



Estimation of trailer off-tracking using visual odometry

Christopher de Saxe & David Cebon

To cite this article: Christopher de Saxe & David Cebon (2018): Estimation of trailer off-tracking using visual odometry, Vehicle System Dynamics, DOI: [10.1080/00423114.2018.1484498](https://doi.org/10.1080/00423114.2018.1484498)

To link to this article: <https://doi.org/10.1080/00423114.2018.1484498>



© 2018 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 18 Jun 2018.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Estimation of trailer off-tracking using visual odometry

Christopher de Saxe ^{a,b} and David Cebon ^a

^aDepartment of Engineering, University of Cambridge, Cambridge, UK; ^bCouncil for Scientific and Industrial Research, Pretoria, South Africa

ABSTRACT

High-capacity vehicles have been shown to be highly effective in reducing emissions associated with road freight transport. However, the reduced manoeuvrability of long vehicles often necessitates the use of active trailer steering. Path-following trailer steering systems are very effective in this regard, but are currently limited to on-highway applications due to the manner in which trailer off-tracking is estimated. In this work, a novel trailer off-tracking measurement concept is introduced which is independent of wheel slip and ground surface conditions, and requires no additional sensor measurements or parameter data from the tractor. The concept utilises a stereo camera pair affixed to the trailer and a visual odometry-based algorithm to calculate off-tracking. The concept was evaluated in detailed simulation and full-scale vehicle tests, demonstrating its feasibility and highlighting some important characteristics. RMS measurement errors of 0.11–0.12 m (3.3–3.6%) were obtained in a challenging visual environment.

ARTICLE HISTORY

Received 8 January 2018
Revised 17 May 2018
Accepted 28 May 2018

KEYWORDS

Off-tracking; trailer steering;
visual odometry; stereo
vision; articulated vehicles

1. Introduction

Logistics efficiency and greenhouse gas emissions are pressing research areas today as populations grow, economies strive to be more competitive, and as the world aims to meet demanding greenhouse gas emissions reduction targets. Freight transport, particularly road freight transport, is fundamental to all of these issues. In the UK in 2014, 185 billion tonne kilometres of inland freight was moved, and of this 74% was transported by road on heavy goods vehicles (HGVs) [1]. Domestic transport accounts for approximately 20% of the UK's total CO₂ emissions, and HGVs account for 21% of this or 4.2% of the total [2]. The case is similar in Europe, with road transport responsible for about 20% of total CO₂ emissions, approximately one quarter of which is attributable to trucks and buses [3].

Both the logistics efficiency and emissions associated with HGVs in performing a given freight task are a direct function of vehicle efficiency. High-capacity vehicles (HCVs) have increased capacity due to increased permissible mass or dimensions relative to conventional HGVs. They are normally longer and have more articulation points than conventional HGVs. The use of HCVs has been trialled and implemented in a number of countries including Australia, South Africa and parts of Europe [4–6], and proved to be

CONTACT Christopher de Saxe  [cgsaxe@csir.co.za](mailto:cdsaxe@csir.co.za)

a highly effective means of addressing efficiency and emission challenges in road freight transport.

The introduction of such vehicles is not without its challenges, chief among which in Europe and the UK is their reduced manoeuvrability. The Longer Semi-Trailer (LST) trial in the UK [7] is studying the benefits of semi-trailers which are 2.05 m longer than conventional semi-trailers. However, all participating vehicles must still meet the UK's 'roundabout test' requirement, and this almost certainly necessitates trailer steering technology.

1.1. Off-tracking and trailer steering

Off-tracking primarily refers to the 'cut-in' behaviour of a truck or trailer as it navigates a turn. Off-tracking is illustrated in Figure 1(a) for a tractor semi-trailer. In this research, off-tracking is measured as the deviation of the rear of the trailer relative to the path of the hitch point, measured laterally to the longitudinal trailer axis.

Conventional trailer axles are unsteered which limits manoeuvrability. A number of passive and active trailer steering technologies exist to address this. Passive systems have limited scope for improving manoeuvrability, and often result in unwanted side-effects such as increased 'tail swing' or degraded stability at high speeds [8,9]. In comparison, active steering systems can approach optimal manoeuvrability performance without unfavourable transient effects and have been shown to improve high-speed stability and safety performance [9–11].

Early work on active steering by Hata et al. [12] and Notsu et al. [13] introduced the idea of the 'path-following' steering concept. Building on this, Jujnovich and Cebon [14] developed an active path-following semi-trailer steering system in which a 'follow-point' at the rear of the trailer follows the path of the lead point (fifth wheel) for all speeds and paths. This is illustrated in Figure 1(b). Their strategy for achieving this was to compare heading angles at the front and rear of the trailer using a non-linear control strategy. A 'model-matching technique' requiring a reference trailer model was used to address problems associated with steering angle saturation at low speeds. The controller required knowledge of vehicle parameters for both tractor and trailer, as well as measurements of speed,




Figure 1. Trailer off-tracking: (a) unsteered trailer, (b) steered trailer with path-following control.




Figure 2. Trailer steering performance in a roundabout manoeuvre (modified from [15]).

articulation angle and trailer yaw rate. The performance of the path-following concept compared with a conventional trailer and command steer is shown in Figure 2, highlighting its superior cut-in and tail swing behaviour.

An alternative (linear) strategy was proposed by Cheng et al. [10,16], using a ‘virtual driver’ located at the rear of the trailer and a simpler proportional-integral-derivative controller to estimate off-tracking at low speeds. The system was also capable of controlling the vehicle at high speeds, minimising a cost function with respect to both high-speed stability and off-tracking. The controller required knowledge of various tractor and trailer parameters, as well as measurements of tractor speed and steer angle, trailer yaw rate and articulation angle in order to estimate trailer off-tracking, and hence calculate the required trailer steer angles. The off-tracking estimation method also assumed zero wheel side-slip at low speeds.

By basing the control on off-tracking error instead of heading angles, Cheng’s approach avoids possible steady-state tracking errors that are possible in Jujnovich’s approach. Further, Cheng’s controller does not require a reference trailer to handle steer angle saturation. However, like Jujnovich’s controller, it requires sensor measurements from the tractor as well as knowledge of tractor parameters for the vehicle model. In addition, the assumption of zero slip limits the application of the system to high friction (on-highway) conditions. This is a reasonable assumption for many applications, but is limiting for icy conditions on the highway and for off-highway applications such as logging, livestock and dairy collections from farms, and military convoys. Off-highway applications are prone to low friction, cambered and inclined roads, and tight corners, giving rise to potentially large wheel slip.

The ability to measure trailer off-tracking, in a manner which is independent of friction conditions and requires no parameter data or sensor measurements from the tractor unit, would improve the commercial prospects of path-following trailer steering systems. Furthermore, such a solution could find application in other areas requiring a versatile means of measuring trailer off-tracking. For example, it could be used to measure the swept

path performance of heavy vehicles for certification with performance-based regulatory schemes, such as those in Australia and South Africa [4,5].

1.2. Trailer off-tracking sensing

Miao and Cebon [17] demonstrated the limitations of Cheng's controller in simulations of roundabout manoeuvres with camber and low friction. Miao [18] then showed how the controller could be reformulated to take an independent measurement of off-tracking, i.e. relaxing the no-slip assumption and subsequently proposed a 'ground-watching navigation system' (GWNS) for independent off-tracking estimation.

The ground-watching concept was demonstrated for a tractor semi-trailer combination using two downwards-facing cameras mounted beneath the semi-trailer, one near the fifth wheel and one near the trailer follow-point. Miao used FAST features [19] and SURF descriptors [20] to detect features in each image of the road surface. Feature matching between front and rear camera images was achieved using the FLANN algorithm [21]. A RANSAC scheme was adopted to calculate the planar homography from the feature matches, and single value decomposition was used to determine planar rotation and translation from the homography. A variation of the system that used a single camera was also presented for instances when off-tracking was too large for the dual camera system to function.

Performance of the system was demonstrated in vehicle tests on dry tarmac with a tractor semi-trailer. Off-tracking measurement errors of 0.05 m were obtained at 10 Hz in open-loop tests (i.e. with non-steered trailer axles and no control). In closed-loop tests, using measurements from the GWNS as inputs to Miao's modified path-following controller, path-following errors of less than 0.1 m were obtained.

