

Hypergraph-Oriented 3D Reconstruction, Interpretation And Analysis Of Air Flows

Denis Klimentjew, Nils Erik Flick, Till Bosselmann, and Jianwei Zhang

Dept. of Informatics, University of Hamburg, Vogt-Kölln-Straße 30, 22527 Hamburg, Germany

{klimentjew, flick, bosselm, zhang}@informatik.uni-hamburg.de

Abstract

The paper introduces a new air flow analysis approach based on Particle Tracking Velocimetry (PTV). One of the special features of the proposed method is that after the tracer particles are detected, matching and tracking are jointly conducted. To this end, we introduce an interpretation module based on a directed hypergraph for 3D curve reconstruction. At first the 2D inter-frame locations are localised and used for the extraction and calculation of 3D keypoints. Through 3D keypoints which are evaluated by the hypergraph together with the time information in several steps, reverse curve matching for path selection can be reconstructed and the resulting trajectories visualised.

In contrast to the preceding works our approach tries to describe the measuring data by 3D trajectories directly instead of estimating 2D trajectories first and then matching afterwards. A higher precision can be achieved also with complicated trajectories. A certain independence of the reflections and lighting conditions is reached by the interpretation. Moreover, the path of particles can also be reconstructed with the minimum number of 3D keypoints under consideration of the path energy minimization.

1 Introduction

According to a Dow Jones Newswires prediction, the number of aircraft built by Boeing and Airbus will double by 2023/2026. The number of passenger aircraft worldwide will be approximately 35,000. Passenger comfort plays an increasingly important role for the marketing of modern aircraft. Passenger comfort is strongly influenced by the structure and intensity of cabin air flow. Most aircraft are different from each other, because the airlines order the cabin layout for a certain purpose and the aircraft are manufactured individually. In this case, each new cabin has to be checked for compatibility with the air conditioning requirements by the manufacturers for the custom layout. More than that, in most cases adaptations and changes of the air conditioning system are necessary.

Therefore, a computer-aided system which evaluates the measurements and visualises the air-streams inside the cabin would be enormously helpful. The resulting system could assist the developers not only in the prototypes, but also in the assembly of the aircrafts to be modified.

2 Related Work

Over the course of years, several measuring technologies have been discovered. One of the most well-known ones is based on the use of anemometers which are placed at different places in the aircraft. Further information about the measurement technology can be found in [1], measurement and assessment can be found does in [2]. For any realistic air-flow model, several anemometers should be placed and calibrated. Moreover, only one dimension of the air-flow is measured and afterwards all data must be merged. As a result the procedure takes a lot of staff, time, and is cost-intensive.

Furthermore, there are three general categories of methods which measure the velocity of air-flow with the help of tracer particles. Particle Image Velocimetry (PIV) uses a setup consisting of cameras and a laser scanner [3]. As particles pass through the laser illuminated plane, two images are taken within a short period of time. On these images, a cross correlation of image regions is performed, in order to obtain the direction and velocity of the flow. This approach is limited by the area which can be illuminated by the laser, as well as the acquisition and transfer rate of the cameras. The main difference between PIV and Hot-wire anemometry respectively Laser Doppler velocimetry techniques is that PIV produces two dimensional vector fields [4].

The principle idea behind Particle Streak Tracking (PST) is to control the scene illumination through a shutter system [5]. It generates a train of pulses, during which the camera is exposed to the reflections of the tracer particles in the scene. As in PTV, only a single plane is illuminated, and the pulses are synchronised with the exposure time of the camera, so that exactly one pair of pulses fall into each captured frame. In each pair, one pulse is longer than the other. This separates the image of each particle trajectory into two components with different length. Knowing this, it is possible to determine the direction of each particle. For an application of PST for 3D velocity measurement of flows on a single plane see [6]. The PST methods support only a low number of particles. Therefore the resulting information density is rather low. This is a serious handicap for the measurement of turbulent air-flow.

Particle Tracking Velocimetry (PTV) methods identify single particles, and track them from frame to frame [7][8]. This makes PTV more suitable for experiments with a low number of particles. The lower number of particles allow depth reconstruction for each particle, provided correspondence can be established.

The disadvantages of the methods that deliver a two dimensional vector field are significant. One of the velocity components (depth direction) is disregarded, some particles and turbulences will not be sufficiently recognised or be totally absent. In other words the resulting visualisation is not complete and accurate. In addition, most resulting applications need elaborate setups and are time-consuming.

Many of the above methods need many hours or even days for the evaluation of measurement data and the visualisation of air-flows.

In this paper, we propose a new approach for the evaluation of measurement data and the visualisation of air-flows inside the aircraft cabin based on PTV. The setup consists of only two synchronised cameras with long exposure time, similar to the setup described by Bosbach et al. [9] as illustrated in fig. 1. The helium filled soap-bubbles that are injected through the air-intakes are detected via the stereo camera system. Moreover, the results for the soap bubbles are traced. First, the 2D inter-frame locations of the bubbles are localised and then used for the calculation and extraction of 3D keypoints. The biggest difference from the other methods is to describe the measuring data directly by 3D trajectories instead of estimating the first 2D trajectories and then matching them to the 3D trajectories. A higher precision can be achieved even with

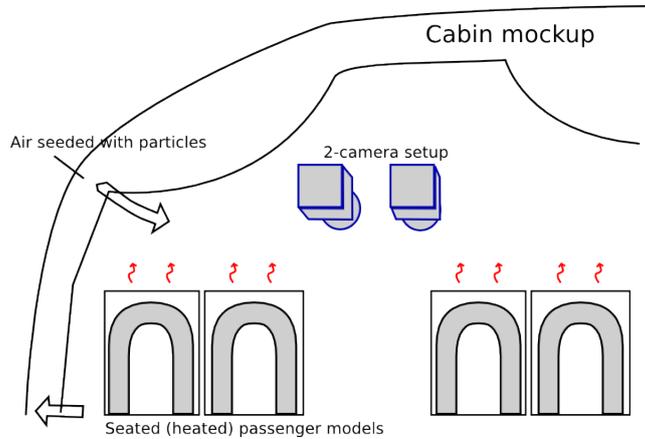


Figure 1: Cross section of a full scale aircraft cabin mock-up, as used in our experiments

complicated trajectories. Our method uses reverse curve matching for path selection, so a directed hypergraph is used to evaluate a data-fit term on a curve obtained from a few points obtained from an earlier step.

