Computer-Aided Civil and Infrastructure Engineering **18** (2003) 264–274

CAD-Based Photogrammetry for Reverse Engineering of Industrial Installations

Johan W. H. Tangelder,* Pierre Ermes,† George Vosselman‡ & Frank A. van den Heuvel

Department of Geodesy, Faculty of Civil Engineering and Geosciences, Delft University of Technology, Thijsseweg 11, NL-2629 JA Delft, The Netherlands

Abstract: *When designing an industrial installation, construction engineers often make use of a library of standardized CAD components. For instance, in the case of a servicing plant, such a library contains descriptions of simple components such as straight pipes, elbows, and T-junctions. A new installation is constructed by selecting and connecting the appropriate components from the library. This article demonstrates that one can use the same approach for reverse engineering by photogrammetry. In our technique, the operator interprets images and selects the appropriate CAD component from a library. By aligning the edges of the component's wire frame to the visible edges in the images, we implicitly determine the position, orientation, and shape of the real component. For a fast object reconstruction the alignment process has been split in two parts. Initially, the operator approximately aligns a component to the images. In a second step a fitting algorithm is invoked for an automatic and precise alignment. Further improvement in the efficiency of the reconstruction is obtained by imposing geometric constraints on the CAD components of adjacent object parts.*

1 INTRODUCTION

Maintenance, performance analysis, and revamping of an industrial installation are facilitated by accurate information in the form of a CAD model that describes the position, orientation, and shape of all the components of the industrial installation. For many installations the as-built situation has to be reconstructed from measurements because the plant design maps are no longer available, not up to date, or not accurate enough. Traditionally, these measurements are done by tacheometry or classical photogrammetry (Cooper and Robson 1996). In both cases this requires the application of a large number of targets that are stuck to the objects.

In this article we present a new and more efficient measurement approach using digital photogrammetry and CAD models. This reverse engineering approach resembles the design process of industrial installations. For the design of new plants one nowadays makes use of catalogs of standardized components like straight pipes, elbows, T-junctions, and more complex parts of various sizes and materials. By selecting components from the catalog and connecting them to previously selected components in a workspace, the CAD model of a plant can be designed incrementally. In the reconstruction process of the plant models the same kind of catalogs are being used. The operator starts by interpreting the images and selects the appropriate component from the catalog (Figure 1a). By interactively shifting, scaling, rotating, and deforming the CAD component, the component is aligned to the object's edges that are visible in the images. Once the pose (position and rotation) and shape parameters of the object are approximately known (Figure 1b), digital image analysis is used to accurately fit the component's contour edges onto the object edges in the images (Figure 1c). The task of the operator is, therefore, limited to supplying approximate values for the pose and shape parameters. This semi-automatic approach improves both the speed and the accuracy of the reconstruction.

Clearly, the pose and shape parameters of adjacent components are related to each other. To further enhance the speed and accuracy we therefore incorporate

*Now at Department of Computer Science, Utrecht University.
†Now at Philips Medical Systems.
‡To whom correspondence should be addressed. E-mail: *g.vosselman@geo.tudelft.nl.*

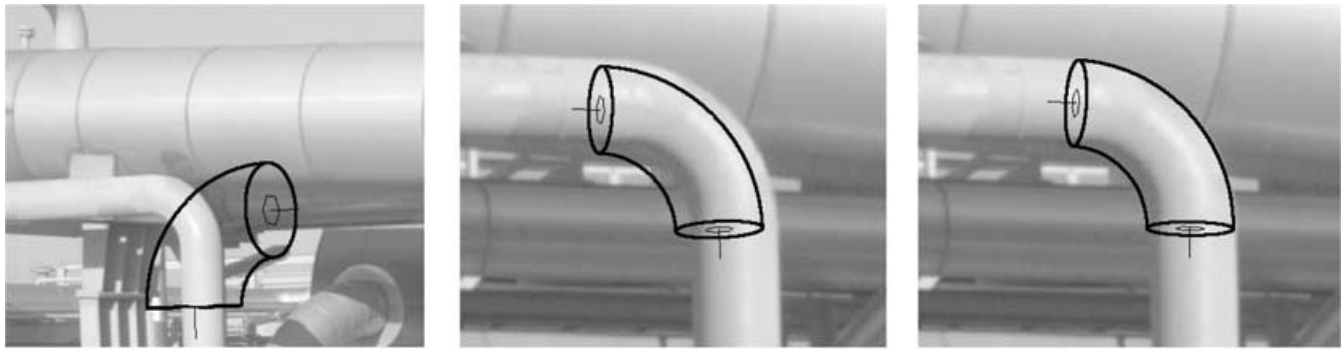

Fig. 1. Interactive alignment of a CAD model to an image: (a) initial position of a model in an image, (b) position after approximate alignment by the operator, and (c) precise position after fitting to image gradients.

geometric constraints in the estimation of the object parameters. This leads to a very significant reduction in the number of measurements that need to be performed by the operator.

With the current state of the art in the CAD-based vision, a completely automated reconstruction does not seem to be feasible. Images of industrial scenes, as shown in Figure 2, are often characterized by a very high complexity, large numbers of (partial) occlusions, and many objects that are not relevant for the reconstruction. Specular reflections and dirt on the installations further complicate an automatic scene analysis. For these reasons, we took a semi-automatic approach, which allows us to take advantage of the image interpretation skills of a human operator.

In Section 2, we first discuss previous work related to object reconstruction from images. In Section 3, we discuss the various object model representations and constraints that are used in the developed approach. The Section 4 describes the measurement process and elaborates on the mathematical equations underlying the manual and automatic alignment of the object models to the images. Experimental results of the implementation are shown in Section 5. The main features and results of the developed method are summarized in the Section 6.

2 RELATED WORK

To reconstruct a 3D model from a collection of 2D images of an object, Rockwood and Winget (1997) define an energy function between the object's images and corresponding images of an articulated mesh representing the 3D object. The energy function measures the differences between the images of the object and the equivalent images of the articulated mesh. Then the 3D model is reconstructed by repeated adjustment of the mesh using simulated annealing to minimize the energy function.