Although the system addressed the issue of wheel slip dependence, it assumed an unchanging and planar road surface with static features. In conditions where wheel slip is likely (i.e. where there is mud, standing water, ice, snow), it is also likely that the road surface is non-planar and displaces with the passing of the vehicle, and so may be unsuitable for the ground-watching concept. Furthermore, cameras mounted beneath a trailer in these conditions would also be very prone to dirt and water splash which could significantly affect the performance of feature extraction and matching.

This paper develops an off-tracking measurement concept which is independent of tractor measurements or parameter knowledge and wheel slip conditions, and which makes no assumptions regarding the visual properties of the ground surface. We present a novel camera-based concept which builds on and addresses the limitations of Cheng's path-following controller [10] and Miao's ground-watching navigation system [18].

2. Off-tracking estimation concept

The off-tracking estimation concept utilises a stereo camera pair mounted on the roof of a trailer which captures stereo images of the surrounding landscape. These images are processed using visual odometry to determine the motion of the trailer relative to the surroundings. Cameras can be either mounted along the side or rear edge of the trailer, though rear-facing cameras cannot be used if there is an additional 'link' trailer. The open-source VISO2-S visual odometry algorithm [22] was adopted for the image processing task

and was coupled with a buffer-based off-tracking calculation algorithm. The combined VISO2-S and off-tracking calculation algorithm is denoted as ‘VISO-OT’.

Target off-tracking errors of 0.10 m RMS and 0.15 m maximum are sought. This is in line with previous work on path-following control [14,23]. These are deemed to be reasonable upper bounds for obstacle avoidance and control stability. For comparison, in the Australian performance-based standards scheme for heavy vehicle certification, vehicle swept path measurements are permitted an error of up to 0.1 m [4].

2.1. Visual odometry

Visual odometry is the estimation of the pose and motion of a camera through a 3-D scene. Advances in visual odometry algorithms have resulted in its widespread use for autonomous road vehicles and mobile robotics. Compared to other odometry systems such as wheel speed sensors and GPS, ¹ visual odometry offers high precision, low-cost hardware and independence from traction conditions. Although vehicle-based visual odometry is commonplace in autonomous vehicles (see, for example, [22,24]), little work has been done with heavy vehicles, with the exception of recent work by Miao, Harris, de Saxe and Cebon [18,25–27].

The VISO2-S algorithm [22] assumes that the stereo cameras have been calibrated and the images rectified. The ‘stereo baseline’ (lateral separation between left and right cameras) must also be known. Details of the VISO2-S algorithm may be found in [22] and can be summarised into the following steps:

- (1) A stereo image pair is obtained and ‘corner-like’ features are detected in each image.
- (2) ‘Circular’ feature matching is performed, comparing features between left and right images (normal stereo matching) as well as between current and previous image pairs. Features are accepted if matching succeeds through the entire ‘circular’ loop of four images.
- (3) A ‘bucketing’ process [28] divides the images into a rectangular grid, and each ‘bucket’ may only store a maximum number of features. This ensures a good distribution of features around the image, minimising the effects of bias and of moving objects.
- (4) Ego-motion is estimated by minimising reprojection errors through Gauss-Newton optimisation with respect to the rotation matrix R and translation vector T .
- (5) The ego-motion estimation incorporates a RANSAC strategy to remove outliers.
- (6) A constant acceleration Kalman Filter is used to minimise noise.

In performance tests on the KITTI data set, the VISO2-S algorithm was shown to yield 2.44% translation error and $0.0114^\circ/\text{m}$ rotation errors on average with a 0.5 m stereo baseline [29]. The algorithm runs at 20 fps on a single processing core using the recommended parameter settings. Using a representative trailer length of 14 m (from fifth wheel to follow-point), translation drift of 2.44% would result in approximately $0.0224 \times 14 \text{ m} = 0.3136 \text{ m}$ of lateral off-tracking between the fifth wheel and the rear of the trailer.

It is expected that the accuracy of VISO2-S can be improved nearer to the target maximum error of 0.15 m by using a larger stereo baseline. Increasing the baseline of a stereo camera pair can improve depth accuracy and hence odometry accuracy [30] but at the expense of increased difference in perspective. The distance to scenery in this application

was deemed suitably large for relatively large baselines to be considered without the need for specialised wide-baseline stereo vision algorithms (see, for example, [30–32]). The maximum practical baseline would be approximately 2.5 m for rear-facing cameras, constrained by maximum vehicle width. However, other practical limitations on maximum baseline exist, such as the ease of stereo camera calibration and the rigidity of the camera mounting.

2.2. Off-tracking calculation

Figure 3(a) illustrates step-by-step (or frame-by-frame) yaw-plane motion of a tractor semi-trailer combination at frames i , $i-1$, $i-2$, $i-3$ and so on. The camera origin is arbitrarily assumed to be at distance a behind the fifth wheel and b ahead of the follow-point.




Figure 3. Off-tracking calculation process from visual odometry data. (a) Camera trajectory from visual odometry. (b) Fifth wheel trajectory. (c) Rotated trajectory in i th reference frame. (d) x-y co-ordinates in i th reference frame.

Raw visual odometry data are in the form of a rotation matrix, R , and translation vector, T , at each frame, where R and T are relative to the prior vehicle location, and in that vehicle's reference frame. The yaw-plane components of these data are shown in the figure as Δx , Δy and $\Delta\psi$ (incremental translation and yaw angle). Pitch and roll motions were assumed to be negligible relative to motion in the yaw plane. By making use of a data buffer or shift register, these data were used to calculate off-tracking.

First, depending on whether the cameras are rear-facing, side-facing, or at some intermediate angle between, translational data from the cameras must be rotated and aligned with the vehicle co-ordinate frame. To do this the camera-to-vehicle rotation parameter is defined as ψ_{c2v} . For perfectly rear-facing cameras, $\psi_{c2v} = 180^\circ$, and for cameras pointing to the right, $\psi_{c2v} = 90^\circ$. The co-ordinates can be rotated according to

$$\begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_i = \begin{bmatrix} \cos(\psi_{c2v}) & \sin(\psi_{c2v}) \\ -\sin(\psi_{c2v}) & \cos(\psi_{c2v}) \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_{cam,i}. \quad (1)$$

The motion data from the cameras can then be transformed into fifth wheel motion (denoted by the subscript F for 'front') as follows:

$$\begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_{F,i} = \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_i + a \cdot \begin{bmatrix} \sin(\Delta\psi_i) \\ \cos(\Delta\psi_i) - 1 \end{bmatrix}, \quad (2)$$

where a is the distance from the hitch to the reference point of the cameras. This is illustrated in Figure 3(b). These data are stored in a buffer which initially grows with each new frame.

To find the fifth wheel trajectory *relative to the current vehicle reference frame*, previous data in the buffer must be rotated at each new frame. This can be done by first summing $\Delta\psi$ from the current frame to each previous frame as follows:

$$\psi_{i-k}^{(i)} = - \sum_{j=1}^k \Delta\psi_{i-j}^{(i-j)}, \quad k = 1, 2, 3, \dots \quad (3)$$

The counter, k , denotes the entry in the buffer, increasing from the current frame ($k=1$) backwards. The superscript (i) denotes the co-ordinate frame (i.e. that of vehicle i).

For buffer entry k , and hence frame $i-k$, a rotation transformation from frame $i-k$ to frame i can then be applied to obtain Δx and Δy in the current reference frame:

$$\begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_{F,i-k}^{(i)} = - \begin{bmatrix} \cos(\psi_{i-k}^{(i)}) & \sin(\psi_{i-k}^{(i)}) \\ -\sin(\psi_{i-k}^{(i)}) & \cos(\psi_{i-k}^{(i)}) \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_{F,i-k}^{(i-k)} \quad (4)$$

which is illustrated in Figure 3(c).

The x and y locations of each data point in the current vehicle reference frame (centred on the fifth wheel) may then be found by simple summation of Δx and Δy data:

$$\begin{bmatrix} x \\ y \end{bmatrix}_{F,i-k}^{(i)} = \sum_{j=0}^{k-1} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}_{F,i-k}^{(i)}. \quad (5)$$

This result is illustrated in Figure 3(d).

