2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

Visual Localisation in Outdoor Industrial Building Environments

Stephen Nuske*†, Jonathan Roberts† and Gordon Wyeth*

*School of Information Technology and Electrical Engineering
University of Queensland, St Lucia, Queensland 4072 Australia

†Autonomous Systems Lab, CSIRO ICT Centre
P.O. Box 883, Kenmore, Queensland 4069, Australia
Email: stephen.nuske@csiro.au

Abstract— **This paper presents a vision-based method of vehicle localisation that has been developed and tested on a large forklift type robotic vehicle which operates in a mainly outdoor industrial setting. The localiser uses a sparse 3D-edge-map of the environment and a particle filter to estimate the pose of the vehicle. The vehicle operates in dynamic and non-uniform outdoor lighting conditions, an issue that is addressed by using knowledge of the scene to intelligently adjust the camera exposure and hence improve the quality of the information in the image. Results from the industrial vehicle are shown and compared to another laser-based localiser which acts as a ground truth. An improved likelihood metric, using per-edge calculation, is presented and has shown to be 40% more accurate in estimating rotation. Visual localization results from the vehicle driving an arbitrary 1.5km path during a bright sunny period show an average position error of 0.44m and rotation error of 0.62deg.**

I. INTRODUCTION

In recent years, heavy industry has begun to investigate the use of automated mobile equipment to address productivity and personnel safety issues. The work reported here is part of a larger push to fully automate large ground vehicles operating in outdoor environments and we present a vision-based method of vehicle localisation. The technique has been developed and tested on a large forklift type robotic vehicle [1] which operates in an outdoor industrial setting (Fig. 1) handling large loads in the steel and aluminium industry. These vehicles operate for long periods outdoors during dynamic and non-uniform lighting conditions.

The operating conditions of the vehicle requires a vision system that is real-time, accurate and robust to illumination conditions. As well as the outdoor lighting conditions, the vehicles also move through sheds and therefore experience the challenge of the indoor-outdoor transition, which is discussed in [2].

The ultimate aim is to create a dependable fully autonomous vehicle. Robustness of the localisation system is a key element required to achieve dependability. One way of achieving robustness is through the use of multiple independent localisation systems that can be cross-checked and compared in order to automatically detect failure or degradation of one of the systems. To date we have developed a scanning laser localisation system [3]. The vision-based localisation system presented in this paper acts as an independent system to increase the dependability of the system.

Fig. 1. The large forklift vehicle in its operating environment.

The configuration of the buildings is permanent and have been surveyed to create a map which can be localised from using a camera. Recently, it has become popular to autonomously build visual maps comprising of image point features, some examples are [4], [5]. However, there are issues with autonomously building maps with image features, one problem is that the features are not view-point invariant in non-planar scenes as Vedaldi and Soatto [6] explain. Perhaps an even more significant problem is that the features do actually vary with lighting conditions in 3D environments, which is acknowledged by Lowe [7]. Another difficulty for autonomous map building is created by transient objects arriving and leaving the dynamic industrial environment. There are no guarantees in the image-feature extraction process that the features extracted and placed into the map are permanent. Therefore a localisation system based on a map with temporary features has the potential to fail.

The problems discussed above leads to the conclusion that autonomous map building from image-features is not a suitable solution for the high robustness required for the targeted application area. Thus a 3D-edge-map approach, which uses manual map surveying, has been logically selected for this work in order to meet the reliability requirement imposed by the industrial application.

978-1-4244-1647-9/08/$25.00 ©2008 IEEE.

Manual surveying of the map has many benefits for the industrial application presented in this paper. The application environment contains large buildings and sheds that are to be repetitively navigated. A map of the environment can be surveyed in just one day, a comparatively short period for a map that will be useful for countless hours - if not decades - of productive operation. Manual map surveying ensures that only permanent objects are included in the map, and that the description of the environment is explicitly separate from any specific lighting condition. All of which allow for reliable long term operation.

The 3D-edge localisation technique in this paper is based on Klein and Murray's [8] particle filter method which can be run in real-time computation on a Graphics Processing Unit (GPU). Here the technique is applied with a modified likelihood metric aimed at increasing accuracy. The outdoor lighting conditions are dealt with by a novel exposure control algorithm. The intelligent algorithm uses knowledge of the scene to control the exposure of important areas in the image while ignoring other areas. The algorithm shows significant improvements over the conventional exposure control algorithms which are naive and are susceptible to overcompensation. Results of robust localisation are presented in bright and highly-non uniform lighting conditions.

