The Effect of Lighting in Nighttime Vision-Based Localisation

Ashley Tews and Paulo Borges
CSIRO, Australia
firstname.lastname@csiro.au

In an effort to create high reliability autonomous ground vehicles capable of operations 24 hours a day, redundancy plays a key role. In particular, localisation systems need to have multiple separate systems to allow for position validation and backup in case of failure of the primary localiser. These systems need to provide high performance in day and night conditions. In this paper, we investigate the onboard and offboard lighting required to allow a vision-based localisation system to provide high performance during night operations in an industrial environment. The comprehensive discussion on sensors and setup and the quantitative evaluation presented provide a relevant discussion towards using vision systems 24 hours a day.

1 Introduction

For autonomous industrial vehicles, the localisation system needs to be robust to the operational and environmental conditions, and provide the required accuracy for the intended operations, 24 hours a day, in various environment and weather conditions. It also needs to have relatively inexpensive hardware components to be feasible to industry, and ideally require as few additional computational resources as possible for satisfactory operation.

Our goal is to develop a robust localisation system for autonomous industrial vehicles that has levels of redundancy to allow the vehicle to continue operations despite failures and be able to detect and manage degradations in performance. The system needs to be reliable and ideally require as little as possible additional resources to allow them to operate. Currently, we have a reliable primary localiser based on retro-reflective beacons and lidars. This system has been in use for several years on our autonomous Hot Metal Carrier (Figure 1) and demonstrated high accuracy and robustness under varying environment and weather conditions.

In this paper, we focus on evaluating a secondary, vision-based localisation system, that has demonstrated high accuracy under daylight conditions [Nuske *et al.*, 2008; 2009]. A vision-based approach offers a different modality to the lidar system which improves reliability under primary system failure. One of the issues with using images from the visible spectrum is its performance variation with changes in lighting conditions. Low-light and high-dynamic range cameras alleviate these issues to an extent but they still suffer in dark environments. In this paper, we test the effects of different common lighting methods on performance of a vision-based localisation system, when compared with a known accurate lidar-based system used as a ground truth. The main contribution is the evaluation of the lighting and sensor required in night-time operations to allow a vision-based localiser to perform to a high enough standard to be considered as a viable secondary localiser for autonomous industrial vehicles.

Figure 1: The autonomous Hot Metal Carrier undertaking payload handling operations in a duststorm.

The remainder of this paper is arranged with the related work next in Section 2. Section 3 outlines the methodology including the cameras and lighting used. Section 4 shows the results of comparison between the ground truth and vision-based localisation system under different lighting conditions. A discussion is also included to summarise the performance of the different combinations. Conclusions are presented in Section 5.

2 Localisation Using Edge Features

Most of the related literature using vision solutions assume that good quality images are available. Under these circumstances, a common approach for vision-based localisation is to build visual maps consisting of image point features, such as SIFT [Lowe, 2004] or SURF [Bay *et al.*, 2006] applied to outdoor scenarios [Weiss *et al.*, 2007; Tabuse *et al.*, 2011]. Although very good results can be achieved in specific environments, a significant concern is the fact that the features are not invariant to non-planar scenes [Vedaldi and Soatto, 2008]. The main challenge with point features, however, is that they change significantly depending on lighting conditions [Lowe, 2004], which can make them inadequate for repeatable longer term operations [Nuske *et al.*, 2009]. When precise localisation is not necessary, Valgren et al. [Valgren and Lilienthal, 2010], for example, have indicated that SURF can be applied for topological localisation across different seasons. A sequential learning matching method for place recognition [Milford and Wyeth, 2012] has also shown very good results, although continuous precise localisation is still a challenge. Other attempts to address illumination changes include combining different SLAM algorithms [Glover *et al.*, 2010], the generation of dynamic maps [Biber *et al.*, 2005], and emphasis on high dynamic range [Irie *et al.*, 2011]. Still, these methods deal with relatively minor perceptual change caused by sun movement during the day, or constrained changes in an indoor office environment.

