photogrammetric modelling system by inclusion of laser scan data. They proposed that the range data could be used to solve for parameters of Facade primitives directly or be used to detail the simple proxy via displacement maps. Our approach here differs in that we have detailed the proxy directly by addition of scan data and point-based primitives.

2. Method of Scene Acquisition.

Figure 1 shows the execution pipeline that leads to scene acquisition. The modelling environment chosen for the implementation is VTK, The Visualization

Toolkit [11]. VTK offers an inexpensive, powerful, extensive and extensible system for 3D modelling. Many of the tools employed in this work utilise the toolkit's powerful input, filtering, rendering and interaction frameworks. We now describe in detail those steps associated with acquiring a scene.

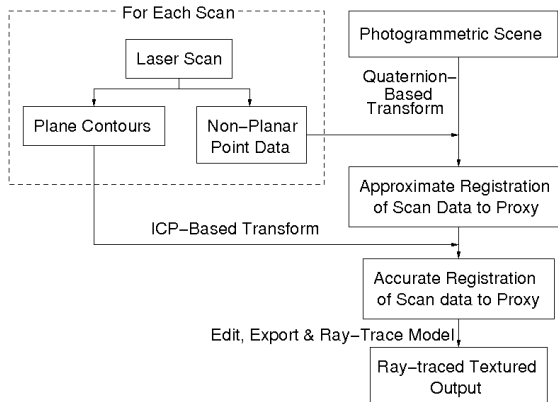


Figure 1. Hybrid Rendering pipeline

2.1. Plane Extraction.

Input to the system is in the form of a Façade modelled environment and range data from a laser scanner.

Each Façade environment output describes a simple block set of polyhedra combined with the set of camera positions, calibration parameters and image files used to construct the block model of the scene. Our adaptation of a class of 3D polygonal reader object in VTK allows facade scene files to be represented inside the toolkit as polyhedral data. Light sources are added to the scene and double as repositories for the camera calibration and image file details for texture mapping later in the process.

The scanner data is provided by a 3rdTech time-of-flight infra-red laser scanner. This equipment allows the collection of up to approximately 20 million points in a 360 degree capture with a resolution of 150 microns. The maximum range is approximately 20 metres. A rotating mirror deflects the beam to increasing elevations in each column of scan while the whole housing rotates in well constrained azimuthal increments. This leads to some variation in the number of samples taken over all the columns. Certain (particularly dark) surfaces are found to give very low or spurious return. **Figure 2(a)** shows a range image from our scanner.

Initially our idea was to register the data sets using the ICP algorithm and then loop through the scan data and remove points which lie within a threshold distance from the corresponding Façade polygons. A

limitation of the ICP algorithm however, is that it performs poorly when one or other of the source and target data sets is more or less complicated or extensive than the other (as is invariably the case with a simple geometric model versus a highly detailed scan). This necessitates the approach described herein where the salient features seen in the simple model are extracted from the scan data and used to compute an accurate transform. Following [12], a range image is divided into a grid of square point clusters and the spherical sample set transformed into xyz space. A best-fit plane is then computed for each square and superclusters (scene-scale planes) of squares are grown where neighbouring squares fall within a threshold orientation and their centroids lie within a threshold distance in space to each other. We extract each supercluster contour and discard all internal supercluster planar points while retaining all other non-supercluster data.

Due to the variation in sample distribution between columns we first resample the scanned range data into a structured grid and perform linear interpolation to fill in any holes between adjacent values in a column. The structured range data is divided up into a grid of small squares of side n samples and we try to fit a plane to all the samples in any square which contains more than 70% of points with a return. In [12] Jacobian iteration was used to compute a local plane, but we found this method to select a significant number of incorrect planar orientations. Using a Random Sampling Consensus (RANSAC) approach [6] gave us much better results. RANSAC operates by choosing a number of test-case planes to fit to a sample set and calculates which (if any) of these tests fits the samples best. To do this, 3 points are chosen at random for each test case and these points form the test plane. The number of members of the sample set which lie within a threshold distance from the test plane is then computed along with the cumulative distance of all such points from the plane. If any test plane lies close enough to a threshold number of samples, then the plane with the lowest cumulative distance from its samples is used as the best fit and the square is marked as 'locally planar'. Next we loop through all locally planar squares examining neighbouring squares for a similar planar orientation and proximity in space. Squares which fulfil these criteria are added to a list of superclusters of squares. A final list of superclusters is created based on a threshold minimum number of square members and we evaluate the average planar orientation of all squares in a supercluster. **Figure 2(b)** shows the planes extracted from the image of **figure 2(a)**. The boundary contour for each supercluster is extracted using the Marching Squares algorithm [11] in spherical theta-phi space. By the nature of this algorithm, the contour points often lie just outside of the supercluster and so the closest