The remainder of the paper is structured as follows. In Sec. II, an edge-based localisation technique is described that can be used to estimate the position and heading of the mobile vehicle given a sparse 3D-edge-map of the doors in the environment. Sec. III deals with the issue of camera exposure and shows how knowledge of the scene can be used to intelligently adjust the camera exposure to improve the quality of information in the image. Results from experiments on a vehicle in a building site are presented in Sec. IV. Finally, conclusions are given in Sec. V.

II. Edge-Map localisation

There are many examples of 3D edge localisation in robotics and computer vision. Kosaka and Kak [9] developed one of the initial edge-based localisation techniques for navigating indoor hallways using a 3D-edge-map of the doors and walls. This type of technique has for the most part been applied in indoor environments. A more recent real-time technique developed by Drummond and Cipolla [10] has been applied outdoors by Reitmayr and Drummond [11]. The outdoor results presented by Reitmayr and Drummond are during overcast weather, whereas in this paper localisation is demonstrated in more extreme sunny conditions.

The previous 3D-edge-based techniques calculate only a single pose estimate each iteration, which is susceptible to error. Multi-modal techniques such as the recent particle filter method developed by Klein and Murray [8] offer improved robustness by maintaining many pose estimates per frame. Klein and Murray show how to compute the multiple estimates in real-time on a Graphics Processing Unit (GPU).

We base our work on the particle filter method of Klein and Murray [8]. Klein and Murray showed the technique applied to the tracking of a single object, such as a printer,

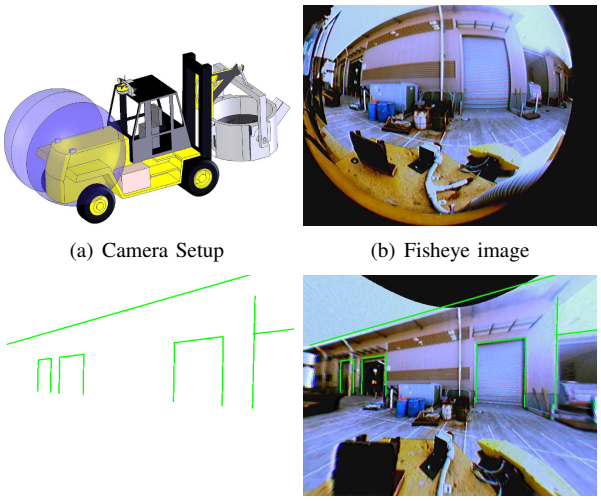

(a) Camera Setup (b) Fisheye image

(c) 3D-edge-map of buildings (d) Un-distorted image with projected 3D-edge-map

Fig. 2. Examples of the fish-eye camera setup and calibration. Two fish-eye cameras are placed at the front of the vehicle facing sideways. The blue hemispheres represent the field of view of the cameras.

from a range of just a few metres in a regular indoor environment. In our work, the same technique is applied using a fish-eye cameras mounted on a vehicle tracking the doors, walls and roof lines in a large outdoor industrial building environment. This application area also presents new challenges with the more extreme nature of the lighting conditions, which we deal with in Sec. III.

For the experiments shown in this work, a 3D-edge-map of the industrial buildings is used. The map is sparse and comprises of around 20 large industrial buildings and was surveyed in one day. Only permanent parts of the buildings are included - which is difficult to guarantee in an automated map building system. An example of the map is in Fig. 2(c).

Once the map is created a vehicle moving through the environment can be localised by matching edges in the map with edges extracted from the camera images. The comparison between map is calculated for each pose hypothesis in a particle filter and provides a likelihood measure, discussed in Sec II-B. Before the likelihood can be calculated, the camera lenses need to be calibrated, so that the 3D-edge-map can be projected to the image plane. A dual fish-eye camera setup is used, seen in Fig. 2(a), the calibration is discussed in the following section.

A. Fish-Eye Cameras

Fish-eye lenses that have a 185 degree FOV are used in this paper to capture a wide view of the environment, which is essential when the vehicle is near walls. The specialised lens model adopted for calibration is from Geyer and Danilidis [12], Ying and Hu [13] show that this model can be applied to fish-eye cameras. The model assumes that all pixels in the fish-eye image map onto a sphere located in front of the image plane. The model consists of four parameters, m, the distance from the center of the sphere

to the image plane, C_x and Cy, the center of projection on the image plane, and, l, the distance from the center of the sphere to the intersecting focus point of the light rays and the center of projection line.

