Computational Understanding of Multiple Images

Reactive LTR Project 21914 — CUMULI
Month 36 Progress Report

January 10, 2000

INRIA Grenoble (F)
INRIA Sophia-Antipolis (F)
University of Lund (S)
Fraunhofer-IGD (D)
Image Systems AB (S)
IMETRIC (CH)

http://www.inrialpes.fr/CUMULI

Esprit LTR 21914 CUMULI

Month 36 Progress Report

The CUMULI Consortium

Introduction

The LTR project CUMULI aims to develop automated methods for accurate 3D industrial measurement from multiple images, with particular emphasis on the geometric aspects of the problem. The consortium is inter-disciplinary, with a mixture of computer vision and geometric reasoning researchers (Grenoble, Sophia, Lund) and applications-oriented partners in close range industrial photogrammetry (IMETRIC), motion measurement from high-speed image sequences (Image Systems) and 3D modelling for augmented and virtual reality applications (Fraunhofer IGD). CUMULI tries to capitalize on the significant recent advances in multi-image geometry in the computer vision community, and apply them in mature, real world visual measurement problems. We also hope that this will lead to some fruitful cross-fertilization between the vision, photogrammetry, reasoning, and modelling communities.

A brief overview of the project is as follows (see the Project Programme for more details). There are three main work packages, each aimed at a different aspect of 3D measurement from multiple images. WP 1 and 2 concentrate on core topics in geometric vision and photogrammetry: projective and uncalibrated structure; constraints, calibration and Euclidean structure; and measurement of non-point-like geometric features (lines, conics, curves, surfaces…). WP 1 focuses on discrete images and the algebraic aspects of these problems. It is directly linked to the IMETRIC demonstrator whose main goal is more flexible methods of recovering multi-camera orientation. WP 2 concentrates on continuous image sequences, tracking, and motion estimation. It is linked to the Image Systems crash-testing demonstrator, in which the non-rigid 3D structure and motion of a car crashing into the side of a stationary car will be recovered from high speed image sequences.

WP 3 is more speculative. Geometric models are more than just coordinates, they are complex networks of incidence relations, constraints, *etc*. To measure or build such models from uncertain data, we need methods of recognizing and mobilizing these constraints to ensure consistency, reduce uncertainty, improve registration, generate new matching hypotheses, …. Hence, geometric reasoning under uncertainty will be a key topic in future systems. WP 3 is an initial attempt to adapt existing geometric theorem proving technology to this sort of application. It is linked to the Fraunhofer IGD demonstrator, which focuses on model-image calibration, registration and occlusion problems for augmented reality applications.

For each topic and site, we give man-months (mm) used this year (months 25–36 of CUMULI). A table at the end shows total man months used to date (months 0–36). At the request of the consortium, CUMULI has received a 6 month unpaid extension and will now finish at the end of February 2000 (month 42).

Note: The partner 'Innovativ Vision Image Systems AB' is now independent of the parent company Innovativ Vision AB, and has thus shortened its name to 'Image Systems AB'.

WP 1: Multi-Camera Geometry, Discrete Images

Work package 1 concerns the theoretical and practical aspects of deriving very accurate estimates of 3D geometrical structure from a discrete set of images of a static scene. The practical work focuses on the needs of IMETRIC, who also supplied the test data. The theoretical and algorithmic results are used not only in this work package, but also form a basis for work package 2 and 3. Hence, Image Systems and IGD have also followed the work in this package and used some of the results.

Task 1.1: Projective Multi-camera Geometry

Grenoble (0.5mm) and Lund (1mm) have collaborated on a study of minimal cases for projective reconstruction from incomplete point correspondences [30]. When all points are visible in all images the minimal cases are well-known, but in practice correspondences in some of the images are often missing. The study introduces a framework for the missing data problem, derives all the minimal cases for 3 and 4 images, and gives practical algorithms for them. The methods are a useful foundation for RANSAC-style approaches to camera initialization.

Grenoble (0.5mm) has specialized its multi-image matching tensor formalism to the case of plane + parallax (images homographically warped to align a common reference plane) [39]. This greatly simplifies the geometry, and allows an efficient and stable multi-image projective reconstruction method based on a rank-1 factorization of a residual parallax matrix into point depths and camera centre vectors (rather than a rank 4 factorization into 3D points and projection matrices).

Task 1.2: Euclidean and Affine Multi-camera Geometry

Lund (1mm) and Grenoble (0.5mm) have continued their collaboration on autocalibration from partial knowledge of the camera intrinsic parameters [21, 19, 22]. Computer algebra and geometric arguments are used to derive the critical motions (motions for which the autocalibration constraints are intrinsically too weak to allow all calibration parameters to be recovered) for several calibration constraints including vanishing skew, known aspect ratio and full internal calibration modulo unknown focal lengths.

Grenoble (0.5mm) have developed a new method for Euclidean structure from motion from three affine views [33]. The method is based on intrinsic three-view properties, in particular an equivalence between 2-D affine cameras and 1-D projective ones operating on the plane at infinity [29]. We show that the relative camera orientations are entirely encoded by the 1-D trifocal tensor of the plane at infinity, and from this derive two new algorithms for SFM from three views. One method uses just the minimal 1-D trifocal tensor, the other the full affine three-view constraints. In contrast, previous three view affine SFM algorithms use only the two-view constraints. The algorithms have been demonstrated on real image sequences.

Grenoble (1mm) has developed several new quasi-linear methods for camera initialization from a single image of a few known 3D points [37]. These lie in the family of n-point quasi-linear 'pose + X' methods, where 'X' is some combination of calibration parameters. Existing methods give pose of a calibrated camera ($n = 4$, X empty), and pose + 5 parameter internal calibration (Direct Linear Transformation, $n = 6$). The new methods give pose + focal length ($n = 4$) and pose + focal length + principal point ($n = 5$). One of the main motivations was to provide useful camera initialization methods for the WP2 and WP3 demonstrators.

