

Visual Servoing of a Car-like Vehicle – An Application of Omnidirectional Vision

Kane Usher, Peter Ridley

School of Mechanical, Medical &
Manufacturing Engineering
QUT, Brisbane 4000.
kane.usher@csiro.au

Peter Corke

CSIRO Manufacturing &
Infrastructure Technology
P.O. Box 883, Kenmore 4069.
peter.corke@csiro.au

Abstract

In this paper, we develop the switching controller presented by Lee et al. [Lee *et al.*, 1999] for the pose control of a car-like vehicle, to allow the use of an omnidirectional vision sensor. To this end we incorporate an extension to a hypothesis on the navigation behaviour of the desert ant, *cataglyphis bicolor*, which leads to a correspondence free landmark based vision technique. The method we present allows positioning to a learnt location based on feature bearing angle and range discrepancies between the robot's current view of the environment, and that at a learnt location. We present simulations and experimental results, the latter obtained using our outdoor mobile platform.

1 Introduction

A key skill for the autonomous navigation of mobile robots is the ability to find particular locations in a workspace. Further to this, in order to perform useful tasks, a mobile robot requires the ability to servo to particular poses in the environment. For the nonholonomic, car-like vehicle used in these experiments, Brockett [Brockett, 1983], showed that there is no smooth, continuous control law which can locally stabilise such systems.

Insects in general display amazing navigation abilities, traversing distances far surpassing the best of our mobile robots on a relative scale. To do this, evolution has provided insects with many 'shortcuts' enabling the achievement of relatively complex tasks with a minimum of resources in terms of processing power and sensors [Weber *et al.*, 1999]. In particular, the high ground temperatures encountered by the desert ant, *cataglyphis bicolor*, eliminates pheromones as a potential navigation aid, as is used by ants in cooler climates [Lambrinos *et al.*, 2000]. The desert ant navigates using a combination of path integration and visual homing.

Visual homing, also known as visual piloting, is the process of matching an agent's current view of a location in a

distinctive locale to a (pre-stored) view at some target position. Discrepancies between the two views are used to generate a command which drives the agent closer to the target position. The process enables the agent to 'find' positions in distinctive locales. These distinctive locales can then be linked to generate paths through an environment [Hong *et al.*, 1991; Kuipers and Byun, 1991], eliminating the need for complex map-like representations, instead embedding this knowledge in terms of what the agent's sensors can 'see'.

When applied to nonholonomic mobile robots, the constraints of Brockett's theorem prevent the insect-based strategies from completely resolving the pose stabilisation problem; they enable servoing to a position but cannot guarantee a particular orientation (see e.g. [Lambrinos *et al.*, 2000; Weber *et al.*, 1999] – it should be noted that these vehicles could spin on their own axis, and thus orientation errors on reaching 'home' could be resolved with a simple switching strategy – the problem is more complex for a car-like vehicle).

In the control community, the problem of stabilising a mobile robot to a specific pose has generally been approached from two directions; the open-loop and closed-loop strategies. Open-loop strategies seek to find a bounded sequence of control inputs, driving the vehicle from an initial position to some arbitrary position, usually working in conjunction with a motion planner (e.g. [Latombe, 1991; Murray and Sastry, 1993]). Finding this sequence of control inputs is difficult for the nonholonomic case and, in the event of disturbances, a new plan has to be formulated.

The closed-loop strategies consist of designing a feedback loop using proprioceptive and exteroceptive sensors to provide estimates of the vehicle's state (see e.g. [Luca *et al.*, 1997]). Feedback control systems are generally more robust to uncertainty and disturbances when compared to their open-loop counterparts. All real mobile robots and sensors are subject to noise and uncertainty — feedback control would thus seem essential. However, due to the well-known limitations presented by Brockett, there is no smooth, continuous control law which can locally stabilise closed-loop nonholonomic systems to a pose [Brockett, 1983]. These

limitations can be side-stepped by either relaxing the constraints on desired pose (i.e. stabilising to a point without a guarantee on orientation, as for the insect inspired approaches), using discontinuous control techniques (see e.g. [Aicardi *et al.*, 1995; Conticelli *et al.*, 1999; Astolfi, 1995; Lee *et al.*, 1999]), or by using time-varying control (see e.g. [Samson and Ait-Abderrahim, 1991]). Much of the literature does not address what has been found in this study to be a significant limitation of many control algorithms for the pose stabilisation of car-like vehicles — input saturation.