These calibrated parameters enable the fish-eye image to be transformed into an undistorted image. Examples of distorted/undistorted images can be seen in Fig. 2. The corrected image is generated as follows; for each pixel coordinate $[U_u, U_v]$ in the corrected image (Fig. 2(d)), with centre $[C_x, C_y]$, the corresponding distorted coordinate $[D_u, D_v]$ in the fish-eye image (Fig. 2(b)) is calculated by;

$$D_u = Rcos(atan(\frac{U_v}{U_u})) + C_x \tag{1}$$

$$D_v = Rsin(atan(\frac{U_v}{U_u})) + C_y \tag{2}$$

where

$$R = \frac{sin(\theta)(m+l)}{cos(\theta)+l} \tag{3}$$

and

$$\theta = atan(\frac{\sqrt{U_u^2 + U_v^2}}{f}) \tag{4}$$

where f is the effective focal length of the undistorted projective image, measured in pixels from the center of the sphere. f is also used in the projection model to project the 3D-edge-map onto the image plane and can be selected according to the effective FOV that is required. The resulting undistorted image is converted into an edge-image using Canny's algorithm [14] with a 3×3 kernel.

B. Likelihood Metric

The likelihood measure for each particle is generated through comparison with the edge-image and the 3D-edge-map. The 3D-edge-map is projected according to each particle's pose onto the image plane, so a direct comparison can be made. Using a regular CPU this per-particle comparison would be very time consuming. Klein and Murray [8] present a fast method to perform this computation on a graphics processing unit (GPU).

The process begins by placing the undistorted edge-image into the GPU's texture memory at the beginning of the iteration. The 3D-edge-map is called to be rendered by a custom fragment shader program. The program allows the counting the visible edge pixels of the 3D-edge-map that align with edge-pixels in the undistorted edge-image. The custom fragment shader program only permits pixels to pass through the pipeline that align with edge-pixels in the edge-image. The pixels that pass this custom fragment shader program are counted using the OpenGL occlusion query extension [15].

Klein and Murray present the likelihood measure of the particle, P_χ, as a ratio between the count of aligning edge-pixels (a) and the total number of visible edge-pixels (v), calculated as follows:

$$Likelihood(P_\chi) \propto exp(\kappa \frac{a}{v}) \tag{5}$$

where κ is a constant that weights the likelihood measure.

Klein and Murray show this metric can successfully track objects, but the simple ratio of pixel counts leads to the situation where large edges, such as the roof-lines of the buildings, dominate other smaller edges, such as the door edges. This is simply because the majority of pixels are in the roof-edges. Smaller edges provide important localisation information and should have more consideration in the likelihood metric.

This paper presents a modified likelihood metric which calculates per-edge measurements, instead of just a sum over the whole image. The new metric calculates the ratio of aligning-to-visible edge pixels for each edge, j. This is calculated using occlusion queries for each edge, giving the two measurements a_j and v_j. The first component of the new metric is the original Klein and Murray global ratio. The second component is a sum of per-edge ratio's calculated as follows;

$$Likelihood(P_\chi) \propto exp(\kappa \frac{a}{v} + \lambda \frac{\sum_{j=0}^{n} \frac{a_j}{v_j}}{n}) \tag{6}$$

where n is the number of edges. The second component of this equation treats each edge equally regardless of its size. This penalises hypotheses that have smaller edges that do not align, even if the overall count of aligning pixels is high. This will enable the filter to maintain better track of the door edges in the environment. The new per-edge component has its own constant λ and this has to be tuned in conjunction with κ, striking a balance between the global and per-edge components.

C. Occlusions

The particle filter will be robust to occlusions caused by unknown objects, if there are more visible portions of the buildings than occluded portions.

Self-occlusions, where one building occludes another (known occlusions), can be dealt with by the depth buffer. Klein and Murray present a real-time technique using a sub-sampled depth buffer. The technique is to render faces of the buildings to the depth buffer, then only edges that are in front of the faces will pass through. The depth buffer is limited in resolution, which leads to the problem of a surface blocking its own edges. To avoid this issue the surfaces are recessed back a distance from the edge. The offset distance between surface and edge needs to be larger than the resolution of the buffer at that depth.