Sophia-Antipolis (3mm) have continued to refine their formalism for detecting special types of camera displacement and scene structure in projective, affine and Euclidean geometry. They have

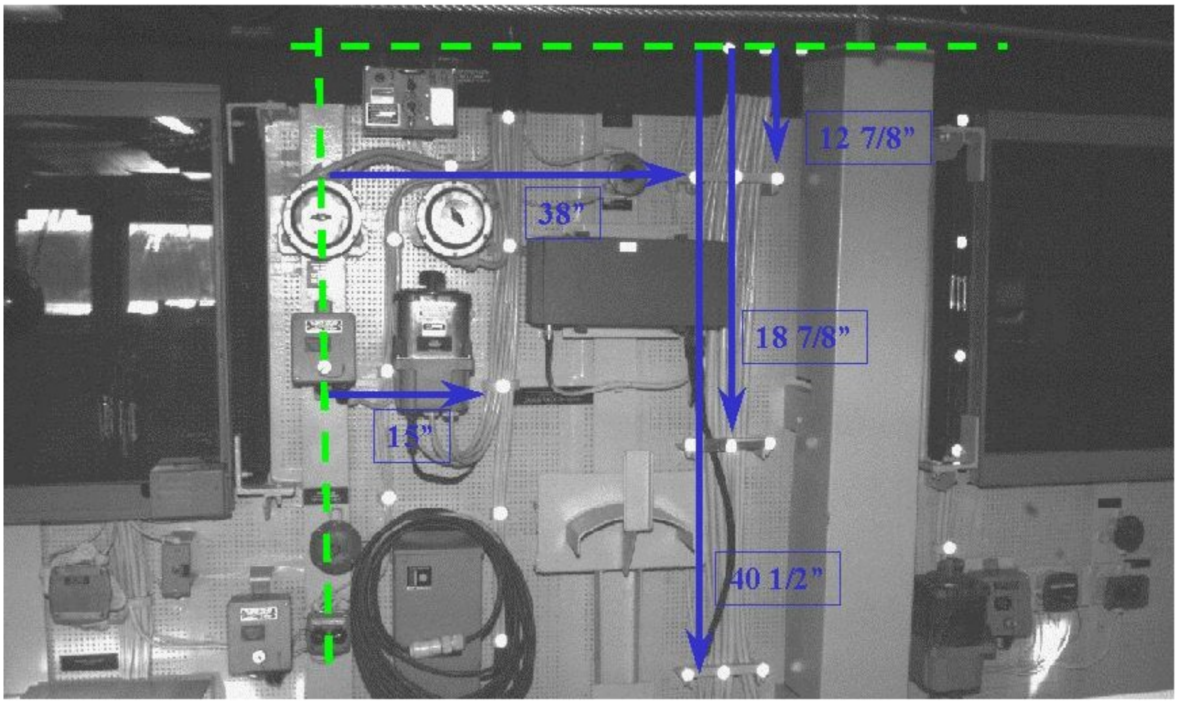

Figure 1: A mock-up of as-built measurement on a ship installation. Currently, reflective targets at known 3D points are needed for camera initialization. We are developing more flexible initialization methods based on unmeasured 3D points and natural non-point scene features (lines, rectangles, cylinders. . .).

derived the combined model and camera singularities of uncalibrated monocular image sequences, for which the fundamental matrix correspondence model degenerates to a homographic one [25, 26, 27]. The code has been tested on CUMULI data and made available to the consortium.

Task 1.3: Geometric Features and Uncertainty

Lund (2mm) has continued its work on minimal structure and motion problems. In a recent study the cases of 5 lines in 4 affine views and 6 lines in 3 affine views were studied and solved [5]. Work has also continued on curve reconstruction [10, 11] and flexible calibration using little information [3].

Grenoble (1mm) have continued their study of the photogrammetric theory of bundle adjustment, and begun an investigation of sparse and iterative numerical algorithms for the linear update prediction step in this [40]. The aim is partly educational — owing mainly to an unfamiliar literature and terminology, the vision community is still unaware of some of the basic photogrammetric developments, which is causing a good deal of duplication of effort — and partly to improve the performance of bundle methods on large and difficult problems.

Task 1.4: Automatic Estimation of Camera Pose

Figure 1 shows a mock-up of as-built measurement on a ship installation. In this case the scene is instrumented with reflective targets, but IMETRIC would eventually like to be able to initialize the camera orientations from untargeted natural scene features (lines, rectangles, cylinders. . .).

Imetric:
1) Testing and improvement of direct solutions for estimation of camera pose (3 mm). The initial implementation of the direct solutions for estimation of camera pose has been improved and evaluated in a large number of industrial tests. The performance of the direct solution was improved to better than 99.9%. The direct solution was augmented by additional checks and redundant solutions, which select the best solution out of several potential ones. The total system can be made 100% reliable. It refuses only 1 out of 1000 measurements according to our tests. The new solution was tested for example via a 3-month trial run at Boeing, which was part of the certification of a TI^2 system to be used for production at Boeing. This is the first time that Boeing has certified a 3D Image Metrology system for guidance of a CNC-machine. This can be seen as a major success for CUMULI and 3D Vision in general. Currently similar tests are going on with a consortium of 3 European aerospace manufacturers. In another application, two automated systems have been in operation since September 1998. These systems are operating with 100% reliability and are a major success in the engineering industry.

2) Preparation of implementation of direct solutions for relative orientation (0.5mm). As already mentioned in earlier reports, there are a number of situations where the placement of the so called "orientation and scale cross" is cumbersome and/or impossible. IMETRIC has now drawn up the procedures required for the use of crossless orientation techniques in industry, and intends to implement these techniques as a next step.

3) Preparation of specifications and test-data for relative orientation from non-point features (1 mm). Unfortunately, it has proven to be very difficult to get clearance for suitable industrial test data for the non-point feature based pose and relative orientation problems. The industrial (nuclear power plant) partner who originally hoped to provide this test data finally withdrew due to confidentiality considerations. IMETRIC then found another partner who was interested in the project and willing to provide data from as-built ship installations. Unfortunately, when the negotiations were almost complete, that company was purchased by a large US defense contractor, who tightened all their security regulations. An application was placed for a "Technical Assistance Agreement" allowing the release of this data (and some other collaborations, too). This is still under consideration by the US State Department. All parties are pushing hard for this to go through, but at present we are unable to say if or when release of the data will be approved. In the meantime, IMETRIC has used the point-based algorithms mentioned above to develop a target-based method for the orientation of these images. Furthermore, a prioritized list of the most useful combinations of natural features for this problem has been prepared for the academic partners.

