VISION-BASED DOCKING FOR BIOMIMETIC WHEELED ROBOTS [1]

Ian R. Manchester * **Andrey V. Savkin** *

* *School of Electrical Engineering and Telecommunications, University of New South Wales, UNSW Sydney, 2052, Australia.*
Email: ianm@student.unsw.edu.au, A.Savkin@unsw.edu.au

Abstract: We present a new control law for the problem of docking a wheeled robot at a certain location with a desired heading. Recent research into insect navigation has inspired a solution which uses just one environment sensor: a video camera. The control law is of the "behavioral" or "reactive" type, in that no attempt is made to observe the relative pose of robot and target, all control actions are based on immediate visual information. Docking success under certain conditions is proved mathematically, and simulation studies show the control law to be robust to camera calibration errors.

Keywords: Robots, Vision, Biocybernetics

1. INTRODUCTION

It is currently very popular among roboticists to draw inspiration from the animal kingdom Arkin (1998); Bar-Cohen and Breazeal (2003); Franz and Mallot (2000). This trend is termed "biomimetics". Robot navigation strategies thus derived, often categorized as "behavioral" or "reactive" robotics, aim at the construction of simple control strategies which use direct sensory information, rather than a structured environmental model. Such strategies are essential when robots must operate in a complex environment using simple sensors. Frequently, a complete environment model would be difficult or impossible to reconstruct with the sensed information.

In this paper we propose one such strategy for the problem of positioning a wheeled robot at a certain location with a certain heading, i.e. *docking*, using information provided by a video camera. The kinematics of the robot are non-holonomic, so standard techniques of visual servoing (see, e.g., Hutchinson et al. (1996)) cannot be directly applied. We introduce a change of variables and a camera space regulation condition which allow solution of the problem via a relatively simple nonlinear control law.

This paper draws on previous work in precision missile guidance Manchester and Savkin (2004, 2002); Manchester et al. (2003); Savkin et al. (2003). This work involved missile guidance with an impact-angle constraint, and was built on a combination of geometrical considerations, and recent work in robust control and filtering theory Petersen et al. (2000); Petersen and Savkin (1999).

The remarkable ability of honeybees and other insects like them to navigate effectively using very little information was a source of inspiration. In particular, the work of Srinivasan and his co-authors, explaining the use of optical-flow in honeybee navigation (Srinivasan et al. (2000) and references therein). His work has previously prompted work on helicopter navigation Barrows

[1] This work was supported by the Australian Research Council

et al. (2003) and missile-guidance systems Manchester et al. (2003).

A bee's eyes are immobile and fixed focus, and are not sufficiently separated for stereopsis to be of any real use. With these sensors, and such minute brains for processing, it seems that accurate estimation of distances is quite beyond them. However, they still manage to make smooth landings on surfaces, and find their way to and from the hive, for example. Recent studies by Srinivasan and others have indicated that sensing optical flow is one of their most useful navigation tools.

Many studies have been done, but one particularly striking example is the algorithm a bee uses to land on a flat surface. The bee looks down at the ground, and measures its optical flow, or angle-rate. It is straightforward to show that, if it keeps this optical flow constant, and keeps its vertical velocity a constant proportion of its horizontal velocity, it will make a smooth landing on the surface. We refer the reader to Srinivasan et al. (2000) for details.

Nature has evolved many such simple "vision-space" strategies, keeping certain angles and optical flows constant, which lead to effective behaviour in physical space (or configuration space). The docking strategy we present is directly modelled after these.

Most studies of this problem can be roughly grouped into two approaches. One focuses on the robot's "configuration space", i.e. the relative positions and angles of the robot and target, and perhaps obstacles, in the plane. All these relations are assumed to be available to the control law, and from them it chooses some desirable path. Examples are found in de Wit and Sørdalen (1992); Laumond (1998); Souères and Laumond (1996); Kelly and B.Nagy (2003) and references therein.

The method described in de Wit and Sørdalen (1992) is similar in its approach to the method presented in this paper, in that the aim is to follow to a circular path. The main differences are that, firstly, they assume a slightly simpler kinematic model (often termed the *unicycle* model), and secondly, they are able to prove exponential stabilization to the desired final location, but at the expense of a control law which is more complicated and requires more information.