D. Propagation Model

Motion measurements from the vehicle can be formed into a model that propagates the particle filter. The uncertainty in the motion model m_t is defined by a Gaussian distribution as follows:

$$m_t = \varphi \tag{7}$$

φ is the Gaussian uncertainty distribution on the measured motion, δ. Here wheel encoders and steering encoders from the vehicle form δ as a 2D translation, t_x, t_y, and also a rotation around the vertical axis, r_z. Alternative types of motion measurements are presented by Klein and Drummond [16]

and by Machand et al. [17]. The ground in the environment is not perfectly flat and therefore slight vertical translations, t_z, and roll and pitching's, r_x, r_y, of the vehicle must be taken into account. There are no measurements in these additional degrees of freedom. So these unknown degrees of freedom are included in the propagation model by small perturbations across all 6 degrees of freedom, defined by the constant α.

The uncertainty in pose estimation is greater when the velocity increases and also is greater along the direction of the motion. This is modeled in the distribution by having a perturbation component that is proportional to δ, according to the constant β. The uncertainty distribution, φ, is defined by the mean, μ, and variance, σ^2, as follows:

$$\mu = \delta \tag{8}$$

$$\sigma^2 = \beta\delta + \alpha \tag{9}$$

III. Intelligent Exposure Control

The exposure control algorithms in most cameras use a grey-world assumption. These algorithms aim to control the mean intensity over the whole image to a predefined intensity level, regardless of the content of the scene. In our work there are specific areas of interest in the scene (i.e. the doors on the buildings) that must be correctly exposed. The conventional approach to exposure control causes overcorrection, resulting with an image that has incorrectly exposed areas. An example of overcorrection is shown in Fig. 3(a) which shows a lens flare that runs down the image even an abnormal sensor response. But more importantly the figure shows that the standard exposure control algorithm has over corrected for the sunny sky, leaving the buildings underexposed.

This paper presents a novel intelligent-exposure-control algorithm that adjusts the exposure according to known areas in the image. The algorithm aims to maximise the strength of image-edges corresponding to 3D-map edges, while ignoring all non-essential areas of the image. The algorithm first samples the intensity of pixels near the tracked edges. The 3D-edge-map is projected into the image according to the current pose estimate and short scans of pixels are taken along the normal of the edges. Fig. 5 shows which pixels are sampled from the scans. The control algorithm is as follows:

$$E_t = E_{t-1} + (1.0 - (.pdfilon I_r)) \tag{10}$$

where E_t is the exposure level at time t. The IEEE1394 IIDC [18] digital camera used in experimentation has two exposure parameters available, an analog to digital gain and a digital shutter time. These two parameters are scaled between 0-1, combined and represented by E. $.pdfilon$ is a damping constant and I_r is the ratio between the mean intensity, I_m from the pixel scans (Fig. 5) and the goal intensity I_d:

$$I_r = \frac{I_m}{I_d} \tag{11}$$

$I_d = 180$ has been determined based on the measurements made in Fig. 4 which show edge strength is at a maximum when the intensity of the sampled pixels is high but not saturated. The damping constant $.pdfilon$ has been set to