In view of this, the remaining work on the IMETRIC demonstrator will focus mainly on relative orientation from unknown 3D points (currently in this application, known 3D points are required to find initial camera poses). For the problem of orientation from non-point features, only the industrial motivation and proof-of-concept MATLAB implementations will be shown in the demonstrator. However, IMETRIC does intend to implement these methods fully after the end of CUMULI.

Grenoble (1mm) has continued its work on calibrated camera pose from known 3D points, as improved methods for this well-studied but rather delicate problem would be useful for all 3 industrial demonstrators. We have refined [32] last year's quasi-linear 4- and 5-point methods [31], and introduced a new family of quasi-linear 4-point methods, and a new type of nonlinear 3-point method based on eigendecomposition of a 5×5 matrix [1]. So far, all of the quasi-linear methods we have tested are less accurate than the 3-point nonlinear ones. We are currently working to understand this problem and rectify it, as accurate 4 point (*i.e.* unique solution) methods would be useful in all 3 demonstrators.

Lund (0.5mm): Lund's main goals for the final 6 months of WP1 are:

1) Establish a final implementation of computer vision algorithms to allow the IMETRIC application to work without the orientation cross. (This is a cross bar with coded markers used to calculate initial camera poses. Eliminating it from the process would have important economic and safety advantages). Within CUMULI, the following algorithms have been or will be transferred to the IMETRIC system

- 6 point relative orientation assuming 2 images with known internal calibration. (This was supplied with deliverable 1.1/24+2.1/24)
- 4 point relative orientation assuming 2 images with known internal calibration and assuming that the points are coplanar. (This was supplied with deliverable 1.1/24+2.1/24).
- 8 point relative orientation assuming 2 images with known internal calibration. (This is the standard 8 point algorithms and has been supplied recently.)
- If time allows, model selection algorithms will also be supplied.

2) Demonstrate the usefulness of alternative features. This will be an off-line final demonstrator. It will use the test data supplied by IMETRIC if this is available by the end of the project.

WP 2: Image Streams and 3D Motion

This work package closely parallels WP 1, but emphasizes image streams and motion rather than discrete images of static scenes. Virtually all of the scientific work listed under WP 1 is also applicable to WP 2, but to avoid duplication we have listed it *only* under WP 1. Please keep this in mind and refer to the WP 1 section as necessary. Here we describe only work aimed specifically at image streams and motion estimation.

Task 2.1: Incremental Projective Camera Geometry

Lund (0.5mm) have developed a theory that generalizes the multilinear constraints to the case of a continuous stream of images. In this framework, the epipolar constraint, as well as the trilinear, quadrilinear and all higher order ones have been generalized to the continuous case. A key result is that the third order continuous constraints are needed to reconstruct the scene and the camera motion, *i.e.* optical flow is not sufficient [2].

They have also developed an iterative and recursive algorithm that solves for projective structure and motion from closely sampled image sequences. This algorithm calculates the camera matrices and reconstructs the scene incrementally [18]. The work on iterative and recursive reconstruction algorithms is valuable for the car-crash application, because the vast number of images implies that recursive methods have to be used rather than batch ones.

Grenoble (1mm) has developed [38] an alternative approach to differential multi-image matching constraints for image streams, based on finite difference expansion rather than the Taylor expansion used by Sophia and Lund [42, 2]. The new approach gives much simpler tensors and equations, and has a very much clearer relation to the discrete case. We develop the theory systematically, and consider several special cases including 'tensor tracking' — the updating of a matching tensor along a sequence of images. We relate this to nonlinear optimization iterations for refining the estimate of a fixed tensor.

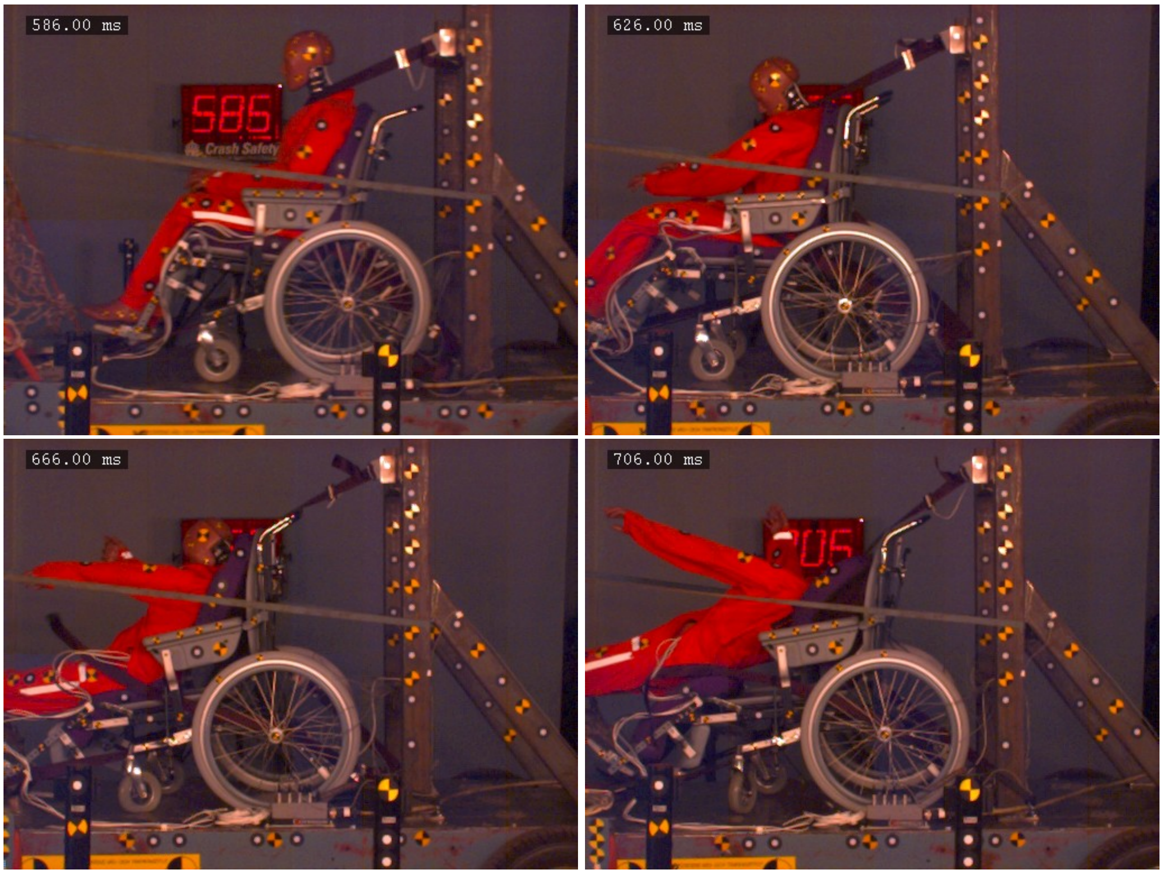

Figure 2: Images 20,40,60,80 from Image Systems new “Wheelchair” test sequence.