Mitigating these lighting issues, the use of edge features has been proposed [Nuske *et al.*, 2008; Rofer and Jungel, 2003; Zhou *et al.*, 2003], illustrating that edges from infrastructure (buildings, posts, trees, etc) are a very strong environmental feature that are fairly robust to illumination changes. Among methods that exploit edge features, additional robustness can be achieved by comparing an edge-filtered version of the input from the camera with a predefined 3D edge map of the environment. Following this approach, Drummond and Cipolla [Drummond and Cipolla, 2002] and Reitmayr and Drummond [Reitmayr and Drummond, 2006] developed a real-time edge localisation technique that can be applied indoors and outdoors. One limitation of those works is that the 3D edge-based techniques calculate only a single pose estimate per iteration, reducing the robustness of the system as it becomes susceptible to isolated errors. Multi-modal alternatives such as the particle filter method proposed by Klein and Murray [Klein and Murray, 2006] provide better performance by maintaining several pose estimates at each frame. In an extension of Klein and Murray's algorithm, a novel observation function for the particle filter was developed by Nuske et al. [Nuske *et al.*, 2009] (we refer to this method as particle filter edge-based localisation (PFEL) in this paper), further improving the localisation accuracy. In addition, their algorithm also uses an intelligent exposure control that improves the quality of the relevant edge information in the image. The authors show results outdoors over sunny and rainy weather, illustrating significant robustness in challenging and changing illumination conditions. Exploiting the fact that in structured environments edges are often straight lines, edge features can be further amplified by applying the Hough transform to the high-pass image, filtering in lines that are more likely to be part of the predefined 3D edge map of the environment [Borges *et al.*, 2010].

As mentioned earlier, the goal of this paper is to evaluate the lighting required in night-time operations to allow a vision-based localiser to perform to high enough standard to be a viable secondary localiser for autonomous industrial vehicles. Considering the techniques above, the one proposed by Nuske et al. [Nuske *et al.*, 2009], presents high robustness in our industrial site, where the architectural configuration of the buildings is fixed, and provides a good set of edge line features from which to localise. For this reason, we choose PFEL as the basis for our low-light operation evaluation. More details about this localisation algorithm are given in Section 3.2.

3 Method

A primary motivation for our research is to have it deployed in industrial environments and gaining industry acceptance. One of the implications of this is that we need to keep the costs of system implementation as low as possible, but highly effective. Based on this motivation and the performance of the vision-based localiser [Nuske *et al.*, 2008], our approach is to determine the level of normal industrial lighting around infrastructure and on the vehicle that provides high performance.

Figures 2 and 3 show the test environment during day and night from a building webcam. The environment has been used for autonomous vehicle development for several years and the laser-beacon localisation system [Pradalier *et al.*, 2008] has demonstrated consistently accurate and reliable performance in this area.

Each of the major aspects to be considered are described in the remainder of this section.

3.1 Lighting

As discussed in Section 2, systems that rely on features can have differently matched features across a very short period of time. However, some features, such as building and shed door outlines may still be strong enough

Figure 2: The test environment at daytime. The reflective beacons used in the laser beacon localisation system can be seen as white tape attached to bollards. The environment measures approximately 50m by 35m.

Figure 3: The test environment at night, illuminated by exterior shed lights, and through the two open doors of the robot shed. While it may appear there is a substantial amount of lighting, the building camera has good low visibility performance, and the illumination through the shed doors was not used in the tests.

to allow for good localisation. This is illustrated in Figures 4 and 5, which show the same images as above with Sobel edge-detection applied. Apart from the major differences in scenes including the barrels and trees, the main building and doorway edges can be seen in both figures. In the test area, site lighting consists of lights above personal door entries to sheds, and flood lights mounted to the outside of sheds. These provide sufficient illumination for people to move around the area during the night.

The other main source of illumination is from the lights onboard a vehicle. These can include driving lights illuminating the road in front and behind the vehicle, and side floodlights to allow better situational awareness for the vehicle operator. For our purposes, we used the forward and rear vehicle lights, and floodlights on the side of the vehicle to illuminate the scene where the cameras were facing as shown in Figure 6.

Figure 4: The daytime image with a Sobel filter applied.