At a given frame i , (x, y) data of fifth wheel path history points can be used to find the off-tracking at the rear of the trailer. Assuming the dimension $a+b$ to be known, the most recent path history point for which

$$y_{F,i-k}^{(i)} \leq -(a+b) \quad (6)$$

denotes the point beyond which the fifth wheel trajectory is past the rear of the trailer.

The value of k at which this condition is met is denoted k_b or the buffer length. The buffer only needs to retain k_b data points, with older data being discarded and the most recent data added at each new frame in the data buffer.

The off-tracking, e_{tr} , may be calculated by linear interpolation as follows:

$$e_{tr} = x_{F,i-(k_b-1)}^{(i)} + \left(x_{F,i-k_b}^{(i)} - x_{F,i-(k_b-1)}^{(i)} \right) \left(\frac{y_{F,i-(k_b-1)}^{(i)} + (a+b)}{y_{F,i-(k_b-1)}^{(i)} - y_{F,i-k_b}^{(i)}} \right). \quad (7)$$

If at least one trailer length has passed, the above calculation of off-tracking will be possible. If not, the buffer will not contain sufficient data for off-tracking to be calculated. In this case the next frame is obtained, and the buffer simply grows in size until sufficient data are available and off-tracking can start to be calculated. Thereafter, the buffer size will vary with the speed and path of the vehicle, only keeping data old enough to calculate off-tracking. When a vehicle stops, it is possible to store the contents of the buffer. So in theory, the calculation of a semi-filled buffer should only occur the first time a new vehicle combination is launched.


Integration drift is inherent in visual odometry systems as a result of summing incremental motion data in the above fashion. In typical automotive and mobile robotics applications, these data are integrated indefinitely with time to obtain global position estimates, so the integration errors can grow unbounded. However, in this application integration is performed only from the fifth wheel to the rear of the trailer and so integration drift is bounded by the length of the trailer and will not grow indefinitely with time. Further, the effects of any outliers in the visual odometry data will be removed from the buffer after one trailer length has passed.

3. Simulations


3.1. Simulation overview

As a proof-of-concept, and to assess the theoretical accuracy of the VISO-OT system, simulations were carried out in a virtual 3-D environment using Autodesk Inventor [33]. Autodesk Inventor includes functionality to generate animations using a simulated perspective camera, which represents an idealised pin-hole camera with no lens distortion and with accurately known parameters. This functionality has proven to be useful in evaluating computer vision algorithms in a simulated environment under ‘perfect’ conditions, thus determining an upper bound on the achievable performance in the field [18,25]. A virtual stereo camera pair was made to travel through a visually representative road and roadside environment, simulating the motion of an off-tracking trailer.

The ‘virtual environment’ is shown in Figure 4(a). Road width was set to 5 m (within the UK rural road design guidelines), and the size of roadside objects and textures were



(a)



(b)

Figure 4. CAD simulation environment: (a) Perspective view, (b) plan view of constant side-slip manoeuvre.

chosen to be representative. Soft ambient lighting and shadows were incorporated and all scenery was stationary. A stereo camera pair was made to follow a straight path along the road at a given slip angle, representative of a trailer moving with constant off-tracking. This is illustrated in Figure 4(b). Stereo images from the cameras were rendered at 10 fps as the trailer travelled 100 m along the path at a constant speed of 5 m/s. VISO-OT was then used to process the image sequences and estimate off-tracking.

Images were captured at a resolution of 1344×391 using virtual cameras with a 100° field of view. The cameras were located 3 m above ground level with zero tilt or roll angles relative to the ground. The simulated stereo cameras adhered to the simple pin-hole camera model with no distortion, and so could be easily calibrated from knowledge of the image size and field of view.

The default input parameters for VISO2-S were used, including 50 RANSAC iterations per optimisation, outlier flow and disparity thresholds of 5 pixels, a bucket size of 50×50 pixels and a maximum of two features per bucket. For off-tracking calculations, the trailer length from fifth wheel to follow-point ($a+b$ in Figure 3) was taken to be 14 m. This is approximately representative of a UK LST. Given the trailer side-slip angle and knowing




Figure 5. Stereo images at a slip angle of 10° and a baseline of 2.5 m. (a) Rear-facing cameras. (b) Side-facing cameras.

the trailer length, the ‘ground-truth’ off-tracking was known accurately in each case for comparison.

Investigations into the sensitivity of the system were carried out using batch simulations, with regard to variations in stereo baseline, camera location (side/rear), trailer slip angle and scenery density. All CAD and visual odometry processing was performed on a 3.2 GHz 6-core desktop computer.

3.2. Results

Example views from the stereo camera pair are shown in Figure 5 for both rear and side-facing camera cases. Locations of matched features are shown in each image. The example case shown is for a trailer side-slip angle of 10° and a baseline of 2.5 m.

3.2.1. Stereo baseline and camera location

The effect of stereo baseline was studied for a reference scenario of 14° trailer slip and a baseline of 0.5 m. A range of baselines was considered from 0.5 to 5.0 m in 0.5 m increments. Rear and side-mounted camera scenarios were considered. (Baselines of > 2.5 m are possible for side-mounted cameras, but are not practical for rear cameras.)

Results are shown in Figure 6. The metrics shown include RMS and mean off-tracking measurement errors, measured over the duration of each simulation run, as well as the mean number of feature matches and the percentage mean inlying feature matches from VISO2-S, indicating the relative performance of the underlying visual odometry measurements.

A clear reduction in feature matching is evident with increasing baseline, as expected due to the increasing difference in perspective of the two cameras. However, little effect on errors is evident until a baseline of approximately 2.5 m. Until this point, RMS errors are in the region of 0.01–0.04 m (0.3–1.2%), and mean errors in the order of 0.01 m.

The reduction in the number of feature matches and increase in errors are more pronounced for the side-facing cameras. This is expected due to the smaller average feature

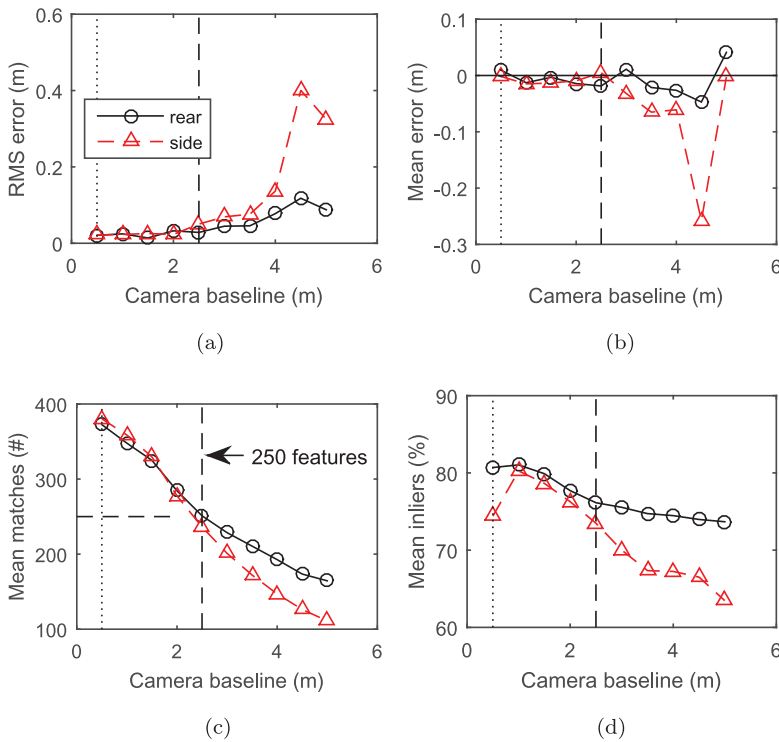


Figure 6. Sensitivity to the stereo baseline for 14° slip (reference = 0.5 m). (a) RMS errors. (b) Mean errors. (c) Mean no. of matches. (d) Mean no. of inliers.

depth which is less suited to wider baselines. The point at which errors grow (at a baseline of 2.5 m) appears to correlate with the number of feature matches dropping below 250 (shown in the figure).

At reasonable baselines (below 2.5 m), the errors for rear and side cameras are comparable. However, for higher baselines, the side cameras experience a significant increase in RMS errors compared with the rear-facing cameras. The average feature depth for the side cameras is lower, and an increased baseline yields more difficult feature matching for less distant features. So an increased baseline will have more effect on the side-facing cameras. Rear cameras therefore offer better accuracy and consistency over the anticipated operating range of off-tracking.