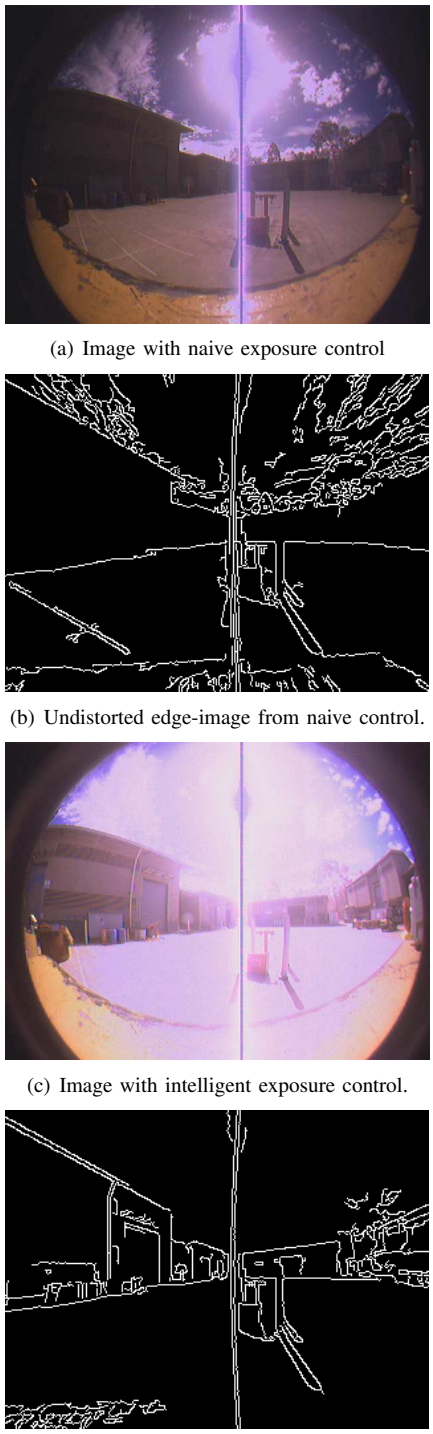

(a) Image with naive exposure control

(b) Undistorted edge-image from naive control.

(c) Image with intelligent exposure control.

(d) Undistorted edge-image from intelligent exposure control.

Fig. 3. Example of the bright lighting conditions. The sun causes a flare in the fish-eye lens and a dark line down the image due to errors in the sensor's response. Overcompensation for the sunlight can occur using naive exposure control found on most cameras. 3(a) shows an example of overcompensation where the buildings are under-exposed. 3(b) is the the corresponding edge-image where no edge-features are detected on the doors of the buildings. 3(c) shows that with the use of intelligent exposure control algorithm developed in this paper, the buildings are correctly exposed. As a result, the edges are detected on the doors and on the other areas on the buildings in 3(d).

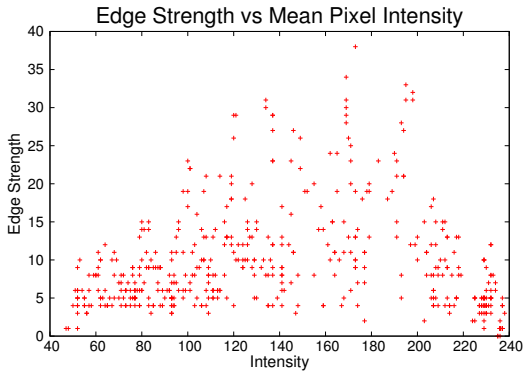

Fig. 4. The edge strength sampled from the doors of buildings, plotted against mean 8-bit pixel intensity of the sample. Edge strength is defined as the average intensity difference in a 3×3 pixel neighbourhood. This graph was recorded over a period of time as the exposure of a stationary camera was incrementally increased.

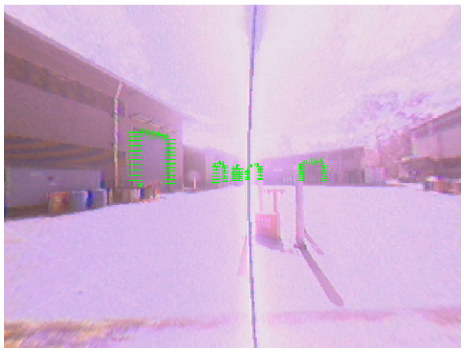

Fig. 5. Image showing the pixels that are sampled for the exposure control algorithm.

0.02, after empirical tests showing this value to provide a balance between quick response to lighting changes and stable control. Fig. 3(c) shows a typical result of this intelligent exposure control algorithm, where the buildings are properly exposed. The intelligent algorithm presented in this paper gives drastic improvements in edge-detection, edges on the doors and other areas of the buildings are detected in Fig. 3(d) and not in Fig. 3(b).

IV. Results

Consumer grade digital cameras with fish-eye lenses are mounted facing sideways on the vehicle capturing 640×480 images. The camera setup is seen Figure 2(a). The intrinsic and extrinsic camera parameters are manually calculated and verified by projecting the edge model into the image plane using the ground truth and ensuring the edges are aligned correctly with the recorded video stream. Fig. 2 shows the undistorted image using the calibrated fish-eye model. For the experiments shown in this section the lens parameters from (Eq. 2,3,4) used were focal scale $f = 175$ pixels, centre of projection of $C_x = 293, C_y = 305$, and fish-eye parameters $l = 2.81$, $m = 811$. Initially the various parameters of the particle filter need to be calculated. The approach used to calculate these parameters is to record a short sequence of video, odometry and ground truth pose data from the vehicle traveling around the environment covering most areas and orientations. The particle filter can be run off-line several times through the same two minute sequence of recorded data to optimise these parameters. This is achieved by using different parameters each cycle and comparing the average pose estimate error. The ground truth pose was collected from laser scanners using the method described in [3]. The number of particles is an important parameter to achieve accurate localisation, ideally the filter will have a large number of particles, but this comes with a larger computation cost. The goal operating rate for the system is 10Hz, and the particle filter can be operated at this rate using 1000 particles.

