

Camera-Based Articulation Angle Sensing for Heavy Goods Vehicles

Christopher de Saxe and David Cebon

Abstract—Articulation angle sensing is an essential component of manoeuvrability and stability control systems for articulated heavy goods vehicles, particularly long combination vehicles. Existing solutions to this sensing task are limited by reliance on trailer modifications or information or by measurement accuracy, or both, restricting commercial adoption. In this paper we present a purely tractor-based sensor concept comprising a rear-facing camera and the parallel tracking and mapping (PTAM) image processing algorithm. The system requires no prior knowledge of or modifications to the trailer, is compatible with planar and non-planar trailer shapes, and with multiply-articulated vehicle combinations. The system is validated in full-scale vehicle tests on both a tractor semi-trailer combination and a truck and full-trailer combination, demonstrating robust performance in a number of conditions, including trailers with non-planar geometry and with minimal visual features. Average RMS measurement errors of 1.19, 1.03 and 1.53 degrees were demonstrated for the semi-trailer and full-trailer (drawbar and semi-trailer) respectively. This compares favourably with the state-of-the-art in the published literature. A number of improvements are proposed for future development based on the observations in this research.

Index Terms—articulated vehicles, articulation angle, computer vision, sensors

I. INTRODUCTION

THE movement of domestic freight is fundamental to the functioning and growth of both developed and developing nations [2]. Domestic freight transport in most countries is largely dominated by road haulage, with the majority of this freight being moved on heavy goods vehicles (HGVs). Although HGVs are favoured for most domestic freight movement due to the speed and flexibility of service, it is an expensive and carbon-intensive mode of transport compared to other modes such as rail.

The growth of freight transport demand on limited infrastructure, coupled with ambitious CO₂ emission reduction targets, has driven technology and policy development to investigate and implement measures to improve the efficiency of road freight transport. Of the technological and policy measures explored, the use of ‘Long Combination Vehicles’ (LCVs) has proven to be one of the most impactful, with relatively minor technological and infrastructural barriers.

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

C. de Saxe and D. Cebon are with the Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK. C. de Saxe is also with the School of Mechanical, Industrial and Aeronautical Engineering, University of the Witwatersrand, 1 Jan Smuts Ave, Johannesburg, 2000, South Africa.

LCVs (also known as ‘road trains’ or ‘high capacity vehicles’), are truck combinations with two or more trailers, able to carry more freight per truck than traditional truck combinations. Successful implementations and pilot projects in (amongst others) Australia, Finland, Sweden, Netherlands and South Africa have demonstrated significant potential for fuel use and carbon reduction, averaging 12% in the relatively modest South African pilot [3], and with savings of up to 30% in Australia [4].

An illustration of seven common HGV combinations is shown in Figure 1. The last three would be commonly deemed to be LCVs, and the last four possess two points of articulation (at either a pintle hitch, turntable, or fifth wheel). Although LCVs can make use of existing infrastructure (though may be constrained to certain routes due to their length), wider adoption is limited by reduced manoeuvrability and potentially compromised high-speed stability due to the number of articulation points. Emerging technologies such as active trailer steering [5], [6], [7], autonomous reversing [8], combined braking and steering control [9], and anti-jackknife control [10] have been shown to improve performance in this regard, but wide-spread uptake has been limited by practical commercialisation constraints.

A particular challenge for the roll-out of such technologies is the requirement for additional vehicle sensors and, for LCVs particularly, the requirement for articulation angle sensing. Assuming conventional trailers, it is necessary for articulation angle sensing systems to be based on the tractor unit, where control processing and actuation signals originate. Currently however, these technologies rely on either trailer-based articulation angle sensors with non-standard or experimental communication links, or estimation techniques which require knowledge of the trailer states and other non-standard sensors.

Camera-based systems have been demonstrated in the literature which have sought to address the challenge of a fully tractor-based, remote sensing solution. However, the image processing algorithms adopted have either not demonstrated acceptable measurement accuracy in full-scale tests, not fully addressed the practical constraints of earlier non-camera-based systems, or both. Depending on the algorithm used, these systems variously retain some trailer dependence (i.e. requiring known trailer markers), require some knowledge of vehicle properties which can vary in practise or are limited to specific trailer types or shapes. Additionally, no existing systems have been shown to be compatible with vehicle combinations comprising more than one articulation point, in which each articulation angle must be measured simultaneously.

In this paper, we present a new camera-based system for

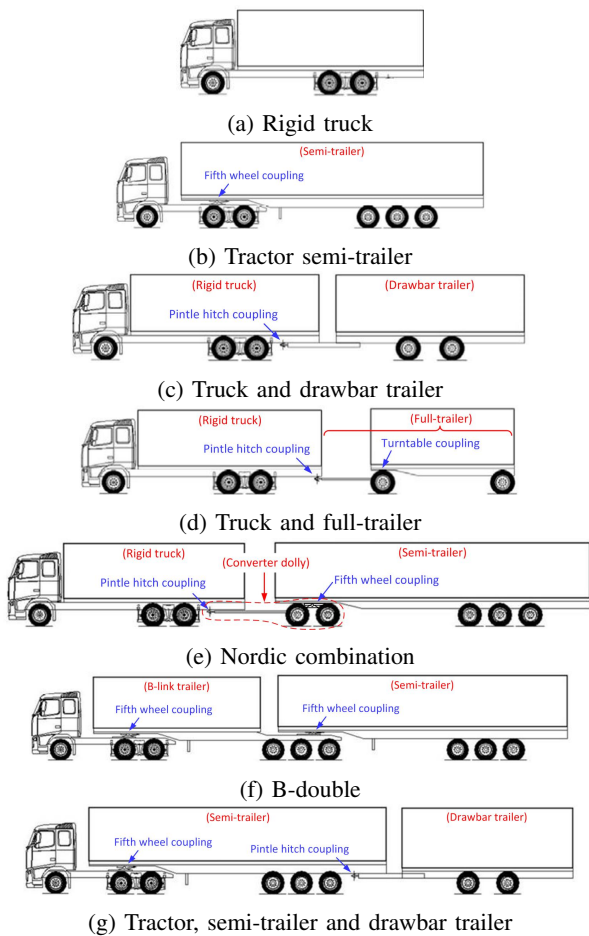


Fig. 1: Common HGV and HCV combinations (original illustrations from [12])

articulation angle measurement and assess its quantitative and qualitative performance under a range of conditions. The contributions of the paper are the novel application of the efficient and versatile parallel tracking and mapping algorithm (PTAM) [11] to the articulation angle sensing problem. In doing so, we have contributed a solution which addresses a number of the standing challenges for this task, including: fully tractor-based sensing, improved accuracy against the state-of-the-art, no reliance on known vehicle dimensions, no reliance on trailer-based markers or indicators, simultaneous measurement of more than one articulation angle, and functionality with a range of trailer types and shapes.

A. Related work

Existing trailer-based articulation angle measurement systems include a kingpin-mounted sensor from Vehicle Systems Engineering B.V. Netherlands (the ‘VSE sensor’) [13], the ‘Orbisense’ magnetic sensor by AB electronic Ltd [14], a custom string potentiometer solution [15], and the custom articulation angle sensor incorporated by TRIDEC (Netherlands) as part of their active trailer steering system [16].

In addition to being trailer-based, the above systems are all ‘contact’ type sensors (with the exception of the Orbisense sensor). This constrains the flexibility for use with various

tractors and trailers and trailer types. For tractor-based sensing, a good case can be made for non-contact sensing. In the literature, non-contact sensing solutions for this task have largely utilised state estimation methods (making use of other sensor information and models) or cameras.

State observer methods have been demonstrated by Bouteldja *et al.* [17] (using a state observer and an Extended Kalman Filter), Chu *et al.* [18] and Ehlgen *et al.* [19]. Moderate measurement accuracies were obtained, but only Ehlgen *et al.*’s system was demonstrated in field tests, achieving a maximum measurement error of 5° . Ziaukas *et al.* [20] utilised an Extended Kalman Filter, requiring speed and trailer yaw rate measurements and knowledge of several vehicle parameters. Physical tests on a tractor semi-trailer demonstrated maximum errors of up to around 10° .

More recent work has focussed on the use of cameras, arguably for their versatility and significant advances in the performance of computer vision algorithms. In 2009, Schikora *et al.* [21] present a system for tractor semi-trailers, comprising an upwards-facing camera on the tractor unit, viewing an encoder plate fitted to the underside of the semi-trailer (hence not strictly a tractor-based solution). Another system for rigid drawbar trailers was also proposed using a rearward-facing infrared camera on the towing unit, and a series of infrared diodes fitted to the trailer (again not strictly tractor-based).

In 2013, Caup *et al.* [22] presented a system for rigid drawbar trailers using a rear-facing camera. The system requires the hitch location and drawbar length (a feature of the trailer) to be known, and utilised a template-matching image processing method. The system was validated in field tests. Also in 2013, Harris [23] presented another rear-facing camera solution, for tractor semi-trailer combinations. Three image processing methods were assessed including homography decomposition, stereo vision (with the addition of a second camera), and a template-matching method which showed improvements over the first two methods in a limited evaluation. All three methods required the trailer front to be planar (flat).

Fuchs *et al.* [24] presented a similar rear-facing camera concept in 2014, which requires markers in a known location to be fitted to the trailer, together with knowledge of vehicle geometry and camera location (and so not strictly trailer-based). With perfect knowledge of camera location, a sub- 1° accuracy was obtained in simulation, but errors rose significantly when uncertainty in camera location was incorporated. In 2015, Fuchs *et al.* added a Kalman Filter to their system, improving the accuracy [25].