Task 2.2: Continuous Constraints and Euclidean Structure

Lund (1mm) have worked on the problem of obtaining Euclidean structure and camera parameters from weak assumptions on the camera geometry [3]. The latest result shows that it is possible to auto-calibrate the camera and obtain Euclidean reconstruction from one constant (but unknown) intrinsic parameter. These auto-calibration methods are of great importance for the car-crash application, where the cameras are not pre-calibrated and thus need to be calibrated automatically.

Task 2.3: On-line Calibration and 3D Motion from Image Streams

This task is devoted to work directly associated with the Image Systems demonstrator: rigid and non-rigid motion estimation for points on two crashing cars during a safety test.

Image Systems (6.4mm): Image Systems have produced an additional test data set, consisting of images of a crash test with a dummy in a wheel chair mounted on a car chassis. The test was filmed with three non-synchronous cameras at a frequency of 500 Hz. This data set differs from the previous ones primarily in that there is a larger number of points with known 3D coordinates. The data has been distributed to the academic partners in PNG format on a CD-ROM. Figure 2 shows a few images from one camera of this sequence.

Discussions within Image Systems and with the academic partners have led to a refined plan for the CUMULI demonstrator. The demonstrator will be based on Image System’s existing TrackEye software system. The main demonstration will use the wheel chair data set. Features (primarily points) will be identified manually and tracked in each camera view, using continuity constraints. Images from moving prism cameras will be stabilized using fixed points in the lab.

Initial values of the camera parameters will be obtained using a variety of calibration algorithms, such as Grenoble's DLT-like methods [37]. The important goal is that they should use as little a priori information as possible, although it is acceptable to require known 3D coordinates for a small number of points. A final 3D reconstruction will then be obtained by a bundle adjustment. The demonstrator will also demonstrate calibrated pose and relative orientation algorithms using the same data; since (unfortunately) no independent calibration of the cameras has been performed, this demonstration will either use the internal parameters determined in the first demonstration or nominal values.

Implementation of the CUMULI calibration algorithms in C++ and integration into Image System's products is in progress. Image Systems are evaluating the performance of these algorithms on data typical for their applications. The 4-point DLT-like algorithms will be used in existing products to provide initial estimates for a non-linear, iterative optimization, and thus considerably simplify their use. This will be delivered to customers in early 2000. The Grenoble or Lund quasi-linear calibrated pose algorithms are also strong candidates for implementation in commercial products during the next year.

The Lund bundle adjustment algorithm will also be incorporated into Image System's existing products within the near future. Currently, these products (aimed at the flight test market) contain functionality for calibrating the cameras from control points in the scene, and then performing 3D reconstruction by an intersection algorithm. A full bundle adjustment will provide increased accuracy for these customers. The ability to use lines and conics in addition to points is valuable in these applications, as well as in the car crash application.

Image Systems have also worked on improving their tracking algorithms for image sequences similar to the car crash test data set, and have started implementing CUMULI algorithms for track continuity. An increased interest in 3D calculations from the crash test customers, as well as the practical work with obtaining the test data sets, has led to an improved appreciation of the practical demands of the industry and the constraints of the test environment.

Grenoble (0.5mm): For camera initialization in the WP2 (and also WP3,1) demonstrators, Grenoble have delivered quasi-linear coplanar 4-point and 4- and 5-point DLT-like algorithms for camera pose and partial calibration, and also several nonlinear 3 point and quasi-linear 4 point methods for pose only. Nonlinear refinement routines for these problems have also been developed. Image Systems are currently selecting, implementing and testing the best few of these methods. Grenoble are also collaborating with Image Systems on the tracking algorithms.

Lund (0.5mm) have been in close contact with Image Systems, discussing which types of auto-calibration methods will be implemented in the final demonstrator. Preliminary investigations and evaluations of different auto-calibration methods have been made on the wheel-chair sequence. Especially, the case of known intrinsic parameters except for different focal lengths and unknown y-component of the principal point seems promising. Further actions will be made to test out this method and transfer the algorithms to Image Systems.

Sophia-Antipolis (3mm) previously developed algorithms for segmenting scenes into planes under various types of rigid displacements. These have now been tested on the data sets supplied by Image Systems. The formalism developed in CUMULI will also be used in the new IST Project 12224 CarSense, on vision applications for the automotive environment at low driving speeds. Sophia has also made available an existing point matching module as an additional algorithm for the final demonstrator.

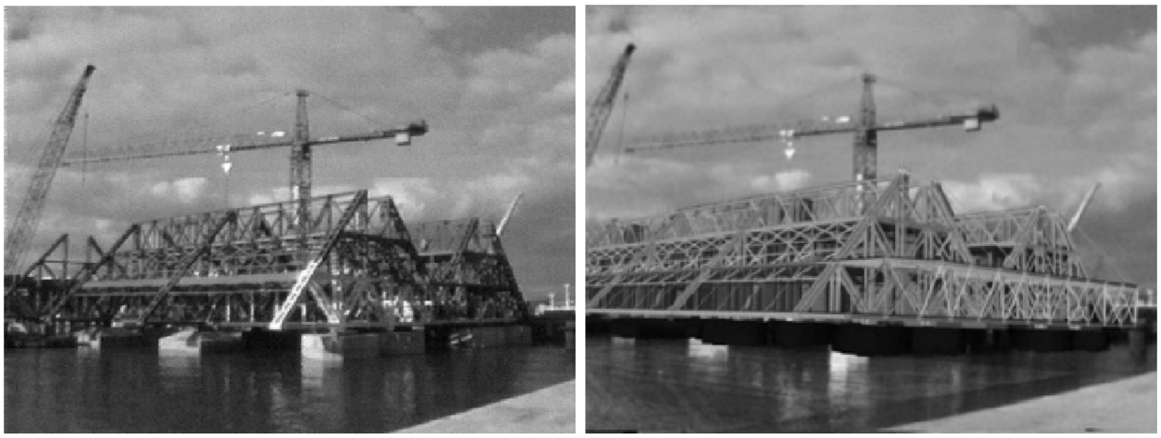

Figure 3: Augmentation of a building image based on IGD's calibration tool.