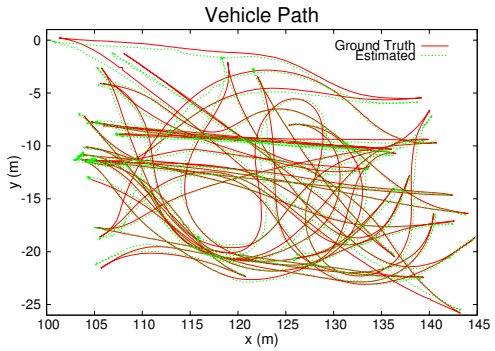

(a)

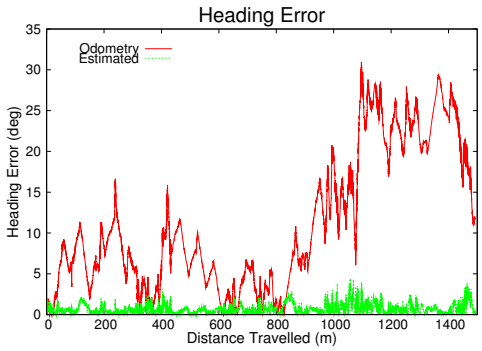

(b)

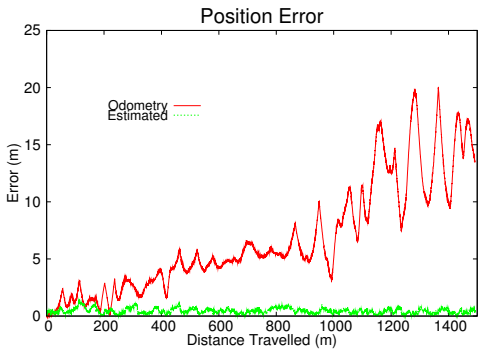

(c)

Fig. 6. Results from extended operation of the forklift type vehicle.

A. Extended Operation

Once the system has been implemented and tuned, it is put through a rigorous evaluation of its abilities through an extended period of operation around 2pm on a sunny day. At this time the lighting conditions are challenging because the sun is bright and on an angle. When the vehicle turns around the exposure control algorithm has to adjust quickly according to the direction each camera is facing. The vehicle was driven along an arbitrary path for a total distance of 1.5km during the experiment. The vehicle travelled through a wide range positions and orientations ensuring the system is well tested. Fig. 6(a) shows the path travelled during the extended operation experiment.

Fig. 6(b) presents the heading estimate error of the vision system which is maintained at an average error of 0.62deg, as opposed to the accumulated odometry error which drifts to a maximum error of 30deg. Fig. 6(c) shows the position error of the visual localisation. The vehicle's position was correctly estimated to an average error of 0.44m over the 1.5km run. The maximum error at one stage crept out to 1.4m and 4.4deg, this can be seen right at the top left of Fig 6(a), in the section of the path the vehicle is far from the buildings in the camera view. Though the average errors indicate that the system is sufficiently accurate for autonomous navigation around the site.

A video of the results of this experiment accompanies this document and a sequence of images is presented in Fig. 8. The sequence shows correct tracking of the buildings from a wide range of poses and with vehicle transitions from shadow to full sunlight. Some images have curved lens flares from the sun making some areas of the buildings not visible, in this situation the matching between model and image will not be perfect. The particle filter is robust to this scenario as it does not require a perfect match, rather the filter will re-sample particles with the best match. The fish-eye camera view has enough visible area for the tracking to not be affected by the lens flare. In some images the lens flares are dark, this is presumably due to an abnormal sensor response to the direct sunlight.

B. Likelihood Metric

A new likelihood metric is presented by this paper (Eq 6) and is compared with the accuracy of Klein and Murray's metric (Eq 5). The two metrics are evaluated against the laser-based localiser on the same sequence of data and their recorded errors are shown in Fig 7. There is only minimal improvement in position accuracy of the proposed metric over Klein and Murray's, both metrics give average position errors of slightly less than 0.5m. However there is a 40% improvement in heading accuracy. The proposed metric gives a average heading error of 0.62deg whereas Klein and Murray's is 1.01deg. This significant improvement in heading estimation is attributed to the proposed likelihood metric giving a greater importance to aligning the door edges. Klein and Murray's simpler likelihood metric will give only a fraction of the importance to the door edges as they only have a fraction of the pixels of the roof edges.