The other main approach focuses on "camera space" or "visual space". It is no longer assumed that the robot has access to the full configuration, but only the image of the target (and obstacles) as the camera sees them. Typically it also knows how they ought to look if the goal is achieved. From this information a control law is assigned which drives the appearance of the target towards its goal. That is, dynamics are examined in camera space. Examples of this approach are found in the papers Santos-Victor and Sandini (1997); Hashimoto and Noritsugu (1997); Lee et al. (1999); Conticelli et al. (1999); Zhang and Ostrowski (2002); Cárdenas et al. (2003) and references therein.

Our paper can be seen as a blend of the two approaches. A simple camera-space condition is defined which, if kept, leads to desirable configuration-space trajectories.

2. PROBLEM STATEMENT

Our aim is to design a control law by which a car-like vehicle may dock to a target point. The information available to the control law is consistent with the use of a video camera as the main sensor.

We now described the kinematic model of the robot, the measurements available to it, and finally give a complete definition of the problem statement.

The relative position of vehicle and target is given in polar form (see Figure 1). The vehicle's position is an extension-less point in the plane, and is identified in a physical system with the mid-point of the rear axle. The scalar quantity r is the range between the vehicle and the target, and the angle ε is the angle between the desired heading, and the line-of-sight from the car to the target. These two quantities can be thought of as polar coordinates, placing the vehicle with respect to the fixed target frame.

Two more angles are required to completely characterize the state of the system. These are the heading of the vehicle, and the angle of its steering wheels. The angle λ is the angle between the vehicles current heading and the line-of-sight. The angle ϕ is the angle of the steering wheels, with respect to the centerline of the car, and is controlled with the input u. The forward speed is controlled with the input v.

The reason for this unusual representation of the state will become clear later in the paper, when the *CNG Principle* is described, and the control law derived.

The state-space of the car-target system is then the manifold $\mathbb{R} \times \mathbb{T}^3$ of states $(r, \lambda, \varepsilon, \phi)$, where $\mathbb{T}$ is the circle group: $\mathbb{R} \mod 2\pi\mathbb{Z}$.

The equations of motion on this manifold are given by the following differential equations. These are given for a front-wheel-drive car. To make our control law independent of the forward-velocity of the car, the dynamics are derived with respect to path length, not time.

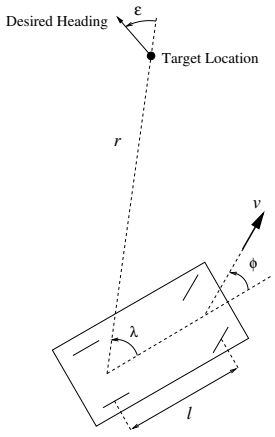

Fig. 1. System geometry

The change of variables $ds = v\cos\phi dt$ allows us to pass from one representation to another

Hereafter, x' denotes derivative of a variable x with respect to path length s. The dynamics of the states in this form are given below:

$$\lambda' = \frac{\sin\lambda}{r} - \frac{\tan\phi}{l}, \quad \varepsilon' = \frac{-\sin\lambda}{r}, \tag{1}$$

$$r' = -\cos\lambda, \quad \phi' = \frac{u}{v\cos\phi}. \tag{2}$$

Where l is the distance between the front wheels and rear wheels.

We now discuss the measurements available. Inspired by the elegant instinctual behavior of insects, and the practical need for controlling vehicles with simple sensors, we use a measurement model consistent with a single video camera mounted on the robot, and a simple optical flow algorithm.

The main restriction felt with this model is that the range to the the target, r, is not directly measurable. Furthermore, in certain situations it is unobservable, or weakly observable, from the measurements we do have. For this reason we do not use this quantity in our control law.

The angular position of the dock-target in the field of view is the angle λ. The derivative of this variable is the optical flow of the image. A simple algorithm such as Srinivasan (1994), as was used in Manchester et al. (2003), can calculate this value with very little computation.