In 2019, De Saxe and Cebon presented a further development on the rear-facing camera and template matching approach of Harris [26], incorporating an Unscented Kalman Filter to smooth the noisy template-matching measurements. The system relied on some knowledge of the trailer geometry (making the solution not fully tractor-based), but sensitivity to this was shown to be small. High accuracy was demonstrated in simulations, while field tests on a tractor semi-trailer showed reasonable accuracy compared to the state-of-the-art, with RMS errors of 0.8 – 1.8° . A large contributing factor to the experimental errors was relative pitch and roll motion between tractor and trailer, which was not included in the simulations.

Also in 2019, Dahal *et al.* [27] presented a neural network-based approach to trailer detection and angle measurement called ‘DeepTrailerAssist’, using a combined convolutional network and long short-term memory network. The system was trained on 11 different trailer types, each of which had an inherent straight front edge. The authors do not present the measurement accuracy in detail, other than to state that ‘87% of the time the estimation was accurate within a tolerance of 1°’.

As a benchmark for the current research, results from the above literature were investigated to extract the maximum measurement errors ϵ_{\max} demonstrated, and maximum articulation angles assessed (Γ_{\max}). RMS errors would arguably serve as a better benchmark, but are not easily available in all the studies. It is clear from where simulation tests have been followed by full-scale testing that simulation results are a relatively poor indication of performance. It is hence crucial to carry out full-scale testing.

B. Research gap and objective

The rear-facing camera configuration, coupled with a suitable computer vision algorithm, is the clear candidate for further development, given its versatility, scope for being entirely tractor-based, and the prevalence of powerful modern computer vision algorithms. The described vision-based methods were shown to be limited to planar trailers, with the exception of [22] which was limited to drawbar-type trailers. Further, they were not demonstrated to be compatible with two or more articulation points—an important consideration for LCVs—or not strictly tractor-based requiring known markers or knowledge of trailer dimensions.

For this work, it was deemed fundamental to develop a system which was independent of trailer dimensions, type or planarity, such that the system could be compatible with a wide range of trailer types in use today, including box trailers, flat-bed trailers, tanker trailers, car-transporters etc without modification. Given the prevalence of LCVs with two or more articulation points, it was also deemed important for the system to be adaptable for the measurement of multiple articulation angles.

In summary, the research objective was to develop a new algorithm for camera-based articulation angle sensing, which improves on the accuracy and practical limitations of earlier systems, and to investigate the performance of the system in full-scale vehicle tests. The system should meet the following design requirements:

- 1) It should be entirely tractor-based, with no reliance on measurements from or geometric information about the trailer.
- 2) It should not rely on the existence of known visual features on the trailer.
- 3) It should be compatible with planar and non-planar trailer shapes.
- 4) It should be compatible with two or more articulation points.
- 5) It should offer acceptable measurement accuracy for control applications: in the order of 1 degree.
- 6) It should measure articulation angles up to 90 degrees.




Fig. 2: Rear-facing camera arrangement (shown for the case of a tractor semi-trailer combination)

II. SYSTEM OVERVIEW

Based on previous work, a rear-facing camera configuration was adopted. The camera was mounted to the tractor unit and faced the front of the trailer. This is shown in Figure 2. The example of a single-articulation tractor semi-trailer vehicle combination is depicted, but the system will be shown later to be compatible with other combination types.

A. PTAM-based image processing

For the image processing algorithm, a Simultaneous Localisation And Mapping or ‘SLAM’ approach was adopted. Primarily adopted in robotics applications, SLAM pertains to the problem of simultaneously determining the position and orientation of a moving agent and constructing a map of its three-dimensional environment. This is similar to the problem of trailer motion sensing; in this case the agent and environment are interchangeable with the tractor and trailer respectively. The objective is to determine the orientation of the trailer relative to the tractor, and to build a map of the trailer to ensure compatibility with various planar and non-planar trailer shapes. This interchangeability is possible provided the trailer occupies the majority of the camera field of view.

Many solutions to the camera-based or visual SLAM problem (vSLAM) exist, some of which are optimised for single cameras, multiple cameras, or other sensor inputs, and include FastSLAM [28], FastSLAM 2.0 [29], Graph-based SLAM [30], Topological SLAM [31], LSD-SLAM [32], ORB-SLAM [33] and ORB-SLAM2 [34].

For this work, the Parallel Tracking And Mapping (PTAM) algorithm by Klein and Murray [11] was adopted. PTAM is a novel and highly efficient implementation of mono-camera vSLAM, in which the localisation and mapping processing tasks are formulated in parallel. The PTAM algorithm has been optimised for localised as opposed to exploration applications, meaning that the system expects to operate with a single scene and not to constantly move into new scenes requiring recurring map extensions. Being optimised as a localised system, PTAM has superior pose estimation efficiency and accuracy versus exploratory systems, with obvious benefits for the articulation

angle sensing task. The name derives from the parallel nature in which the localisation (tracking) and mapping tasks are carried out, on two separate CPU threads.

The algorithm as implemented for the articulation sensing task is summarised in Figure 3 [35]. For an unseen trailer, initially the algorithm requires two images of the trailer front to create the initial three-dimensional trailer map (creating a point cloud from matched features in the stereo image pair). These are called ‘keyframes’. The system must then be zeroed, which is achieved through a brief straight driving manoeuvre.

In this work, the initialization and zeroing processes have been performed manually, however these processes can reasonably be automated in future work. For example, a simple initialisation manoeuvre could be carried out whenever a new trailer is hitched to the tractor. This might involve an initial steering input to generate the stereo keyframes, followed by following a straight line to zero the system. The system could be programmed to detect once a sufficient steer angle has been input for a sufficient time to capture the stereo keyframes, followed by detecting when zero steer angle has been observed for a sufficient time to assume that zero articulation has been achieved in order to zero the system. The initialisation and zeroing data could be stored in the system per individual trailer so that the manoeuvre need not be repeated, and previously used trailers can be recognised using a simple image recognition system if a new trailer is hitched.

The stereo initialisation step is important, as this defines the initial map upon which all pose updates are based. A small translation of the scene between the two initial keyframes is required to complete the stereo mapping process. The exact size of this translation is not prescribed and does not need to be known. For typical hand-held camera applications, Klein and Murray arbitrarily assumed the size of this translation in world co-ordinates to be 10 cm in order to generate the initial map using the five-point algorithm [36]. For the trailer sensing application the exact magnitude is not known and can vary, meaning the resultant Cartesian map is only accurate up to a scale factor (*i.e.* relative dimensions and translations are accurate, but the precise scale is unknown). As we are only concerned with rotations (*i.e.* articulation angle), this is of no consequence.

Ongoing orientation measurement is achieved through matching features in the current image with features in the existing trailer map (tracking). To do this, PTAM detects corner features using the FAST algorithm (Features from Accelerated Segment Test) [37] and utilises a simple 8×8 pixel patch around each corner as the feature descriptor. This approach was favoured for its computational efficiency. The current location of the patch relative to the original keyframes is predicted using a decaying velocity motion model (similar to an alpha-beta filter [38]), and the new location is finalised through template-matching the original patch in a circular region around this point. The patch is warped according to its predicted new location to improve matching correlation. A perspective projection pin-hole camera model is used for patch warping and tracking. The pin-hole model is summarised in Equation 1 [39], [40].

$$\tilde{\mathbf{w}} = \mathbf{K} \left[\mathbf{R} \mid \mathbf{T} \right] \tilde{\mathbf{X}} \quad (1)$$

$$\begin{bmatrix} \xi u \\ \xi v \\ \xi \end{bmatrix} = \begin{bmatrix} k_u f_x & 0 & u_0 \\ 0 & k_v f_y & v_0 \\ 0 & 0 & 1 \end{bmatrix} \left[\mathbf{R} \mid \mathbf{T} \right] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

X, Y, Z are the location of a feature point in world coordinates, u and v are the pixel coordinates of the feature in the image plane, \mathbf{R} and \mathbf{T} are the 3×3 rotation matrix and 3×1 translation vector of the relative camera motion, \mathbf{K} is the camera calibration matrix containing the ‘intrinsic’ camera parameters, f_x and f_y are the focal lengths in the horizontal and vertical directions, u_0 and v_0 represent the location of the optical axis, k_u and k_v are scaling parameters for rectangular pixels, and ξ is the scaling factor (unknown in this case due to the scale ambiguity of using a single camera).

Lens distortion is accounted for using the simplified FOV-model [41]. In this case, pixel coordinates are assumed to be distorted by a purely radial distortion which is described using a single distortion coefficient s as follows (where X_C, Y_C, Z_C are the location of the feature point in camera-centred coordinates):

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u_0 \\ v_0 \end{bmatrix} + \begin{bmatrix} f_x & 0 \\ 0 & f_y \end{bmatrix} \frac{r'}{r} \begin{bmatrix} X_C/Z_C \\ Y_C/Z_C \end{bmatrix} \quad (2)$$

where

$$r = \sqrt{\frac{X_c^2 + Y_c^2}{Z_c^2}} \quad r' = \frac{1}{s} \arctan \left(2r \tan \frac{s}{2} \right) \quad (3)$$

In parallel to the tracking thread, the mapping thread checks if new features (and hence potential new keyframes) are evident. This would happen in practise if the trailer angle increased to a point where a new side of the trailer was seen for the first time (if not captured in the initial two keyframes). If a new keyframe is detected, it is added to the existing map and the map is updated. If no new keyframes are detected, the existing map is refined through bundle adjustment [42].

B. Articulation angle extraction

The C++ source code for PTAM is open source [43] and was modified for application to the articulation sensing task. The algorithm provides two primary outputs: a 3-D map of feature point locations, and the camera pose in the form of a 3×1 translation vector \mathbf{T} and a 3×3 rotation matrix, \mathbf{R} (see Equation 1). During stereo initialisation, the rotation matrix is initialised relative to a dominant plane found in the first two keyframes, and the translation vector represents the motion of the camera relative to its original location.