For the task at hand, we argue the case for the feedback control methods, with vision used as the primary sensor. However, it is clear that a measure of open-loop planning is also usually required in order to prevent deadlock situations from occurring. Here, we are concerned only with the feedback component for the task of servoing to a position and orientation. The use of visual feedback is finding increasing application in the solution to this problem (see e.g. [Hashimoto and Noritsugo, 1997; Murrieri *et al.*, 2002; Conticelli *et al.*, 1999; Das *et al.*, 2001]), but the use of omnidirectional camera systems has not yet been fully explored.

The remainder of this paper is arranged as follows: Section 2 details the switching control technique and some of the adaptations we have made to cope with noisy vision data and to ensure stability; Section 3 describes our experimental system and briefly outlines some preliminary results; Section 4 concludes the paper and presents some directions of future interest.

2 Control Strategy

In this section, the switching control strategy presented by Lee *et al.* [Lee *et al.*, 1999] is developed. We then show how a derivative of the insect inspired Average Landmark Vector model of navigation presented by Lambrinos *et al.* [Lambrinos *et al.*, 2000] can be used to provide the required quantities to the switching controller.

2.1 Kinematics

Our vehicle is (approximately) car-like in its kinematics. Referring to Figure 1, the kinematics in Cartesian space of our experimental vehicle are:

$$\begin{aligned} \dot{x} &= v \cos \theta \\ \dot{y} &= v \sin \theta \\ \dot{\theta} &= v \frac{\tan \phi}{L} \end{aligned} \quad (1)$$

where v is the vehicle's forward velocity (measured at the centre axle of the rear wheels), L is the vehicle's length, ϕ is the steering angle, and the point (x, y) refers to the centre of the front axle. (It is more usual to define the coordinates of a rear-wheel driven vehicle with respect to the centre of the rear axle. However, on our vehicle, the sensors which provide position information are mounted over the front axle and thus it makes more sense to use this as the reference point — it is a simple transformation between the front and rear points on the vehicle.)




Figure 1: The vehicle and the coordinate system used. All angles are counter-clockwise positive.

2.2 Control Law Development

The technique presented by Lee *et al.* [Lee *et al.*, 1999] operates in Cartesian space and sidesteps Brockett's theorem via a discontinuity in the control law. The control law operates in two distinct stages. Without loss of generality, consider the goal pose to be $(x, y, \theta) = (0, 0, 0)$. The initial stage involves minimising y and θ ; i.e. the vehicle converges to the x -axis with an orientation of 0. The second stage then moves the vehicle along the x -axis to the desired point. Each controller is described in the following sections.

Control Law — First Stage

The first stage of the control law takes the vehicle to the x -axis with the correct orientation. This is achieved via the minimisation of the Lyapunov-like function [Lee *et al.*, 1999]

$$V = \frac{1}{2}k_1y^2 + \frac{1}{2}\theta^2$$

This function is radially unbounded and positive semi-definite. Its derivative is:

$$\begin{aligned} \dot{V} &= k_1y\dot{y} + \theta\dot{\theta} \\ &= k_1yv \sin(\theta) + \theta v \frac{\tan(\phi)}{L} \end{aligned}$$

With the following choice for the steering angle,

$$\phi = \arctan \left[-\frac{L}{v} \left(k_2\theta + k_1v \frac{\sin(\theta)}{\theta} y \right) \right]$$

Lee *et al.* [Lee *et al.*, 1999] showed that y and θ converge to zero.