3.2.2. Trailer slip angle

In addition to the reference case of 14° trailer slip (about 3.5 m off-tracking), a wide range of other slip angles and hence magnitudes of off-tracking were investigated. Slip angles of between 0° and 60° were considered. Although 60° of side-slip yields an unrealistic off-tracking of 24 m (as per the definition in Figure 1(a)), these higher slip scenarios were included for possible additional insights.

The results are shown in Figure 7 for a rear-mounted camera configuration. Results for three baselines (0.5 m, 1.5 m, 2.5 m) are shown for comparison. For realistic slip angles

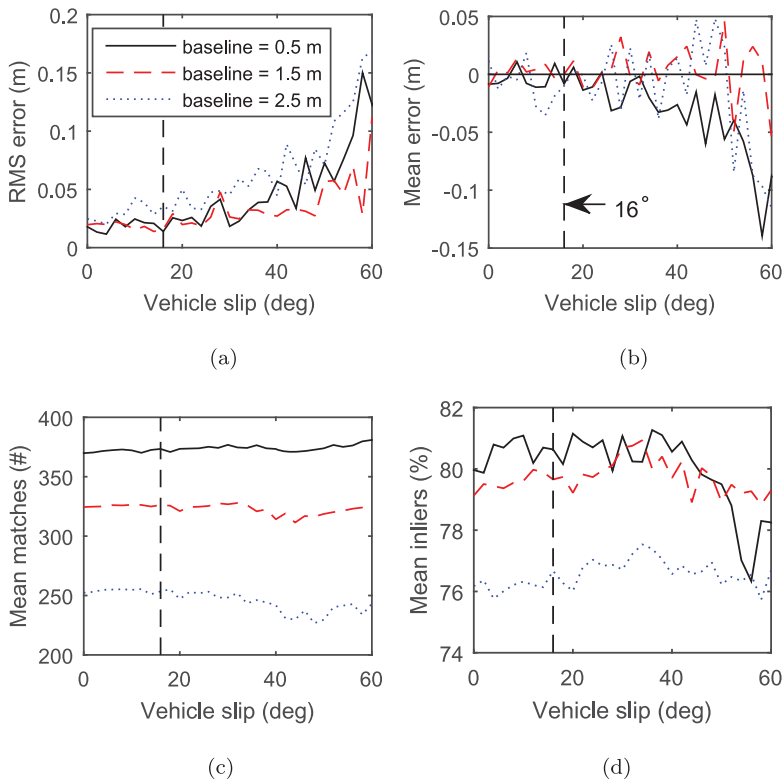


Figure 7. Variation in performance for rear-facing cameras as a function of vehicle slip angle and stereo baseline. 16° trailer slip (4 m off-tracking) shown as dashed line. (a) RMS errors. (b) Mean errors. (c) Mean no. of matches. (d) Mean no. of inliers.

below 16° (4 m off-tracking), the effectiveness of the feature matching is seen to be approximately insensitive to slip angle, but clearly sensitive to baseline. RMS errors for the 0.5 and 1.5 m baselines are comparable, and range from 0.01 to 0.03 m. 2.5 m baseline errors are consistently higher between 0.02 and 0.04 m. Mean error trends are similar, with a maximum bias of 0.01 m for the 0.5 m and 1.5 m baselines, rising to up to 0.035 m for the 2.5 m baseline case.

The mean number of feature matches for the 0.5m and 1.5 m baselines are approximately 380 and 325, well above the threshold of 250. However, for the 2.5 m baseline, the number of matches is on the limit of this threshold, sometimes dropping below it. This explains the larger RMS errors for the 2.5 m baseline, compared to those for the 0.5 and 1.5 m cases.

These results are promising, given the sub-0.15 m target accuracy. However, they are only an indication of the limit of performance in an idealised environment, with perfectly calibrated cameras, and consistent lighting and scenery conditions.

3.2.3. Scenery density

Sensitivity of the measurements to the density of the surrounding 3-D scenery was assessed in order to give insights into real world performance under low feature environments.

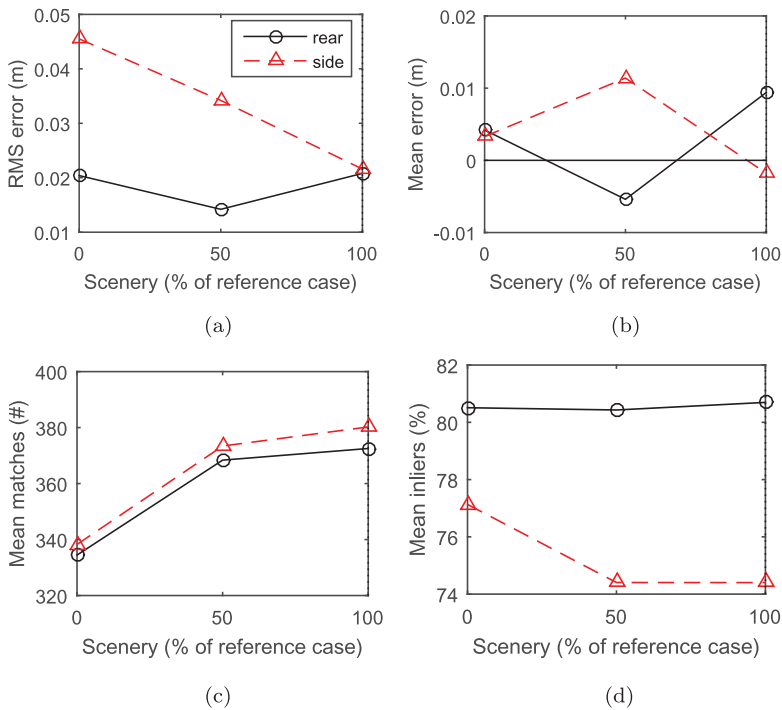


Figure 8. Sensitivity to the density of the roadside scenery (reference = 100%). (a) RMS errors. (b) Mean errors. (c) Mean no. of matches. (d) Mean no. of inliers.

Three scenarios were considered. The reference case was deemed to have ‘100%’ of the available scenery. A 50% case was then considered, wherein 50% of the 3-D objects were removed from the scene, at an approximately even distribution in the scene. Lastly, a 0% scene was considered in which all the 3-D scenery was removed, leaving only the ground textures.

The results for the three scenery scenarios are shown in Figure 8. As expected, the number of features drops with a reduction in scenery, for both rear and side cameras. The effect on errors, however, is distinctly different between the rear and side cameras. The rear cameras appear almost insensitive to the changes, which is a significant finding. This suggests that in a typical off-highway scene such as large crop fields or pastures, the algorithm is still able to detect and match sufficient features to maintain accuracy. This suggests that the most important feature points are ground-based, and so having a good view of the ground and road surface will be beneficial for performance.

In the case of the side cameras, however, there is a twofold increase in RMS errors when the scenery is removed. Because the number of features in both rear and side cases is above 250, performance here is largely dictated by the average depth to feature (depth errors are smallest for closer features). With side cameras, the average feature depth changes more when the scenery is removed than in the case with rear cameras. For the side cameras, the scenery is closer to begin with and more dominant in the field of view. So the effect of removing it is more severe, resulting in the observed rise in errors.




Figure 9. (a) Side-camera off-tracking results, 5° slip, 1.5 m baseline; (b) visual odometry data for (a). (The camera optical axis is along the Z_c -axis, which is in the direction of off-tracking measurement for side-facing cameras.)

3.2.4. Bounded drift

Instances of temporary error drift were observed in some results, for example, in Figure 9(a) (side cameras, 5° slip, 1.5 m baseline). Drift develops in region ‘A’, resulting in a constant error in region ‘B’. The source of this can be seen in Figure 9(b), in the visual odometry data in the camera Z_c -direction, where Z_c in this case is in the direction of off-tracking measurement (as is the case with the side cameras). While the data exhibit predominantly zero-mean noise, in region ‘A’ there is a distinct sequence of biased outliers (circled) relative to the dashed reference value.

The sum of the magnitudes of these outliers equates to a cumulative error of about 0.08 m, which is comparable to the observed off-tracking error in region ‘B’. The effect disappears after approximately one trailer length (14 m) has passed as the corrupting data points are discarded from the buffer.