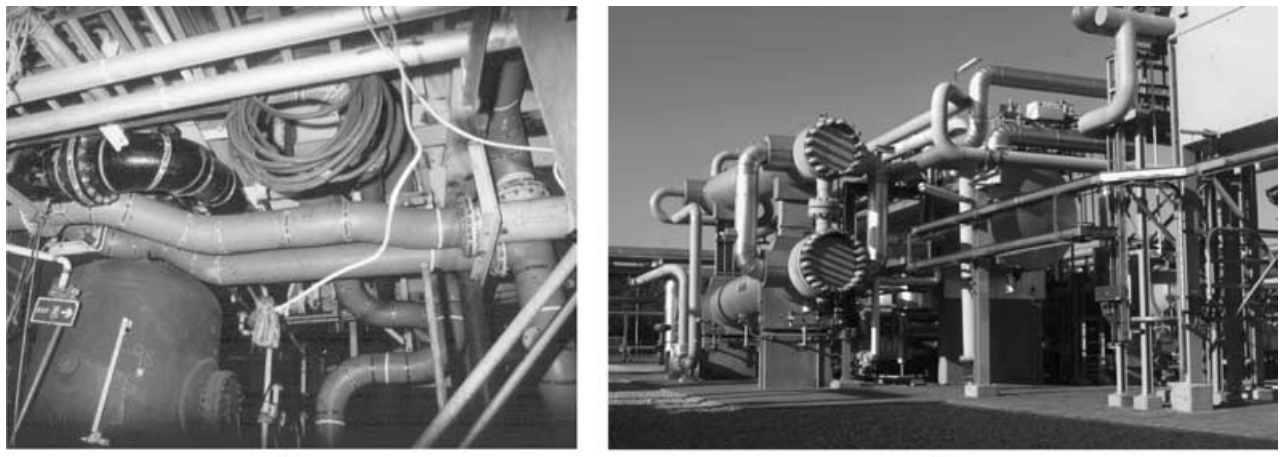

Fig. 2. Complex industrial environments: (a) a ship's engine room and (b) a gas-drying plant.

Fua (1996) uses an energy function to fit a polyhedral object model to a collection of 2D images of an object. This energy function is defined as the negative sum of the absolute gray value gradients along the edges of the object in all the images. With the steepest gradient method the positions of the corners of the polyhedral model are adapted such that the energy function is minimized. Lowe (1991) describes a least-squares algorithm that fits the projected edges to edge pixels. These edge pixels are defined as the pixels with a gray value gradient above some preset threshold. The perpendicular distances of the edge pixels to the nearest edge quantify the errors in the approximate alignment of the object model. The fitting algorithm estimates the changes to the values of the pose and shape parameters that minimize the square sum of these distances. Vosselman and Veldhuis (1999) compared the performance of Fua and Lowe's approach and found that the latter approach is much faster due to the comparatively small number of required iterations. This is one of the reasons why we also selected a least-squares approach for our object reconstruction method.

Debevec et al. (1996) describe a system for modeling and rendering architecture from photographs. The operator instantiates parameterized geometric primitives such as boxes, prisms, and surfaces of revolution. The primitives are located by marking edges in the images that correspond to edges of the model. Thompson et al. (1999) describe a system for feature-based reverse engineering of mechanical parts from 3D points acquired with a noncontact sensor. To model a feature the user selects a feature type and a view in which the parameterized feature model can be seen on the object. The user then specifies enough points on the displayed image to indicate the approximate location and shape of the feature. Finally, the parameterized feature model is fitted to the sensed 3D point data. Our approach improves on that of Debevec and Thompson because it allows the measurement of contour edges of curved object models like cylinders.

Several authors developed environments for the reconstruction of cylinders from the measured straight contour edges in multiple images (Jones et al. 1996; Benning 1997; Hilgers et al. 1998). The approach described in this article, however, is not specific to the reconstruction of cylinders, but can deal with all kinds of objects with a parameterizable model.

Werghi et al. (1999) describe a framework, whereby constraints are added and integrated in the model reconstruction process, resulting in an optimal trade-off between minimization of the shape-fitting error and the constraint tolerances. The range data are given as sets of 3D measurement points representing quadric surfaces belonging to the object to be measured. Their method estimates the different surface parameters, taking into account the geometric relationships between these surfaces. The sum of the least-squares residuals of the measurement points is minimized taking into account constraint values within a predefined tolerance. In our approach both the measurements and the constraints are handled in the same way as explained later in this article.

The functions available in the CAD packages for interactive manipulation differ from the functions needed for model-based measurement. For instance, van Emmerik (1990) describes a system to translate, scale, or rotate a 3D-object model by dragging control points in a 2D window displaying the object model. In this approach the positions of the displayed edges and vertices change. In the measurement process, the operator needs another kind of interactive manipulation in which edges are dragged one by one. Lang and Förstner (1996) describe an approach dedicated to model-based measurement. By dragging corner points of an object model to the corresponding position in an image, one or two parameters of the object model can be adapted. The parameters that are adapted depend on the selected corner point. In our approach, both corner points and points on the edges of the hidden-line projections can be dragged to their correct location. This feature enables us to also reconstruct objects without corner points such as cylinders. Furthermore, we do not require a direct relationship between the points selected by the operator and the adapted parameters. As a consequence, the operator does not need to have knowledge about the object parameterization.

Although the research described in this paper focuses on the reconstruction from images, it should also be noted that laser-scanning devices placed on tripods receive increasing attention (Varady et al. 1997; Chapman and El-Hakim 2000; Son et al. 2002). Both data sources have their strengths. Photographs are often easier to interpret and more accurate with respect to the location of object edges, whereas point clouds derived by laser scanners may offer a higher degree of automation. In our ongoing research, we will focus on the integration of both data sources for a more efficient reconstruction of plant models.

3 REPRESENTATION OF OBJECT MODELS

Object models can be described in various ways (Mortenson 1997). The two most important types are briefly discussed, because the way of representing an object model affects the way in which the parameters of the model can be manipulated.