The hypergraph connects several 3D keypoints through adjacent frames and contains time information. In the next step the subpaths are combined. The accruing curves are enhanced with a B-spline in consideration of the energy, so-called active contour models. An explanation of the theory can be found in sections 4.2 and 4.3.

Independence of the lighting conditions and reflections is reached by the interpretation to a certain degree. Moreover, the way of particles can also be reconstructed with the minimum number of 3D keypoints under consideration of the path energy minimization. We do not assume to have perfect control of the experimentation environment and our method can deal with a variable number of helium filled soap bubbles. Furthermore, the investigated volume is considerably larger than in most other PTV applications. Based on the theoretical foundations, we implemented the prototype and tested it in the full scale aircraft cabin mock-up. Our application calculates and visualises the results online without lengthy delays.

The rest of this paper is structured as follows. In the next section we present the theoretical foundations of our method. In it we describe the applied model, different pre-processing steps, particle segmentation, as well as the curve fitting and interpretation. Then, in section 5, we describe the experimental setup. The implementation and experimental results are described and discussed in detail. Finally, we present our conclusion in section 6.

3 Model Description And Abstraction

PTV is about obtaining information on fluid flows by means of digitized images. Consequently, one must employ image processing techniques.

Starting with the physical model, PTV methods are based on the assumption that trajectories of small particles closely follow pathlines (eq. 1) of the fluid flow being observed, which allows the reconstruction of aspects of the time dependent velocity field.

In fact, the measurements amount to a 1-dimensional sampling of the time-dependent velocity field.

$$\frac{d}{dt}C(t) = \vec{v}(C(t), t) \quad (1)$$

Our method can deal with long exposure times, provided the intra-frame portions of trajectories are sufficiently well described by low-order terms.

We do not make any further assumptions about the flow, but if the smaller scales of turbulence [10] cannot be resolved, that is a question of the resolution attainable by the observation of tracer particles; no amount of image processing can extract information which was lost in the image formation process.

Motions which are not distinguishable should be treated with a kind of smoothness assumption (in the spirit of multi-scale analysis, one likes to separate the main flow patterns from details); one must introduce the least spurious detail possible, as per the maxim known as *pluralitas non est ponenda sine necessitate* (“plurality should not be posited without necessity”) [11].

Now a video camera is a system composed of a lens system and a sensor array. The role of the lens is to obtain a two-dimensional geometrical image of a three-dimensional scene, and its intended effect is accurately modeled by projective geometry. The sensor array is a device which measures illuminance integrated over a certain exposure time, and at the same time performs a digitization and therefore allows the continuum to be processed.

In our simplest image formation model, we consider an image to be a sampled version of a continuous function. There are theoretical reasons for establishing images as elements of a Schwartz space $\mathcal{S}(\mathbb{R}^n)$ [12]. Notably, one can define convolutions and other LTI operators. Digitized images, by virtue of the sampling process, are associated with a certain resolution. The resolution of an image is arbitrary but limits the inner scale of the observations [13]. Noticing this leads immediately to scale space theory, viewing an image at all possible scales (i.e. from coarse to fine). Florack et al. [14] derive Gaussian scale space as the unique possibility from some general assumptions about low-level vision without *a priori* knowledge. In the PTV application, it is not necessary to build an entire scale-space representation because particles appear at a limited range of scale.

Each frame is composed of background, particle traces and noise. We assume a simplified model, wherein particles are transparent. In any case, the detection does not depend on observed intensity, except for the requirement of sufficient contrast; assuming that particles are opaque also do not impact the detection process.

This image formation model can thus be formulated by describing the frame as a sum of three signals: illuminance gathered from the particle traces, scene background, and pixel noise (as indicated in eq. 3).

It is advantageous to abstract from signal theory and think of an image as a continuous function or scalar field [13] and describe its properties in terms of isophotes (subject to scale-space considerations) rather than pixels.

The essential point is that particle traces appear at a certain scale to be ideal lines (1D piecewise submanifolds of the plane), thickened by a unimodal, radially symmetric PSF. This fact is not invalidated by the complex nature of the actual reflections (tracer particles are often highly non-Lambertian scatterers); higher frequencies are smoothed out at the correct scale.

The PSF’s support does not really matter, but the apparent diameter ρ of a particle can be defined by noting where the signal-to-noise ratio makes detection of a

well-contrasted particle impossible (e.g. by introducing suitably “local” features and observing error bounds in classification).

For practical purposes, ρ is about the particle’s projected diameter: the experimental setup should be designed such that it is of the order of at least a few pixels, to ensure sufficient spatial sampling.

The effect of the finite aperture is often described simply as a linear filter [15], disregarding depth. In our model this effect is subsumed, together with the finite extent of the image of a particle and its specular reflection, into a compound PSF, which does not need to be analyzed further because the only important properties for our purposes are that it be unimodal and radially symmetric.

The endeavour of accommodating long exposure times entails that the temporal discretization can no longer be usefully approximated by a Dirac comb; instead, one obtains an integration process over a finite, and non-neglectable time, which excludes, for instance, optical-flow based methods of motion estimation. Another factor diminishing the utility of consecutive frame differences is that with a motion in the direction of an optical ray, the particle would escape detection, which means a systematical error.

4 Theoretical Description Of The Originating Method

4.1 Preprocessing

The PTV system we developed is based on two synchronised cameras with long exposure time. We used Zhang’s calibration method [16] to obtain the camera parameters. As a result of the single camera calibration, we got its intrinsic parameters and the lens distortion. This was done separately for both cameras. After that, we used the semi-automatic calibration method to calibrate the stereo camera system. The transformation matrix between the two cameras’ image planes also needs to be determined. There are a wide array of algorithms for this problem, e.g. Bouguet’s [17]. To determine the rectification matrices based on intrinsic and extrinsic camera parameters, we implemented a procedure based on the work of Fusiello, Trucco and Verri [18].