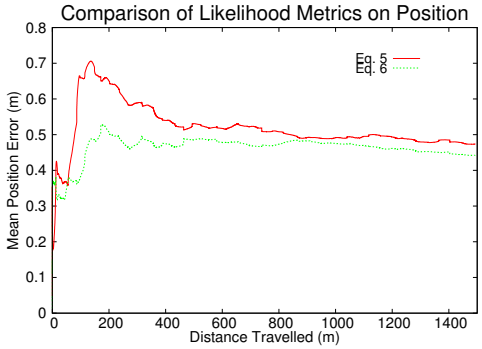

(a)

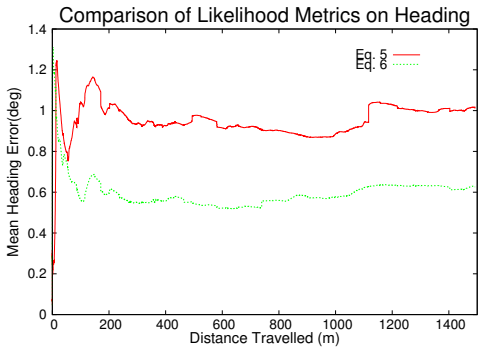

(b)

Fig. 7. Comparison between recorded errors from Klein and Murray's likelihood metric (Eq 5) and this paper's proposed likelihood metric (Eq 6). There is only minimal improvement in position accuracy of the proposed metric over Klein and Murray's, however there is a 40% improvement in heading accuracy.

V. Conclusions

This paper has described a vision-based localisation system for a ground vehicle operating in an outdoor industrial setting. The vision localiser is based on Klein and Murray's approach [8] with an improved likelihood metric and is applied outdoors using a sparse 3D-edge-map of the building environment. The technique is processed at 10Hz by utilising a particle filter and a standard GPU.

An intelligent exposure control algorithm is presented to enable operation in the dynamic and non-uniform outdoor lighting conditions. The algorithm uses knowledge of the scene to adjust the camera exposure and hence improve the quality of the important information in the image.

Finally, the performance of the localisation system is demonstrated on an industrial vehicle while it is driven around a test site for an extended period where the vehicle covered a total distance of 1.5km. The pose estimates from the vision-based localiser were compared to a laser-based localiser which acted as a ground truth. The pose of the vehicle was estimated by the visual localiser to an average position error of 0.44m and average rotation error of 0.62deg over the 1.5km path. The improved likelihood metric, which uses per-edge calculations, proved to be 40% more accurate in estimating the heading of the vehicle.

ACKNOWLEDGMENTS

This work was funded in part by CSIRO ICT Centre's Autonomous Ground Vehicle project, CSIRO's Light Metals Flagship project and the School of Information Technology and Electrical Engineering at the University of Queensland. The authors gratefully acknowledge the following members of the Autonomous Systems Lab's team for their assistance with the setup, calibration and experimental work conducted during this project: Cedric Pradalier, Ashley Tews, Peter Hansen, Paul Flick, Polly Alexander, Felix Duvallet and Felix Ruess.

REFERENCES

[1] J. Roberts, A. Tews, C. Pradalier, and K. Usher, "Autonomous hot metal carrier - navigation and manipulation with a 20 tonne industrial vehicle," in *Proceedings of IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 2007, pp. 2770–2771, video paper.

[2] S. Nuske, J. Roberts, and G. Wyeth, "Extending the dynamic range of robotic vision," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Florida, USA, May 2006, pp. 162–167.

[3] A. Tews, C. Pradalier, and J. Roberts, "Autonomous hot metal carrier," in *Proceedings of IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 2007, pp. 1176–1182.

[4] N. Karlsson, E. D. Bernado, J. Ostrowski, L. Goncalves, P. Pirjanian, and M. E. Munich, "The v-slam algorithm for robust localization and mapping," in *Proceedings of IEEE International Conference on Robotics and Automation*, 2005.

[5] P. Newman, D. Cole, and K. Ho, "Outdoor slam using visual appearance and laser ranging," in *International Conference on Robotics and Automation*, 2006.