Constructive solid geometry (CSG) is widely used for the object modeling. With this technique, an object is

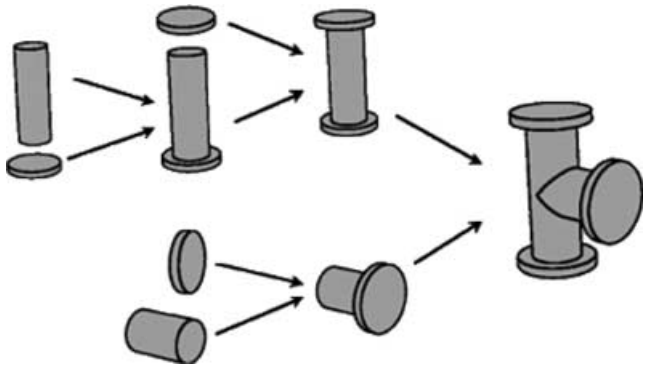

Fig. 3. CSG tree of a flanged T-junction.

composed by taking unions and intersections of several primitive shapes such as rectangular boxes, spheres, cylinders, and cones. Each primitive has its own shape parameters and its pose (position and orientation) parameters. For example, a cylinder has two shape parameters, radius and length; three position parameters, x, y, and z; and two parameters that define its orientation (because it is symmetric about its axis). With a CSG tree, these primitive elements are combined to represent more complex, composite objects such as flanged pipes and T-junctions. Finally, one CSG tree that combines the CSG trees of its primitive elements can describe the industrial installation itself. Figure 3 shows the construction of the CSG tree of a flanged T-junction.

Boundary representations (B-rep) of objects describe the geometry of the points, edges, and surfaces of the object boundaries together with the topological relationships between these points, lines, and surfaces. B-rep modeling is more generic than the CSG modeling but it is also more verbose. For instance, for the representation of a cube, the positions of eight corners are needed as well as a topological graph that connects the surfaces to the edges that are in turn connected to the vertices. Because, in general, a B-rep model requires more parameters to describe a shape than a CSG model, we selected CSG as a representation scheme to model parts of industrial installations. Besides, CSG modeling is commonly used for plant design. Within our measurement process we directly estimate the pose and shape parameters of the CSG primitives. The B-rep models are also used, but only as an intermediate step in the computation of the visible edges of the wire frames projected into the images.

To obtain useful object models with CSG trees, constraints on the parameters have to be specified. For composite objects we distinguish the following three types of constraints. The mathematical formulation of these constraints is elaborated in Ermes et al. (1999).

1. *Internal constraints.* Internal constraints describe restrictions on the parameters of the composing primitives of the CSG tree. A T-joint, for instance, is composed of two cylinders that are connected at a 90° angle. Such conditions are already specified in the CAD-model catalog.
2. *External constraints.* External constraints describing relations between components described by the CSG trees have to be specified. Most often this relates to the connection of adjacent components. These constraints are indicated while modeling an installation.
3. *Standard constraints.* Standardized components can be described by models with fixed values of one or more shape parameters. For instance, often the radius of pipes is known or restricted to a limited number of values. These kind of constraints are again specified in the CAD-model catalog.

These constraints not only assure a consistent modeling of a site but also have a large impact on the efficiency of the reconstruction. Each constraint that can be specified reduces the degrees of freedom of the model. As a consequence, it also reduces the number of measurements that are required to obtain an approximate alignment of the model to the images. The extensive usage of constraints can lead to substantial reductions in the number of measurements.[1]

4 THE MEASUREMENT PROCESS

The operator measures a complete industrial installation in an incremental fashion using the CSG tree data structure described in the previous section. After the first component is measured, the operator selects the object model of an adjacent component from the catalog. The measurement of this component takes into account the component's internal constraints (as specified in the catalog) and the external constraints with the adjacent previously measured components. The confidence in a constraint is specified through weights in the process that estimates the CSG parameters. If it is known that standardized components are used, then high weights are assigned to the standard constraints used in these components. If the operator is unsure that two pipes are perpendicular, an external constraint with a low weight is preferred.

Images are assumed to be taken with a calibrated camera (i.e., known focal length, principle point, and lens distortion). During the interactive modeling, it is assumed that the viewpoints of the images (the so-called exterior orientations) are known. In Section 5.1, it will be explained how the estimation of the image orientations

can also be included in the reconstruction process. To simplify the description of the principles we will assume orientations discussed in this section.

Using orientation information, the visible edges of the wire frame of the selected model can be calculated and projected into the images. This projection also computes the contour edges of curved surfaces. Because the parameters of the object are yet unknown, the model is initially instantiated with an arbitrary pose and shape within the field of view of a selected image (Figure 1a).

The actual measurement of the component is now performed in two steps: (1) interactive manipulation of its hidden-line projections in multiple images and (2) precise alignment to the component boundaries in the image using image analysis. In both steps of the measurement procedure, weighted observation equations are formulated that establish the relation between the position of the edges in the images and the parameters of the CAD model. Also, the constraints are formulated as weighted observation equations. In general, a set of these nonlinear observation equations is too complex to solve analytically. Therefore, these equations are linearized and solved using an iterative weighted least-squares algorithm.

4.1 Manual alignment of the model to the images

The operator adapts the model by selecting a point in the vicinity of an edge of the projected model with a mouse. For the selected point a so-called observation equation, as discussed in the next section, is set up. This observation equation relates incremental changes to the pose and shape parameters of the model to the distance between the mouse pointer and the edge. New pose and shape parameters are computed by a least-squares approach using all the observation equations set up so far, and a new hidden-line projection is computed. Next, the operator moves the pointer to the component boundary in the image. The observation equation is updated, taking into account the new distance to the edge. This is repeated in a loop, using the current pointer position in each iteration, until the operator stops measuring by releasing the mouse button. Experience shows that when the iterations are computed fast enough (less than 1 second per iteration, approximately) the operator gains the impression that he or she is dragging the model edge, along with the pointer, to the object boundary in the image.