supercluster member point must be found for each contour point. With the contour thus constrained, each contour point spherical pair is projected on to the supercluster average plane and a single plane boundary obtained. We maintain a list of supercluster boundaries and discard all other superfluous points internal to the superclusters. We then create a list of all other points from the clusters that were locally non-planar. In this manner we discard up to 80% of the original scan data and retain the non-planar detail from the scene.

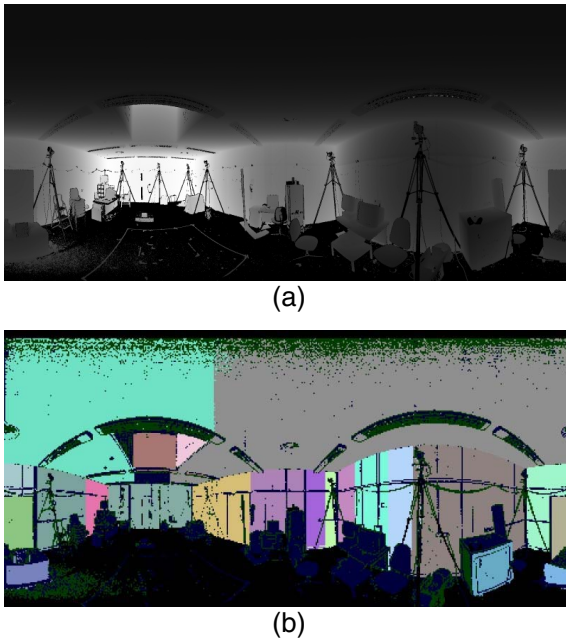


Figure 2. (a) Range image from our scanner. (b) Planar features extracted from the range data using a cluster size of 5 by 5 pixels. The pastel coloured areas are the superclusters (= planes) discussed in the text, dark blue areas represent locally planar areas not extensive enough to form superclusters, dark green represents a non-locally planar cluster and black (as in the range image) represents insufficient sampling density.



2.2. Approximate Alignment

With the Facade block model and laser range data read in to VTK the next task is to register the data sets into alignment. The ICP algorithm is used to compute an accurate transform but this first requires an initial approximate alignment to succeed.

As the Facade model contains the information to texture map the scene, this model's coordinate system must remain static and the scan data be registered to it. In order to calculate an approximate transformation between the two data sets, a number of

correspondences between the models must be specified. As the Façade model is a very simple block representation and the range scan highly detailed, automation of this part of the process is highly unlikely. For this reason, the initial registration stage requires user-defined point correspondences called landmark points to be selected. Four well-chosen landmarks on each model can suffice in evaluating a transform and the interactive point-picking functionality in VTK is employed in allowing the user to specify point matches across the models. A tolerance value in picking means that only vertices in the Façade model can be chosen so no inter-polygonal points are allowed. This anchors the point choosing procedure to a high degree of accuracy across the models. Accurately chosen correspondences in the range data can be specified as a result of the employment of VTK's interaction framework. This allows the user to orient the range cloud and zoom in on any single point out of the millions of points in the dataset to register as close to where the Façade vertex occurs as possible. The algorithm developed by Horn [8] which computes the approximate transform between the two sets of landmark points is used. This 'Landmark Transform' algorithm is a non-iterative closed-form solution that uses unit quaternions to calculate the rotation component of absolute orientation. A minimum of three points is required to calculate the rotation between the left and right coordinate systems. Triad structures are defined for each system. A line joining the first point to the second point becomes the triad x-axis. A line perpendicular to the x-axis and in the plane formed by the third point becomes the y-axis, and the z-axis is perpendicular to these two in a right-handed sense. The rotation between the two underlying Cartesian systems is then the transformation between the two sets of triads. Column unit vectors representing these axes are defined in terms of left and right systems and are adjoined to form left and right 3x3 matrices. The desired rotation is found to be the product of the left matrix with the transpose of the right matrix. The translation component of transformation is given by the difference of the centroid of one system and the rotated and scaled centroid of the other. The scale factor is the ratio of the RMS deviations of the respective centroids. Unfortunately it is unlikely that three measured points will be exact in both systems. More likely is that residual errors will be introduced in the measurement and so using more points to calculate the transform will increase the accuracy of the result. Thus the problem becomes one of minimising residual errors. Horn [8] finds that this is achieved when the sum of the dot products of corresponding coordinates in the right system and the rotated coordinates in the left is maximised. This can be achieved using unit quaternions, and the best result is the eigenvector