[6] A. Vedaldi and S. Soatto, "On viewpoint invariance for non-planar scenes," UCLA, Tech. Rep., 2006.

[7] D. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[8] G. Klein and D. Murray, "Full-3d edge tracking with a particle filter," in *British Machine Vision Conference*, 2006.

[9] A. Kosaka and A. Kak, "Fast vision-guided mobile robot navigation using model-based reasoning and prediction of uncertainties," in *Proc. Int. Conf. on Intelligent Robots and Systems*, 1992.

[10] T. Drummond and R. Cipolla, "Real time visual tracking of complex structures," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2002.

[11] G. Reitmayr and T. Drummond, "Going out: robust model-based tracking for outdoor augmented reality," in *International Symposium on Mixed and Augmented Reality*, 2006, pp. 109–118.

[12] C. Geyer and K. Danilidis, "Catadioptric projective geometry," *International Journal of Computer Vision*, vol. 45, no. 3, pp. 223–243, 2001.

[13] X. Ying and Z. Hu, *Computer Vision - ECCV 2004*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2004, vol. 3021/2004, ch. Can We Consider Central Catadioptric Cameras and Fisheye Cameras within a Unified Imaging Model, pp. 442–455.

[14] J. Canny, "A computational approach to edge detection," *IEEE Transactions Pattern Analysis Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.

[15] *NVIDIA OpenGL Extension Specifications*, NVIDIA Corporation, February 2007.

[16] G. Klein and T. Drummond, "Tightly integrated sensor fusion for robust visual tracking," in *Proceedings of British Machine Vision Conference*, 2002.

[17] E. Marchand, P. Bouthemy, F. Chaumette, and V. Moreau, "Robust real-time visual tracking using a 2d-3d model-based approach," in *IEEE International Confernece on Computer Vision*, 1999.

[18] *IIDC 1394-based Digital Camera Specification*, 1st ed., 1394 Trade Association, Regency Plaza Suite 350, 2350 Mission College Blvd., Santa Clara, CA 95054, USA, July 2000.

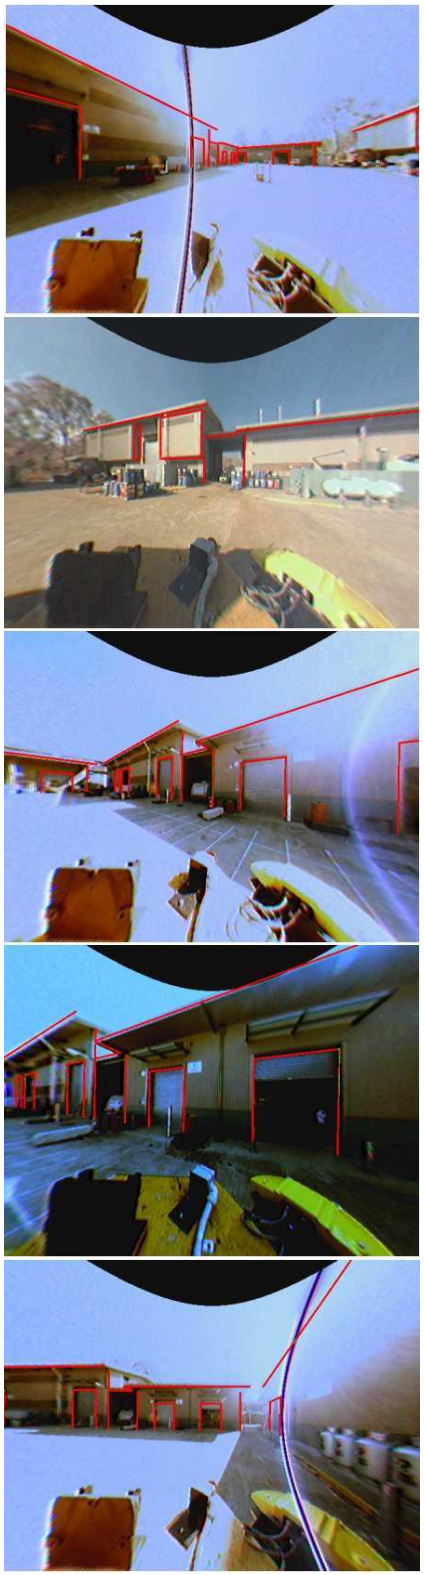

Fig. 8. Undistorted images taken from the experiment, overlaid in red with the 3D-edge-map projected from the estimated pose.