4.2 The observation equations for the manual alignment

For all points measured by the operator, observation equations are set up. These equations relate the distance between a measured point and the nearest wire frame edge to the changes in the parameters of the CSG model. For each measured point the observation equation

$$\Delta u_j = \sum_{i=1}^{n} \left(\frac{\partial u}{\partial p_i} \right)_j \Delta p_i, \qquad W\{\Delta u_j\} = w_j \tag{1}$$

is set up. In this equation, Δu_j is the signed distance from the jth measured point to the nearest projected model edge, p_i ($1 \leq i \leq n$) are the n model parameters, $(\partial u/\partial p_i)_j$ is the partial derivative of the distance with respect to p_i, and w_j is a user-defined weight. Ermes et al. (1999) describe an approach to compute the partial derivatives analytically, which is far more efficient than a numerical approach. Solving the system of all observation equations results in an estimation of changes to the parameters of the CSG model such that the distances between the measured points and the wire frame edges are minimized. The measurements of points on the object edges result into a set of linearized equations as shown above (Equation (1)). Let this set be denoted as

$$\Delta u = A\,\Delta p \tag{2}$$

with Δu as the vector of measured distances, A as the so-called design matrix with the above-mentioned partial derivatives, and Δp as the vector of changes in the parameters of the object model. If the number of measurements (size of Δu) is larger than or equal to the number of parameters (size of Δp), the changes in the model parameters can be estimated by the least-squares method, that is,

$$\widehat{\Delta p} = (A^T W A)^{-1} A^T W \Delta u \tag{3}$$

with W as a diagonal matrix. The diagonal elements represent the weights of the observations. More precisely, W is the inverse of the assumed covariance matrix of the observations. For the manual measurements we assumed a standard deviation of 0.5 pixel. Hence, the weights on the diagonal of W equal 4. When all observations are equally weighted, the weights have no effect on the parameter estimation. They are, however, of importance for the estimation of the precision of the estimated parameters.

When starting the measurement of a component, the number of observations will, of course, be smaller than the number of parameters. To avoid a singularity in solving the equation system, we add the approximate (i.e., current) values of all component parameters as observations to the equation system. In the linearized equations, this implies that the value 0 is used as an observation for the change in a parameter value. This set of n equations

$$0 = \Delta p \tag{4}$$

(with 0 as the null vector of size n) is added to the set observation equations (Equation (2)) with a low weight. The parameter changes are estimated by Equation (3).

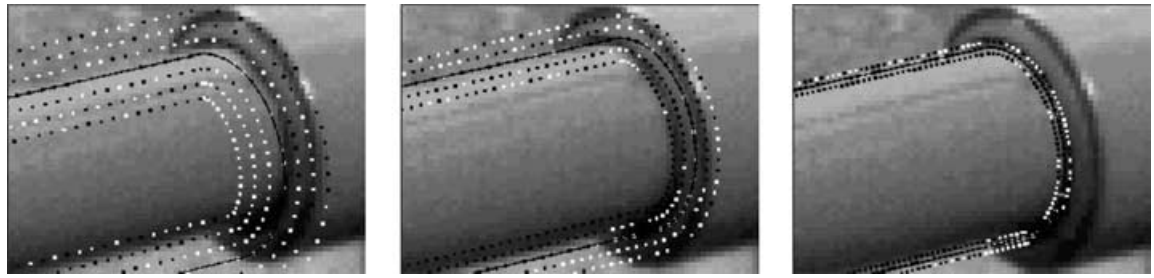

Fig. 4. Pixel-based fitting of pixels to wire frame edges. From left to right: shrinking profile lengths during the iterations. Black profile points were used in the calculation. White profile points were eliminated because their gray value gradients were not perpendicular to the wire frame edge.

The effect of the low weights is such that the additional equations do not contribute significantly to the estimation of a specific parameter if the value of this parameter can be derived from the measured points. However, if the performed measurements are insufficient to estimate a parameter, the least-squares estimation is determined by the additional equations, as a result of which the estimated parameter change will be zero. In general, the behavior of the observation equations is such that the distances between the measured points and the wire frame edges of the changed component as well as the size of the parameter changes are minimized. This procedure allows a stable dragging of the wire frames over the images.

While moving the mouse pointer to the correct position, many mouse events are generated. At each mouse location, only one iteration of the parameter estimation is performed. For an accurate parameter estimation with linearized equations one usually needs to iterate. Because the distances between the subsequent mouse locations are very small when precisely positioning the mouse pointer in the image, the approximate values obtained by the estimation at the previous mouse location are already accurate. Therefore, one iteration at each mouse location suffices for an accurate parameter estimation.

4.3 Fitting the model to image gradients

In the second step of the measurement procedure, digital image analysis is applied to accurately fit the projected model onto the component in the images. The fitting procedure consists of a pixel-based fitting stage and an edge-model-based fitting stage. Both fitting stages are implemented by an iterative-weighted least-squares algorithm.

In the pixel-based fitting stage, the pose and the shape of the object is computed by fitting the edges of the object model to pixels with high gray value gradients using an iterative least-squares estimation. For a number of pixels in a buffer around the projected model edges, linearized observation equations are set up. The observation equation of a pixel is weighted by the squared gray value gradient at the pixel. Setting up an observation equation for each pixel in the buffer would make the fitting algorithm time consuming. Therefore, we subsample the buffer with profiles perpendicular to the edges and shrink the buffer after each iteration of the pixel-based fitting stage by reducing the length of the profiles as illustrated in Figure 4. Observation equations are set up only for points on these profiles. The point density is increased each time the buffer size is reduced. If the profile length is reduced to 2.0 pixels the last iteration of the pixel-based fitting stage is performed.

Next, the edge-model-based fitting stage is initiated. In the edge-model-based fitting stage, the gray values in the buffer are matched with the gray values predicted by a Gaussian-smoothed step edge, and the standard deviations of the measured parameter values are estimated. Again for each pixel in the buffer, a linearized observation equation is set up as described below. The iterations of the edge-model-based fitting stage are repeated until the adjustment of each parameter value divided by its estimated standard deviation is smaller than a predefined error bound (e.g., 0.01).