Of interest here is the rotation matrix and particularly the yaw angle of the camera relative to the observed scene (*i.e.* the trailer). In order to extract the articulation angle, the rotation matrix is decomposed into a combination of sequential rotations about each axis, known as Euler angles: roll (ϕ), yaw (ψ) and pitch (θ). This is summarised as follows:

$$\mathbf{R} = \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{bmatrix} = \mathbf{R}_z(\phi)\mathbf{R}_y(\psi)\mathbf{R}_x(\theta) \quad (4)$$

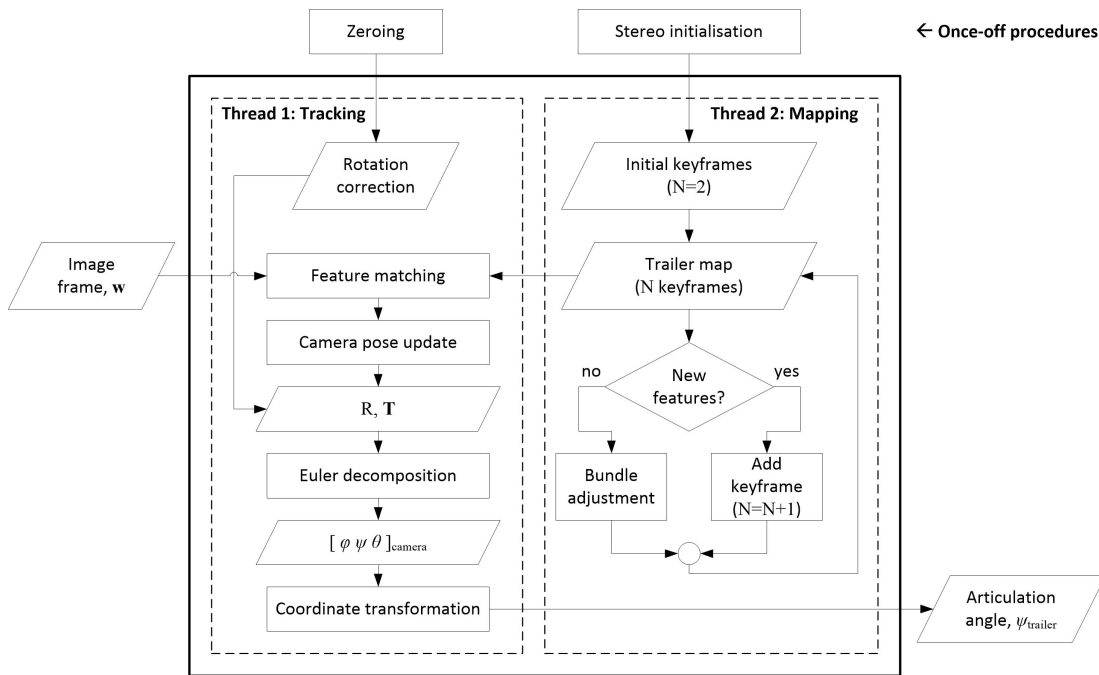


Fig. 3: Flowchart of the PTAM-derived articulation angle measurement method [35]

where

$$R_z(\phi) = \begin{bmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$R_y(\psi) = \begin{bmatrix} \cos \psi & 0 & \sin \psi \\ 0 & 1 & 0 \\ -\sin \psi & 0 & \cos \psi \end{bmatrix}$$

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}$$

These Euler angles may be extracted from the rotation matrix using a combination of the above definitions and logic to reject multiple solutions. The method of [44] was used for this purpose. These Euler angles are in the camera coordinate frame however, and must be translated into the trailer coordinate frame as follows:

$$\begin{bmatrix} \phi \\ \psi \\ \theta \end{bmatrix}_{trailer} = \begin{bmatrix} \cos \psi & 0 & \sin \psi \\ 0 & 1 & 0 \\ -\sin \psi & 0 & \cos \psi \end{bmatrix} \begin{bmatrix} \phi \\ \psi \\ \theta \end{bmatrix}_{camera} \quad (5)$$

The articulation angle is then:

$$\Gamma = \psi_{trailer}. \quad (6)$$

C. Implementation

The modified PTAM algorithm was implemented in C++ on a 64-bit Linux laptop computer with a dual-core 2.4GHz Core i3 processor and 4GB of RAM. Steady frame processing rates of 20 fps were achieved. The built-in camera calibration tool available with PTAM was used to calibrate cameras.

III. SIMULATIONS

A. Vehicle model

As an initial validation, a simple 3D computer model of a tractor semi-trailer was created in Autodesk Inventor, with a single degree of freedom for articulation angle. The model is shown in Figure 2. A virtual camera was fixed behind the tractor cab as shown. This functions as a simple ‘pin-hole’ camera with no distortion. Photo-realistic visual texture was added to the trailer based on an available test unit owned by the Cambridge Vehicle Dynamics Consortium (CVDC).

Representative articulation angle signals were generated using a three degree-of-freedom ‘bicycle model’ of a tractor semi-trailer vehicle combination. (More details can be found in [26].) The model was subjected to a square wave steer input, yielding an approximately sinusoidal articulation angle signal. The articulation angle signal was then fed into the CAD model to generate sequences of image data from the virtual camera, representative of the articulation angle signals.

B. Results

Two tests were carried out (one run each), one reaching a maximum articulation angle of 50°, and another reaching 90°. The 90° case served to test performance when the side of the trailer came into view, forcing the need for new keyframes to be captured. Results are shown in Figure 4. Γ_{lim} represents the angle beyond which the front of the trailer disappears from view, and only the trailer side is visible (around 70°).

In both tests the errors increase approximately in line with increasing articulation angle, potentially due to the increasing oblique camera view relative to where the stereo initialisation took place. No obvious degradation of accuracy is evident beyond Γ_{lim} . Accuracy is good relative to the the published

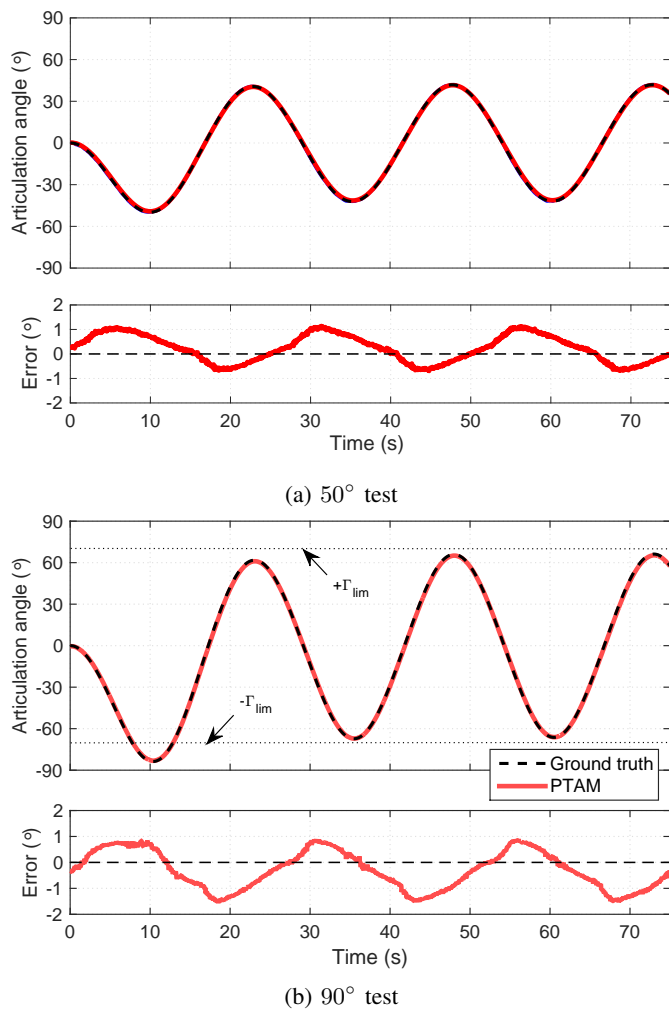


Fig. 4: Example simulation results

literature, while the targeted practical shortcomings of the state-of-the-art have been overcome.

The maximum errors do not appear to align with maximum articulation angles, but rather seem to peak at different values of Γ between tests, close to 0° in the 90° test. This may seem to suggest that the algorithm exhibits some specific sensitivity near these angles, but in fact the variation appears to be a relatively smooth function of articulation angle. The root cause of this behaviour is deemed most likely to be variations in the initialisation process. A more detailed investigation into this behaviour can be found in [1] and later in this paper.

Figure 5 shows details of the features tracked during the simulation. Features on both the trailer sides were added to the map successfully as they become visible at higher values of Γ . In the third frame, it can be seen that the front of the trailer has rotated out of view at high negative Γ . This is in accordance with $-\Gamma_{lim}$ being exceeded as seen in Figure 4.

IV. FIELD TESTS: TRACTOR SEMI-TRAILER

A. Vehicle and instrumentation

Initial full-scale vehicle tests were carried out on a tractor semi-trailer combination, with one articulation point. The test

vehicle and instrumentation used is illustrated in Figure 6. A Point Grey Flea3 USB 3.0 camera was fitted behind the tractor cabin, with a vari-focus lens set to $f = 2.8$ mm. Images were captured in greyscale at 640×480 and at 20 fps. The ‘ground truth’ articulation angle measurement came from a VSE rotary potentiometer articulation angle sensor [13] fitted to the trailer kingpin.

Testing was carried out at Bourn airfield near Cambridge. Two trailer front faces were considered: a planar front (a predominantly flat surface) and a non-planar front comprising a mock refrigeration unit (a common addition to trailer fronts).