WP 3: Algebraic Symbolic Reasoning

IGD (7.5mm) has continued work on its Augmented Reality (AR) tool. The goal is to be able to flexibly augment any number of images with as little external 3D and camera information as possible. For this reason, IGD is enriching the tool with many different kinds of computer vision algorithms, rather than focusing on a single sub-problem.

For single images, several calibration and pose estimation algorithms have been implemented, including the Direct Linear Transformation, Weng's method [1], Grenoble's 4-point DLT method [37], and the 4 point homography based method. Depending on the number and geometry of the known 3D points, the user can select the most appropriate algorithm to augment the image.

When more views are available, several additional calibration and optimization algorithms can be used. If no 3D information about the scene is available, augmentation can use structure from motion methods, setting the virtual object interactively in the reconstructed scene. For stereo image pairs, a corner detector followed by a matching routine and robust estimation of the epipolar geometry allows semi-automatic to automatic structure from motion.

Finally, from video sequences, feature tracking based on correlation techniques allows a complete initial camera path to be reconstructed. To improve the results and reduce any jitter, a robust optimization technique (L-estimator) can then be applied over the whole sequence.

Following this approach of algorithmic diversity, IGD have also proposed a new "calibration propagation" algorithm that is in-between calibrated structure from motion and auto-calibration. Starting from a calibrated image, the calibration is propagated to other, uncalibrated ones without using 3D knowledge. Results and possible industrial applications are presented in [36]. Figure 3 shows an building image augmented using the IGD calibration tool.

The next problem of AR after defining the camera parameters is the handling of occlusion between real and virtual objects. Here again computer vision algorithms are very useful. Simple models of the scene are constructed using several views. These are used as occlusion masks of the real scene, to allow correct rendering of the augmented image. The reconstruction is done by selecting polygons in different views. The covariance of the reconstructed points is computed and visualized, to give the user an estimate of the quality of the reconstruction.

Finally, the CUMULI AR-Tool also contains imaging functionalities like mosaic building and compensation for interlacing.

INRIA Sophia (8mm) has improved its algebraic geometric reasoning system for computer vision.

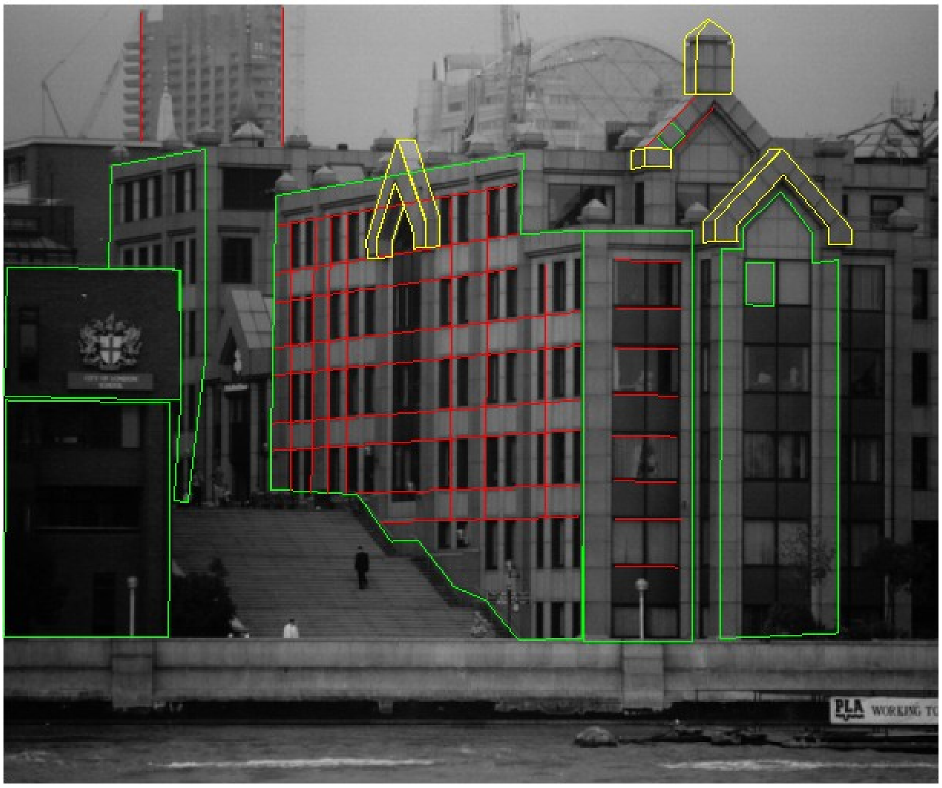

Figure 4: Constrained parametrization of a building frontage from the IGD 'London' test data, using the Sophia software.

This provides automatic minimal parameterization of 3D constrained environments given as sets of primitives (points, lines and planes). The system is now more reliable, faster, and consumes less memory. This has been achieved by improving the algorithm that computes the constructive order of the geometric set of primitives, and by using sparse techniques in the reduction algorithm [12]. Figure 4 shows a parametrization calculated using this software.

INRIA Sophia has also focused on the camera + map calibration problem. Their stratified approach makes use of some of the constraints (parallelism and orthogonality in vertical planes) between a map and an image at all the stages of calibration. This allows the recovery of a calibration matrix for the image that depends on only 2 parameters. The method has been tested on CUMULI data with very satisfactory results. Other Euclidean constraints can then be used to recover the last two parameters through numerical optimization. These tools will be transfered as executable binaries. Work on improving the current implementation is continuing.