4.4 The observation equations for the model fitting

The pixel-based fitting process relates the estimation of the pose and shape parameters to the gradient values of the pixels in the vicinity of the projected model edges. For this purpose, we modified the fitting approach by Lowe (1991). Lowe introduces an observation equation for each pixel, which fulfills two conditions: (1) the gray value gradient should be above some threshold, and (2) the pixel should be within some range of a projected edge. We follow Vosselman and Veldhuis (1999) in dropping the first condition. This means that for all pixels within some range of a projected edge, the observation equation is introduced. To ensure that the pixels with the higher gradients dominate the parameter estimation, the squared gray value gradients of the pixels are used as

weights to the observation equations, that is,

$$\Delta u_j = \sum_{i=1}^{n} \left(\frac{\partial u}{\partial p_i} \right)_j \Delta p_i, \qquad W\{\Delta u_j\} = \left(\left(\frac{\partial g}{\partial u} \right)_j \right)^2 \tag{5}$$

where Δu_j $(1 \leq j \leq m)$ is the signed distance from the jth pixel to the projected model edge, and $(\partial g/\partial u)_j$ is the gray value gradient at the pixel perpendicular to the model edge. The partial derivatives $\partial u/\partial p_i$ are again computed analytically, according to Ermes et al. (1999). Using the weights instead of the first condition by Lowe has the advantage that one does not need to specify a threshold for the selection of edge pixels. It also results into a higher accuracy. The incremental changes in the parameters are again determined by least-squares estimation (Equation (3)). The weight matrix W is a diagonal matrix with the gray value gradients $W\{\Delta u_j\}$ as the diagonal elements.

Note that in the pixel-based approach, the observation Δu_j is not stochastic and that the weights do not represent measurement accuracies. Therefore, it is not possible to estimate the standard deviations of the obtained parameter values by propagation of the standard deviation of the observations. In the edge-model-based fitting approach presented next, the standard deviation of the observations can be estimated and propagated to the parameter values. Experiments showed that the edge-model-based approach requires better approximate positioning by the operator and more iterations than the pixel-based approach. We therefore first iterate with the pixel-based approach and switch to the edge-model-based approach for the last few iterations.

The edge-model-based fitting procedure matches the gray values in a buffer around the projected edges of an object model with the gray values predicted by the Gaussian-smoothed step edge. We apply the principle of differential matching described by Förstner (1993) to estimate a shift of the model edges and to propagate the standard deviation in the gray values to the standard deviation of the obtained parameter values. At the transition between the object and the background, there is a discontinuity in the gray values that can be described by a step function. Due to the image formation process this transition loses its sharpness in the images. This is modeled by smoothing the step edge with a Gaussian as illustrated by Figure 5. The obtained Gaussian step edge h predicts the gray values at the pixel by

$$h(u) = h_{\text{bg}} + (h_{\text{obj}} - h_{\text{bg}}) \frac{1}{\sigma_{\text{s}}\sqrt{2\pi}} \int_{-\infty}^{u} e^{-t^2/(2\sigma_{\text{s}}^2)}\, dt \tag{6}$$

where u is the signed distance of the pixel to the transition between the object and the background. h_{obj} and h_{bg} denote the gray values of the object and the background,

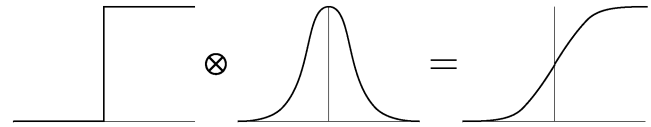

Fig. 5. Convolution of a step edge by a Gaussian kernel as model for a gray value profile across an edge.

respectively. The parameter controlling the amount of smoothing is σ_{s}.

For each pixel within some range of the projected model edges an observation equation is set up as follows. First, the shift Δu_j between the projected model edge and the transition between the object and the background is related to the difference of the observed gray value and the predicted gray value at the pixels divided by the first derivative of the smoothed step edge, that is,

$$\Delta u_j = \frac{h(u_j) - g(u_j)}{h'(u_j)} \tag{7}$$

where u_j $(1 \leq j \leq m)$ is the signed distance at the jth pixel to the projected model edge, and $g(u_j)$ is the gray value at the pixel.

The linearized equation (5) relates this shift Δu_j to the adjustments to the parameter values Δp_i. In this equation, $(\partial u/\partial p_i)_j$ is the partial derivative of the signed distance at the jth pixel to the projected edge with respect to the parameter p_i. By combining Equations (5) and (7), we obtain the observation equation

$$\Delta g_j = h(u_j) - g(u_j) = \sum_{i=1}^{n} \left(\frac{\partial u}{\partial p_i} \right)_j h'(u_j) \Delta p_i \tag{8}$$

that directly relates the difference of the actual and predicted gray value at a pixel to the required changes in the parameter values. After each iteration step the standard deviation σ_g of Δg_j is estimated from the residuals of the least-squares fit by

$$\hat{\sigma}_g = \sqrt{\frac{\sum_{j=1}^{m} (h(u_j) - g(u_j))^2}{m - n}} \tag{9}$$

For each pixel within some range of the projected model edges, such an observation equation (8) is set up with weight $1/\hat{\sigma}_g^2$, where $\hat{\sigma}_g$ is computed at the previous iteration step. Furthermore, the constraints are included in the adjustment as weighted linearized observation equations as described by Ermes (2000). The constraint equations can contain several different types of variables. All the parameters (pose and shape) are directly related to the components in the CSG tree.

The parameter values are updated with the weighted least-squares estimates of Δp_i. The projections of the model edges are recomputed, and the weighted least-squares adjustment described above is repeated until the

adjustment of each parameter value divided by its estimated standard deviation is smaller than a predefined error bound or the relative improvement of the estimated standard deviation $\hat{\sigma}_g$ is below a predefined threshold.

5 RESULTS

In this section, we illustrate the application of our photogrammetric modeling technique with a particular example: modeling a part of a gas-drying plant. We first focus on the manual part of the modeling process in Section 5.1. In Section 5.2 results of the automatic model-fitting procedure are presented, showing the effect of applying external constraints. Examples of other projects and a short description of the method are available through our website at http://www.geo.tudelft.nl/frs/piping/.