The velocity of the vehicle is chosen according to its initial orientation with respect to the goal. If the vehicle is facing the goal, it is given a constant positive velocity, else it is given a constant negative velocity, i.e.

$$v = \begin{cases} k_3 & \cos \phi_{initial} < 0 \\ -k_3 & \cos \phi_{initial} \geq 0 \end{cases}$$

where ϕ is the initial orientation of the goal relative to the vehicle.

The velocity v could be switched at any time, depending on obstacle layout etc.. Lee addressed local minima problems via the use of ‘intermediate’ points. At this stage in our work, we assume there are no obstacles, and we simply set the velocity equal to a constant, k_3 , with the sign assigned as explained above. However, we set a limitation on the distance the vehicle can be from the goal, and reverse the velocity if this distance is exceeded.

Control Law — Second Stage

On reaching the x-axis and $\theta = 0$, the following control law is invoked:

$$v = -k_3x$$

$$\phi = \arctan \left[-\frac{L}{v} \left(k_2\theta + k_1v \frac{\sin(\theta)}{\theta} y \right) \right]$$

bringing the vehicle to the desired pose of $(x,y,\theta) = (0,0,0)$. In the original work of Lee et al., the steering angle was set to zero in the second stage of control. We have found that in practice, due to the imperfections of ‘real’ systems, it is necessary to control the steering angle in the second stage of control. An analysis of the Lyapunov function for this case shows that convergence is still guaranteed with the above choice of steering angle.

In combination, these control laws stabilise the vehicle to the desired pose from any initial condition. Again in practice, the imperfections of ‘real’ systems requires that switching to the second stage of control occurs when the y-axis distance and the error in orientation (θ) drop below pre-specified thresholds.

The control technique presented above assumes that the distance and angles to the target location can be measured, from which the quantities required for the controller can be calculated. In the next section, we detail a landmark-based vision technique which yields this information without explicit knowledge of the robot’s pose; sensing is provided by both an omnidirectional camera and a compass.

2.3 Ant navigation

The desert ant, *cataglyphis bicolor*, is unable to use pheromones to navigate due to the high ground temperatures found in its habitat. Thus, it relies on a combination of visual piloting and path integration, enabling it to find nest openings of less than a few millimetres after foraging journeys of several hundred metres [Wehner *et al.*, 1996; Lambrinos *et al.*, 2000]; see Figure 2 for an example of a foraging journey. On this journey, the ant has travelled a round trip distance of over two hundred metres, returning to a nest with an opening of less than 5mm, equating to a drift rate of less than 0.0025% in its navigation system. Comparing this to high-end commercial inertial navigation units for land based navigation, which have drift rates of the order 0.1% of distance travelled, demonstrates the effectiveness of this insects navigation system.

Here however, we are interested in the visual piloting component of the ants navigation system. When visual piloting, ants take a rather unprocessed view of the target location and match it with a current view, using the discrep-

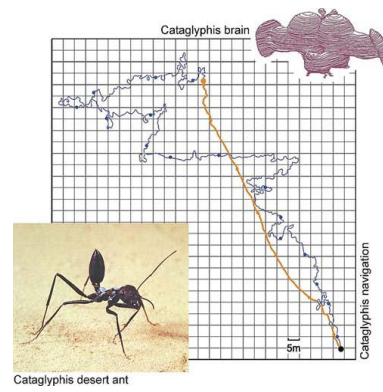


Figure 2: An example of the amazing navigation feats of the humble ant. Diagram courtesy www.neuroscience.unizh.ch/e/groups/wehner00.htm, Rudiger Wehner.

ancies to derive a direction of movement. Matching characteristics used are the differences in landmark bearings, apparent size and apparent height [Lambrinos *et al.*, 2000].