Six tests with the planar trailer front were carried out, comprising three periodic ‘square wave’ steering tests with articulation angles up to 30° (denoted ‘per30’) and three ‘general’ driving manoeuvres (denoted ‘gen’). Six tests with the non-planar trailer front (refrigeration unit) were carried out. Three runs were carried out for each test. The maximum articulation angle was also increased: two periodic tests up to 30° (‘per30_3d’), two periodic tests up to 50° (‘per50_3d’), and two general tests up to 55° (‘gen_3d’). Vehicle speed was approximately 5 km/h for the step steer tests and in the range 0–10 km/h for the general driving tests.

B. Initialisation and zeroing

An example of the stereo initialisation and zeroing procedure is given in Figure 7, showing a time history of PTAM and ground truth measurements. As some lateral movement of the semi-trailer or drawbar begins, the first keyframe is captured. Once sufficient additional lateral motion has occurred, the second keyframe is captured (shown at approximately 12 seconds), defining the initial stereo map. Hereafter tracking and measurement takes place, but the co-ordinate system is not necessarily aligned with the vehicle, as is evident in the measurement bias.

At approximately 72 seconds, when it was observed that the trailer was straight, a zeroing command was sent to the algorithm, applying a zeroing correction to all subsequent measurements. At this point the instantaneous rotation matrix, R , is recorded (see Equation 4). Subsequent rotation matrices were post-multiplied by this reference matrix to adjust the co-ordinate frame accordingly, before performing Euler decomposition to extract the articulation angle.

C. Results

Example measurement time histories are shown in Figure 8, including tests with a planar and non-planar trailer front. The non-planar tests at larger Γ exhibited greater errors than the planar tests at smaller Γ ; it appears that there is no sensitivity to the planarity of the trailer. Measurements also appear to be robust to transient pitch and roll variations. In the non-planar tests, keyframes from the trailer sides are added, with no discernible effect on performance.

Similar peak error behaviour to that discussed in Section III-B can be seen here. In [1], this behaviour was studied in detail and the primary root cause was concluded to be variations in the initialisation process. The shape of the error function with respect to articulation angle was found




Fig. 5: Simulated semi-trailer feature map

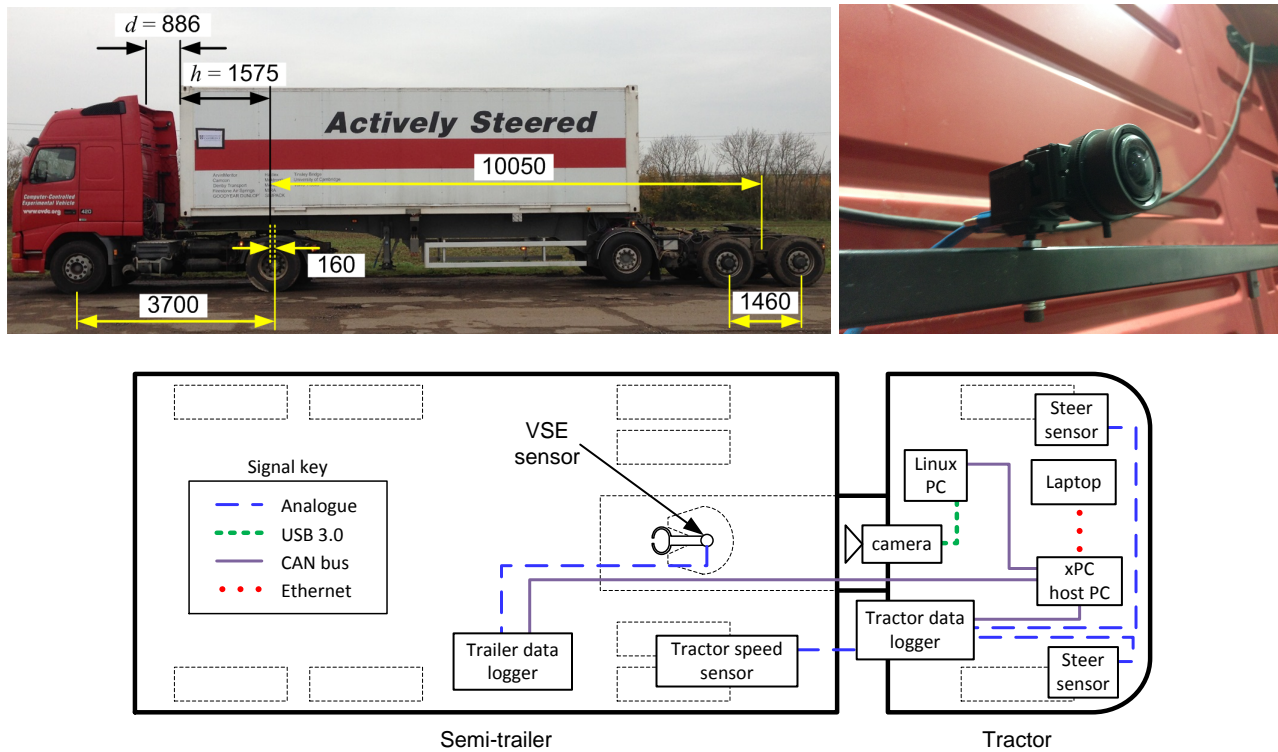


Fig. 6: Tractor semi-trailer test vehicle and instrumentation (all dimensions in mm) [26]

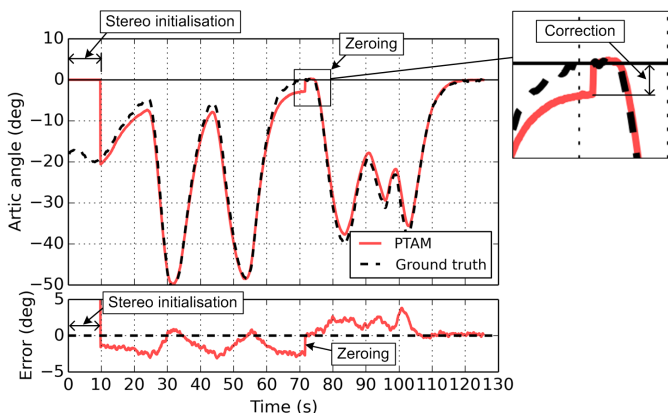


Fig. 7: Stereo initialisation and zeroing in field tests

to be consistent in shape, but varied from test to test in its precise location which gave rise to the maximum errors manifesting at different articulation angles between tests. The initialisation process was carried out manually and as such was difficult to reproduce consistently between tests, especially with the full-scale vehicle tests. Further investigation into and improvements to this process would undoubtedly improve performance. However, even given this behaviour, the algorithm has been demonstrated to perform well in terms of accuracy and practicality when compared to the current state-of-the-art. Additional considerations into the behaviour of errors and their potential sources are given in Section VII.

Feature-tracking details are provided in Figure 9 both planar and non-planar tests. For the non-planar test reaching 55° ,

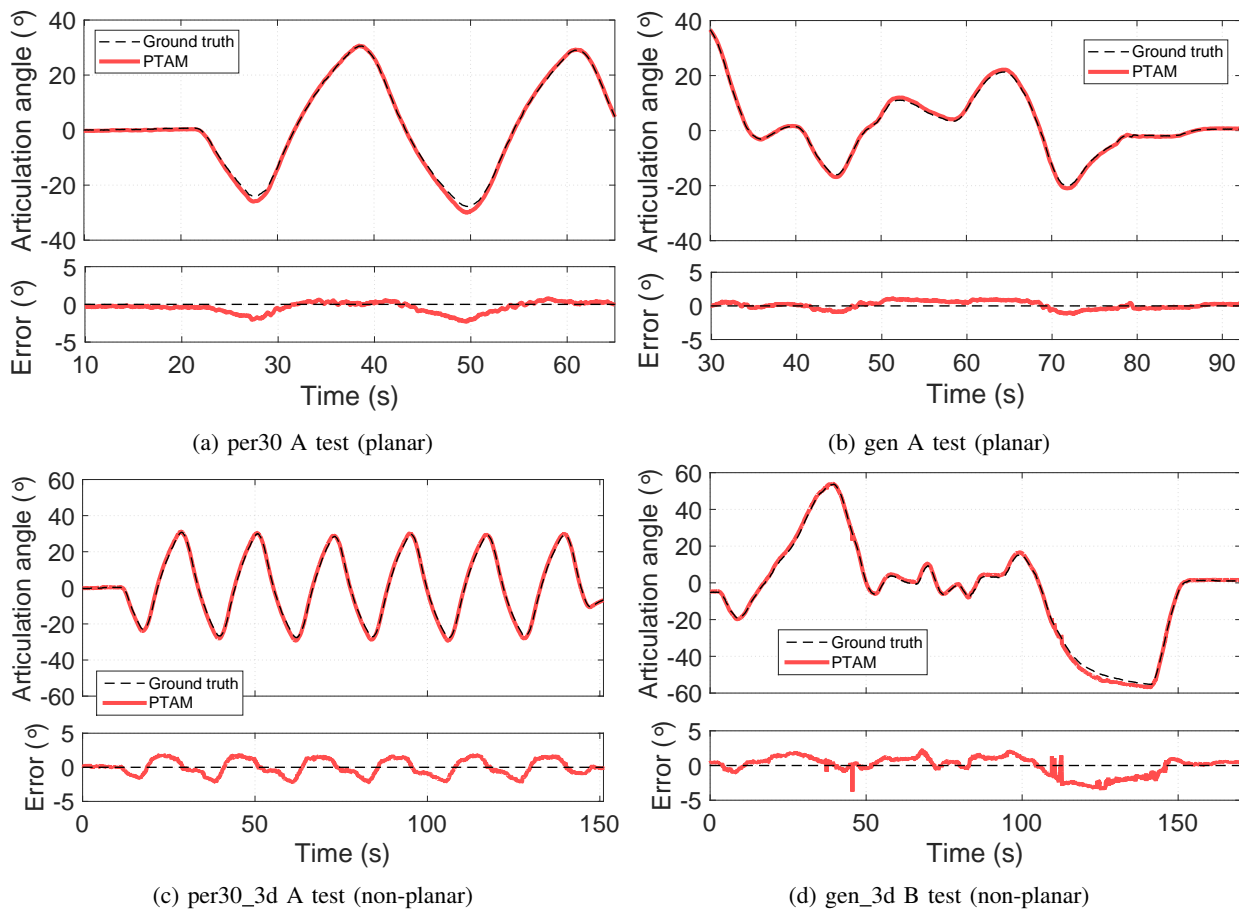


Fig. 8: Vehicle test time histories

it can be seen that features on the side of the trailer are successfully detected and tracked when the front face is no longer visible. Overall, the performance is good and the PTAM signal compares well with that of the ground truth. As expected, the errors are noticeably larger than those in the simulations, but still within good performance limits for this application. (Note when comparing to the simulation results that the error axes in Figure 8 are set at -5° to $+5^\circ$.)