Although ideally no drift would be present, the temporary or ‘bounded’ nature of the drift here is a significant benefit over using visual odometry methods for global positioning estimates. In this case, the drift was well below the maximum acceptable error of 0.15 m, but only in these idealised conditions. Depending on the tolerances of the application, it would be desirable to try to detect or otherwise minimise drift in practical applications of the system.

3.3. Summary

The simulation experiments have illustrated the theoretically achievable accuracy of the off-tracking measurement concept, with RMS errors around 0.01–0.04 m for representative conditions with ideal cameras. They have also highlighted some useful characteristics

including sensitivity to stereo baseline and camera location, robustness to low density scenery and drift behaviour which is bounded by the data buffer as a function of trailer length.

4. Field tests

4.1. Experimental setup

Full-scale field tests were conducted on a tractor semi-trailer vehicle combination. The combination is shown in Figure 10, and consisted of a 3-axle ‘B-link’ semi-trailer and a Volvo FH12 6×2 tractor. The tag axle of the tractor was permanently lifted during testing, as well as the front axle of the semi-trailer, making the combination effectively a 4×2 tractor with a 2-axle semi-trailer.

Two Point Grey Flea3 USB 3.0 cameras with wide-angle lenses were used for the stereo visual odometry, mounted to a custom machined aluminium mounting frame. The mount allowed for multiple camera baselines to be used, from 0.1 m up to 1 m in increments of 0.1 m, and was mounted to the top of the semi-trailer using strong magnets. The magnets attached the mounting frame to the top of the vehicle, enabling easy removal of the mount to locate it at the rear or side of the trailer, or to adjust the baseline of the cameras. When mounted to the top of the semi-trailer, the cameras were approximately 4.2 m above ground. The mounting locations of the cameras are shown in Figure 10.

The instrumentation and communication systems are shown in Figure 11. A Linux-based computer was mounted inside the trailer for image capture and processing, and connected to the cameras via USB 3.0. A laptop computer in the cabin of the tractor was connected to the image processing unit via ethernet. This was used to remotely access the Linux computer, to send commands to start and stop test runs, and to change camera and run parameters as needed.




Figure 10. Test vehicle combination with cameras (dimensions in mm).




Figure 11. Vehicle instrumentation layout (cameras shown in rear-facing configuration).

A trailer-mounted RT3022 inertial and GPS navigation unit was used to obtain measurements of global trailer position and heading angle, from which a ‘ground-truth’ measurement of off-tracking was calculated. The RT3022 contains precision accelerometers, gyroscopes and GPS receivers, with dual GPS antennae to improve heading measurement accuracy. A GPS base station (the ‘RT GPS-Base-2’) was used to provide differential corrections, improving position measurements to a quoted accuracy of 20 mm RMS.

The RT3022 unit was rigidly mounted to the floor inside the trailer, aligned with the longitudinal axis of the trailer but offset laterally from the centre line due to practical constraints. The primary and secondary antennae were mounted on the roof of the trailer with a longitudinal separation of 2.342 m, also limited by practical constraints. CANbus messages from the RT3022 were logged via a PCAN-USB on the Linux computer. Ground-truth off-tracking calculations using RT3022 measurements were carried out using a similar buffer-based approach to the VISO-OT algorithm. Maximum RMS errors using this method were estimated to be 0.03 m. Details on the processing of RT3022 data may be found in [34].

The VISO-OT algorithm was compiled to the Linux computer from C++ in much the same form as used for the simulations. Minor additional functionality was added to account for lens distortion, laterally offset cameras, RT3022 CAN message logging and adjustable camera and image parameters. The algorithm ran at 10 fps. Individual cameras were calibrated according to the methodology of Zhang [35] and the distortion model of Heikkilä [36]. Stereo calibration was carried out using the method of Fusiello et al. [37]. The OpenCV library [38] was used for this purpose.

Tests were conducted at Bourn Airfield near Cambridge. Figure 12 shows an aerial view of the test site, including camera views in the north-east, north-west and south-west directions. Visually, the test site consisted of an area of tarmac 40 m wide with various imperfections and potholes, surrounded by large grain fields. The grain fields provided a challenging scenario for the visual odometry algorithm. Such a scenery might be

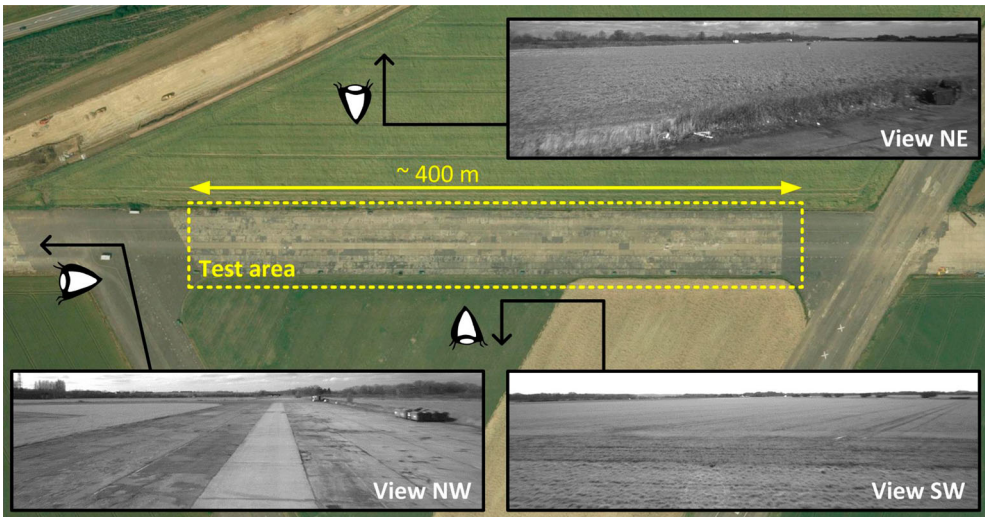


Figure 12. Aerial view of the test site at Bourn Airfield, showing camera views of the scenery in three directions.

experienced in a challenging off-highway application, and so provides a good assessment of the system's robustness.

Three types of manoeuvre were conducted, namely: 'figures-of-eight' and left and right roundabout turns (three full turns per manoeuvre). Rear camera ($\psi_{c2v} = 180^\circ$) and side-camera (left-facing, $\psi_{c2v} = 270^\circ$) configurations were assessed at baselines of 500, 700 and 900 mm. Manoeuvres were conducted at relatively low speeds, in the range of 15 km/h (during cornering) to 40 km/h (during some straight sections). Potholes and other surface features were not avoided by the driver, so as to include these disturbances in the assessment of the system's performance. Two tests of each manoeuvre were performed for each configuration.

Initial straight-line manoeuvres were carried out to determine any biases in both RT3022 and camera measurements due to mounting misalignment, and these were corrected for in subsequent test results.

4.2. Results

Example stereo image pairs overlaid with stereo feature matches are shown in Figure 13. Off-tracking measurements as a function of time are given in Figure 14 for the 500 mm stereo baseline. Results are shown for all three manoeuvres and for both rear and side-camera configurations. Results for one of the two tests per manoeuvre are shown. Included beneath each plot is the time history of the number of successful feature matches per image pair (after outlier rejection), as well as the calculated error signal between the VISO-OT and ground-truth measurements.

At the beginning of each plot, there is a period of between 50 and 100 frames without data ($\approx 5\text{--}10$ s). This is the period during which one trailer length of data was being accumulated in the buffer (for both VISO-OT and RT3022 results), and so off-tracking is not yet defined.




Figure 13. Example rear and side-camera views with detected features. (a) Rear cameras. (b) Side cameras.

Overall, off-tracking measurement accuracy was reasonable, with errors generally less than 0.5 m. However, some instances of large errors up to 1 m were observed. The number of features matched per image pair was in the region of 400–800 in most manoeuvres, but as low as 50 features in one instance (Figure 14(c)).

The number of feature matches are well above the 250 threshold identified in the simulations, below which errors were expected to rise steadily. The number of matches drops to about 200 features in some cases, without an observed effect on errors. In Figure 14(c), however, the number of feature matches drops to well below this to approximately 50 features, although there is no evidence of this having affected the measurement errors.