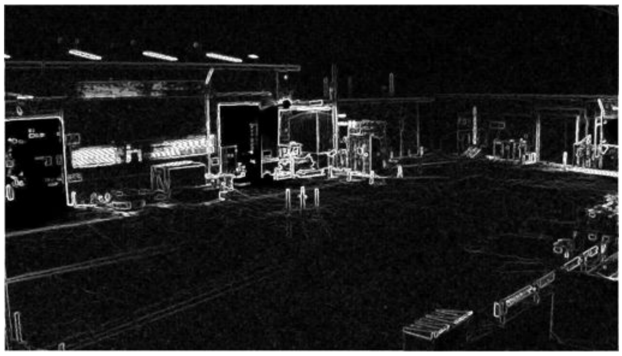

Figure 5: The night image with a Sobel filter applied.

3.2 Edge-based Localisation

Based on the lighting aspects discussed above and the literature review in Section 2, PFEL is the vision-based localisation method employed in this work. It uses cameras mounted on a vehicle tracking linear features such as building edges, doors, and roof lines in a large outdoor industrial building environment. For this task, a sparse 3D edge map of the site is utilised, consisting of around 20 large industrial buildings. This map can be generated via professional surveying, or acquired automatically with laser range sensors [Borges *et al.*, 2010]. Examples of the surveyed map for the test environment are given in Figure 7. 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 onboard camera images (Figures 8(a) and 8(b)). The comparison between the image and the map is calculated for each pose hypothesis in a particle filter and provides a likelihood measure for that particle. In our case, where the images were dark, we performed histogram equalisation to improve the edge detection. Although some of the images looked quite dark to the human eye (e.g. Figure 10), because of the high signal-to-noise ratio in the sensor (discussed in Section 3.3 below), a histogram equalisation was able to improve significantly the edge detection without causing excessive noise artifacts.

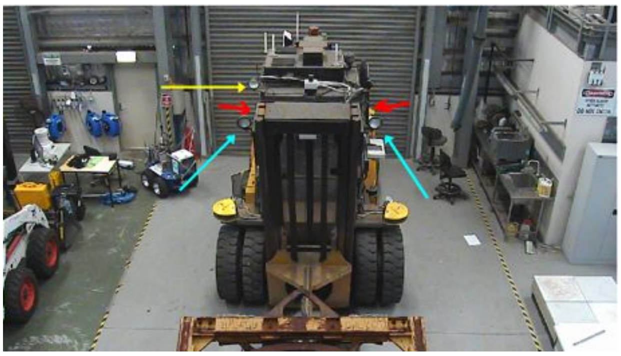

Figure 6: The lights onboard the vehicle. There are two forward driving lights, two side lights (yellow arrow), and two rear lights (cyan arrows) since the vehicle reverses in to pick up the payload. In this figure, one of the side lights is facing rearwards towards the vehicle's pickup hook. The other is not visible. The red arrows indicate the approximate positions of the two cameras, mounted above the left and right mudguards.

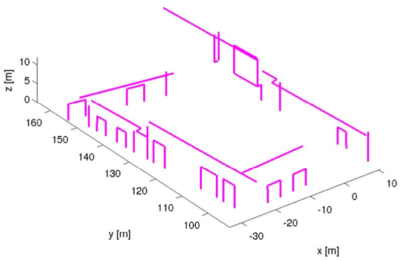

Figure 7: 3-D edge map of the test area generated from surveyed data.

Full details on the particle filter implementation can be found in the original publication [Nuske *et al.*, 2009]. This reference also provides additional information about the system, such as robustness to occlusions and intelligent exposure control to deal with challenging outdoor lighting conditions.