V. FIELD TESTS: TRUCK AND FULL-TRAILER

In the previous section, the PTAM method was shown to be effective on tractor semi-trailer combination. These combinations are common throughout the world, and so it is important that the articulation angle sensor is compatible with them. As discussed previously however, it is important to demonstrate compatibility with LCVs with more than one articulation point as well. These are increasing in use in Europe and are already common in countries including Sweden, Australia, New Zealand and South Africa. In this section, the applicability of the sensor concept is extended to a vehicle combination comprising a rigid truck, converter dolly, and semi-trailer. These have been widely adopted in Sweden and Finland and are hence known as the ‘Nordic combination’, typically 25.25 m in length. This is depicted in Figure 1 (the fifth vehicle).

In order for the sensor system to measure two articulation angles, still using a single camera, both articulating trailer bodies must be visible in the camera’s field of view. Two regions of interest can then be defined, one for each of the drawbar and semi-trailer. Each region can then be processed independently using the same PTAM-method as before, yielding independent articulation angle measurements. The only modifications required are the positioning of the camera to ensure the correct field-of-view, and the separation of the view into separate regions of interest for independent processing.

In practice, for a fixed mounting location of the camera near the top of the tractor chassis, it would be reasonable to set fixed regions of interest which would work for any drawbar semi-trailer (or full-trailer) combination. This could comprise simply of one vertical delineation separating the field-of-view into a lower third (for the drawbar) and an upper two thirds (for the semi-trailer). Otherwise, it is possible that a simple optical flow algorithm could be included as part of an initialisation procedure to separate the field-of-view into any number of regions of interest depending on the number of rotating bodies in the field of view, for any trailer type. In this work the field of view was set manually.

A. Vehicle and instrumentation

A Nordic combination vehicle was made available for testing by Volvo Group Trucks Technology (GTT), Sweden.

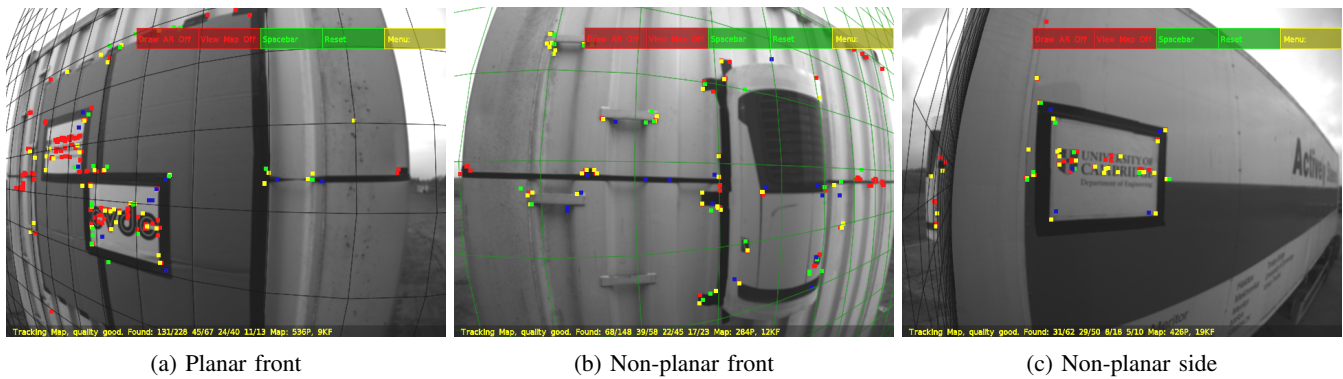


Fig. 9: Semi-trailer feature maps

The combination consisted of a rigid truck, converter dolly, and semi-trailer. The available truck was bare of any superstructure, but the addition of typical superstructures would not be expected to obscure the camera field of view. The semi-trailer had a planar front, though as demonstrated previously this is not a requirement for accurate operation of the system. The vehicle and instrumentation is illustrated in Figure 10. The camera was mounted to the chassis of the truck unit, near the location of the existing reverse-aid camera already fitted (the location of which would also suffice).

The same Point Grey Flea3 USB camera with Fujinon lens used previously was fixed to the rear of the truck chassis as shown, approximately along the longitudinal axis of the truck. The camera was calibrated using the built-in PTAM calibrator as before. From this location the field of view of the camera included both the drawbar of the dolly and the front of the semi-trailer. Regions of interest in the images were chosen manually, and are illustrated in Figure 11. The sizes of the partitions were 440×270 and 560×180 pixels for the semi-trailer and drawbar respectively, and were centred laterally.

Testing was carried out at the Hällered proving ground near Gothenburg. Manoeuvres were carried out around the trailer storage yard next to a workshop which provided the visual scenery and manoeuvres representative of truck loading and unloading areas. Arbitrary visual markers were added to the otherwise bare trailer front using pieces of duct tape, which were then removed during testing to compare performance. Six manoeuvres were carried out: a left turn loop ('left turn'), a right turn loop ('right turn'), two arbitrary routes around the yard with both left and right turns ('arbitrary A' and 'arbitrary B'), a short manoeuvre over an uneven unpaved surface ('uneven'). The 'left turn' manoeuvre was then repeated with the trailer markers removed ('left turn, no markers'), for comparison of performance with and without markers. Three runs of each manoeuvre were carried out.

The 'ground truth' articulation angle measurements were obtained via a VSE sensor at the fifth wheel, and a custom sensor provided by Volvo at the pintle hitch. Stereo initialisation and zeroing was carried out in the same manner as the previous tests, with the semi-trailer and dolly regions of interest being processed separately.

B. Results

Example time histories of articulation angle measurements are shown in Figure 12 (semi-trailer angle, left turn, marker and no marker tests). PTAM measures the total relative pose between camera and environment, and so in the case of the semi-trailer this is the *sum* of drawbar and semi-trailer articulation angles. Therefore, in this case the sum of pintle hitch and VSE sensor measurements were used as the ground truth.

The manoeuvres conducted provided a large variation in background scenery, and the results show no discernible sensitivity to this. There is also no clear difference in performance between the left turn results and left turn with no markers, indicating that the markers are not necessary, and that the system performs well with visually bare trailers. Unusually high errors (not in line with the other tests) are evident in the manoeuvres on uneven ground. These tests were intended to demonstrate sensitivity to transient pitch and roll motion of the trailer. Anecdotal evidence gained by observing the bumpy vehicle motion in the video sequences, and comparing the results with smoother runs, suggests that there is little if any sensitivity to this. However zeroing these tests was difficult given that there was not an obvious period of straight-line driving. This likely resulted in an inaccurate zeroing of articulation angle yielding a small bias, combined with a potentially non-zero pitch and or roll, yielding an inaccurate pitch and roll component to R. (See Equation 4.)

Example initialised feature maps are shown in Figure 13, for the semi-trailer (marked and unmarked), and the drawbar. It is clear that the markers introduced additional feature points compared with the unmarked semi-trailer, but that the unmarked semi-trailer still provided sufficient features for initialisation and tracking. In both cases there was a concentration of features around the service connections and along the bottom edge of the semi-trailer around rivets, trailer name plates and trailer identification code. The drawbar possessed sufficient natural features for initialisation and tracking, in the form of edges, bolts, and connector points.