For the background removal, as stated in the model description, an image I can be seen as the sum of several components:

$$F_{i,f}(\vec{x}) = \alpha \cdot BG_{i,f}(\vec{x}) + N_{i,f}(\vec{x}) + T_{i,f}(\vec{x}) \quad (2)$$

or simplified

$$F_i = T_i + BG_i + N_i \quad (3)$$

where F denotes the image, T_i the foreground region containing the traces, BG_i the background, and N_i the camera noise at frame number i .

In a static scene, BG_i is affected only by changes in illumination (and reflections). The main of BG_i ’s intensity can be removed by calculating a median image over several frames. This template is then subtracted from each input image. If an effort was made to keep illumination constant, the remaining BG_i intensity is low enough to not affect further segmentation steps significantly, so that we can assume $F_i = T_i + N_i$, which simplifies the following segmentation steps.

4.2 Particle Segmentation

Keypoint extraction may be too computationally expensive to apply it to entire images. In order to obtain regions of interest in which to search for keypoints, region based

segmentation methods are useful.

Structuring element with hysteresis: Our first approach for segmentation was to apply a morphological operator to test for four-connectivity on the pixel grid by eroding the binary image with a cross-shaped structuring element. Testing for four-connectivity seemed like a suitable, if unrefined, indicator for the presence of traces, as the pattern is unlikely to occur at random ($\frac{1}{32}$ if BG_i has been removed and N_i consists of white noise). The detection rate, however, is strongly dependent on the threshold used to obtain the binary image. The regions gained were then used as starting points in order to complete the segmentation result by hysteresis on the original image.

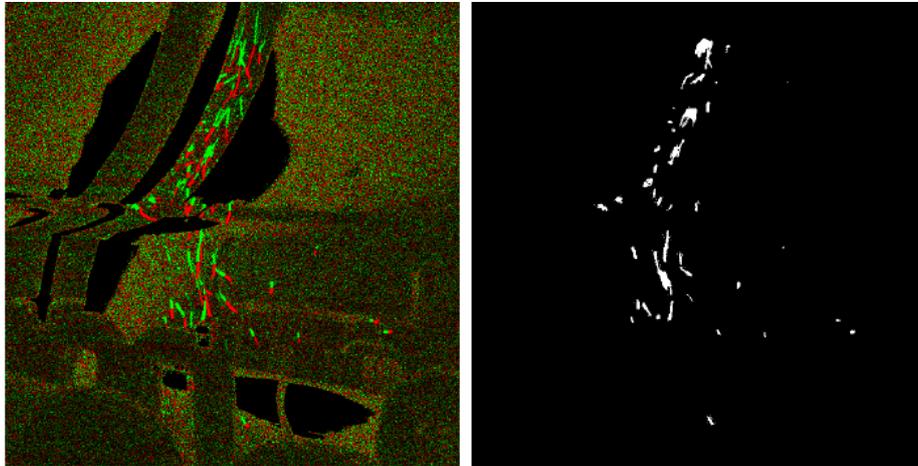
Gradient Magnitude: Particle traces, at an appropriate spatial scale, show a higher gradient magnitude than the surrounding area. For an image $F(x,y)$, it is defined as presented below (Eq. 4).

$$|\Delta F(x,y)| = \sqrt{\left(\frac{\delta F(x,y)}{\delta x}\right)^2 + \left(\frac{\delta F(x,y)}{\delta y}\right)^2} \quad (4)$$

This was exploited for segmentation by calculating the gradient magnitude at each pixel and filtering through a threshold Tr , which was set to a multiple n of the square root of the empirical variance of gradient intensities V : $Tr = n * \sqrt{V}$

We found that changing n does not change the detection yield significantly. While the algorithm is suitably quick, and gradient estimates can be reused for detection, it breaks down on images with low contrast, thus low overall variance, leading to too many spurious regions of interest.

MeanShift: If F is viewed as feature space, the MeanShift algorithm [19] can be used for segmentation. It operates by following the density gradient with a kernel in



(a) Original difference image

(b) Mask image

Figure 2: Performance of the Mean Shift based mask generator with kernel radii of 8. Not all traces are distinguishable as separate segments, which has no bearing on the keypoint extraction. Some traces are partially segmented, making a dilation or blurring advisable before the mask is used. The detection rate does not change significantly with different kernel values.

the feature space until a maximum has been reached. Pixels sharing the same gradient

maximum are assigned the value of that maximum. All other pixels are assigned the density estimate of the kernel centered at their position. If the kernel is chosen to be suitably large, and the noise is uniform and white distributed, the method is noise resistant.

Although by no means perfect, the MeanShift based segmentation showed the best performance of all the tested methods, see fig. 2. Due to its suitability to varying classes of input images. All methods can be sped up by using a suitable representation which allows operating on multiple scales, such as [13].

4.3 Curve Fitting And Interpretation

First, we will exploit known image structure to obtain a combinatorial representation as a starting point for further analysis. In most points, the projected curve can be approximated as in Eq. 5.

$$\begin{aligned} \Pi(t_a + t) &= h \circ \vec{\zeta}(t_a + t) = \\ &= \Pi'(s) = \vec{\Pi}(t_a) + \frac{\delta \Pi'}{\delta s}(0) \cdot s + \frac{1}{2} \frac{\delta^2 \zeta'}{\delta s^2}(0) \cdot s^2 + \mathcal{O}(s^3) \end{aligned} \quad (5)$$

where ζ is the space curve, h the de-homogenization, Π the projected curve and Π' its unit-speed version (parametrized by arc length, since t is unobservable in practice).