Lund (1mm): Aiming at integrating geometric knowledge into computer vision methods, Lund has concentrated on numerical approaches to geometric reasoning under uncertainty. This work relies heavily on basic tools from WP1 and WP2, optimized for the particular demands on model building techniques made by VR and AR. Specifically, the focus is on numerically enforcing affine and Euclidean constraints on the reconstruction. For instance, as shown in [35], the formalism allows prior information to be handled efficiently, *e.g.* the coplanarity of certain points, or the parallelism of a given line to another line or a plane such as the ground plane. The methods allow any number of points to be handled simultaneously in any number of images. Another advantage, making it possible to cope with occlusions, is that all points need not be detected in all images. In [35], it is also shown that the geometry of the scene can be recovered more accurately by taking information from maps of the scene into account. Complications are caused by the fact that the images often deal with scenes that are nearly flat, and with all viewpoints from the same vicinity.

Recently, Lund has also prepared a set of commented algorithms and implementations that are easy for IGD to handle for integration into the demonstrator.

Grenoble (1mm) has further demonstrated the success of applying Clifford algebra and term-rewriting to automated geometric reasoning and computer vision. Software packages in Objective CAML and Maple V implementing the proposed method and techniques are capable of proving a number of geometric theorems, as well as some geometric properties studied in computer vision in both 2D and 3D [14, 15]. Research carried out in this task also leads to an inductive approach for automated proving of Clifford algebraic identities [43].

WP 4: Management, Deliverables and Dissemination

The management of the project is going smoothly. The collaboration is working well and all partners are pulling their weight in the project.

Test Data

As noted above under WP1.4, despite significant efforts by IMETRIC, the industrial pose-from-features test data has again been delayed and may not now be available by the end of the project. We stress that the problem of pose and relative orientation from non-point features is of strong interest for IMETRIC and a number of their industrial clients. But unfortunately, most of the applications are commercially highly sensitive and it has proven very difficult to get real industrial test data cleared for release. As noted under WP1.4, the IMETRIC demonstrator will not be delayed, but the orientation from features component will now be shown only as an industrial motivation and proof-of-concept MATLAB implementations. These methods will however be implemented by IMETRIC after the end of CUMULI.

As described under WP2.3, Image Systems have recently (Oct'99) delivered an additional set of test data involving a dummy in a crashing wheel chair. The academic partners are currently studying this data.

Articles and Reports

European Community support for CUMULI has aided and been acknowledged in (at least) the following 36 publications over the last 12 months [24, 12, 20, 17, 18, 6, 4, 22, 16, 10, 5, 3, 9, 8, 19, 21, 30, 7, 28, 23, 11, 36, 27, 25, 26, 14, 15, 43, 38, 37, 32, 29, 33, 39, 1, 34]. A selection of some of these papers is included as an appendix.

As agreed at the last review, we will not produce a dedicated paper for the deliverable 'Overview article: Discrete and continuous multi-image Euclidean reconstruction' (1.2/30(i) + 2.2/30(i)). Our experience with the previous overview article (on projective reconstruction — deliverable 1.1/24(i) + 2.1/24(i)) showed that the result would probably not be very useful within the project, and would be too disparate to publish owing to the wide range of work done in CUMULI! However a number of individual papers have been published on this theme under CUMULI, including [30, 21, 22] (which are collaborations between Lund and Grenoble), and [27, 26, 34, 33, 3, 5, 10, 20].

For deliverable 3.2/36 'Article: automated geometric reasoning for database coherence and uncertainty' we have nominated the following papers: [35, 13, 12, 15].

This document itself is the month 36 Periodic Progress Report.

Libraries and Demonstrators

As agreed, the WP3 demonstrator 3.2/30 will now be shown at the end of the project, together with the WP1 and WP2 demonstrators.

Most of the remaining deliverables are software libraries. These are internal deliverables, designed to aid the technology transfer to the industrial partners by providing example implementations of the algorithms developed in CUMULI. As discussed at the last review, they have been revised and reorganized to reflect the actual needs of the industrial partners. In particular, there is an increased emphasis on camera initialization and pose methods, which have proved very useful to all three industrial partners. Since the last review, Grenoble have delivered an updated version of their pose and calibration library, including improved versions of the 4 and 5 point DLT pose+calibration routines, new routines for nonlinear 3-point and quasi-linear 4-point camera pose (several routines each), and new routines for pose from various combinations of features. Lund have prepared revised versions of their pose, relative orientation and structure from motion routines, and are preparing routines for autocalibration. Sophia have delivered C/C++ code for feature correspondence, and for recognizing special types of camera motion, and for segmenting scenes into planes under them.

Industrial Dissemination

The CUMULI AR results have been shown in many demonstrations to industrial people at IGD. The potential of this technology lead to the creation of a new industrial consortium of major companies based in Germany. The consortium applied for German government funding and was granted the project ARVIKA (http://www.arvika.de).

In this framework, IGD has organized an industrial AR forum to present and discuss industrial solutions, mainly in the area of prototyping, assembly, construction, maintenance and service. The participants are companies like Volkswagen, Daimler-Chrysler, Airbus, DASA, Ford, Siemens, BMW and several machine constructors like Ex-Cell-O.

In parallel, one car manufacturer has already planned to extend its Virtual Reality system with Augmented Reality functionality, and has started a pilot project with IGD in which several CUMULI results like camera calibration and pose estimation are applied.

The formalism developed by Sophia during CUMULI will also be used in the new IST Project 12224 CarSense, on vision applications for the automotive environment at low driving speeds.

The CUMULI project was presented at an industrial session at the annual Swedish workshop SSAB, arranged by the Swedish Association for Automated Image Analysis, in Gothenburg, March 1999. Presentations were made by Image Systems (Anders Kalldahl) and Lund (Gunnar Sparr), and also also included a video demonstration from IMETRIC. The audience was about 100 people, about 30 of whom were industrial.

The CUMULI consortium still plans to hold an industrially-oriented seminar in Paris just after the end of the project. The provisional date is 30 March 2000.

Workshops and Exchanges

This year IGD again co-organized the (second) International Workshop on AR (IWAR'99). This event was in the USA, but the workshop will be held in Europe next year, and promoted to a Symposium (ISAR 2000, International Symposium on Augmented Reality, October 5-6, 2000, Munich Germany).

Grenoble were the principal organizers of a very successful workshop *Vision Algorithms: Theory and Practice* [41] at the International Conference on Computer Vision, 21-22 September 1999, in Corfu, Greece. This was the largest workshop at ICCV with around 100 attendees. Although not explicitly linked to CUMULI, it embodied many CUMULI themes.