We can conclude from this that the system does not require any artificial trailer markers, specially-designed or otherwise, and performs well with visually plain trailers. Compared to marker-based systems, this helps to ensure that the system

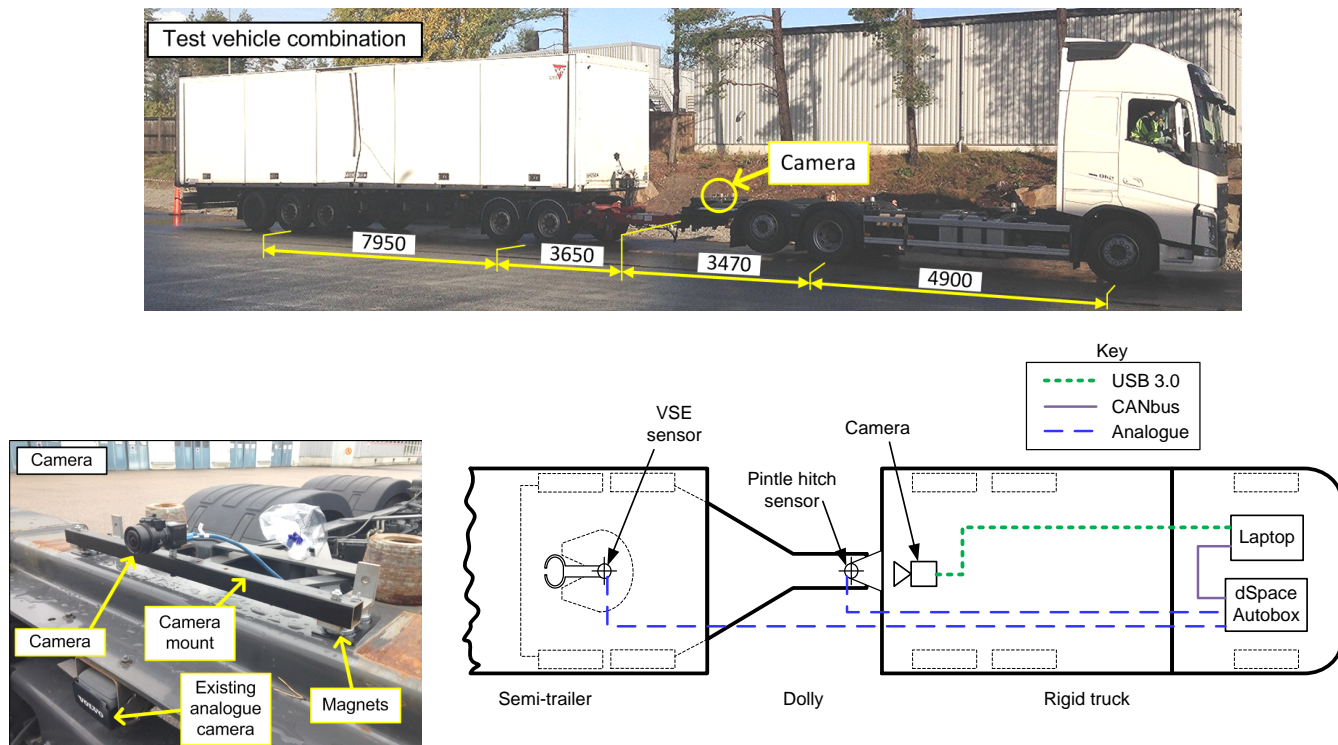


Fig. 10: Nordic combination test vehicle and instrumentation (all dimensions in mm)




Fig. 11: Image regions of interest used for processing (zero and non-zero articulation angles shown)

is entirely independent of the trailer and does not require modification when new trailers are hitched or due to trailer wear and tear.

It was evident during testing that a larger rotation was required during the drawbar stereo initialisation process. If the motion was too small, the initialisation would either fail or yield obviously erroneous and/or erratic measurements. This is due to the fact that the camera was located near the centre of rotation of the drawbar (at the pintle hitch), meaning that relative motion between the camera and observed body was dominated by a rotational element, with little translation as required for the stereo initialisation. In practice, it would be

straight-forward to detect a failed initialisation, or to otherwise ensure sufficient motion between keyframes in an automatic fashion. Adding a lateral offset between the camera and hitch axis would help reduce the angle required for initialisation, by increasing the translation component of the relative motion (\mathbf{T}).

VI. PERFORMANCE COMPARISON

A summary of the performance of the system is given in Table I, together with the available performance information of comparable systems discussed in Section I-A. The table

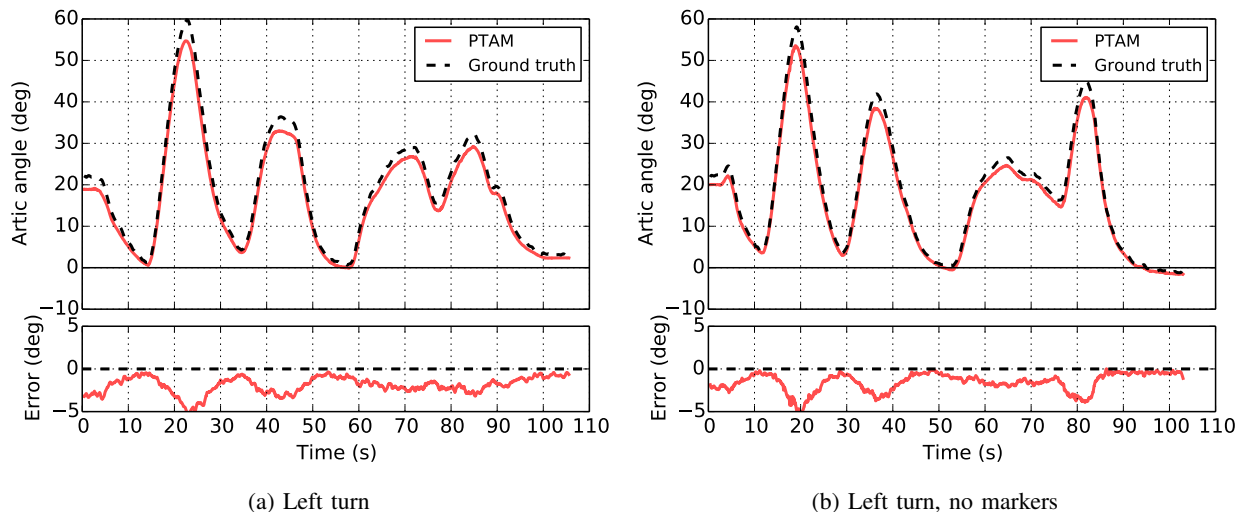


Fig. 12: Time histories, semi-trailer

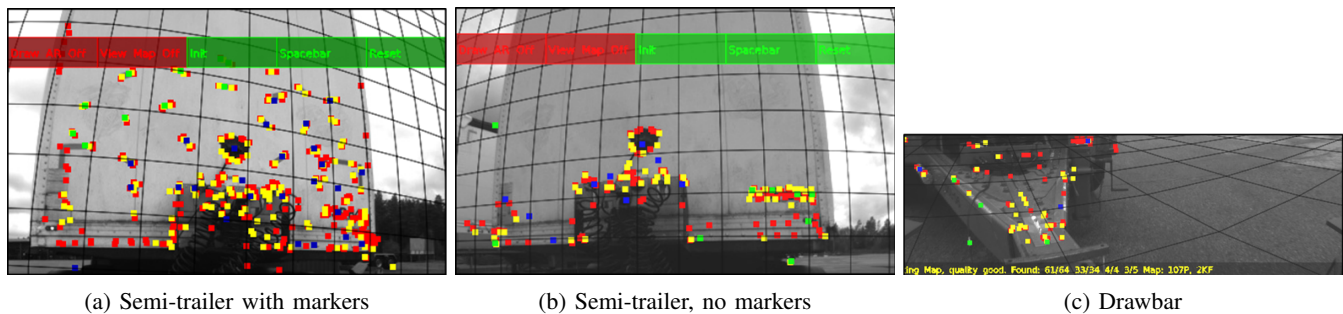


Fig. 13: Full-trailer feature maps

includes the maximum articulation angle tested, Γ_{\max} , the maximum observed error over all tests, ϵ_{\max} , and the RMS error over all tests (where available), ϵ_{RMS} . The encoder and infrared systems of [21] have not been included as they require physical trailer modifications other than just markers. Information on RMS errors was not available from the published literature, but has been included for the PTAM system. For field tests of the PTAM system where multiple runs per test were carried out, the RMS and maximum errors have been averaged over the runs. The table has been separated into systems which have been assessed only in a simulation environment, and those which have undergone full-scale field tests.

The PTAM system performs well in simulation compared to other simulation-tested systems. The marker tracking systems perform better, but rely on specially-located markers placed on the trailer. The template matching system also performs better, but is limited to trailers with a planar front, and is restrained to a maximum articulation angle. In comparison, the PTAM system was tested up to 84° . Field test results provide a more accurate view of expected in-field performance. In field tests, the PTAM system exhibits the lowest maximum errors, with a significantly smaller variation in performance compared to some. The PTAM system was also tested to higher articulation angles than the others, and is theoretically unconstrained in this respect. Also, the PTAM system has been demonstrated to function with trailers with non-planar

fronts (the state observer and stereo vision methods should theoretically manage this too), and has been demonstrated to function effectively for a bi-articulated trailer (which none of the other systems have). Overall RMS errors are shown to be in the region of 1° . Provided the maximum errors can be reduced, through filtering or further investigation into the error sensitivities, this indicates that the PTAM system has the potential for incorporation into several articulated vehicle control systems in future.

Quantitative performance was not the only performance objective for this work. Additional important qualitative properties were stipulated in Section I-B, to ensure that the system was commercially viable and practical. A comparison of the qualitative performance of the PTAM system and the state-of-the-art against these requirements is given in Table II. Where adherence with a given metric was not explicitly demonstrated but was deemed achievable with the described method, the metric has been deemed to be met. It is clear from the table that state-of-the-art methods fail to meet at least two of the required performance metrics, while the PTAM method has been demonstrated to meet all requirements.