Many methods use templates or intensity maxima for detection and tracking. Going by intensity maxima is not an option for data with long exposure times, because they do not come with a well defined time stamp; while one could envisage using some formulation of templates (usually contour-based), it seemed more promising in terms of efficiency and also for the sake of parsimony, to continue thinking in terms of thickened lines as opposed to regions.

4.3.1 Keypoints

We introduce inter-frame locations of particles as a choice of 2D keypoints, which can be detected on each view individually; the 2D keypoints introduced above are then combined to 3D keypoints, all while respecting the epipolar constraint¹.

During our initial research, the idea of searching directly for 3D keypoints was rejected because searching, e.g. by masking regions of low contrast and then applying Newton's method to a response function in two dimensions, is less costly in two dimensions than it is in three and, more to the point, there is no essential advantage in doing so.

All available *image* information being essentially two dimensional because it stems from two 2D views, the fundamental ambiguity of epipolar geometry is not resolved by choosing another representation.

Subpixel measurements are introduced naturally in the context of modern image analysis, but we would like to point out that it is not directly responsible for increased precision in our setup; that can be a by-effect, but the main advantage of subpixel-methods is that they allow to deal more meaningfully with image information. The gain in precision, which is smaller than the apparent radius ρ of a tracer particle, is in practice smaller than calibration errors, which may be of the order of several pixels even in a carefully calibrated setup.

¹Actually, strictly enforcing the epipolar constraint at this point would mean no 3D keypoints at all, because of subpixel localization!

Inter-frame locations are a very good candidate for anchoring further steps because they are well localized spatially as well as temporally. Furthermore, they are detectable in the image series. At the same time, they can be said to have a standard appearance.

The suspicion that inter-frame locations can be detected with high accuracy was confirmed by a test we carried out to investigate whether different *regions* corresponding to background, endings and interiors of thickened lines could be classified with a max-margin linear classifier, and which turned out successful even within a non-adapted feature space, taken from another experiment – as long as contrast was sufficient.

There is more to the choice of such *2D keypoints* as essential significant feature:

- they are very robust against the choice of exposure time, because their aspect, except for inevitable overlaps, depends only on first order of arclength of the original trajectory, and not on curvature or other aspects of its shape;
- no matter how long the duration of the integration of *intra*-frame information, these keypoints are always equally well localized (localization of observations in time and space being a cornerstone of optical velocimetry).

For the software prototype, the choice fell on Harris' and Stephens' combined edge and corner detector [20] over other possibilities because of excellent localization and high specificity (empirical, synthetic and actual image series), and especially for its avoidance of multiple detections (which would not be good for the accuracy of measurements).

This detector is based on the gradient second-moment tensor, or structure tensor, M , whose components are image derivatives taken at an inner scale. It is defined as $hs = \det(M) - k \cdot \text{tr}(M)$ and is a smooth function of location (and scale).

When is calculated on each (background-subtracted) frame separately, then corroborated between neighboring frames, it indicates probable intra-frame locations of particles.

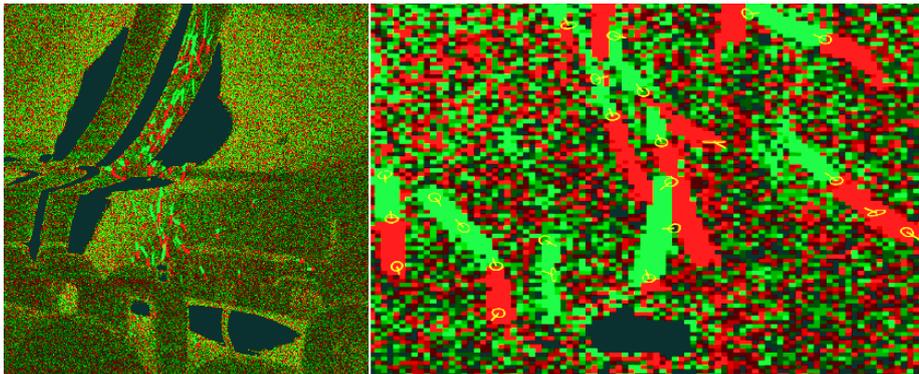


Figure 3: Example of correctly detected inter-frame locations, superposed onto a color-coded two-frame difference image.

The fact that it consistently showed better detection rates (lower false-positive, lower false-negative) than an isophote-curvature derived criterion and a detector based on critical points of a certain function calculated on the time-series remains somewhat surprising if one looks at the images.

Because it is a gradient-based detector, it good at capturing properties which do not change much with exposure time – which are also related to first-order properties, i.e. gradients – and we found that the choice of scales does not have appreciable influence on the localization of the inter-frame position, except in crowded areas. This is because any wandering of the spatial maxima is offset by the requirement that the feature be strong on both frames – distribution of the gradient is locally symmetrical at the inter-

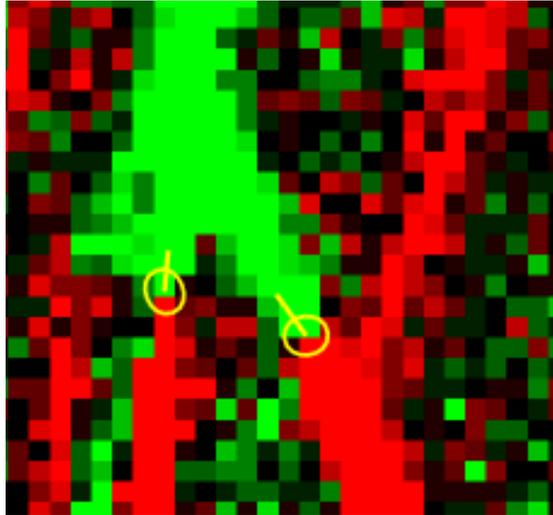


Figure 4: Example of correctly detected inter-frame locations, superposed onto a color-coded two-frame difference image.

frame location, as predicted by the model. Figures 3 and 4 show several examples.

4.3.2 Interpretation

As mentioned before, the epipolar constraint is a soft constraint in the analysis of actual stereo image data, in the sense that some tolerance must be allowed for.