3.3 Cameras

The camera used plays an essential role in the quality of the localisation results. Using standard consumer cameras, we have performed tests with a number of CMOS and CCD sensors. Examples include the Basler sca750-60gc, the Axis 233d, the Unibrain Fire-i 1.2, the Pointgrey Dragonfly 640×480C, and the Basler A312fc

The experiments indicated that in low light, the camera's sensors discussed above are not sensitive enough (even with some amount of auxiliary external lighting, as discussed in Section 4) and are not suitable for the localisation algorithm, considering the necessary frame rate and resolutions. Among the tested cameras, however, one class of sensors - the Sony ICX285 series - showed very good low light performance, providing satisfactory input images at 15 frames per second. This sensor has a quantum efficiency around 64% with a low dark current (8 electrons/pixel/second). The quantum efficiency is defined as the percentage of the generated electronic charges by the incoming photons, whereas the dark current is the current produced when no photons are reaching the sensor. As the ICX285 is a relatively popular sensor, this was our primary choice. Therefore, most of the results reported in Section 4 were obtained using the Basler sca1400-gm camera, which uses the ICX285 of resolution 1280×960. A small number of preliminary tests with more recent sensors (the Sony ICX674) with a quantum efficiency of 68% have also presented satisfactory images. This indicates that sensors with a quantum efficiency above 60% and low dark noise are potential candidates for mobile robotic platforms operating outdoors in relatively low light, when considering algorithms that rely on strong features like edges. In contrast, most of the low performance sensors tested had a quantum efficiency of 50% or less. The lenses used with these cameras were the Kowa LM5JC10M with a 5mm focal length

Apart from the Basler sca1400-gm, we also report results from the usage of a thermal camera, which in our case is the Thermoteknix Miricle 307KS camera, of resolution 640×480.

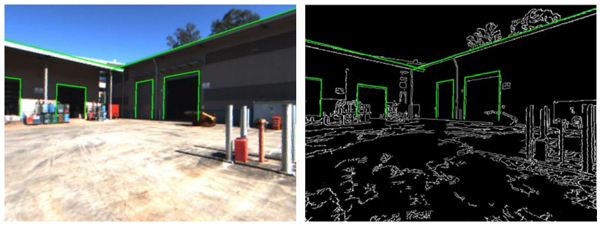

(a) Image after undistortion, overlaid with edge model (in green).

(b) Edge image after undistortion, overlaid with edge model (in green).

Figure 8: Original and edge images with models overlaid.

3.4 Test Procedure

The testing procedure involves evaluating the PFEL method under varying lighting conditions at night for the Basler sca1400-gm and the Thermoteknix Miricle thermal cameras mounted to the vehicle. For each combination of lighting conditions and cameras, the same path is driven around the environment. The autonomous vehicle is used for this purpose as it has demonstrated high repeatability and accuracy in hundreds of demonstrations in the test area. The ground truth localiser used for the autonomous vehicle is based on a particle filter laser beacon system. Independent tests of the system demonstrate accuracies between 5cm to 20 cm depending on the local beacon density. We have found these accuracies to be adequate for undertaking autonomous materials handling tasks.

Since the internal map used by the vision-based localiser consists of external infrastructure edges (Figure 7), the vehicle starts its mission outside its parking shed to allow the vision system to localise. The vehicle then conducts a mission involving reversing to picking up its payload, circling the test environment, reversing in to drop off the payload, and then lining up to reverse into the shed to park and shutdown. We show results from the start location to the park location which is inside the building. We let the localiser run inside the building to examine the drift in the localiser once it has lost its references.

4 Results

The lighting parameters tested for the thermal and visible cameras were:

- no site or vehicle lights
- vehicle lights only
- site lights only
- both vehicle and site lights

For each combination of lights and cameras, a set of figures is presented in this section. These show - the path comparison between the ground truth localiser (shown in red), edge-based localiser (shown in green) and wheel odometry (shown in blue); the error between the edge-based localiser and ground truth; histogram of the errors over the run; and a sample image from the camera. For the path comparisons, the start locations are shown where the blue circles appear and the end is at the top left of the figures.

To provide a reference to the integrity of the vision-based localiser, a trial run was conducted during a sunny day when illumination conditions were considered favourable. Figure 9 shows the resulting path comparison to the ground truth localiser.

4.1 No Site Lights

We arranged for the building and security lights to be turned off for the tests in this section. The only constant light source for all tests was from a screen mounted on the side of the vehicle used for monitoring the vehicle's systems. Figure 10 shows the environment from a building webcam with site lights off and the vehicle lights on. The light at the front (left) is an orange system indicator light and appears bright in the figure due to the webcam overexposing the image. The onboard cameras cannot see these dim lights and they have a negligible effect on environment illumination. The higher light is the forward facing driving light and the lower is the computer screen.