VII. ERROR SENSITIVITY

In both the simulation and field tests, some error sensitivity as a function of articulation angle was observed. This suggests

TABLE I: Comparison of the quantitative performance of articulation angle estimation methods

	Method	Validation	Γ_{\max}	ϵ_{\max}	ϵ_{RMS}
[17]	State observer	Simulation	90°	8°	-
[18]	State observer	Simulation	3–20°	0.3–10°	-
[23]	Homography decomposition [†]	Simulation	17–90°	3.2–8.4°	-
[24]	Marker tracking [†]	Simulation	50°	0.5°	-
[25]	Marker tracking [†]	Simulation	15–30°	0.6°	-
[26]	Template-matching [†]	Simulation	50°	0.7°	-
	PTAM (semi-trailer) [†]	Simulation	50–84°	1.1–1.6°	0.7°
[19]	State observer	Field tests	48°	5.4°	-
[22]	Template-matching [†]	Field tests	20–55°	5.5–7.6°	-
[23]	Stereo vision [†]	Field tests	17–52°	3.3–18°	-
[26]	Template-matching [†]	Field tests	37°	2.3–6.8°	-
[20]	State observer	Field tests	50°	10°	-
[27]	Deep learning [†]	Field tests	-	-	-
	PTAM (semi-trailer) [†]	Field tests	23–37°	1.2–2.6°	0.7°
	PTAM (semi-trailer, non-planar) [†]	Field tests	28–50°	1.7–3.8°	1.2°
	PTAM (Nordic, trailer) [†]	Field tests	54–61°	3.5–5.1°	1.5°
	PTAM (Nordic, drawbar) [†]	Field tests	30–40°	2.5–4.6°	1.0°

[†] Computer vision method

TABLE II: Comparison of the qualitative performance of articulation angle estimation methods

Method	No trailer sensor inputs	No geometric knowledge	No trailer markers	Non-planar trailers	Multiple articulations	90° articulation	
[17] [18] [19] [20]	State observer	✗	✗	✓	✓	✗	✓
[22]	Template-matching [†]	✓	✗	✓	✗	✗	✓
[23]	Hom. decomposition [†]	✓	✓	✓	✗	✗	✓
[23]	Stereo vision [†]	✓	✓	✓	✗	✗	✓
[24] [25]	Marker tracking [†]	✓	✗	✗	✗	✗	✗
[26]	Template-matching [†]	✓	✗	✓	✗	✗	✗
[27]	Deep learning [†]	✓	✗	✓	✗	✗	✗
[11]	PTAM [†]	✓	✓	✓	✓	✓	✓

[†] Computer vision method

some shared error mechanism. The degree of sensitivity appears to be consistent between the tests within the same test scenario, but does not appear to be consistent between each different scenario. In some results there the sensitivity is positive (increasing over-estimation with increasing articulation angle) and in others it is negative (increasing under-estimation with increasing articulation angle). In some the sensitivity is larger, and in others, smaller. Here we try to consolidate this sensitivity and explain its possible sources.

In the tractor semi-trailer tests (Section IV), increasing Γ led to increasing *positive* errors (Γ is over-estimated), while in Sections III and V, increasing Γ led to increasing *negative* errors (Γ is under-estimated). This is demonstrated in a comparison of maximum error and maximum articulation angles, shown in Figure 14. The results for all runs of all manoeuvres (both simulation and experiment) are included, and a linear least squares fit has been added for each set of tests. (A zero-intercept has been forced, consistent with the observed behaviour with correct zeroing.)

It is clear that the sensitivity is unique for each configuration of the camera-trailer system. The drawbar and semi-trailer measurements in the Nordic tests display a similar behaviour, and these share the same camera setup and environment. It is also observed that the theoretically ideal simulation tests also exhibit this sensitivity, but to a smaller extent. In this case the

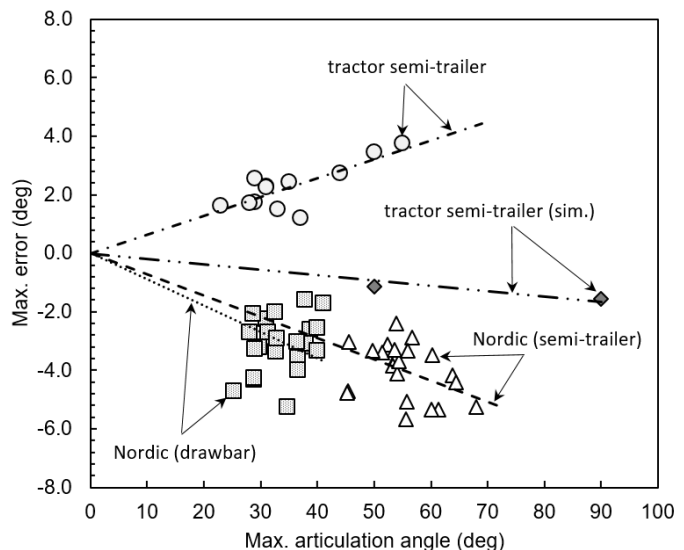


Fig. 14: Maximum errors vs. max. articulation angle, all tests

camera is theoretically perfectly aligned and calibrated, with no distortion, and the system has no pitch or roll degrees of freedom. Therefore it is concluded that an underlying systemic sensitivity exists which is shared by all tests, and that an

additional environmental and physical sensitivity exists in the field tests.

The underlying cause of these errors could be one or a combination of several factors:

- Inaccuracies in the initialised stereo map: This would lead to increasing errors with increasing Γ as the scene rotates further away from the initial keyframes. The extent of this error could depend on the distance between keyframes or the number and vertical distribution of features in the scene.
- The relatively simplistic camera model used: Only radial distortion is assumed, and only one distortion coefficient is used to model this. In reality lenses often give rise to 'barrel' distortion, which is a function of both radial and tangential distance, requiring four or more coefficients. We would expect that inaccuracies in the distortion model would yield higher errors in the periphery of the field-of-view, which is synonymous with larger Γ .
- Inaccuracies in the zeroing process: If zeroing is carried out at non-zero Γ , a measurement bias would result. If a non-zero pitch or roll angle exists during zeroing, this would give rise to an error as a function of Γ .
- Camera positioning: Inaccuracies in the camera mounting could yield unexpected measurement errors, likely a function of Γ if there is an inherent pitch or roll offset. Correct zeroing should take care of a yaw angle offset in the camera mounting.

VIII. CONCLUSION

A concept for articulation angle estimation for articulated HGVs has been presented, utilising a rear-facing camera, and image processing based largely on the PTAM algorithm of Klein and Murray [11]. The system is entirely tractor-based, non-contact, does not require knowledge of or modification to the trailer unit/s, and runs in real-time at 20 fps (20 Hz). The system has been validated in full-scale vehicle trials on both a tractor semi-trailer combination (single articulation) and a Nordic combination (two articulations), and shown to yield good performance relative to the published literature. The system is theoretically expandable to other vehicle combinations with two or more articulation points; this will require optimisation of the camera location and field of view segmentation. Repositioning of the camera, or the use of more than one camera, may be required in some instances, especially in combinations with two or more trailers where the lead trailer is a semi-trailer such as a B-double (see Figure 1, and [35]).

This concept presents a solution to the task of tractor-based non-contact articulation angle sensing for articulated vehicles, which can assist in the commercialisation of vehicle control systems, which ultimately supports the uptake to long combination vehicles for the benefit of economies and the environment. It is expected that performance of the system can be improved in future work, through optimisation and automation of the stereo initialisation and zeroing processes, and through a more detailed investigation into sensitivity to pitch and roll motion. Additionally, recent advances in computational power may allow for more precise feature detection and tracking algorithms to be explored to further improve accuracy.

ACKNOWLEDGEMENT

The authors would like to thank Volvo Group Trucks Technology (GTT) for their support with field testing at Hällered, especially Johan Eklöv and Gustav Neander. Thanks are also extended to the CUED technical staff and colleagues in the Cambridge Vehicle Dynamics Consortium for their assistance with vehicle instrumentation and field testing (Will Midgley, Matthew Miao, Ayelet Ashkenazi, Graeme Morrison, Amy Rimmer, Leon Henderson, Yanbo Jia, Richard Roebuck).

FUNDING

This work was funded by the Cambridge Commonwealth, European and International Trust (CCEIT), the Council for Scientific and Industrial Research (CSIR, South Africa), and the Cambridge Vehicle Dynamics Consortium (CVDC). At the time of writing the Consortium consisted of the University of Cambridge with the following partners from the heavy vehicle industry: Anthony Best Dynamics, Camcon, Denby Transport, Firestone Industrial Products, Goodyear, Haldex, MIRA, SDC Trailers, Tinsley Bridge, Tridec, Volvo Trucks, and Wincanton.

REFERENCES

- [1] C. C. de Saxe, "Vision-based trailer pose estimation for articulated vehicles," PhD thesis, University of Cambridge, 2017.
- [2] I. Feige, "Freight transport and economic growth – two inseparable twins?" in *Transp. trade Econ. growth - coupled or decoupled?* Berlin Heidelberg: Springer, 2007, pp. 9–73. [Online]. Available: https://doi.org/10.1007/978-3-540-68299-8_2
- [3] P. Nordengen, R. Berman, C. de Saxe, and J. Deiss, "An overview of the performance-based standards pilot project in South Africa," in *Proceedings of the 15th International Symposium on Heavy Vehicle Transport Technology (HVT15), Oct. 2–5, 2018*. Rotterdam: International Forum for Road Transport Technology (IFRTT), 2018.
- [4] NHVR, "Performance Based Standards -An introduction for road managers," 2019. [Online]. Available: <https://www.nhvr.gov.au/files/201810-0924-pbs-a-guide-for-road-managers.pdf>
- [5] C. Cheng, A. Odhams, R. L. Roebuck, and D. Cebon, "Feedback control of semi-trailer steering at low speeds," *Submitt. to Proc. Inst. Mech. Eng. Part D J. Automob. Eng.*, 2011.
- [6] B. A. Jurnovich and D. Cebon, "Path-Following Steering Control for Articulated Vehicles," *J. Dyn. Syst. Meas. Control*, vol. 135, no. 3, p. 031006, 2013. [Online]. Available: <https://doi.org/10.1115/1.4023396>
- [7] R. L. Roebuck, A. Odhams, K. Tagesson, C. Cheng, and D. Cebon, "Implementation of trailer steering control on a multi-unit vehicle at high speeds," *J. Dyn. Syst. Meas. Control*, vol. 136, no. 2, pp. 021 016–021 016–14, dec 2013. [Online]. Available: <https://doi.org/10.1115/1.4025815>
- [8] A. J. Rimmer and D. Cebon, "Implementation of reversing control on a doubly articulated vehicle," *J. Dyn. Syst. Meas. Control*, vol. 139, no. 6, pp. 061 011–061 011–9, apr 2017. [Online]. Available: <https://doi.org/10.1115/1.4035456>
- [9] G. Morrison and D. Cebon, "Combined emergency braking and turning of articulated heavy vehicles," *Veh. Syst. Dyn.*, vol. 55, no. 5, pp. 725–749, may 2017. [Online]. Available: <https://doi.org/10.1080/00423114.2016.1278077>
- [10] L.-K. Chen and Y.-A. Shieh, "Jackknife prevention for articulated vehicles using model reference adaptive control," *Proc. Inst. Mech. Eng. Part D J. Automob. Eng.*, vol. 225, no. 1, pp. 28–42, jan 2011. [Online]. Available: <https://doi.org/10.1243/09544070JAUTO1513>
- [11] G. Klein and D. Murray, "Parallel tracking and mapping for small AR workspaces," *6th IEEE ACM Int. Symp. Mix. Augment. Real.*, nov 2007. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4538852>
- [12] J. Aurell and T. Wadman, "Vehicle combinations based on the modular concept," Nordic Road Association, Tech. Rep., 2007. [Online]. Available: <http://www.nvfnorden.org/lisalib/getfile.aspx?itemid=260>
- [13] VSE, "Product information for ETS trailer," Veenendaal, The Netherlands, 2009. [Online]. Available: <http://www.v-s-e.com/uploads/documents/leaflet-ets-trailer-en.pdf>