An elegant, correspondence free, homing method developed from hypotheses on how these ants might use visual piloting is the Average Landmark Vector model. An ALV for any particular position in the workspace is found by summing unit vectors towards all currently visible landmarks, and dividing by the number of landmarks. By matching the current ALV with a pre-stored ALV of the target location, a homing vector can be formed which drives the agent (robot) towards the target location [Lambrinos *et al.*, 2000]. In order to consistently add the vectors in the ALV model, an absolute reference direction is required, and, unless apparent size information is incorporated, a minimum of three landmarks is needed.

The nature of our sensor, an omnidirectional vision system, led us to investigate improvements to the ALV method. In its original form, the ALV method required the bearings to landmarks only. Range information could be incorporated by including landmark apparent size, and slight improvements to the performance could be made. However, we have found that by including range information directly, a significant improvement is made and in fact, the distance and angle to the goal are yielded directly. In addition, the minimum required landmarks is reduced to one. An example of the IALV method is shown in Figure 3. In essence, the IALV method is equivalent to finding a position relative to the centroid of the landmarks in the workspace.

As with the ALV method, the IALV method is purely sensor-based. Landmark bearings are readily ascertained with an omnidirectional camera. If a flat-earth assumption is made, range information can be derived from an omnidirectional camera image through the geometry of the camera/mirror optics. Alternatively, optic flow techniques could be used to determine landmark range (see e.g. [Chahl and Srinivasan, 1997]).

One of the advantages of the ALV, and hence the IALV method, is that knowledge of a target location is contained within a single quantity. This reduces the need for complex




Figure 3: Illustration of the the IALV method for two landmarks in a workspace. The IALV’s are found by adding the vectors to the individual landmarks, and dividing the resulting vector by the number of landmarks. The home vector is then calculated by subtracting the target IALV from the currently IALV.

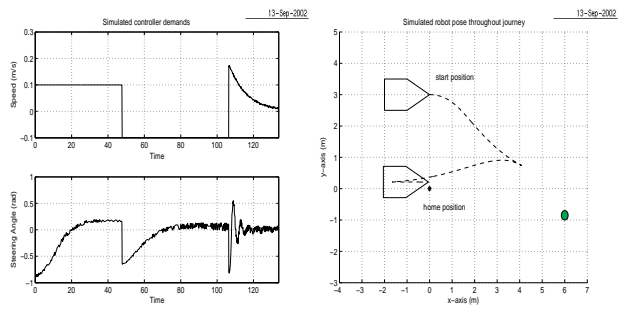
map-like representations of the environment and is well suited for a topological navigation method, (see e.g. [Levitt and Lawnton, 1990; Gaspar *et al.*, 2000; Mataric, 1992; Kuipers and Byun, 1991]). Additionally, landmarks need not be unique, and the need for landmark correspondence is also bypassed. Many of the other homing algorithms require that the landmarks in the current image be matched with those at the target location, usually by minimisation of the sum of the bearing differences (see e.g. [Weber *et al.*, 1999]). If landmarks are occluded or missing, these methods can fail. Of course, like all sensor-based techniques, this method has a finite catchment area, limited by the omnidirectional sensor’s range and, in addition, has the potential to suffer from perceptual aliasing, or in a similar sense, the local minima problem.

The homing vector provided by the IALV method can be used to drive the agent towards home but does not provide a means of guaranteeing a final orientation. However, the quantities derived from the IALV can easily be converted to the states required by the switching controller developed earlier.

2.4 Simulations

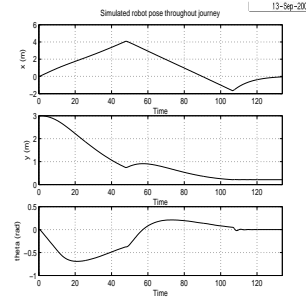
In this simulation, landmarks are modelled as points, consistent with the requirements of the IALV model which merely requires the range and bearing to one or more non-unique landmarks in the workspace. An IALV is taken at a target location, nominally the origin, with this target IALV being matched with the agent’s current IALV throughout the vehicle’s journey.