This circumstance can even be leveraged to accommodate for imperfections in calibration; indeed, for every matching which is made between 2D keypoints, the image is subsequently locally subjected to an affine warping, according to the average offset normal to the epipolar lines, which approximates the local effects of a re-calibration “as if the matchings were confirmed”.

In this way, intra-frame observations are not lost even if calibration is less than perfect, and moderate calibration errors do not result in non-detection, only in a predictable error in the estimation of velocities.

Keypoints appearing on consecutive frames might belong to the same particle. However, which 3D keypoints actually correspond to particle locations is not known beforehand because of the epipolar matching ambiguity, and which continuations correspond to particle trajectories is not known either.

For that reason, an immediate interpretation is not possible; it is necessary to take intra-frame observations into account at this point. These are not timestamped and we know only that they were accumulated between the instants $T \cdot i$ and $T \cdot (i + 1)$, and furthermore we know the temporal ordering because we regard a particle trace as a blurred line.

It is true that the 3D keypoints are already furnished with a tangent direction for each view except when the motion is exactly along an optical ray, which occurs almost never, and which determines the direction in 3D except for a motion inside an epipolar plane (i.e., each keypoint is a point in $\mathbb{R}^4 \times \mathbb{R}P^1 \times \mathbb{R}P^1$), which makes a preferred search direction plausible, but extrapolation, with moderate exposure times, is not possible especially in the case where turbulent motion is resolved and even so, it can only serve to reduce the number of possible candidates.

In the following, we always assume that all inter-frame locations which are detectable with sufficient confidence have been detected.

While it might be possible to achieve good results by suitably tracing out the trajectories directly on the image, we think it preferable not to proceed in this direct way but instead to generate hypotheses and filter them according to how well they explain the frame content. This circumvents most problems local path-following methods have with ambiguous situations, especially in the presence of occlusion.

The advantages to indirect analysis via hypotheses are manifest:

- Thanks to a global, top-down, view, results are much less likely to be influenced by local fluctuations, noise and difficult situations (i.e. crossings of traces) than local approaches.
- There are enhanced opportunities for a true probabilistic interpretation of the image series, as explained below
- One can unambiguously fit simple curves to just a few keypoints, which nevertheless remain accurate to high order if desired (B-splines lend themselves to it).

The third point above can be understood as an instance of Occam’s Razor; moreover, it renders optimization over curve spaces completely unnecessary.

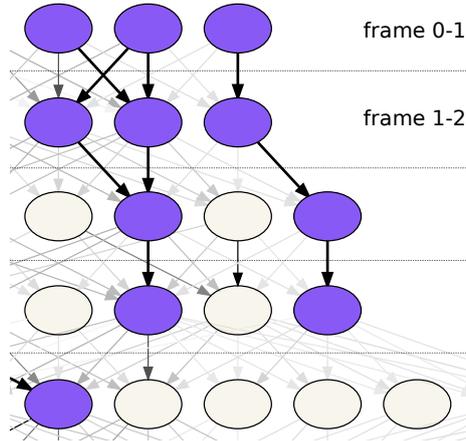


Figure 5: Directed graph of continuations.

All possible immediate continuations form an acyclic, directed graph (fig. 5). We extend it to a directed hypergraph by considering all paths up to an arbitrary number of frames f . f can be small, of the order of about 5 frames, and longer range dependencies are disregarded, because the particle motion can be described locally. By

cutting hyperedges (e.g. bottom-up, by removing edges first and enforcing transitivity under these constraints), one can partition the hypergraph of continuations into non-branching segments. These partitions are possible *explanations*, or *interpretations*, of the image evidence g and should be assigned a probability.

An application of the Bayes theorem: in eq. 6, β represents the curve parameters and g the image evidence. The prior distribution $p(\beta)$ can be picked on physical grounds; for example, one should favor interpretations which do not require excessive kinetic energy to realize. $p(\beta|g)$ would be read as the probability of a single trajectory being supported by the image evidence; the probability of the whole hypergraph partitioning depends on the individual trajectories' probabilities, which are independent except for the interdiction of crossings and for the handling of subchains.

$$p(\beta|g) = \frac{p(g|\beta)p(\beta)}{p(g)} \quad (6)$$

In our implementation of the framework, $p(g|\beta)$ is replaced by any of a set of criteria for good fit, ideally a monotonically increasing function of the actual probability, while the hypergraph partitioning itself is carried out by a greedy algorithm, which orders the possible trajectories by their score and discards all but the best one in ambiguous situations (proceeding recursively from longer chains to subchains). $p(\beta)$ was chosen as to penalize curves with abnormally high acceleration and $p(g)$ is irrelevant, because one is mostly interested in relative probabilities; A unified way of determining a threshold for any of the scores so as to reject bad explanations (in case no path is viable) would be a topic for further research.

The problem of interpreting the images is thus reduced from a problem involving properly an integration over continuous probability distributions to the purely combinatorial question of finding the chains of keypoints which correspond to actual particle traces, reducing the search space to a realistic size by arguing that all sensible curves which contribute most to the probability of a path looks about the same and can be reduced to a point estimate.

4.3.3 B-Splines

Each path candidate (chain) from the hypergraph is converted to a smooth trajectory candidate by means of B-splines [21].

As alluded to above, B-splines play a simplification role. On the one hand, detail which cannot be resolved when measuring is not spuriously resolved in the description (since B-spline spaces with few knots do not allow for it), and over fitting does not occur; on the other hand, they approximate data well (the approximation of a smooth function is accurate to high order, see [21] for details).

Note that B-splines are not suited for extrapolation. This is not a disadvantage in our application.

5 Prototype Development And Evaluation

5.1 Experimental Setup

Frames $\{F_{i,f}\}$ captured by the camera number i at frame number $f \in \{0, \dots, f_{max}\}$ are images obtained by the cameras sensors integrating from $\frac{f}{framerate}[s]$ to $\frac{f}{framerate} - gap[s]$ along the time axis. The gap during which the sensor is "not exposed" should be

as short as possible; it is always disregarded in the remainder of the text. In our setup, the camera positions are fixed with respect to an immobile scene. The backgrounds overall brightness may vary, because of oscillations in lighting intensity (indeed, fluorescent lamps might serve as light sources), which flicker at a frequency usually out of tune with the imaging system. This should not throw off our system, and indeed in the experiments it did not, even without more sophisticated background subtraction.