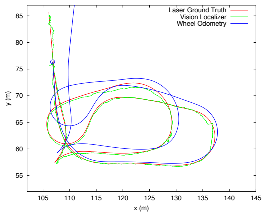

(a) Path comparison

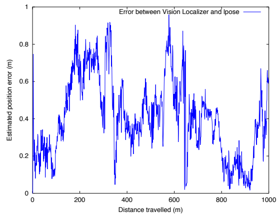

(b) Path error

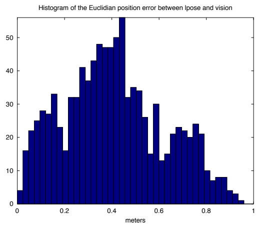

(c) Error histogram

Figure 9: Performance of the vision localiser during a cloudy day with relatively uniform lighting across the test environment. This is used as a benchmark for comparing the night trials with. Note, this was undertaken with manual driving. The blue circle in (a) represents the starting point of the test run.

Figure 10: The test environment with vehicle lighting and no site lighting.

Visible Camera Results

The results in this section are for the visible camera's localisation performance compared to the laser-beacon localiser. Figure 11 was undertaken with no site or vehicle lights and Figure 12 is with only vehicle lights.

Thermal Camera Results

The thermal camera results are included to evaluate how well the vision localiser can process the different modal-

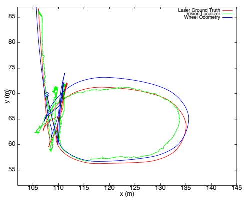

(a) Path comparison

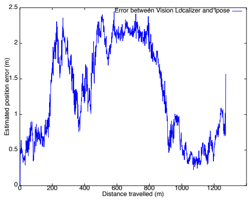

(b) Path error

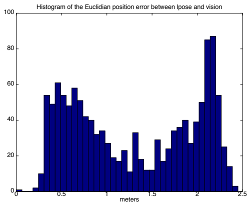

(c) Error histogram

(d) Camera view

Figure 11: Performance of the vision localiser with no site lights or vehicle lights.

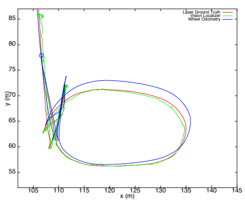

(a) Path comparison

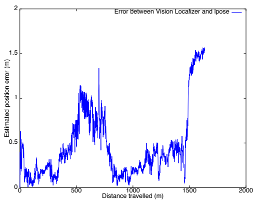

(b) Path error

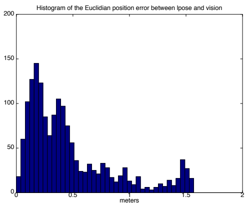

(c) Error histogram

(d) Camera view

Figure 12: Performance of the vision localiser using only vehicle lighting.

ity and for completeness (Figures 13 and 14). No additional pre-processing steps were undertaken on the thermal image stream before generating position estimates.

4.2 Site Lights

In these tests, the site lights were activated. These consist of standard external shed lights and personal door lights. Figure 15 shows a view from a building webcam with the site and vehicle lights on. Note the vehicle's side lights illuminating the shed.

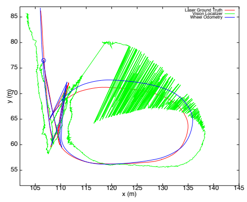

(a) Path comparison

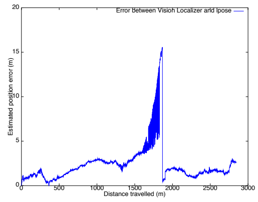

(b) Path error

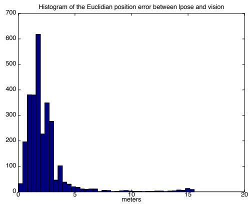

(c) Error histogram

(d) Camera view

Figure 13: Performance of the thermal camera localiser with no lighting.