associated with the most positive eigenvalue of a symmetric 4x4 matrix derived from sums and differences of the relative coordinates. **Figure 3(a)** shows two data sets after approximate alignment.



(a)
(b)

Figure 3. (a) Approximately registered photogrammetric model and scan (without planes). (b) Registered data sets after accurate alignment.

2.3. Accurate Alignment.

Once the approximate relative orientations of the data sets are found, the range scan points and planes

are then subjected to this transformation and rendered with the Facade model. The second stage of registration fine-tunes the alignment by taking the approximately transformed data and passing it to the ICP algorithm [3]. This algorithm is an iterative and computationally efficient method for the accurate registration of 3D forms whether they are points, lines, segments, polygons, splines or implicit surfaces. This is useful for registering or measuring shape similarity between two different representations of similar models. A reasonable initial approximation to the registration is required for the algorithm to operate successfully whereafter a mean square distance metric can be used as termination condition of the algorithm. Each point in the unaligned source data set is examined to find the closest point on the target data set (this may be a new internal surface point if the point did not previously exist). All source point/closest point pairings are passed to the Landmark Transform algorithm described above to compute the transform that maps between the data sets. The source points are then transformed accordingly and the relative mean square distance between the sets is calculated to test for termination. The process is repeated iteratively until this condition is met or a maximum number of iterations is reached. The algorithm is found to cause the expeditious and highly accurate convergence of the data sets with a matter of only a few tens of points supplied as the source sample where a good initial approximation has been supplied. **Figure 3(b)** shows 2 data sets after accurate alignment.

By aligning the scan data to the Facade proxy in this manner it allows any number of scans to be registered into the same scene. With the simple proxy acting as reference frame, multiple scans can be registered on the basis of the corresponding planes in the data sets. **Figure 4** shows multiple scans registered to the same proxy. This technique is similar in essence to that employed in [15] for multiple scan alignment, with the proxy acting as a global coordinate system for alignment.

2.4. Model Editing, Exporting and Rendering.

With the data sets properly aligned, it is now possible to texture map the original Facade model formative images back onto the registered combined data to increase the level of detail of the final output. Before this is done it may be desirable to perform steps that improve the level of detail shown in the combined data. Consider the images shown in **Figure 5**. **5(a)** shows the representation of a block model from Facade; **5(b)** shows the laser scan point cloud of the same model as tiny spheres. Superposition of the two data sets is shown in **5(c)** and one can see how the concave details in the church doorway and windows

have been lost due to the simplistic blocky nature of




Figure 4. Multiple scans can be registered to the same proxy (shown coloured)

the Facade model. A solution to this occlusion problem is to allow the user to remove all or parts of polygons in the Facade data. Removing parts of polygons allows the range data behind to be viewed and consequently texture mapped. In such cases with the detailed range data registered over the Facade model, it might seem advantageous to remove the whole Facade model and leave only the fine detail behind. While this is certainly the case where the range data provide complete coverage of the model area, those areas where scanning has been occluded will show up as holes and it is thus useful to retain the underlying model to fill in these holes (Figure 5(d)). We have implemented various clipping functionality to allow the data sets to be cleaned up and to remove any occluding regions resulting from the Façade model's simplistic nature. Erroneous points may also be present in the range data, especially due to specular reflections of the laser beam. These can be efficiently removed by clipping. This point in the pipeline has the two data sets registered accurately and the Facade model manipulated to show maximum detail with minimum point count. All that remains is to output the scene so that it can be rendered into high quality images. In order to resolve occlusions from the projected camera images we ray-trace the output scenes using BMRT - the Blue Moon Rendering Tools [7], a high quality RenderMan compliant system which incorporates global illumination algorithms. The camera images are projected onto the model by the use of custom light source shaders in RenderMan. A VTK