The angle ε must be known, as it is not an environmental variable, but part of the problem statement. Two possibilities of how it might be calculated are: (a) Visual analysis of the dock-target image may allow us to judge the angle between the line-of-sight and the target-heading. Or (b) If the target-heading is defined as an abstract bearing, the heading of the vehicle could be dead-reckoned and from this and the angle λ, ε could be calculated. Detailed considerations of this issue are beyond the scope of this paper.

Further to the information from the video camera, we need some knowledge of the internal state of the vehicle. Specifically, we assume knowledge of the forward speed v, the angle of the steering wheels ϕ and the distance between the axles l.

2.1 Complete Problem Statement

Our complete problem statement is this. To find a control law of the form

$$u = f(l, \phi, v, \varepsilon, \lambda, \dot{\lambda}) \tag{3}$$

such that range and angle error at final time, i.e. $r(T)$ and $\varepsilon(T)$, are minimized Corresponding to this, we make the following definition:

Definition 1. A docking manoeuvre is considered perfect if there exists some finite time T such that

$$r(T) = 0,$$
$$\lim_{t\to T} \varepsilon(t) = 0.$$

A limit is used in the above definition because if $r = 0$ the angle ε is undefined.

3. CONTROL LAW

From the optical flow measurements, we can cancel the component due to the robot's rotation ($= v\sin\phi/l$), and retain only the component due to the relative motion of robot and dock-target. We denote this remaining flow $\mathcal{O}_f$, so:

$$\mathcal{O}_f := \dot{\lambda} + \frac{v\sin\phi}{l} \tag{4}$$

The control input u is then chosen as:

$$e_h := \lambda - \varepsilon, \quad e_c := \frac{2\mathcal{O}_f}{v\cos\phi} - \frac{\tan\phi}{l}, \tag{5}$$

$$u := lv\cos^3\phi(ae_c + be_h). \tag{6}$$

Here we can think of e_h as the heading error, and e_c as the curvature error, as the car describes a path toward the target.

The gains a and b should both be positive, and can be chosen with the following guidelines:

- The dynamics of the linear system $e_h'' + ae_h' + be_h = 0$ should represent suitable regulation to the desired path,
- The range $r_0 := 2/a$ should be small enough that divergence from the desired path within this region of the target is acceptable.

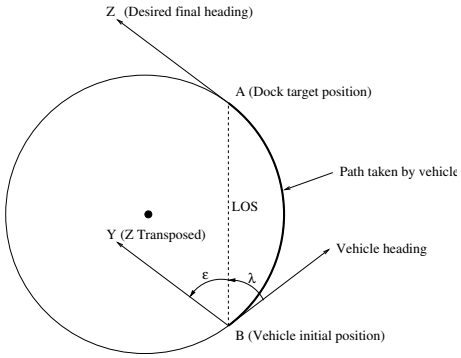

Fig. 2. Geometry for Theorem 1

A discussion of the reasoning behind this control law, and the tuning guidelines, is presented over the next two sections.

4. CONTROL LAW DERIVATION

The method with which we arrived at the above control law is slightly different than most previous approaches. The control objective is to reach some final state, but rather than trying to derive a controller which provides some type of stability to this state, our approach has two stages.

Firstly, simple geometry allows us to pass from the terminal condition to a condition on the instantaneous configuration of the vehicle, this is what we call the *CNG Principle*. Secondly, from this instantaneous condition we derive a feedback-control law using methods similar to feedback linearization.

The following theorem forms the basis of our control law, and was proved in Manchester and Savkin (2002):

Theorem 1. (Circular-Navigation-Guidance Principle) Introduce the circle uniquely defined by the following properties: The initial and final positions of the vehicle lies on the circle; The desired final-heading vector at the target's position is a tangent to the circle.

Suppose that a controller of the form (3) is designed such that the angles λ and ε are kept exactly equal over the full docking manoeuvre, then the vehicle's trajectory will be an arc on this circle. Furthermore, this will result in a perfect docking manoeuvre, as defined in Definition 1 [2] □

This is visualized in Figure 2, where the point A is the dock target position, and B the vehicle's initial position. BA, then, is the line-of-sight, and let AZ (equivalently BY) be the desired final-heading vector.