5.2 Implementation

In order to evaluate the system, a prototype was programmed in C++, using data structures and algorithms from the VIGRA library for image processing. Aside from an implementation of the detection and tracking algorithm, a GUI was build using the QT library, containing a visualisation widget based on the VTK library. The prototype can visualize detected particle trajectories, and a model of the experimental environment in 3D. The overall structure of the prototype can be seen in figure 6.

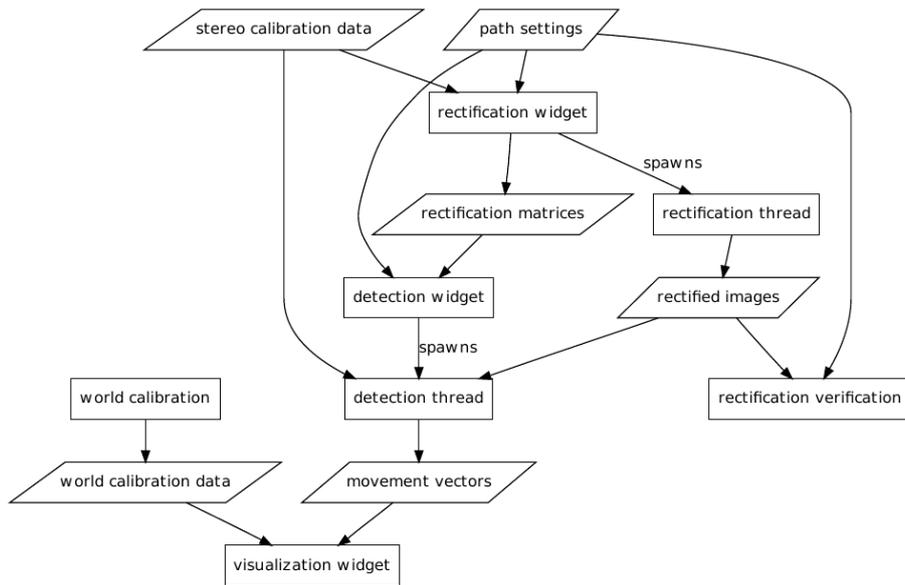


Figure 6: Illustration of the implementation of the system. Data is denoted by parallelograms, and processing components by rectangles.

Each feature is implemented as a separate widget. The detection and rectification steps are assigned their own threads, in order to keep the GUI operational while they are processing images. While the system can work on wholly rectified images, as indicated in figure 6, it is also able to rectify (and free from distortion) single points. Considering that neither rectification nor freedom from distortion are necessary for the particle segmentation and keypoint extraction, this yields a considerable performance gain. Rectification is, however, crucial for efficient correspondence search, by facilitating the search for correspondences on the scanlines of image pairs. The prototype is equipped with a widget to allow the user to visually evaluate the quality of the rectification (and in extension of the stereo calibration).

Our PTV system calculates the positions of the particle trajectories relative to one

of the cameras, and not in the world coordinate system. As knowledge of the latter position is required for a meaningful interpretation of the measurements, a widget is provided to assist the user in obtaining the necessary transformation between the two. This is done by calculating the absolute orientation between two sets of points from either system [22]. The user is required to manually establish the correspondences.

5.3 Experimental Results

The implemented prototype comprises temporal information and the reconstructed 3D particle trajectories. In spite of the foundations described in 4.3 the temporal information can be determined through the synchronized cameras and particle traces (blurred lines).

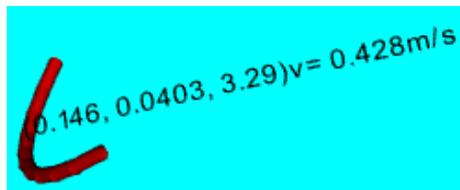


Figure 7: Visualisation of air-flow velocity with numerical indicators placed directly on the reconstructed 3D trajectories.

To compute the velocity of single particles we use two different algorithms, the above mentioned physical properties or the first derivative of the resulting curvature.

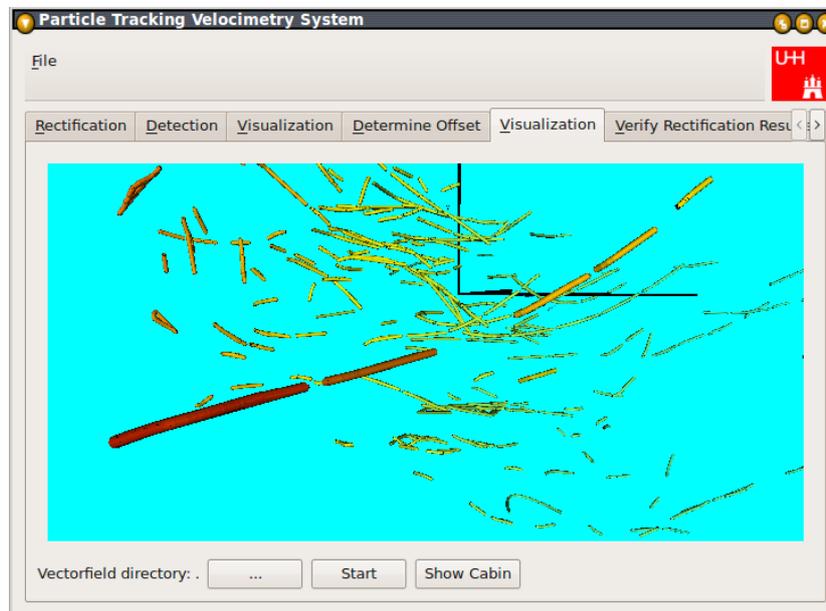


Figure 8: Color and hue coded velocities of reconstructed 3D trajectories. The slowest particles are green, the fastest are red.