The vehicle model used is simply the kinematic equations (as given in equation 1). Also included in the model are input saturation and first-order delays in the steering and velocity loops. A simulation of the omnidirectional sensor (with random noise) provides the IALV, from which the required states are derived and fed to the switching controller. Figure 4 shows the path generated, pose, and control inputs for a starting pose of $(x,y,\theta) = (0,3,0)$ and a goal pose



(a) Demanded inputs.

(b) Vehicle’s path.



(c) Pose throughout journey.

Figure 4: Results of the simulation for a starting pose of $(x,y,\theta) = (0,3,0)$.

of $(0,0,0)$. Gains were set to $(k_1,k_2,k_3) = (0.35,0.1,0.1)$ for these simulations, based on an analysis of the linearised system. The method works for all starting and goal poses.

The method has been extensively tested in simulation and found to be quite robust to input saturation and noise. The next section presents some experimental results which validate our simulations.

3 Experiments

Experimental validation of this visual servoing technique was conducted on our outdoor research testbed. In these experiments, artificial landmarks in the form of red traffic cones (witches hats) were used. The following sections give a brief overview of the vehicle and outline the image processing used to extract the landmarks.

3.1 Robotic testbed

The experimental platform is a Toro ride-on mower which has been retro-fitted with actuators, a control system, and a computer, enabling control over the vehicle’s operations. All control and computing occurs on-board. The vehicle is fitted with an array of sensors including odometry, GPS, a magnetometer, a laser range-finder (for collision avoidance only) and an omnidirectional camera (see Figure 5 for a photograph of the vehicle). For the experiments cited here, the primary sensor used is the omnidirectional cam-



Figure 5: The experimental platform. Note the omnidirectional camera mounted over the front wheels and the box at the rear which houses the control and computer system.

era with the magnetometer providing an absolute reference direction.

3.2 Image processing

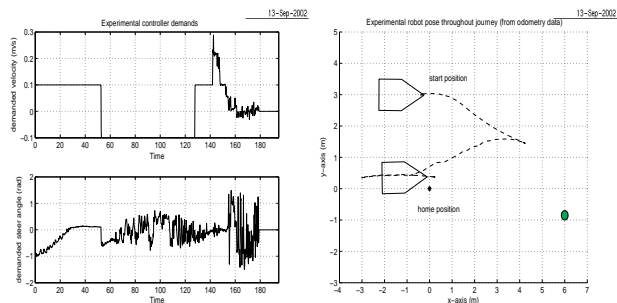
The distinguishing feature of the landmarks used in this experiment is their colour. Hence, we use colour segmentation to track them; not a trivial task in an outdoor environment with no control over lighting conditions. Our frame-grabber provides a YCrCb signal, and, to reduce processing time, we work directly in this colour space.

We use a bivariate histogram on the Cr and Cb of the objects we wish to track, creating a two-dimensional lookup table on colour. As each image is acquired, pixels that fall within the histogram are flagged as belonging to a landmark. Blob extraction is then performed, the objects sorted by size, and very small and very large objects eliminated. The image coordinates of the centroid of each blob are then converted to a polar form, giving the radial distance and bearing of each blob with respect to the centre of the image. At no stage do we attempt to ‘unwrap’ the image; we believe this to be a waste of valuable processing time and instead work directly with the radial image.

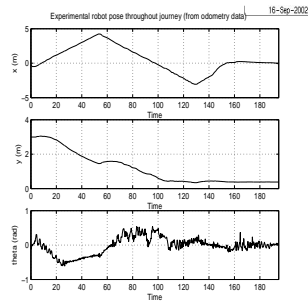
The relative landmark bearings are then combined with the robot’s orientation obtained from the magnetometer to give bearings with respect to the reference direction. Landmark range however is not directly available; we use a flat-earth assumption and knowledge of the camera/mirror optics to derive an estimate of range given a landmarks radial pixel distance from the centre of the image [Horswill, 1993; Das *et al.*, 2001].