Note that in the case where $\lambda = \varepsilon = \pi$ and $\phi = 0$, the car is heading away from the target, and will continue to do so forever. In a sense, the car is following a circle of infinite radius: a straight line.

It is only in this case, corresponding to just a one-dimensional line in a four-dimensional manifold, where a perfect docking manoeuvre will not occur. Since this is a "thin" set, and would be simple to overcome in practice, we do not consider it further.

In order to regulate λ to be equal to ε, we consider two errors: $\lambda - \varepsilon$ and $\lambda' - \varepsilon'$. The second of these can be expanded as follows, from Equations (1,2, 4):

$$\lambda' - \varepsilon' = \frac{2\sin\lambda}{r} - \frac{\tan\phi}{l}, \tag{7}$$

$$= \frac{2\mathcal{O}_f}{v\cos\phi} - \frac{\tan\phi}{l}, \tag{8}$$

giving us Equation (5).

This can also be interpreted in the following way: Given any position of the car in the plane, relative to the dock target, there exists a unique circle it should follow. To follow this circle, it must have a certain instantaneous heading and curvature. There are then two errors worth considering: heading error and curvature error. e_h is obviously the heading error, and e_c is the curvature error.

This follows, since the curvature of the circle defined in Theorem 1 is given by the function $2\sin\lambda/r$, and the instantaneous curvature of the vehicle is given by the function $\tan\phi/l$.

If both of these errors are zero, then the vehicle will follow a circular path to the dock target. We can think of these error functions as describing a two-dimensional *target sub-manifold* of the four-dimensional state-space:

$$M := \{(r, \lambda, \varepsilon, \phi) : e_h = 0 \text{ and } e_c = 0)\}.$$

Viewed like this, our objective is similar to that of sliding-mode control: to regulate the system to a particular sub-manifold on which it is known to behave well.

So we have transformed the terminal-state control problem into an instantaneous-state control problem, i.e. the regulation of e_h and e_c. This is reminiscent of the way a honeybee can land on a surface by regulating certain visual cues. We now tackle this regulation problem in a way similar to input-output linearization (see, e.g., Khalil (1993), Chapter 13), and analyze the resulting control law using Lyapunov theory.

[2] In Manchester and Savkin (2002) the definition of perfect intercept was slightly different. However, in the case we consider here it is equivalent to Definition 1.

Let us choose the heading error, $e_h = \lambda - \varepsilon$, as an output function, and attempt to regulate it using input-output linearization.

Differentiating e_h with respect to path-length, we obtain:

$$e_h' = \frac{2\sin\lambda}{r} - \frac{\tan\phi}{l} = e_c,$$

We differentiate this again, obtaining

$$e_h'' = \frac{2\cos\lambda}{r} e_c - \frac{\sec^3\phi}{lv} u. \tag{9}$$

In this equation we note that the control appears explicitly, so a natural approach would be introduce a fictional control input $\bar{u}$ and set

$$u = lv\cos^3\phi\left(-\bar{u} + \frac{2\cos\lambda}{r} e_c\right) \tag{10}$$

rendering the dynamics from $\bar{u}$ to e_h linear, in fact just a double integrator. However, since the range r is unknown to the controller, we cannot do this.

We then "almost feedback linearize" the system, and treat the first term in (9) like an uncertainty. The second term is canceled with the nonlinear control law:

$$u = lv\cos^3\phi(ae_c + be_h)$$

as given in Section 3, then we have

$$e_h'' + \left(a - \frac{2\cos\lambda}{r}\right) e_h' + be_h = 0. \tag{11}$$

If r is large, this is "almost" like the linear system

$$e_h'' + ae_h' + be_h = 0,$$

and it is clear that, by choosing a and b, both the errors e_h and $e_c = e_h'$can be made to converge in any desired fashion.

5. CONTROL LAW ANALYSIS

Since our control law only "almost" linearized the system, we need some further analysis to understand how the system will behave.

The following simple theorem says this: if we start with zero errors, we will continue to have zero errors and achieve a perfect docking manoeuvre. Another way to put this is that if, at any time, the state $(r, \lambda, \varepsilon, \phi) \in M$ then it will stay in M.