- [14] C. Cheng, "Enhancing safety of actively-steered articulated vehicles," PhD thesis, University of Cambridge, 2009.
- [15] B. A. Jujnovich, "Active steering of articulated vehicles," PhD thesis, University of Cambridge, 2005.
- [16] "Product overview: mechanical and hydraulic steering systems and axle suspensions." [Online]. Available: http://jic.jost-world.com/static/upload/pdf/FLY/32004_FLY_TRIDEC_Uebersicht_Image_160824_EN_SCREEN.pdf
- [17] M. Bouteldja, A. Koita, V. Dolcemascolo, and J. C. Cadiou, "Prediction and detection of jackknifing problems for tractor semi-trailer," in *IEEE Veh. Power Propuls. Conf.* Windsor: IEEE, sep 2006, pp. 1–6. [Online]. Available: <https://doi.org/10.1109/VPPC.2006.364272>
- [18] L. Chu, Y. Fang, M. Shang, J. Guo, and F. Zhou, "Estimation of articulation angle for tractor semi-trailer based on state observer," in *Int. Conf. Meas. Technol. Mechatronics Autom.*, vol. 2. Changsha City: IEEE, mar 2010, pp. 158–163. [Online]. Available: <https://doi.org/10.1109/ICMTMA.2010.342>
- [19] T. Ehlgren, T. Pajdla, and D. Ammon, "Eliminating blind spots for assisted driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 9, no. 4, pp. 657–665, 2008. [Online]. Available: <https://doi.org/10.1109/TITS.2008.2006815>
- [20] Z. Ziaukas, M. Wielitzka, T. Ortaier, and J. P. Kobler, "Simultaneous estimation of steering and articulation angle in a truck-semitrailer combination solely based on trailer signals," in *Proceedings of the American Control Conference*. Philadelphia, PA, USA: American Automatic Control Council, 2019, pp. 2509–2514.
- [21] J. Schikora, U. Berg, and D. Zöbel, "Berührungslose Winkelbestimmung zwischen Zugfahrzeug und Anhänger," in *Aktuelle Anwendungen Tech. und Wirtschaft*, ser. Informatik aktuell, W. Halang and P. Holleczeck, Eds. Springer Berlin Heidelberg, 2009, pp. 11–20. [Online]. Available: https://doi.org/10.1007/978-3-540-85324-4_2
- [22] L. Caup, J. Salmen, I. Muharemovic, and S. Houben, "Video-based trailer detection and articulation estimation," in *IEEE Intell. Veh. Symp.* Gold Coast: IEEE, jun 2013, pp. 1179–1184. [Online]. Available: <https://doi.org/10.1109/IVS.2013.6629626>
- [23] M. P. Harris, "Application of computer vision systems to heavy goods vehicles: visual sensing of articulation angle," MPhil thesis, University of Cambridge, 2013.
- [24] C. Fuchs, S. Eggert, B. Knopp, and D. Zobel, "Pose detection in truck and trailer combinations for advanced driver assistance systems," in *Proc. IEEE Intell. Veh. Symp.* Dearborn, MI: IEEE, jun 2014, pp. 1175–1180. [Online]. Available: <https://doi.org/10.1109/IVS.2014.6856547>
- [25] C. Fuchs, F. Neuhaus, and D. Paulus, "Advanced 3-D trailer pose estimation for articulated vehicles," in *IEEE Intell. Veh. Symp.* Seoul: IEEE, 2015, pp. 211–216. [Online]. Available: <https://doi.org/10.1109/IVS.2015.7225688>
- [26] C. C. de Saxe and D. Cebon, "Measurement of articulation angle by image template matching," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 2019. [Online]. Available: <https://doi.org/10.1177/0954407019833819>
- [27] A. Dahal, J. Hossen, C. Sumanth, G. Sistu, K. Malhan, M. Amasha, and S. Yogamani, "DeepTrailerAssist: Deep learning based trailer detection, tracking and articulation angle estimation on automotive rear-view camera," in *Proceedings - 2019 International Conference on Computer Vision Workshop, ICCVW 2019*, 2019, pp. 2339–2346.
- [28] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *Proc. 8th Natl. Conf. Artif. Intell. Conf. Innov. Appl. Artif. Intell.*, vol. 68, no. 2. Edmonton, Alberta: AAAI, 2002, pp. 593–598.
- [29] M. Montemerlo, S. Thrun, D. Roller, and B. Wegbreit, "FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," *Proc. Sixt. Int. Jt. Conf. Artif. Intell.*, pp. 1151–1156, 2003. [Online]. Available: <https://doi.org/10.1109/IROS.2006.282528>
- [30] G. Grisetti, R. Kummerle, C. Stachniss, and W. Burgard, "A tutorial on Graph-Based SLAM," *IEEE Intell. Transp. Syst. Mag.*, vol. 2, no. 4, pp. 31–43, 2010. [Online]. Available: <https://doi.org/10.1109/MITS.2010.939925>
- [31] A. Angeli, S. Doncieux, J.-A. Meyer, and D. Filliat, "Visual topological SLAM and global localization," in *2009 IEEE Int. Conf. Robot. Autom.* IEEE, may 2009, pp. 4300–4305. [Online]. Available: <https://doi.org/10.1109/ROBOT.2009.5152501>
- [32] J. Engel, T. Sch, and D. Cremers, "LSD-SLAM: Large-Scale Direct monocular SLAM," in *Proc. 13th Eur. Conf. Comput. Vis.*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds., no. Part 2. Zurich: Springer, 2014, pp. 834–849. [Online]. Available: https://doi.org/10.1007/978-3-319-10605-2_54
- [33] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, oct 2015. [Online]. Available: <https://doi.org/10.1109/TRO.2015.2463671>
- [34] R. Mur-Artal and J. D. Tardos, "ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, 2017. [Online]. Available: <https://doi.org/10.1109/TRO.2017.2705103>
- [35] D. Cebon and C. C. de Saxe, "WO/2019/202317 - Method and system of articulation angle measurement," apr 2019. [Online]. Available: <https://patentscope.wipo.int/search/en/detail.jsf?docId=WO2019202317>
- [36] H. Stewénius, C. Engels, and D. Nistér, "Recent developments on direct relative orientation," *ISPRS J. Photogramm. Remote Sens.*, vol. 60, no. 4, pp. 284–294, jun 2006. [Online]. Available: <https://doi.org/10.1016/j.isprsjprs.2006.03.005>
- [37] E. Rosten, R. Porter, and T. Drummond, "Faster and better: a machine learning approach to corner detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 1, pp. 105–119, jan 2008. [Online]. Available: <https://doi.org/10.1109/TPAMI.2008.275>
- [38] A. W. Bridgewater, "Analysis of second and third order steady-state tracking filters," in *AGARD Conf. Proc. no. 252 Strateg. Autom. Track Initiat.*, S. J. Rabinowitz, Ed. Monterey, CA: AGARD, jun 1979.
- [39] R. Hartley and A. Zisserman, *Multiple View Geometry*, 2nd ed. Cambridge, UK: Cambridge University Press, 2003.
- [40] R. Cipolla, "Computer vision and robotics: projection (4F12 lecture notes, Cambridge University Engineering Department)," 2013.
- [41] F. Devernay and O. Faugeras, "Straight lines have to be straight," *Mach. Vis. Appl.*, vol. 13, no. 1, pp. 14–24, 2001. [Online]. Available: <https://doi.org/10.1007/PL00013269>
- [42] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle Adjustment — A Modern Synthesis," in *Vis. Algorithms Theory Pract.* Springer Berlin Heidelberg, 2000, vol. 1883, pp. 298–372. [Online]. Available: https://doi.org/10.1007/3-540-44480-7_21
- [43] "Parallel Tracking And Mapping for small AR workspaces: Licence and source code download page," 2008. [Online]. Available: <http://www.robots.ox.ac.uk/~gk/PTAM/download.html>
- [44] S. M. La Valle, "Geometric representations and transformations," in *Plan. algorithms*, 1st ed. Cambridge University Press, 2006, ch. 3, pp. 81–126. [Online]. Available: <http://planning.cs.uiuc.edu/>



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 Associate in the Centre for Sustainable Road Freight at the University of Cambridge, and a Visiting Lecturer in the School of Mechanical, Industrial and Aeronautical Engineering at the University of the Witwatersrand.



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.