To assist the developers of air conditioning systems we use different means for the

visualisation of the velocity of the particles. For the visualisation of single particles the direct numerical indicator (shown in fig. 7) can be used.

For a large number of particles, the color-coded visualisation of the velocities is advantageous. Figure 8 shows the color and hue coded trajectories of the particles, the slowest particles are green, the fastest are red (color and hue coded).

Exact evaluation of fluid experiments needs ground truth, which is not always easily obtainable: however, for validation, one can resort to synthetic images, employing ray tracing in order to use the observational model as a generative model.



Figure 9: Trajectory accurately reconstructed from synthetic image series.

Experimental results are encouraging: the detection process, when run on a synthetic image series showing a portion of a circular motion, reports after hyperedge selection only two candidates for motion, one corresponding closely to the true motion (with very low average deviation from it of about 1%), see fig. 9. Of course, more extensive testing is needed.

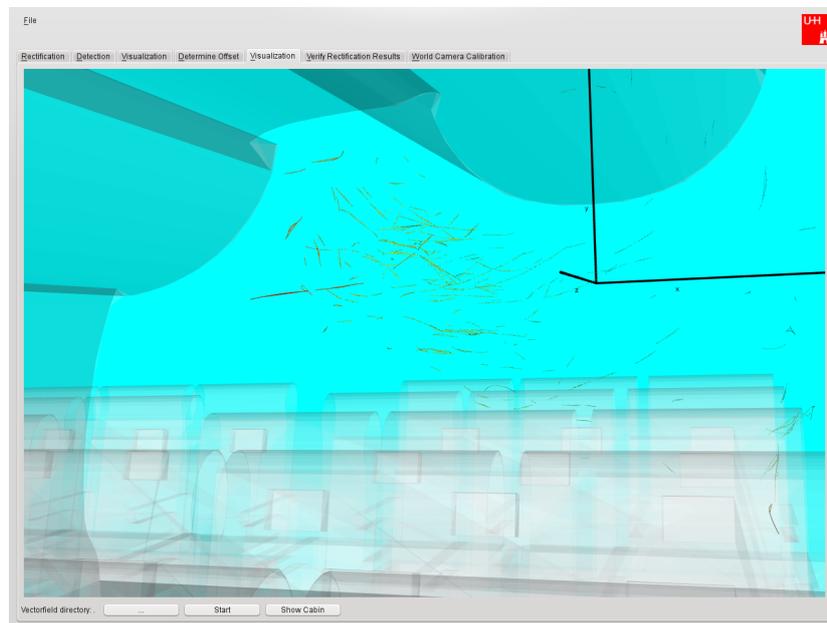


Figure 10: Detected particle trajectories shown in a model of the experimental environment.

The system was also subjected to a series of tests in an industrial environment for which prior measurements with a different system were known. As it was not feasible

to obtain the ground truth, it is impossible to calculate an error. However, the velocity distribution and overall flow orientation matched expectations. The measurement result of one of the test series is shown in figure 10.

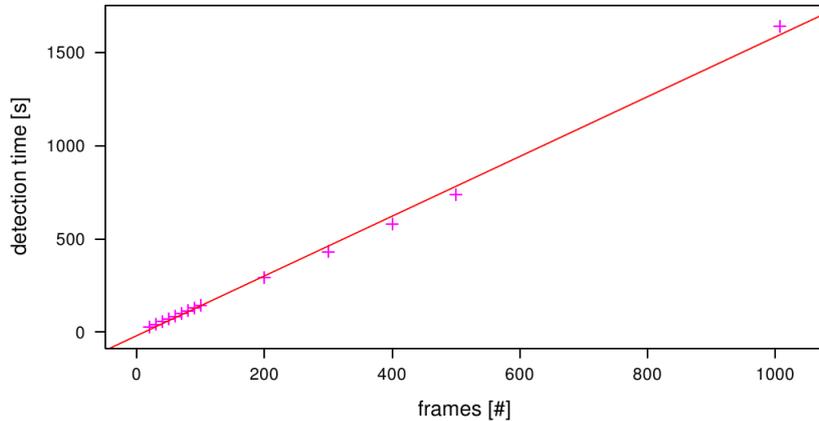


Figure 11: Plot of run time against number of frames average detected keypoints per frame: 50.

An image series of 1000 frames from each camera takes ca. 1500 s on a 1.6 GHz computer. The system scales linearly with the number of frames, as demonstrated in figure 11.

6 Conclusion

The ambition of this work was to develop a depth recovering particle velocimetry system, which can operate under fairly broad conditions, such as requiring only the use of a stereo camera setup and a low sampling frequency. The developed system compose of several modules, which are developed independently from bottom up. First of course well-known and still indispensable camera calibration. The next three steps, namely particle segmentation, key point extraction, matching and tracking constitute the centre and the novel scientific approach of the presented work. Last but not least, the visualisation of the results.

This work demonstrated that depth reconstruction of sparse flow information is practical using a stereo camera setup. Furthermore, tracking of individual particles can be done by generating a graph of possible paths and matching generated curves. After the summary of the results also known problems of the developed method are shown briefly. The selecting a suitable background mask to limit the number of candidate points and prevent the problem from becoming unmanageable from state explosion (super linear growth of possible choices).

The system is sensitive to bad calibration. It degrades performance by forcing wider search windows. Due to the sub optimal conditions in the experimental setup, weak contrast reduced the detection performance of particles. In combination with having to track particles over several frames, this leads to significant detection losses.

As a summary, the direct description of the measuring data by 3D trajectories instead of estimating first the 2D trajectories and this matching afterwards offer oneself

as advantageous. A higher precision can be achieved also with complicated trajectories. The presented system is capable of making essentially one-dimensional measurements of a fully $3 + 1$ - dimensional spatio-temporal phenomenon. The fact that these measurements are subject to a number of uncertain factors is known and noted; the uncertainty, as presented in the main text, can be estimated to a certain extent and that information is available for every individual measurement.