Theorem 2. Suppose the vehicle system (1, 2) has the desired heading and curvature, i.e. $e_h(0) = 0$ and $e_c(0) = 0$, then the vehicle will perform a perfect docking manoeuvre, as per Definition 1. □

Proof of Theorem 2: It is clear from the equation of the system (11) that, $e_h(t) = 0$ and $e_c(t) = 0$ at some time t, then they have been, and will be, zero for all time. This implies, then, that $\lambda = \varepsilon$ for all time, and the claim follows from Theorem 1.
■

Now suppose the state starts outside M, that is, with incorrect heading and curvature. Now we'd like to know something about convergence to the target sub-manifold. The dynamics of (11) are those of a linear system with time-varying coefficients, and can be analysed with Lyapunov theory.

Theorem 3. Consider the function

$$V(e_h, e_c) := be_h^2 + e_c^2. \tag{12}$$

This is a positive-definite quadratic form in the heading and curvature errors, and may be considered as the distance to the target sub-manifold.

Let $[s_1, s_2], s_2 > s_1$ be any path interval over which $V(e_h, e_c, s) \neq 0$ and the following inequality holds:

$$a - 2\cos\lambda/r > 0. \tag{13}$$

Then $V(e_h(s_2), e_c(s_2)) < V(e_h(s_1), e_c(s_1))$. That is, over any interval of non-zero length, the norm of the errors strictly decreases. □

Proof of Theorem 3:

For the proof of this theorem, consider the following linear parameter-varying realization of the system (11), with a state $x = [e_h e_c]^T$:

$$x' = Ax + Bw \tag{14}$$

$$z = Cx. \tag{15}$$

where

$$A = \begin{bmatrix} 0 & 1 \\ -b & 2\cos\lambda/r - a \end{bmatrix},\ B = \begin{bmatrix} 0 \\ 1 \end{bmatrix},\ C = \begin{bmatrix} 0 & 1 \end{bmatrix}.$$

Furthermore, consider the Lyapunov function $V(x) = x^T P x$ where

$$P = \begin{bmatrix} b & 0 \\ 0 & 1 \end{bmatrix}. \tag{16}$$

The derivative of this Lyapunov function with respect to distance travelled reduces to

$$V(x)' = -2e_c^2(a - 2\cos\lambda/r). \tag{17}$$

Now, for any s in the interval $[s_1, s_2]$, it follows from $V(e_h, e_c, s) \neq 0$ that $x(s) \neq 0$. Since x is observable from e_c, it follows that

$$\int_{s_1}^{s_2} e_c(s)^2 ds > 0,$$

and since inequality (13) holds, clearly

$$\delta := 2\int_{s_1}^{s_2} e_c(s)^2 \left(a - \frac{2\cos\lambda(s)}{r(s)}\right) ds > 0.$$

Now,

$$\begin{aligned} V(e_h(s_2), e_c(s_2)) &= V(e_h(s_1), e_c(s_1)) \\ &\quad + \int_{s_1}^{s_2} V(e_h, e_c)' ds, \\ &= V(e_h(s_1), e_c(s_1)) - \delta, \\ &< V(e_h(s_1), e_c(s_1)), \end{aligned}$$

and the theorem is proved.
■

This theorem reflects the following physically meaningful problem: When the vehicle is very close to the desired target location, large gains are required to make it swing around and track the correct path.

It should be noted that if range is measureable, either through some other sensor device, or through vision-processing techniques such as stereopsis, optical flow or image looming, this problem will still be present. Indeed, suppose the control law (10) were used, then as range decreased the gains would be come extremely large, due to the $1/r$ term. The actuator constraints on any real system would thus prevent the exact feedback linearization which is attempted.

6. ROBUSTNESS

It has been mentioned in the literature that a particularly important test of a docking algorithm is the robustness of its terminal positioning precision to imperfect modeling of the kinematics and camera calibration (Cárdenas et al. (2003), Laumond (1998)).

The parameters chosen for the simulation were: $l = 1$m, $v = 1$m/s, $a = 4$, $b = 4.04$. The initial conditions were $r(0) = 7$m, $\lambda(0) = \pi/4$ rad, $\varepsilon = \pi/4$ rad, $\phi = \pi/8$ rad.