3.3 Results and Discussion

The testing arena used for these experiments is a large shed. The vehicle was placed in the middle of the shed defining the goal pose $(x, y, \theta) = (0, 0, 0)$. A ‘landmark’ was then placed at $(x, y) = (5.85, -1)$, and the target IALV was found. The vehicle was then manually moved to $(x, y, \theta) = (0, 3, 0)$, and the control system was activated. Figure 6 depicts the results of the experiment, showing the vehicle servoing to the goal pose based on vision and compass data



(a) Demanded inputs. (b) Vehicle’s path.



(c) Pose throughout journey.

Figure 6: Experimental results for a starting pose of $(x, y, \theta) = (0, 3, 0)$. Note the similarity to the results achieved with the simulation plotted in Figure 4.

alone. Although the vision data is extremely noisy, the system seems quite robust to this.

The similarity to the simulation results plotted in Figure 4 is quite clear, although the experimental system did take longer to stabilise to the desired pose. This could be due to the quite coarse onboard velocity measurements. Also, the rather noisy data provided by the vision system could be in effect here.

4 Conclusion

This paper has described a method of stabilising a car-like vehicle to a target pose based on the discrepancies between a target view of the landmarks in a workspace and the robot’s current view.

The robot’s view of the workspace is summarised by a single quantity, the Improved Average Landmark Vector, which augments the original formulation of Lambrinos *et al.* [Lambrinos *et al.*, 2000] with range information. At each instant, the robot compares the current IALV with that at the target location which yields directly the distance and orientation to the target position. At no stage is there a requirement for landmark correspondence; this represents a key advantage over many other homing methods. This information is then fed into switching controller which stabilises the vehicle to the desired pose based on the information pro-

vided by the vision system. Furthermore, the method uses polar representations of the landmarks, and thus, the need for processor intensive image unwarping is circumvented.

We have presented both simulations and real experimental results showing the validity of our approach even in the face of input saturation and noise. Future work includes extensive experimental validation of the method, along with a deeper understanding of the impact of dynamic elements in the control loops. Finally, it is clear that this technique is ideally suited to topological navigation, and this too represents a direction for further research.

Acknowledgements

The technical support of the CMIT Tractor Team — Peter Hynes, Stuart Wolfe, Stephen Brosnan, Graeme Winstanley, Pavan Sikka, Elliot Duff, Les Overs, Craig Worthington and Steven Hogan — is gratefully acknowledged, while Jonathon O'Brien at UNSW is gratefully thanked for the loan of the Toro ride-on mower.

The first author gratefully acknowledges the funding and technical support provided by CMIT, the funding provided by an APA grant through QUT and would also like to thank Peter Corke and Peter Ridley for their guidance and support throughout this course of research.