There have also been a number of internal exchanges within the consortium. Fredrik Kahl from Lund visited INRIA Grenoble for 3 months at the end of 1998, and during this period Anders Heyden from Lund also visited Grenoble. We are continuing this collaboration, which has so far produced several joint papers [21, 30, 22]. Fredrik Kahl's Lund licentiate thesis [23] consists mainly of his papers published under CUMULI.

Man Months

WORK TASK	MONTH	Gren	Soph	Lund	I.S.	Imetric	IGD	TOT.
1 Discrete Images		17/17	13/12	17.5/17	2/2	11.7/17		61.2/65
1.1 Multi-camera Geometry	0–24	5/4	5/5	5/4	1.1/1	1/1		17.1/15
1.2 Euclidean Structure	0–36	5/4	7/5	4/4	0.3/0	0.6/1		16.9/14
1.3 Features & Uncertainty	0–36	3/4		5/4	0.6/1	0.6/1		9.2/10
1.4 Transfer Imetric	0–36	4/5	1/2	3.5/5		9.5/14		18/26
2 Image Streams		14.5/15	13/11	12/14	11.9/20			51.4/60
2.1 Incremental Geometry	0–24	5/5	5/5	5.5/5	1/1			16.5/16
2.2 Continuous Constraints	6–30	4/5	5/5	4/5	1/1			14/16
2.3 Transfer I.V.	0–36	5.5/5	3/1	2.5/4	9.9/18			20.9/28
3 Reasoning		4/3	24/24	5/7			17.5/22	50.5/56
3.1 Geometric Reasoner	0–18	4/3	12/12	4/4			7/7	27/26
3.2 Visual Applications	12–36		12/12	1/3			10.5/15	23.5/30
TOTAL RESEARCH		35.5/35	50/47	34.5/38	13.9/22	11.7/17	17.5/22	163.1/181
4 Management	0–36	11/13	2.5/3	2.5/3	2/3	0.8/1	1.7/2	20.5/25
TOTAL PROJECT		46.5/48	52.5/50	37/41	15.9/25	12.5/18	19.2/24	183.6/206

Table 1: Man-months consumed in months 0–36 *vs.* total budgeted man-months for entire task, per work task and per site.

The total man-months consumed to date (months 0–36) are listed in table 1, by work task and by site. Overall, there has been a substantial transfer of technology in which all partners have played an active part, through technical discussion at CUMULI meetings, through experiments on the test data provided by the industrial partners, and through the CUMULI deliverables.

References

[1] J. Weng, P. Cohen, and M. Herniou. Camera calibration with distortion models and accuracy evaluation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-14(10):965–980, 1992.

Cumuli References

[1] M.-A. Ameller, B. Triggs, and L. Quan. Camera pose revisited — new linear algorithms. Submitted to ECCV'00, 1999.

[2] K. Åström and A. Heyden. Continuous time matching constraints for image streams. *Int. Journal of Computer Vision*, 28(1):85–96, 1998.

[3] K. Åström and A. Heyden. Flexible calibration: Minimal cases for auto-calibration. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[4] K. Åström and A. Heyden. Stochastic analysis of scale space smoothing. *Advances in Applied Probability*, 30(1), 1999.

[5] K. Åström, A. Heyden, F. Kahl, and M. Oskarsson. Structure and motion from lines under affine projections. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[6] K. Åström and F. Kahl. Motion estimation in image sequences using the deformation of apparent contours. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(2):114–127, 1999.

[7] K. Åström and M. Oskarsson. Solutions and ambiguities of the structure and motion problem for 1d retinal vision. In *Scandinavian Conf. on Image Analysis, Greenland*, 1999.

[8] R. Berthilsson. Densities and maximum likelihood estimation of matching constraints. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[9] R. Berthilsson. Finding correspondences of patches by means of affine transformations. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[10] R. Berthilsson, K. Åström, and A. Heyden. Reconstruction of curves in R^3, using factorization and bundle adjustment. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[11] Rikard Berthilsson. *Extensions and Applications of Affine Shape*. Phd thesis, Centre for Mathematical Sciences, Lund University, December 1999.

[12] D. Bondyfalat, B. Mourrain, and T. Papadopoulo. An application of automatic theorem proving in computer vision. In X.-S. Gao, D. Wang, and L. Yang, editors, *Automated Deduction in Geometry*, number 1669 in LNAI. Springer-Verlag, Berlin Heidelberg, 1999.

[13] Didier Bondyfalat and Sylvain Bougnoux. Imposing Euclidean constraints during self-calibration processes. In R. Koch and L. Van Gool, editors, *Proceedings of SMILE Workshop on Structure from Multiple Images*, Lecture Notes in Computer Science. Springer Verlag, Lecture Notes in Computer Science, 1998.

[14] Thierry Boy de la Tour, Stéphane Fèvre, and Dongming Wang. Clifford term rewriting for geometric reasoning in 3d. In X.-S. Gao, D. Wang, and L. Yang, editors, *Automated Deduction in Geometry — Proceedings of ADG'98*, LNAI 1669, pages 130–155. Springer-Verlag, 1998.

[15] Stéphane Fèvre and Dongming Wang. Combining clifford algebraic computing and term-rewriting for geometric theorem proving. *Fundamenta Informaticae*, 39(1–2):85–104, 1999.

[16] A. Heyden. Tensorial properties of multilinear constraints. *Mathematical Methods in Applied Sciences*, 1999. Submitted.

[17] A. Heyden and R. Berthilsson. Recognition of planar objects using the density of affine shape. *Computer Vision and Image Understanding*, 76(2):135–145, 1999.

[18] A. Heyden, R. Berthilsson, and G. Sparr. An iterative factorization method for projective structure and motion from image sequences. *Image and Vision Computing*, 17(13):981–991, 1999.

[19] F. Kahl. Critical motions and ambiguous euclidean reconstructions in auto-calibration. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.

[20] F. Kahl and A. Heyden. Affine structure and motion from points, lines and conics. *Int. Journal of Computer Vision*, 33(3):163–180, 1999.

[21] F. Kahl and B. Triggs. Critical motions in euclidean structure from motion. In *Conf. Computer Vision and Pattern Recognition, Fort Collins, USA*, 1999.