There is a clear periodic element to the number of feature matches in the left and right roundabout results. There appear to be either three or six cycles per manoeuvre, which is consistent with the number of turns driven per manoeuvre. At the time of testing the sun was setting in the west, and the variations in feature matching can probably be attributed to variations in lighting conditions as the cameras faced towards and away from direct sunlight. These oscillations in feature matches do not seem to have affected any measurement errors directly, as the total number of features matches has remained acceptable. These effects could be reduced in future work through the use of more advanced automatic shutter speed and exposure adjustment, compared to the manual adjustment adopted for these tests.

There is a recurring error trend in the rear-facing camera results, in that positive off-tracking ($+e_{tr}$) generally yields small errors, while off-tracking in the negative direction ($-e_{tr}$) yields a consistent bias. A seemingly similar trend is also evident in the side camera results, but is less consistent, and trends are more difficult to observe as the side-camera results are generally more erratic. The steady bias in the right roundabout test did not originate from zeroing bias, as it only appears at non-zero off-tracking magnitudes. This error trend will be revisited in the next section.

For each test, RMS and maximum errors over the course of the manoeuvre were calculated. Feature matching statistics including the mean number of features matched, and the number of these which survived the outlier rejection process, were also recorded. These results are shown in Figure 15 for the left roundabout manoeuvre, for both side and rear

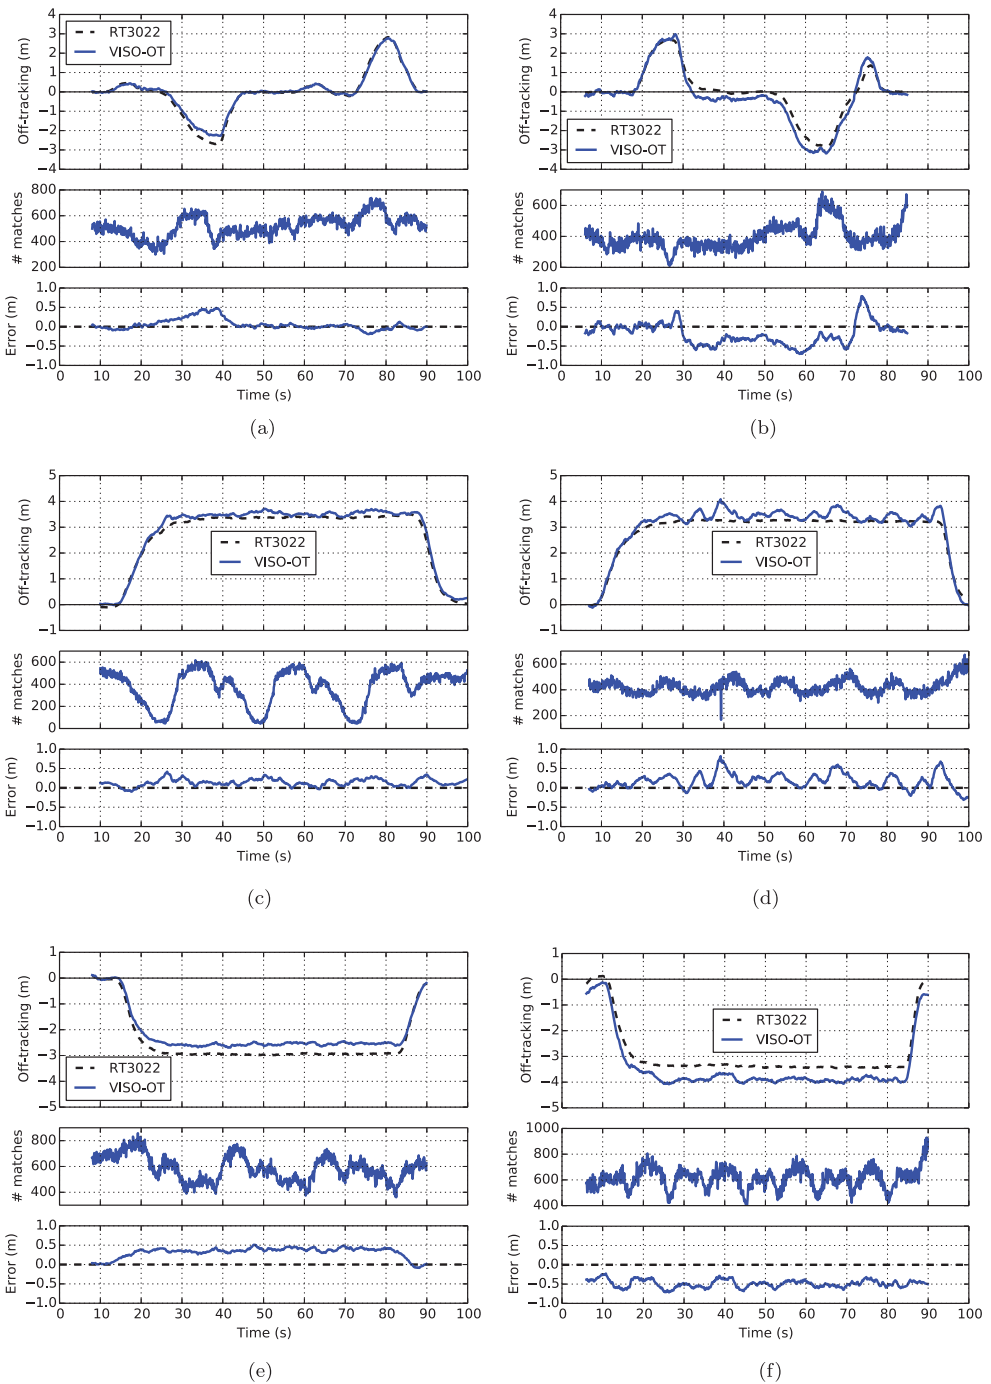


Figure 14. Off-tracking time histories with features and errors, 500mm stereo baseline (speeds in the range 15–40 km/h). (a) Rear cameras, figure-of-eight. (b) Side cameras, figure-of-eight. (c) Rear cameras, left roundabout. (d) Side cameras, left roundabout. (e) Rear cameras, right roundabout. (f) Side cameras, right roundabout.




Figure 15. Off-tracking vehicle test results summary, left roundabout. (a) RMS errors. (b) Maximum errors. (c) Mean no. of matches.

camera configurations and all three baselines. The left roundabout tests were the least affected by bias, and so provides a reasonable picture of the achieved accuracies.

Although the results are erratic as a result of the afore-mentioned oscillations and bias, in general, the rear cameras appear to provide significantly better accuracy than the side cameras. In the left roundabout tests RMS errors of 0.11–0.12 m (3.3–3.6%) and maximum errors 0.31–0.44 m were obtained for the rear cameras, and 0.14–0.59 m (RMS) (4.2–18%) and 0.36–1.16 m (maximum) for the side cameras.

There seems to be an overall increase in errors with increasing baseline for all tests. However, there appears to be a small but significant increase in feature matches with increasing baseline for both rear and side cameras. The scenery in these tests was closest to the ‘0%’ scenery case from the simulations, in which average feature depth is at a maximum. In this case, this should favour larger baselines as is observed. However, the effect is small, and the number of features is already well above the required minimum (≈ 250), and so there does not seem to be a benefit to baselines beyond 0.5 m.

4.3. Systematic errors

The consistency of the errors within each camera configuration for all three manoeuvres suggests an underlying systematic effect. It was hypothesised that these errors may have originated from imperfections in the trailer structure, causing small quasi-static pitch angle and/or roll angle deviations between the left and right cameras. Further, small yaw angle deviations could have resulted when the cameras were attached to the mount during baseline changes. These would have led to small deviations from the stereo calibration, so as to bias the visual odometry measurements in one direction or another, but not so large as to prevent feature matching entirely.

The magnitude and direction of pitch, roll and yaw angle variations may be expected to differ between stereo baselines (as cameras are moved to different positions on the mount) and side/rear camera configurations (as the mount was moved between rear and side mounting locations). However, these were expected to be consistent for all three manoeuvres for a given camera configuration if this hypothesis was correct. Additionally, the trailer structure itself could deform (twist) as a result of moments in the chassis generated during turning. This may have acted to correct for these camera rotations in left turns, but amplify them in right turns.