References

- [Aicardi *et al.*, 1995] Michele Aicardi, Giuseppe Casalino, Antonio Bichi, and Aldo Balestrino. Closed loop steering of unicycle-like vehicles via Lyapunov techniques. *IEEE Robotics and Automation Magazine*, pages 27–35, March 1995.
- [Astolfi, 1995] A. Astolfi. Exponential stabilization of a car-like vehicle. In *Proceedings of the International Conference on Robotics and Automation*, pages 1391–1396, Nagoya, Japan, 1995. IEEE.
- [Brockett, 1983] R.W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. S. Millman, and H. J. Sussman, editors, *Differential Geometric Control Theory*, pages 181–191. Birkhauser, Boston, USA, 1983.
- [Chahl and Srinivasan, 1997] J. S. Chahl and M. V. Srinivasan. Range estimation with a panoramic visual sensor. *Journal of the optical society of America*, 14(9), September 1997.
- [Conticelli *et al.*, 1999] Fabio Conticelli, Bendetto Allota, and Pradeep K. Khosla. Image-based visual servoing of nonholonomic mobile robots. In *Proceedings of the 38rd Conference on Decision and Control*, pages 3496–3501. IEEE, Phoenix, Arizona, USA, December 1999.
- [Das *et al.*, 2001] A.K. Das, R. Fierro, V. Kumar, B. Southall, J. Spletzer, and C.J. Taylor. Real-time vision based control of a non-holonomic mobile robot. In *International Conference on Robotics and Automation*, pages 1714–1719, Seoul, Korea, May 2001. IEEE.
- [Gaspar *et al.*, 2000] Jose Gaspar, Niall Winters, and Jose Santos-Victor. Vision-based navigation and environmental representations with an omnidirectional camera. *IEEE Transactions on Robotics and Automation*, 16(6):890–898, December 2000.
- [Hashimoto and Noritsugo, 1997] Koichi Hashimoto and Toshiro Noritsugo. Visual servoing of nonholonomic cart. In *International Conference on Robotics and Automation*, pages 1719–1724, Albuquerque, New Mexico, USA, April 1997. IEEE.
- [Hong *et al.*, 1991] Jiawei Hong, Xiaonan Tan, Brian Pinette, Richard Weiss, and Edward M. Riseman. Image-based homing. In *International Conference on Robotics and Automation*, pages 620–625, Sacramento, California, USA, April 1991. IEEE.
- [Horswill, 1993] I. Horswill. Polly: A vision-based artificial agent. In *Proceedings of the eleventh national conference on artificial intelligence (AAAI'93)*, Washington DC, USA, July 1993. MIT Press.
- [Kuipers and Byun, 1991] Benjamin Kuipers and Yung-Tai Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 8:47–63, 1991.
- [Lambrinos *et al.*, 2000] Dimitrios Lambrinos, Ralf Moller, Thomas Labhart, Rolf Pfeifer, and Rudiger Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, 30:39–64, 2000.
- [Latombe, 1991] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic, 1991.
- [Lee *et al.*, 1999] Sungon Lee, Manchul Kim, Youngil Youm, and Wankyun Chung. Control of a car-like mobile robot for parking problem. In *International Conference on Robotics and Automation*, pages 1–6, Detroit, Michigan, 1999. IEEE.
- [Levitt and Lawnton, 1990] Tod S. Levitt and Daryl T. Lawnton. Qualitative navigation for mobile robots. *Artificial Intelligence*, 44:305–360, 1990.
- [Luca *et al.*, 1997] Alessandro De Luca, Giuseppe Oriolo, and Claude Samson. Feedback control of a nonholonomic car-like robot. In J.-P. Laumond, editor, *Planning Robot Motion*, chapter 4. Springer Verlag, 1997.
- [Mataric, 1992] Maja J. Mataric. Integration of representation into goal-driven behaviour-based robots. *IEEE Transactions on Robotics and Automation*, 8(3):304–312, June 1992.
- [Murray and Sastry, 1993] Richard M. Murray and S. Shankar Sastry. Nonholonomic motion planning: Steering using sinusoids. *IEEE Transactions on Automatic Control*, 38(5):700–716, May 1993.
- [Murrieri *et al.*, 2002] Pierpaolo Murrieri, Daniele Fontanelli, and Antonio Bicci. Visual-servoed parking with limited view angle. In *Preprints of the International Symposium on Experimental Robotics*, Sant'Angelo d'Ischia, Italy, 2002.
- [Samson and Ait-Abderrahim, 1991] C. Samson and K. Ait-Abderrahim. Feedback control of a non-holonomic wheeled cart in cartesian space. In *International Conference on Robotics and Automation*, pages 1136–1141, Sacramento, California, USA, April 1991. IEEE.
- [Weber *et al.*, 1999] Keven Weber, Svetha Venkatesh, and Mandyam Srinivasan. Insect-inspired robotic homing. *Adaptive Behavior*, 7(1):65–97, 1999.
- [Wehner *et al.*, 1996] R. Wehner, B. michel, and P. Antonsen. Visual navigation in insects: Coupling of egocentric and geocentric information. *Journal of Experimental Biology*, 199:129–140, 1996.