5.1 Modeling a gas-drying installation

A pilot project was undertaken at a gas exploration site of the Nederlandse Aardolie Maatschappij (NAM) in Tjuchem, the Netherlands. The object analyzed was a gas-drying unit, which measures about 30×30 square meters and has a maximum height of about 5 meters. Over a period of 3 hours, 102 images were recorded on the site using a handheld Kodak DCS420 camera (at a resolution of 1012×1524 pixels). For the modeling of this part of the installation 36 images were selected. In Figure 2b one of the images is shown. The interior orientation and lens distortion of the camera were calibrated beforehand. Several dimensions of the installation, such as radii and lengths of piping elements, were tape measured to scale and check the model.

The measurements started with the measurement of a large box-shaped element. The object co-ordinate system is defined by the pose parameters of the box that are constrained to their default values. All the images that show this box could be oriented by interactively adapting the exterior orientation parameters of these images in such a way that the backprojected wire frame matches the image of the box. The exterior orientation parameters of these images are improved by measuring other components of the installation, and adjusting their shape and pose together with the exterior orientation parameters of the images in an overall bundle adjustment.

In fact, the reconstruction process consists of iteratively introducing new model parts and new images in which these parts are visible. The reconstructed model is continuously improved by least-squares adjustment of all the parameters involved. The installation was modeled by 218 CSG models, which were built with a total of 288 geometric primitives. Figure 6a depicts one of the images with the final model of the installation in an overlay. Figure 6b shows the reconstructed model. The primitives and the image orientations are parameterized with a total of 2,706 parameters. The operator measured 2,744 points on model edges and 1,800 constraints were applied, giving a total of 4,544 observation equations. Note that the redundancy is considerable due to the use of the constraints. Without the constraints the operator would have to measure many more edges of the model.

The reconstruction of the model took approximately two weeks, although one must keep in mind that the current application is a not very user friendly first prototype and that a few new primitives (such as a prism and a sphere) needed to be incorporated in the software for this project. It is expected that the need for new primitives will rapidly decrease in the next projects. Note that changes in the software were only required for new shape primitives, not for components that can be composed by set operations on the primitives, or for components with different parameter values as the default components in

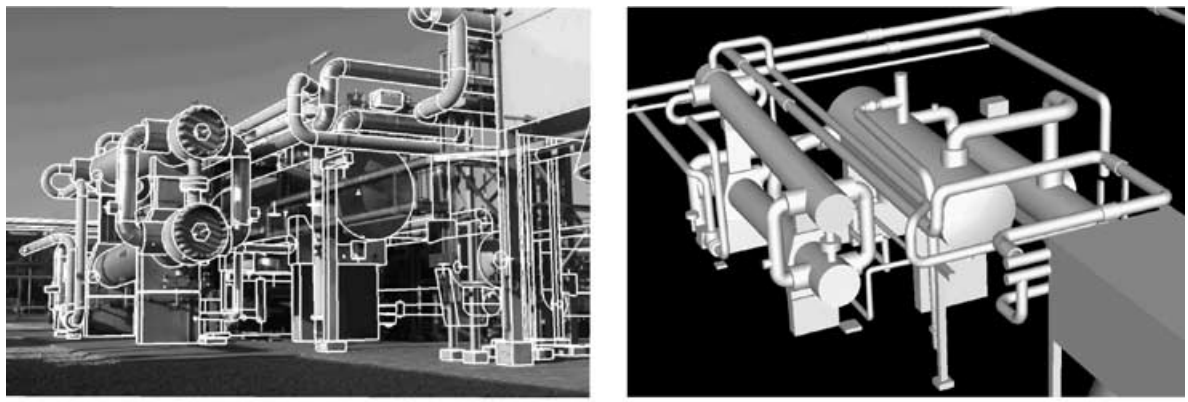

Fig. 6. Reconstruction of a gas-drying plant: (a) wire frame edges of models as overlay on one of the images and (b) resulting 3D model.

the catalog. The latter two kinds of new components can easily be created and inserted in the catalog.

Improvements to the user interface should still be made to increase the efficiency and make this approach more attractive for production work. A comparison with the performance by standard photogrammetric reverse engineering is yet to be done. Parts of the gas-drying plant were reconstructed by a company using standard photogrammetry a few years ago. This reconstruction, however, only consisted of the straight pipes. More complex object parts had not been modeled and constraints were not used. Because of the usage of constraints in our approach we expect a higher overall accuracy when using the same imagery.

Assuming an accuracy of the image measurements made by the operator of 1 pixel, the estimated accuracy in object space is approximately 5–10 mm. We checked the remaining tape measurements against the photogrammetrically obtained model dimensions and the differences were all less than the estimated accuracy.

5.2 Fitting the model to the images

Two images from different viewpoints that show a curved pipe with two straight pipe elements connected to it are used to illustrate the model fitting. This typical example also demonstrates the benefits of the external constraints on the accuracy of the obtained parameter values. We measured the two straight pipes and the curved pipe with the parameterized cylinder and torus models described below.

The orientation of a cylinder is defined by the direction of its axis. This direction is specified by a normalized vector (t_0, t_1, t_2) relatively to a reference coordinate system. Hence, the pose and the shape of the cylinder is described singularity-free by three parameters x, y, and z denoting its position; three parameters t_0, t_1, and t_2 denoting its orientation; and two shape parameters l and r denoting its length and radius, respectively.

A curved pipe element is not symmetric around an axis and therefore its orientation is defined by a normalized quaternion vector (q_0, q_1, q_2, q_3) providing a singularity-free representation of the rotation (Ermes et al. 1999). A curved piping element has the shape of a torus. The centerline of the pipe is a circle with a major radius R. A torus is obtained by sweeping the center of a smaller circle with a minor radius $r < R$ along the centerline of the piping element. The boundary of the torus is the surface swept out by the circle. An angle denotes the part of the centerline of the pipe element that is followed. For instance, a curved pipe between two perpendicular straight pipes is described by a torus with an angle of 90°. Therefore, we model this pipe element by a torus described by three parameters x, y, and z denoting its position; four normalized quaternion parameters q_0, q_1, q_2, and q_3 denoting its orientation; two shape parameters r and R denoting its minor and major radii; and a shape parameter ω denoting the torus angle.