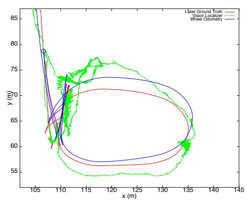

(a) Path comparison

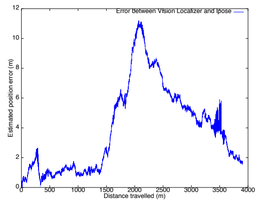

(b) Path error

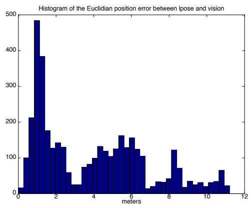

(c) Error histogram

(d) Camera view

Figure 14: Performance of the thermal camera localiser using only vehicle lighting.

Visible Camera Results

The visible camera results are shown in this section in Figure 16 with no vehicle lights and Figure 17 with vehicle lights adding illumination to the scene.

Thermal Camera Results

This section shows the results from the thermal camera with no vehicle lighting in Figure 18 and with vehicle lighting in Figure 19.

Figure 15: The test environment with site and vehicle lighting.

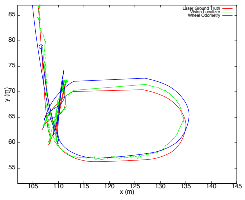

(a) Path comparison

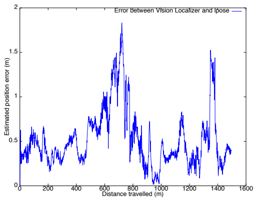

(b) Path error

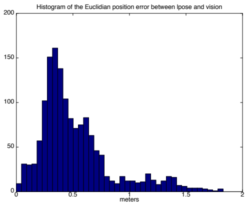

(c) Error histogram

(d) Camera view

Figure 16: Performance of the vision localiser with site lights and no vehicle lights.

4.3 Discussion

The results in Table 1 summarise the localisation performance. We can see that the best results in terms of mean error (defined by the mean absolute Euclidean distance for all points) and standard deviation are obtained when the vehicle lights are on and the site lights are off. Notice that this result is only approximately 20% worse than that obtained for the best day-time run (indicated by the row "Day"). The performance of the visible camera with no lighting except for the screen mounted on the side of the vehicle is interesting. In images with a strong light source such as roadside or building lights, there is a wide contrast in the image from darkness to potentially the bright light source itself. Consequently, edge extraction is inconsistent around the environment depending on the type and coverage of the lighting. Whereas the onboard screen provides diffuse illumination across the

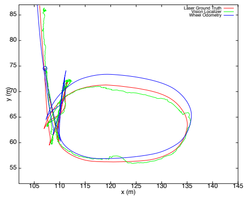

(a) Path comparison

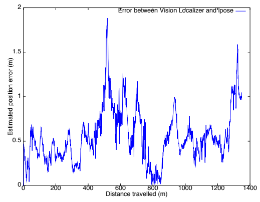

(b) Path error

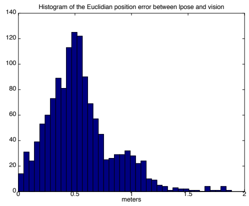

(c) Error histogram

(d) Camera view

Figure 17: Performance of the vision localiser using site and vehicle lighting.

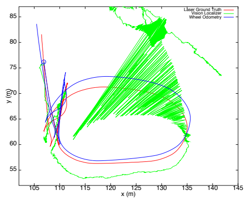

(a) Path comparison

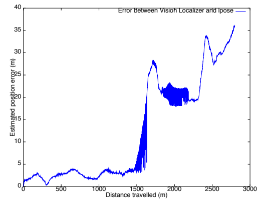

(b) Path error

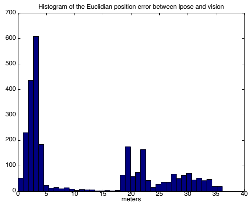

(c) Error histogram

(d) Camera view

Figure 18: Performance of the thermal camera localiser with site lights and no vehicle lighting.

scene, the edges are more consistent in brightness.