4.4. Camera rotation correction

To validate the camera rotation hypothesis, a small investigation was carried out in which artificial pitch, roll and yaw angle transformations were applied to right camera images in post-processing. This would ‘correct’ for the theorised rotation incurred by twisting or bending of the camera mount. If the rotation angles were chosen correctly, the systematic errors in observed off-tracking measurements should diminish.

The right camera was assumed to have rotated only about its optical centre. Roll, yaw and pitch rotation angles were denoted ϕ , ψ and θ . A ‘planar projective transformation’ was applied to the right camera images to simulate 3-D camera rotation as follows [39]:

$$\tilde{\mathbf{w}} = \mathbf{K}\mathbf{R}\mathbf{K}^{-1}\tilde{\mathbf{w}}_0, \quad (8)$$

where

$$\mathbf{R} = \begin{bmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \psi & 0 & \sin \psi \\ 0 & 1 & 0 \\ -\sin \psi & 0 & \cos \psi \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}, \quad (9)$$

and where \mathbf{K} is the camera matrix, $\tilde{\mathbf{w}}_0$ is the image before camera rotation and $\tilde{\mathbf{w}}$ is the image after camera rotation. Images are in homogeneous co-ordinates. This transformation does not assume a planar scene.

The three manoeuvres of the rear-facing 500 mm baseline case (Figure 14(a,c,e)) were used for a small parametric study. Transformations with different combinations of ϕ , ψ and θ were applied to the right camera images in each manoeuvre, image sequences were reprocessed, and the effects of the transformation on off-tracking results were observed. A combination of $\phi = 0.15^\circ$, $\psi = -0.10^\circ$ and $\theta = 0^\circ$ yielded a plausible and consistent set results for all three manoeuvres (suggesting that the right camera was rotated by $\phi = -0.15^\circ$ and $\psi = 0.10^\circ$ during testing). Results are shown in Figure 16.

By comparing Figure 16 with Figure 14(a,c,e), the corrections can be seen to effectively correct for the bias for $-e_{tr}$ values, while not significantly altering errors for $+e_{tr}$. These results confirm that very small ($< 0.2^\circ$) changes in the camera orientation can significantly affect off-tracking measurements, and this is likely to have been the case during these experiments.

4.5. Summary

The full-scale vehicle experiments have yielded a more realistic understanding of the accuracies achievable with the VISO-OT concept, with RMS errors around 0.11–0.12 m for rear cameras in the left roundabout manoeuvre. Systematic errors in the results of other manoeuvres led to the hypothesis that measurements were highly sensitive to the rigidity of the stereo camera mounting, and this was confirmed in a parametric study of the effects of small camera rotations.

5. Conclusions and future work

- (1) A novel concept for measuring trailer off-tracking has been described, using a stereo camera pair, visual odometry and a buffer-based off-tracking algorithm. The concept

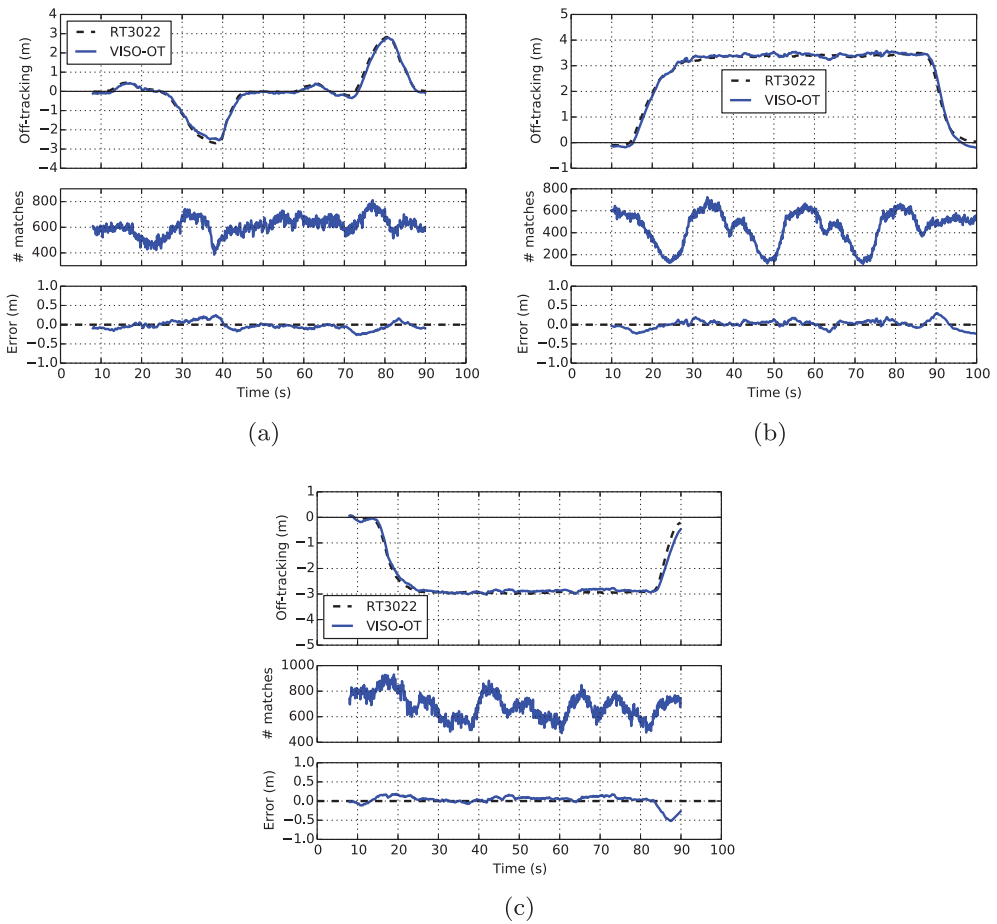


Figure 16. Off-tracking vehicle test results, rear cameras, 500 mm baseline, showing the effect of a simulated rotation of the right camera ($\phi = 0.15^\circ$, $\psi = -0.10^\circ$, $\theta = 0^\circ$). (a) Figure-of-eight, corrected. (b) Left roundabout, corrected. (c) Right roundabout, corrected

addresses the limitations of Cheng’s path-following controller and Miao’s ground-watching system, especially for off-highway applications. It is independent of wheel slip conditions, requires no tractor-based measurements or parameters and does not assume a planar, static road surface.

- (2) The theoretical performance of the system was evaluated in an idealised CAD environment, yielding RMS errors between 0.01 and 0.04 m (0.3–1.2%) under representative operating conditions. A high robustness to scenery density was demonstrated, and drift was shown to be bounded.
- (3) Full-scale field tests were carried out on a tractor semi-trailer combination, in which rear and side-facing camera configurations were evaluated at baselines of 500, 700 and 900 mm. RMS errors of 0.11–0.12 m (3.3–3.6%) were observed for left roundabout tests at a 500 mm baseline, with higher errors in other tests with other baselines.
- (4) Results from other manoeuvres were shown to be negatively affected by small ($< 0.2^\circ$) misalignments in camera mounting. A correction model was proposed to validate

this hypothesis, through which RMS errors were reduced and comparable to the left roundabout tests. Large stereo baselines were found to have a small improving effect on feature matching, but an overall negative effect on errors due to an increased chance of camera misalignment and difficulty in calibration.

- (5) In future work, a rigid stereo camera mount of sub-0.5 m baseline should be used to minimise the effects of trailer frame twist on camera calibration. It is also suggested to conduct some experiments with the sensor ‘in the loop’ with the CVDC’s path-following controller in an off-highway environment. Other suggestions include reducing sensitivity to lighting conditions, incorporating a Kalman filter to smooth off-tracking measurements, and improving side-camera measurement accuracy to enable functionality with multi-trailer combination vehicles.

Note

1. Typical GPS accuracy is in the order of a few metres, insufficient for off-tracking measurement. The target RMS errors of 0.10 m can be reached through the use of a GPS base station for operation in a fixed area, but this is clearly impractical for unrestricted driving. It does, however, serve as an attractive precision measurement tool for experimental validation (see Section 4.1).

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the Cambridge Commonwealth, European and International Trust, UK, the Council for Scientific and Industrial Research (CSIR), South Africa, and the Cambridge Vehicle Dynamics Consortium (CVDC).

Notes on contributors

Christopher de Saxe received the BSc(Eng) and MSc(Eng) degrees in mechanical engineering from the University of the Witwatersrand, South Africa, and the PhD degree in engineering from the University of Cambridge, UK. He is currently a Senior Research Engineer at the Council for Scientific and Industrial Research in Pretoria, South Africa.