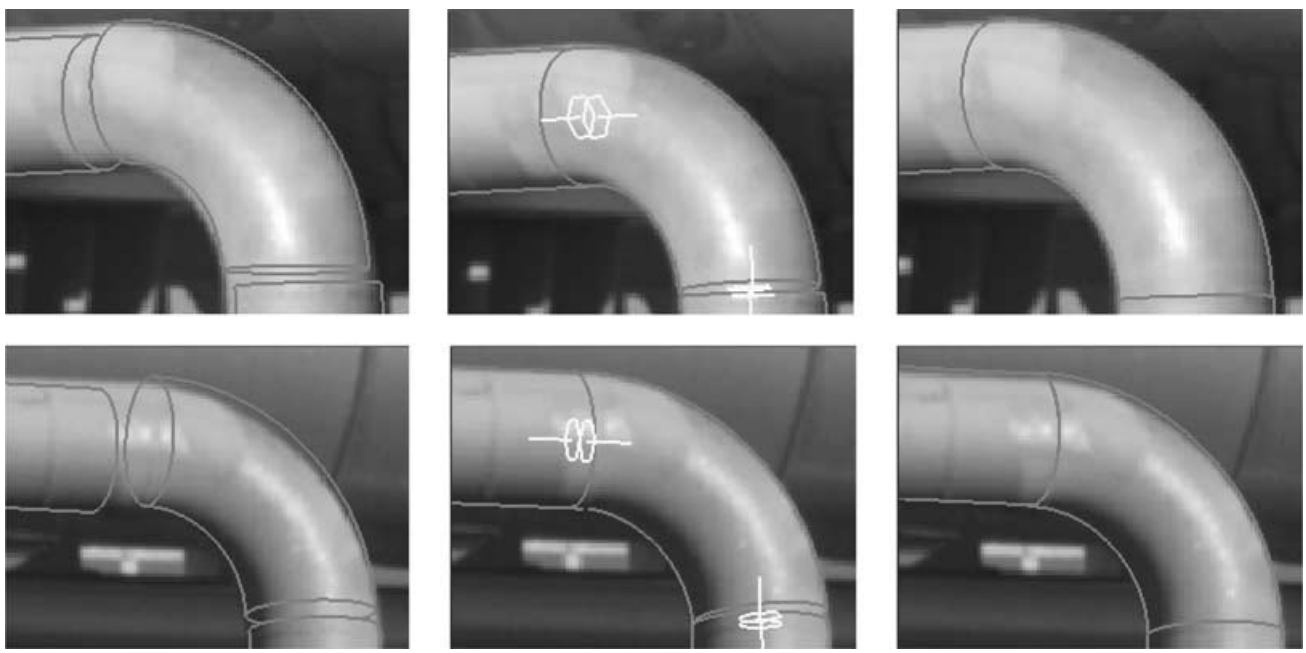

Fig. 7. Effect of using constraints in fitting. Left: Approximate positions. Middle: Fitting results without using constraints. Right: Fitting results using constraints.

One end of the curved pipe is connected to the horizontal straight pipe and the other end is connected to the vertical straight pipe. Therefore, for each curved pipe end the operator has to add a position constraint and a direction constraint both with a high weight. These constraints ensure that the curved pipe ends meet with the straight pipes and have identical orientations. Furthermore, two constraints are introduced to set the radii of the straight pipes equal to the radius of the curved one.

Figure 7 shows the three pipes after the measurements by the operator, after fitting without constraints and fitting with constraints. The iteration process of the least-squares fitting was stopped when the updates of all parameters were smaller than 1% of their formal standard deviation, or when the change in the estimated variance factor became smaller than 0.1%.

Table 1 compares the estimated standard deviations of the parameters resulting from the fitting procedure with and without constraints. Fitting without constraints does not allow all parameters to be estimated accurately, the largest standard deviation of the position parameters being 3.27 mm. There are two reasons for this. First, the geometry of the configuration (i.e., the relative position) of the images is essential for the precision of the parameters. Second, the transitions from curved to straight pipes are smooth and are not clearly defined in the images. As a result, the largest estimated distance between the positions of the pipe ends that are physically connected is 35 mm (Figure 7, middle). Applying the constraints described above reduces the standard deviations of most of the parameter values. Note that this is not the case for the vertical pipe. This is because, as a consequence of the image configuration, a large correlation occurs between the z-position and the length l. This is not improved by the constraints. For the horizontal pipe, a high correlation between the y-position and the length remains (this pipe is oriented in y-direction). The curved element profits most from the introduction of the constraints. The standard deviations of the position parameters reduce by a factor of about 3. The orientation parameters of the element improve with at least a factor of 10. The reduction of the standard deviation of the torus angle ω is most striking. Due to the connection to the two straight pipes this shape parameter is estimated with a high precision.

Table 1
Comparison of the parameter standard deviations (σ_p) of the fitting with and without constraints for the curved (C) pipe connected to the horizontal (H) and vertical (V) straight pipe (σ_p in mm for x, y, z, l, r, and R; ω is in degrees, and the orientation parameters t and q are dimensionless)

Constraints	*No*	*Yes*
H		
x	1.35	0.24
y	2.62	0.93
z	0.29	0.22
t_0	1.2×10^{-3}	0.3×10^{-3}
t_1	0.01×10^{-3}	0.007×10^{-3}
t_2	0.2×10^{-3}	0.2×10^{-3}
l	3.10	1.29
r	0.15	0.16
C		
x	0.67	0.19
y	0.47	0.16
z	2.23	0.91
q_0	3.9×10^{-3}	0.1×10^{-3}
q_1	3.2×10^{-3}	0.06×10^{-3}
q_2	3.4×10^{-3}	0.05×10^{-3}
q_3	0.1×10^{-3}	0.009×10^{-3}
R	1.53	0.96
r	0.17	0.22
ω	0.44	0.01
V		
x	0.23	0.22
y	0.18	0.18
z	3.27	3.40
t_0	0.1×10^{-3}	0.1×10^{-3}
t_1	0.1×10^{-3}	0.1×10^{-3}
t_2	0.007×10^{-3}	0.007×10^{-3}
l	4.20	3.52
r	0.08	0.07