[22] F. Kahl, B. Triggs, and K. Åström. Critical motions for auto-calibration when some intrinsic parameters can vary. *Journal of Mathematical Imaging and Vision*, 1999. Submitted.

[23] Fredrik Kahl. Geometric stratification in computational vision. Licentiate thesis, Centre for Mathematical Sciences, Lund University, April 1999.

[24] G. Klinker, D. Stricker, and D. Reiners. The use of reality models in augmented reality applications. In R. Koch and L. Van Gool, editors, *Proceedings of SMILE Workshop on Structure from Multiple Images*, Lecture Notes in Computer Science. Springer Verlag, Lecture Notes in Computer Science, 1998.

[25] Diane Lingrand. *Analyse Adaptative du Mouvement dans des Séquences Monoculaires non Calibrées*. PhD thesis, Université de Nice - Sophia Antipolis, July 1999.

[26] Diane Lingrand. Perspective projection's approximations: combined model and homographic singularities. Technical Report RR-3682, INRIA Sophia Antipolis, April 1999.

[27] Diane Lingrand. Etude des singularités des homographies dans des séquences monoculaires non calibrées. In *12^{e} Congrès Reconnaissance des Formes et Intelligence Artificielle*, volume 1, February 2000.

[28] M. Oskarsson and K. Åström. Automatic geometric reasoning in structure and motion estimation. In *Scandinavian Conf. on Image Analysis, Greenland*, 1999.

[29] L. Quan. Inherent two-way ambiguity in 2d projective reconstruction from three uncalibrated 1d images. In *Int. Conf. Computer Vision*, pages 344–349, Corfu, Greece, 1999.

[30] L. Quan, A. Heyden, and F. Kahl. Minimal projective reconstruction with missing data. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 210–216. IEEE Computer Society Press, June 1999.

[31] L. Quan and Z. D. Lan. Linear $\geq N$ point pose determination. In *Int. Conf. Computer Vision*, 1998. To appear.

[32] L. Quan and Z.D. Lan. Linear n-point camera pose determination. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(8):774–780, August 1999.

[33] L. Quan and M. Lhuillier. Structure from motion from three affine views. Submitted to ECCV'00, 1999.

[34] L. Quan and B. Triggs. A unification of autocalibration methods. In *Asian Conf. Computer Vision*, January 2000.

[35] G. Sparr. Euclidean and affine structure/motion for uncalibrated cameras from affine shape and subsidiary information. In R. Koch and L. Van Gool, editors, *Proceedings of SMILE Workshop on Structure from Multiple Images*, Lecture Notes in Computer Science. Springer Verlag, Lecture Notes in Computer Science, 1998.

[36] D. Stricker and N. Navab. Calibration propagation for image augmentation. In *Proceedings of Second International Workshop on Augmented Reality*, San Francisco, 1999.

[37] B. Triggs. Camera pose and calibration from 4 or 5 known 3d points. In *Int. Conf. Computer Vision*, pages 278–284, Kerkyra, Greece, September 1999.

[38] B. Triggs. Differential matching constraints. In *Int. Conf. Computer Vision*, pages 370–376, Kerkyra, Greece, September 1999. IEEE press.

[39] B. Triggs. Plane + parallax, tensors and factorization. Submitted to ECCV'00, 1999.

[40] B. Triggs, A. Fitzgibbon, P. McLauchlan, and R. Hartley. Bundle adjustment for structure from motion. In B. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice*. Springer-Verlag, 2000.

[41] B. Triggs, R. Szeliski, and A. Zisserman, editors. *Vision Algorithms : Theory and Practice*, Workshop held during ICCV'99 in Kerkyra (Corfu), Greece, September 1999. INRIA. Final proceedings will appear in Springer LNCS in early 2000.

[42] Thierry Viéville and Olivier Faugeras. The first order expansion of motion equations in the uncalibrated case. *Computer Vision, Graphics and Image Processing*, 64(1):128–146, July 1996.

[43] Dongming Wang. Geometric algebra and reasoning. In *Geometric Algebra: A Geometric Approach to Computer Vision, Quantum and Neural Computing, Robotics and Engineering*. Birkhäuser, Boston, 1999. To appear.

Appendix: Selected Papers

As an example of recent scientific work on CUMULI, we have included a selection of recent papers in this appendix. The following papers have been selected:

- [1] M.-A. Ameller, B. Triggs, and L. Quan. Camera pose revisited — new linear algorithms. Submitted to ECCV'00, 1999.
- [5] K. Åström, A. Heyden, F. Kahl, and M. Oskarsson. Structure and motion from lines under affine projections. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.
- [10] R. Berthilsson, K. Åström, and A. Heyden. Reconstruction of curves in r^3, using factorization and bundle adjustment. In *Int. Conf. Computer Vision, Kerkyra, Greece*, 1999.
- [12] D. Bondyfalat, B. Mourrain, and T. Papadopoulo. An application of automatic theorem proving in computer vision. In X.-S. Gao, D. Wang, and L. Yang, editors, *Automated Deduction in Geometry*, number 1669 in LNAI. Springer-Verlag, Berlin Heidelberg, 1999.
- [22] F. Kahl, B. Triggs, and K. Åström. Critical motions for auto-calibration when some intrinsic parameters can vary. *Journal of Mathematical Imaging and Vision*, 1999. Submitted.
- [24] G. Klinker, D. Stricker, and D. Reiners. The use of reality models in augmented reality applications. In R. Koch and L. Van Gool, editors, *Proceedings of SMILE Workshop on Structure from Multiple Images*, Lecture Notes in Computer Science. Springer Verlag, Lecture Notes in Computer Science, 1998.
- [27] Diane Lingrand. Etude des singularités des homographies dans des séquences monoculaires non calibrées. In *12^e Congrès Reconnaissance des Formes et Intelligence Artificielle*, volume 1, February 2000.
- [36] D. Stricker and N. Navab. Calibration propagation for image augmentation. In *Proceedings of Second International Workshop on Augmented Reality*, San Francisco, 1999.
- [37] B. Triggs. Camera pose and calibration from 4 or 5 known 3d points. In *Int. Conf. Computer Vision*, pages 278–284, Kerkyra, Greece, September 1999.