exporter was developed that exports the geometric scene and the cameras into a RenderMan RIB format.




Figure 5. Hole cutting to reveal occluded detail. (a) The Facade block model. (b) the laser data as tiny spheres. (c) superposition of (a) and (b). Note how detail in the doorway and window recesses is occluded. (d) the same superposition as (c) with holes cut in the Facade model to reveal recess detail. Holes in the wall on the right in (b) are filled by the inclusion of the Facade model.

The contents of the scene in VTK are examined and output in RIB format making use of several powerful features available in BMRT. RIB format allows for the specification of complex polygons and this allows any holes made in Facade polygons to be dealt with. BMRT has no point type primitive to deal with the point cloud, but it does have a sphere primitive so the point set is specified as tiny spheres. This negates any polygon mesh construction problems, were such a strategy followed. Recall that the image and camera details used to construct the Facade model were added to the VTK scene in special light sources. In exporting the scene, these lights are examined and can be specified in the RIB file as special light source shaders that operate like slide projectors. Ray tracing the scene from these light sources removes any requirement for calculation of occlusions within the scene. The camera's calibration and projection parameters define the shader properties and the input image is transformed into a texture which can be projected on to the scene, subject to these shader properties. Specialised surface shading can be specified in BMRT and we have used these to implement view-dependent texture interpolation techniques, which improve the output according to a perception-based metric. This improvement comes from the employment of a

frequency-dependent dual interpolative scheme which operates to minimise the number of perceived Mach bands resulting from poorly modelled geometry or abrupt scene changes. High-frequency components in the input images are modelled with nearest-neighbour interpolation and other frequencies with linear interpolation. Regarding motion between novel scenes, this reduces multiple projections of eye-catching high-frequency areas while exhibiting smooth change across other frequencies [13]. The user can also specify whether to output the range scan or Facade model, with or without texture mapping so that the system doubles as a viewer for any constituent part of the process.

3. Results.

Some scans of the test scene were taken from various locations. A Facade environment output file was created from a representative set of digital photos of the scene taken using a Nikon D1X professional digital camera and introduced into the system. Four corresponding points were selected in each model and the data sets registered on the basis of these and the planar features extracted from the data sets. The sets of registered data were cleaned up and exported into RIB format with the range data in the form of a cloud of tiny spheres. Images from the resulting scenes can be seen in **figure 6**. All scenes are shaded with the frequency dependent interpolation scheme discussed above. Areas not seen by any camera are shaded red.

4. Discussion & Conclusion.

The most obvious result from **figure 6** is the dramatic improvement in the representation of detail between the Facade model and the combined data sets. Although digital image texture mapping does generally give a good result in the Facade system, certain features are poorly represented where the model is overly simple. These areas are particularly time consuming to model in the Facade system but are collected as part of the normal procedure in laser scanning. Areas of occlusion in scanning occur but one can see in the scenes showing the combined data sets that these holes are effectively filled by having the Facade model underneath the laser data. Addition of multiple scans also helps overcome these omissions.

Although some manual effort is required to construct both the Facade model and to perform the alignment we have performed a comparison with the effort required to align the images and scans in [15] and have found our system to be considerable less effort. This is due in part to the ease of performing a single 3D model alignment as compared to multiple 2D image alignments for each camera.

The system is limited in its ability to handle non-architectural type scenes as we are unable to register data without planar features. A further limitation occurs because of the fact that the Facade proxy is approximate in shape leading to a slight misalignment in the scan data after registration. The problem indicates the need for tighter integration between the photogrammetric model and the range data and in particular a better treatment of the errors inherent in each system.