6 CONCLUSIONS

In this article, we presented a new semi-automatic procedure for measuring the as-built situation of industrial installations from a collections of images. The manual interaction is limited to image interpretation, rough alignment of backprojected CAD models with the image content, and constraint specification. The image interpretation leads to the selection of a CAD model from a database with standard components. For the alignment, a graphical user-interface is developed that allows manipulation of shape and pose parameters of the CAD model that is projected onto the images. The method to manually align the model to the images offers a great amount of flexibility. Points, object edges, and contour lines can be dragged with the mouse pointer to their correct location in the image, and as a result the parameters of the model are continuously updated. This is far more efficient than the reconstruction from a large number of three-dimensional points measured separately using a classical photogrammetry approach. Once the pose (position and orientation) and shape parameters of the object are approximately known, digital image analysis is used to accurately fit the hidden-line projections onto the object in the images. The application of constraints between and within CAD models reduces the number of required manual measurements and also the number

of images needed. Furthermore, the precision of the estimated parameters is improved, and a more consistent model is obtained. Many parts of the model are difficult if not impossible to reconstruct without the use of object constraints.

The practical application of this CAD-based measurement method has been demonstrated by reverse engineering of a part of a gas-drying plant. The digital images and a few tape measurements were acquired in about 3 hours. No artificial targets were used. The model was constructed from 36 images and consists of 218 components. Model reconstruction was accomplished with a positional accuracy better than 1 cm for most components. The effects of applying constraints between different components were demonstrated with an example. The resulting improvement in the precision of the shape and pose parameters depends, strongly, on the geometry of the configuration, but exceeded by a factor of 3 for many parameters.

ACKNOWLEDGMENTS

We thank the Nederlandse Aardolie Maatschappij (NAM) for supporting the pilot project at a production cluster of the Groningen gas field. Also, we thank Ingenieursbureau Geodelta for providing the images of the engine room of the ship. This research was supported by the Dutch Technology Foundation (STW) and the IST program of the European Union.

NOTE

1. On http://www.geo.tudelft.nl/frs/piping/overview.html#constraints a short movie demonstrates how, with the help of constraints, a flanged T-junction modeled by five cylinders (i.e., 35 parameters) can be reconstructed by measuring only eight points.

REFERENCES

Benning, W. (1997), PHIDIAS-MS—Eine Digitale Photogrammetrieapplikation unter MicroStation für Nahbereichsanwendungen, *Allgemeine Vermessungsnachrichten*, **104**(1), 16–25.

Chapman, D. & El-Hakim, S. (2000), Short-baseline active triangulation for CAD reconstruction of room-sized industrial environments, *International Archives of Photogrammetry and Remote Sensing*, **33**, part B5, 122–29.

Cooper, M. A. R. & Robson, S. (1996), Theory of close range photogrammetry, in K. B. Atkinson (ed.), *Close Range Photogrammetry and Machine Vision*, vol. 2, Whittles Publishing, Caithness, Scotland, UK, pp. 9–51.

Debevec, P. E., Taylor, C. J. & Malik, J. (1996), Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach, in H. Rushmeier (ed.), *Proceedings of SIGGRAPH 96*, Addison Wesley, Reading, MA, pp. 11–20.

Ermes, P. (2000), Constraints in CAD models for reverse engineering using photogrammetry, *International Archives of Photogrammetry and Remote Sensing*, **33**, part B5, 215–21.

Ermes, P., van den Heuvel, F. A. & Vosselman, G. (1999), A photogrammetric measurement method using CSG models, *International Archives of Photogrammetry and Remote Sensing*, **32**, part 5W11, 36–42.

Förstner, W. (1993), Image matching, in R. M. Haralick and L. G. Shapiro (eds.), *Computer and Robot Vision*, vol. II, Adison-Wesley, Reading, MA, pp. 289–378.

Fua, P. (1996), Model-based optimization: accurate and consistent site modeling, *International Archives of Photogrammetry and Remote Sensing*, **31**, part B3, 222–33.

Hilgers, G., Przybilla, H.-J. & Woytowics, D. (1998), The digital photogrammetric evaluation system PHAUST for as-built documentation, *International Archives of Photogrammetry and Remote Sensing*, **32**, part 5, 226–29.

Jones, M. A., Chapman, D. P., Hamid, A. A. & Deacon, A. T. D. (1996), Close range photogrammetry using geometrical primitives for efficient CAD modelling of industrial plant, *International Archives of Photogrammetry and Remote Sensing*, **31**, part B5, 284–89.

Lang, F. & Förstner, W. (1996), 3D-city modeling with a digital one-eye stereo system, *International Archives of Photogrammetry and Remote Sensing*, **31**, part B3, 415–20.

Lowe, D. (1991), Fitting parameterized three-dimensional models to images, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **13**(5), 441–50.

Mortenson, M. E. (1997), *Geometric Modeling*, 2nd ed., Wiley Computer Publishing, New York.

Rockwood, A. P. & Winget, J. (1997), Three-dimensional object reconstruction from two-dimensional images, *Computer-Aided Design*, **29**(4), 279–85.

Son, S., Park, H. & Lee, K. H. (2002), Automated laser scanning system for reverse engineering and inspection, *International Journal of Machine Tools & Manufacture*, **42**(8), 889–97.

Thompson, W. B., Owen, J. C., de St. Germain, H. J., Stark, S. R. & Henderson, T. C. (1999), Feature-based reverse engineering of mechanical parts, *IEEE Transactions on Robotics and Automation*, **15**(1), 57–66.

van Emmerik, M. J. G. M. (1990), A direct manipulation technique for specifying 3D object transformations with a 2D input device, *Computer Graphics Forum*, **9**(4), 355–61.

Varady, T., Martin, R. R. & Cox, J. (1997), Reverse engineering of geometric models—an introduction, *Computer-Aided Design*, **29**(4), 255–68.

Vosselman, G. & Veldhuis, H. (1999), Mapping by dragging and fitting of wire-frame models, *Photogrammetric Engineering & Remote Sensing*, **65**(7), 769–776.

Werghi, N., Fisher, R., Robertson, C. & Ashbrook, A. (1999), Object reconstruction by incorporating geometric constraints in reverse engineering, *Computer-Aided Design*, **31**(6), 363–99.