7 Further Developments

We propose and plan to develop a modified version of our algorithm where the steps of curve candidate generation and stereo reconstruction are swapped.

If matching of keypoints is required before particles can be tracked, the success of tracking will be critically dependent on the accuracy of calibration, which is an undesirable effect.

Tracking a particle, on the other hand, can be done on a single frame for most of its journey; this requires only the application of the tracking algorithm described in the above paragraphs directly to 2D keypoints. Doing so results in candidates for *projected* curves to be generated.

Parts of curves can be matched because of points which can be unambiguously identified on both views (at this point, camera calibration parameters must be taken into account): epipolar tangencies, as used by Porrill and Pollard [23] for another stereo vision task, are critical points of a suitable parameterization of the 2D curves which do appear simultaneously on both views related by epipolar geometry. It is hoped that such an approach will be less sensitive to miscalibration because the influence of errors in calibration can be exactly determined and conditions for matching set accordingly.

An added advantage will lie in the relaxation of the stringent time synchronization requirement, when matching is no longer dependent on the exact location 2D keypoints.

Still, other difficulties appear with this approach: a number of curves will in general not be matched with any confidence, especially if they lack the aforementioned special points because the particle travels only “up” or “down” during its time of visibility, meaning that more than two cameras will in general be advisable.

8 Acknowledgments

This paper was improved by collaboration with a large number of people and usage of different tools. In particular, we thank the Cabin Innovation Center of Airbus GmbH for their constructive feedback and for the possibility to use a full scale aircraft cabin mock-up and image sequences during our experiments for this work.

References

- [1] P.C. Stainback, K.A. Nagabushana, Review of Hot-Wire Anemometry Techniques and the Range of their Applicability for Various Flows, *Electronic Journal of Fluids Engineering* 119, Transactions of the ASME, 1997
- [2] F. K. Owen, A. K. Owen, Measurement and assessment of wind tunnel flow quality, *Progress in Aerospace Sciences* 44 pp 315-348, Elsevier Ltd., 2008

- [3] G. Guenther, J. Bosbach, J. Pennecot, C. Wagner, T. Lerche, and I. Gores, Experimental and numerical simulations of idealized aircraft cabin flows, *Aerosp Sci Technol* 10, pp 563-573, 2006
- [4] J. Bosbach, M. Kuehn, and C. Wagner, Large scale particle image velocimetry with helium filled soap bubbles, *Experiments in Fluids* 46, no. 3, pp 539-547, 2009
- [5] Y. Zhang, Y. Sun, A. Wang, JL Topmiller, JS Bennett, Experimental characterization of airflows in aircraft cabins, Part II: Results and research recommendations, *ASHRAE transactions* 111, pp 53-59, ISSN 0001-2505, 2005
- [6] D. Mueller, B. Mueller, and U. Renz, Three-dimensional particle-streak tracking (PST) velocity measurements of a heat exchanger inlet flow A new method to measure all three air-flow velocity components in a plane is applied to a steady-state three-dimensional flow, *Experiments in Fluids* 30, no. 6, pp 645-656, 2001
- [7] H.-G. Maas, A. Gruen, and D. Papantoniou, Particle tracking velocimetry in three-dimensional flows, Part A, *Exp. Fluids*, 15, 133-146, 1993
- [8] A. Wang, Y. Zhang, Y. Sun, X. Wang, Experimental study of ventilation effectiveness and air velocity distribution in an aircraft cabin mockup, *Building and Environment*, Volume 43, Issue 3, pp 337-343, 2008
- [9] J. Bosbach, M. Kuehn, C. Wagner, M. Raffel, C. Resagk, R. du Puits, and, A. Thess, Large scale particle image velocimetry of natural and mixed convection, 13th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal, 2006
- [10] Pozrikidis, C., *Fluid dynamics: theory, computation, and numerical simulation*, Springer Verlag, 2009
- [11] Merriam-Webster's Collegiate Dictionary (11th ed.), Occam's razor, New York, ISBN 0-87779-809-5, 2003
- [12] Florack, L., *Image structure*, Kluwer Academic Publishers, 1997
- [13] Lindeberg, T., *Scale-space theory in computer vision*, Springer Verlag, 1994
- [14] Florack, L. M. J., Ter Haar Romeny, B. M., Koenderink, J. J., and Viergever, M. A., Scale and the differential structure of images, *Image and Vision Computing*, Vol. 10, pp. 376-388, 1992
- [15] Koethe, U., *Reliable Low-Level Image Analysis*, Professorial dissertation, 2007
- [16] Z. Zhang, A flexible new technique for camera calibration, In: *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2000
- [17] J.-Y. Bouguet, K. Strobl, W. Sepp, C. Paredes, and K. Arbter, Camera Calibration Toolbox for Matlab, <http://www.vision.caltech.edu/bouguetj/>, 2008
- [18] A. Fusiello, E. Trucco, and A. Verri, A compact algorithm for rectification of stereo pairs, In: *Machine Vision and Applications*, Springer, 2000
- [19] Comaniciu, D., Meer P., Mean Shift: A Robust Approach Toward Feature Space Analysis, *IEEE Trans. Pattern Anal. Mach. Intell.*, Vol. 24, no. 5, ISSN 0162-8828, pp. 603-619, IEEE Computer Society, Washington, DC, USA, 2002
- [20] Harris, C., and Stephens, M., A combined corner and edge detector, *Alvey vision conference*, Vol. 15, Manchester, UK, 1988
- [21] C. De Boor, *A practical guide to splines*, Series: Applied Mathematical Sciences, Vol. 27, ISBN: 978-0-387-95366-3, Springer Verlag, 2001
- [22] B. K. P. Horn and others, Closed-form solution of absolute orientation using unit quaternions, *Journal of the Optical Society of America*, Vol. 4, pp: 629-642, 1987
- [23] Porrill, J., and Pollard, S., Curve matching and stereo calibration, *Image and Vision Computing*, Vol. 9, no. 1, pp. 45-50, Elsevier, 1991