A long-term aim is to improve the registration technique to remove the need for user interaction, although the best approach to this is unclear at present. Other longer-term views on our system include allowing the model to be edited in 3D by manipulating depth map data returned by the laser scanner. Resampling the range data into a structured form and then allowing areas on a surface to be raised or lowered by altering the depth value could achieve this.

5. Acknowledgements.

The authors would like to thank Dr. Paul Debevec for continued access to the Facade modelling system. We would also wish to acknowledge the support of ESPRC for project funding.

6. References.

- [1] E. Adelson and J. Bergen, *Computational Models of Visual Processing*, The MIT Press, Cambridge Mass., 1991.
- [2] N. Amenta, S. Choi and R. Kolluri, "The Power Crust.", *Computational Geometry*, Vol. 19, No. 2-3, pp. 127-153., 2001.
- [3] P. Besl and N. McKay. "A Method for Registration of 3-D Shapes.", *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Vol. 14, No. 2 1992.
- [4] C. Buehler, M. Bosse, S. Gortler, M. Cohen, and L. McMillan, "Unstructured Lumigraph Rendering.", *Proceedings of SIGGRAPH 2001*, 2001
- [5] P. Debevec, C. Taylor, and J. Malik, "Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach.", *Proceedings of SIGGRAPH 1996*, pages 1120, 1996.
- [6] M. A. Fischler and R. C. Bolles, "Random Sampling Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography.", *Graphics and Image Processing*, pages 381-395, 1981.
- [7] L. Gritz and J. Hahn, "BMRT: A global illumination implementation of the renderman standard.", *Journal of Graphics Tools*, 1(3), 1996.

[8] B. K. P. Horn, "Closed-Form Solution of Absolute Orientation Using Unit Quaternions.", *Journal of the Optical Society of America*, Vol. 4, pages 629-642, 1987.

[9] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade and D. Fulk, "The Digital Michelangelo Project: 3D Scanning of Large Statues.", *Proceedings of SIGGRAPH 2000*, 2000.

[10] L. McMillan and G. Bishop, "Plenoptic modelling: An image-based rendering system.", *Proceedings of SIGGRAPH 1995*, 1995.

[11] W. Schroeder, K. Martin and B. Lorensen, *The Visualisation Toolkit*, Prentice-Hall, 1996.

[12] I. Stamos and P. Allen, "3-D Model Reconstruction Using Range and Image Data.", *Proc. of IEEE International*

Conference on Computer Vision and Pattern Recognition, Vol. 1, pages 531-536, 2000.

[13] G. Watson, P. O'Brien & M. Wright, "Towards a Perceptual Method of Blending for Image-Based Models.", *Proceedings of SIGRAD 2002*, 2002.

[14] M. Wright, G. Watson and R. Middleton, "Archetype: Towards the Integration of Photogrammetric and Laser Range Data for Architectural Reconstruction.", *Proceedings of the International Symposium on Virtual and Augmented Architecture*, Pages 255-266, 2001.

[15] McAllister, David K., Lars Nyland, Voicu Popescu, Anselmo Lastra, Chris McCue. *Real-Time Rendering of Real-World Environments*, Rendering Techniques '99, Proceedings of the Eurographics Workshop on Rendering, (Granada, Spain), June 21-23, 1999.




Figure 6. Output from the Hybrid Rendering system.

(a) The Façade geometric proxy with a light source corresponding to each camera used to construct the scene. (b) The same geometric proxy texture mapped with view-dependent texture mapping. Note how the model contains no holes but the texture mapping gives a flat appearance. (c) Point data from the laser scanner displayed as tiny spheres. The planar regions used to align the data sets have been discarded. Other holes in the scene can be seen where the scanner has been occluded, particularly the circular 'footprint' in the foreground. (d) The façade proxy combined with the scan data. Note how the holes in the scan data are filled by the façade proxy. (e) The ability to cut holes in the façade proxy can dramatically increase the level of detail in the model. Note how the wall plane occluding the doorway recess in (d) has been cut to reveal this feature. (f) Textured mapped representation of the model in (e). Red areas are not represented by any images. Incorporation of extra images into the scene will resolve this.