David Cebon received the BE degree in mechanical engineering from the University of Melbourne, Australia, and the PhD degree from the University of Cambridge, UK. He is currently Professor of Mechanical Engineering at the University of Cambridge, and Director of the Cambridge Vehicle Dynamics Consortium and the Centre for Sustainable Road Freight.

ORCID

Christopher de Saxe  <http://orcid.org/0000-0002-9458-809X>

David Cebon  <http://orcid.org/0000-0003-2828-6445>

References

- [1] Domestic road freight statistics, United Kingdom, 2015. London: Department for Transport; 2016. (Report No.: August).

- [2] Gummer J, Kennedy D, Fankhauser S, et al. Meeting carbon budgets – 2013 progress report to parliament. London: Committee on Climate Change; 2013. (Report No.: June).
- [3] European Commission. Strategy for reducing heavy-duty vehicles' fuel consumption and CO2 emissions. Brussels: European Commission; 2014. (Report no. COM(2014) 285).
- [4] National Transport Commission. Performance based standards scheme. The standards and vehicle assessment rules. Melbourne: National Transport Commission; 2008.
- [5] Nordengen PA, Kienhöfer FW, de Saxe CC. Vehicle safety performance improvements using a performance-based standards approach: four case studies. Proceedings of the 13th International Symposium on Heavy Vehicle Transport Technology (HVTT); 2014 Oct 27–31; San Louis: IFRTT; 2014.
- [6] Kraaijenhagen B, Barth T, Kural K, et al. Greening and safety assurance of future modular road vehicles: book of requirements. HTAS EMS project consortium; 2014.
- [7] Evaluation of the longer semi-trailer trial: annual report 2015. London: Department for Transport; 2016. (Report No.: D3150-04).
- [8] Jujnovich BA, Cebon D. Comparative performance of semi-trailer steering systems. Proceedings of the 7th International Symposium on Heavy Vehicle Weights and Dimensions; June. Delft, The Netherlands; 2002. p. 195–214.
- [9] Jujnovich BA. Active steering of articulated vehicles [dissertation]. Cambridge, UK: University of Cambridge; 2005.
- [10] Cheng C. Enhancing safety of actively-steered articulated vehicles [dissertation]. Cambridge, UK: University of Cambridge; 2009.
- [11] Kharrazi S. Steering based lateral performance control of long heavy vehicle combinations [dissertation]. Gothenburg, Sweden: Chalmers University of Technology; 2012.
- [12] Hata N, Hasegawa S, Takahashi S, et al. A control method for 4WS truck to suppress excursion of a body rear overhang. Society of Automobile Engineers; 1989. (Report No.: 892521).
- [13] Notsu I, Takahashi S, Watanabe Y. Investigation into turning behavior of semi-trailer with additional trailer-wheel steering: a control method for trailer-wheel steering to minimize trailer rear-overhang swing in short turns. Society of Automobile Engineers; 1991. (Report No.: 912570).
- [14] Jujnovich BA, Cebon D. Path-following steering control for articulated vehicles. *J Dyn Syst Meas Control*. 2013;135(3):031006.
- [15] Jujnovich B, Roebuck RL, Odhams A, et al. Implementation of active rear steering of a tractor-semi-trailer. In: Jacob B, Nordengen PA, O'Connor A, Bouteldja M, editors. Proceedings of the 10th International Symposium on Heavy Vehicle Transportation Technology. Paris: John Wiley & Sons, Inc.; 2008. p. 358–367.
- [16] Cheng C, Odhams A, Roebuck RL, et al. Feedback control of semi-trailer steering at low speeds. Submitted to the Proceedings of the IMechE Part D Journal of Automobile Engineering; 2011.
- [17] Miao Q, Cebon D. Effects of operating off-highway on the tracking performance of a path-following steering system. FISITA World Automotive Congress. Maastricht; 2014.
- [18] Miao Q. Vision-based path-following control of articulated vehicles [dissertation]. Cambridge, UK: University of Cambridge; 2015.
- [19] Rosten E, Drummond T. Machine learning for high-speed corner detection. In: Leonardis A, Bischof H, Pinz A, editors. Proceedings of the 9th European Conference on Computer Vision Part I (ECCV 2006); 2006 May 7–13; Graz, Austria: Springer Berlin Heidelberg; 2006. p. 430–443.
- [20] Bay H, Ess A, Tuytelaars T, et al. Speeded-up robust features (SURF). *Comput Vis Image Underst*. 2008;110(3):346–359.
- [21] Muja M, Lowe DG. Fast approximate nearest neighbors with automatic algorithm configuration. Proceedings of the International Conference on Computer Vision Theory and Applications (VISAPP); 2009. p. 331–340.
- [22] Geiger A, Ziegler J, Stiller C. StereoScan: Dense 3d reconstruction in real-time. Proceedings of the 2011 IEEE Intelligent Vehicles Symposium (IV); 2011 Jun 5–9; Baden-Baden: IEEE; 2011. p. 963–968.

- [23] Morrison G, Cebon D. Combined emergency braking and turning of articulated heavy vehicles. *Vehicle System Dyn.* **2017**;55(5):725–749.
- [24] Sibley G, Mei C, Reid I, et al. Vast-scale outdoor navigation using adaptive relative bundle adjustment. *Int J Rob Res.* **2010**;29(8):958–980.
- [25] Harris MP. Application of computer vision systems to heavy goods vehicles: visual sensing of articulation angle [dissertation]. Cambridge, UK: University of Cambridge; **2013**.
- [26] de Saxe CC, Cebon D. A visual template-matching method for articulation angle measurement. 2015 IEEE 18th International Conference on Intelligent Transportation Systems; 2015 Sept 15–18; Las Palmas: IEEE; **2015**. p. 626–631.
- [27] de Saxe CC, Cebon D. Visual odometry for trailer off-tracking estimation. In: de Pont JJ, editor. *Proceedings of the 14th International Symposium on Heavy Vehicle Transport Technology (HVTT14)*; 2016 Nov 15–18; Rotorua, NZ: TERNZ Ltd.; **2016**. p. 138–149.
- [28] Kitt B, Geiger A, Lategahn H. Visual odometry based on stereo image sequences with RANSAC-based outlier rejection scheme. *Proceedings of the 2010 IEEE Intelligent Vehicles Symposium*; 2010 Jun 21–24; San Diego, CA: IEEE; **2010**. p. 486–492.
- [29] Geiger A, Lenz P, Urtasun R. Are we ready for autonomous driving? The KITTI vision benchmark suite. *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition*; 2012 Jun 16–21; Providence, RI: IEEE; **2012**. p. 3354–3361.
- [30] Olson CF, Abi-Rached H. Wide-baseline stereo vision for terrain mapping. *Mach Vis Appl.* **2010**;21(5):713–725.
- [31] Konolige K, Agrawal M. FrameSLAM: from bundle adjustment to real-time visual mapping. *IEEE Trans Robot.* **2008**;24(5):1066–1077.
- [32] Matas J, Chum O, Urban M, et al. Robust wide-baseline stereo from maximally stable extremal regions. *Image Vis Comput.* **2004**;22(10):761–767.
- [33] Inventor Professional. **2013**. [cited 2017 June 29]. Available from: <http://www.autodesk.com/products/inventor/overview>
- [34] de Saxe CC. Vision-based trailer pose estimation for articulated vehicles [dissertation]. Cambridge, UK: University of Cambridge; **2017**.
- [35] Zhang Z. A flexible new technique for camera calibration. *IEEE Trans Pattern Anal Mach Intell.* **2000**;22(11):1330–1334.
- [36] Heikkilä J, Silvén O. A four-step camera calibration procedure with implicit image correction. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. San Juan; **1997**. p. 1106–1112.
- [37] Fusiello A, Trucco E, Verri A. A compact algorithm for rectification of stereo pairs. *Mach Vis Appl.* **2000**;12(1):16–22.
- [38] Bradski G. The OpenCV library; **2000**. [cited 2017 July 7]. Available from: <http://opencv.org/>
- [39] Hartley R, Zisserman A *Multiple view geometry*. 2nd ed. Cambridge: Cambridge University Press; **2003**.