On the other hand, when the site lights are on, the vehicle lights do not bring as much benefit, as expected. Nonetheless, a visual analysis shows that the vehicle lights do help in reducing the apparent effect of the point source illumination and overexposure of the site lights on the camera, even though this did not influence significantly the results.

Although the thermal camera provides good images, this modality does not render sharp enough edges for the algorithm. This is illustrated by the poor results in the

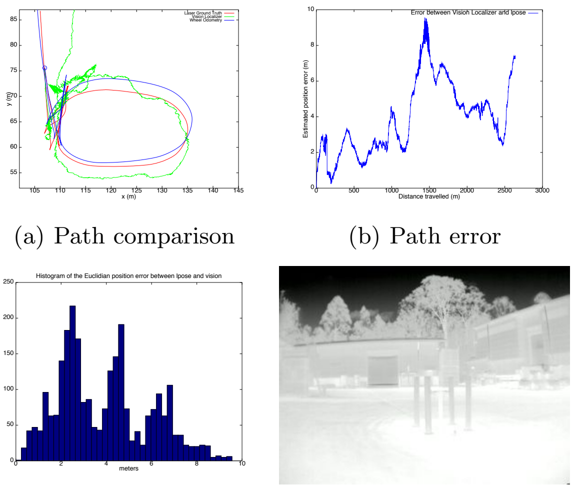

(a) Path comparison (b) Path error

(c) Error histogram (d) Camera view

Figure 19: Performance of the thermal camera localiser with site and vehicle lighting.

thermal rows in Table 1, in particular where the particle filter diverges (e.g. Figure 18). Adjusting the filtering parameters does not aid significantly, as it causes an excessive number of "false" edges, which are not included in the 3D model. It is also clear in the figures that the different lighting methods did not affect the thermal camera images, which is to be expected. This is evident from the high standard deviation and by analysing the cross-track error compared to the ground truth. Even though the values in Table 1 appear as though some lighting scenarios gave reasonable performance, in practice the resultant position estimates were variable since the thermal images were not clear enough for the vision-based localiser.

All figures show that the results are degraded the most whilst the vehicle was cornering. The main reason is due to the amount of time the sensor needs for the incident light to charge each pixel, and under fast rotation the image quality decreases compared to straight driving.

All the other cameras tested (reported in Section 3.3) did not yield any meaningful results, with the algorithm diverging from the start. Therefore, those results are not reported.

5 Conclusions

We have evaluated different types of common lighting available to a vehicle working around an industrial site, for it to be able to localise itself using onboard image-based localisation. A suitable vision-based localisation method that has demonstrated good performance in the test environment under daytime illumination, including sunny, cloudy and dusty days, was used. The input image stream came from an off-the-shelf visible-light monochrome camera and a thermal infrared camera. The different lighting conditions included the faint glow from a computer monitor on the side of the vehicle, vehicle operating lights, and site lights. Tests under different lighting conditions were conducted during night time autonomous vehicle operations at an industrial site. In general, the visible-light camera showed the best performance with only vehicle lights and no site lights. This is mainly due to the more consistent lighting across the scene from the local source, rather than large point sources on infrastructure lighting which effect the camera's exposure compensation across the image which tends to provide more localised bright and dark areas. Depending on the vehicle operation and local environment at the time, this setup could be sufficient for autononomous control of the vehicle. One aspect of future work will be to examine the possibility.

As expected, the camera selection has a critical effect on the algorithm's performance. A low sensitivity camera provides sufficient image information. Before converging to the Basler sca1400-gm, several tests were made with a number of different cameras, most of which were not satisfactory. Our tests indicated that sensors with a relatively high quantum efficiency (above 60%) and low dark noise can potentially be considered candidates for mobile robotic platforms operating in relatively low light outdoors. This, of course, depends on the nature of the algorithm and the task in hand. The thermal camera's performance was noticeably poor which highlights the need to tune the localisation technique to the camera. In this case, either the localisation algorithm needs to be tuned to allow the fuzzier edges appearing in the image to be tolerated, or a preprocessing step introduced to sharpen the edges. These possibilities, and the possibility of fusing the different image streams will also be investigated as part of future work.

References

[Bay *et al.*, 2006] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Computer Vision–ECCV 2006*, pages 404–417. Springer, 2006.

[Biber *et al.*, 2005] Peter Biber, Tom Duckett, et al. Dynamic maps for long-term operation of mobile service robots. In *Robotics: science and systems*, pages 17–24, 2005.

[Borges *et al.*, 2010] Paulo Borges, Robert Zlot, Michael Bosse, Stephen Nuske, and Ashley Tews. Vision-based localization using an edge map extracted from 3d laser range data. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 4902–4909. IEEE, 2010.

Table 1: Summary of the accuracies (in metres) of different lighting conditions and cameras on localisation.

	Camera	Figure	Vehicle Lights	Mean	Median	Stdev
No site lights	Visible	9	Day	0.41	0.40	0.05
		11	None	1.03	1.09	0.31
		12	Vehicle	0.47	0.35	0.15
	Thermal	13	None	3.48	2.7	9.86
		14	Vehicle	3.9	3.6	8.4
Site lights	Visible	16	None	0.52	0.43	0.11
		17	Vehicle	0.55	0.5	0.09
	Thermal	18	None	12.33	3.72	133.21
		19	Vehicle	3.91	3.7	4.0

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

[Glover *et al.*, 2010] Arren J Glover, William P Maddern, Michael J Milford, and Gordon Fraser Wyeth. Fab-map+ ratslam: appearance-based slam for multiple times of day. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 3507–3512. IEEE, 2010.

[Irie *et al.*, 2011] Kiyoshi Irie, Tomoaki Yoshida, and Masahiro Tomono. A high dynamic range vision approach to outdoor localization. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 5179–5184. IEEE, 2011.

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

[Lowe, 2004] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[Milford and Wyeth, 2012] Michael J Milford and Gordon Fraser Wyeth. Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1643–1649. IEEE, 2012.

[Nuske *et al.*, 2008] Stephen Nuske, Jonathan Roberts, and Gordon Wyeth. Outdoor visual localisation in industrial building environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Pasadena, USA, May 2008.

[Nuske *et al.*, 2009] Stephen Nuske, Jonathan Roberts, and Gordon Wyeth. Robust outdoor visual localization using a three-dimensional-edge map. *Journal of Field Robotics*, 26:728–756, September 2009.

[Pradalier *et al.*, 2008] C. Pradalier, A. Tews, and J. Roberts. Vision-based operations of a large industrial vehicle: Autonomous hot metal carrier. *Journal of Field Robotics*, 25(4-5):243–267, 2008.

[Reitmayr and Drummond, 2006] G. Reitmayr and T. Drummond. Going out: robust model-based tracking for outdoor augmented reality. In *International Symposium on Mixed and Augmented Reality*, 2006.

[Rofer and Jungel, 2003] Thomas Rofer and Matthias Jungel. Vision-based fast and reactive monte-carlo localization. In *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, volume 1, pages 856–861. IEEE, 2003.

[Tabuse *et al.*, 2011] Masayoshi Tabuse, Toshiki Kitaoka, and Dai Nakai. Outdoor autonomous navigation using surf features. *Artificial Life and Robotics*, 16(3):356–360, 2011.

[Valgren and Lilienthal, 2010] Christoffer Valgren and Achim J Lilienthal. Sift, surf & seasons: Appearance-based long-term localization in outdoor environments. *Robotics and Autonomous Systems*, 58(2):149–156, 2010.

[Vedaldi and Soatto, 2008] Andrea Vedaldi and Stefano Soatto. Features for recognition: Viewpoint invariance for non-planar scenes. In *IEEE International Conference on Computer Vision*, 2008.

[Weiss *et al.*, 2007] Christian Weiss, Hashem Tamimi, Andreas Masselli, and Andreas Zell. A hybrid approach for vision-based outdoor robot localization using global and local image features. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2007.

[Zhou *et al.*, 2003] Chao Zhou, Yucheng Wei, and Tieniu Tan. Mobile robot self-localization based on global visual appearance features. In *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, volume 1, pages 1271–1276. IEEE, 2003.