These parameters imply that the area in which the path could begin to diverge is approximate $2/a = 0.5$m. Note that in all simulated cases, the terminal positioning error was much smaller than this.

In all the following simulations, the control law is derived as above, as though all parameters were nominal. We then simulate a system where parameters are perturbed by some amount.

6.1 Camera calibration

Here we simulate the effect of incorrect camera calibration. We skew the measurement of λ and the optical flow in a way consistent with an incorrect assumption on the focal length of the camera. We introduce the ratio k_f as the true focal length divided by the assumed focal length.

This parameter was varied from 0.6 to 1.8. In Figure 3 we see graphical plots of trajectories, and numerical data for the final range and final-angle error. It is clear that, although the trajectories throughout the middle stage of the docking manoeuvre vary widely, in all cases the robot docked with less than 1cm positioning error, and less than 10° angle error.

6.2 Control input gain

We now move on to consider errors in kinematic model, specifically, in the steering wheel system. Firstly, we investigate what happens if the relationship between control input and steering-wheel movement is not what we expect. Instead of the assumed relation $\dot{\phi} = u$, we instead simulate the system $\dot{\phi} = k_u u$, where k_u is unknown gain term.

Figure 4 depicts the trajectories and error data as k_u is varied from 0.2 to 5. It is noted that increasing k_u significantly, which means that our control input is stronger than we expect, has little effect on the performance of the vehicle. A larger effect is observed when the control input is weaker than expected, but performance is still very good.

Reducing it significantly (to around 0.2) results in some large oscillations in the trajectory, and errors in both final position and final angle. However, reduction of k_u even to 0.5, meaning our control input is half as strong as we think, does not significantly degrade performance.

6.3 Measurement of steering-wheel angle

The control law (4-6) depends explicitly on our knowledge of the current steering wheel angle, ϕ. In the next set of simulations we consider what happens when this information is wrong. Suppose ϕ is read by a potentiometer which is not tuned correctly, so the resulting measurement is a fixed gain of what it should be. Hence, in the calculation of our control law we replace ϕ with $k_\phi\phi$.

Figure 5 shows the resulting trajectories as k_ϕ is varied from 0.6 to 1.4. Once again, it is seen that the middle stages of the trajectory are strongly affected, but the terminal errors remain quite small. We note that terminal position was very small for all cases, but as k_ϕ got very large, the angle error did increase notably.

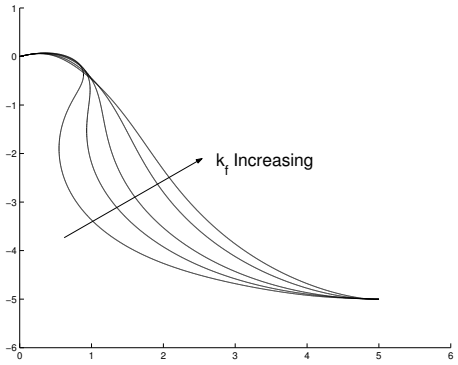

k_f	0.6	0.8	1	1.4	1.8
r(T)($\times 10^{-2}$)	0.44	0.04	0.06	0.23	0.36
$\varepsilon(T)$ (deg)	7.93	2.42	0.12	7.25	9.05

Fig. 3. The effect of incorrect camera calibration

6.4 Steering-wheel angle saturation

In the last set of simulations, we suppose that the steering-wheel angle is restricted to be within some range of angles. This will obviously be true for many practical robotic vehicles, and essentially results in a lower bound on the turning-circle radius.

We represent this with the constraint $|\phi| \leq \phi_s$.

In Figure 6 we depict four trajectories, and four sets of terminal error data. These are for the cases where, firstly, the steering wheel angle is not restricted, and then when it is restricted by $\phi_s = \pi/4, \pi/6$, and $\pi/8$, respectively. In the first three cases the saturation has little, if any, effect on the performance. In the final case, the performance was significantly degraded, simply because the car could not turn around fast enough to get on the right path.

These simulations show that the control law derived above, which does not explicitly account for steering-wheel saturation, does handle sufficiently small levels of saturation without any degradation in performance.

These four sets of simulations show promising prospects for application of our control law when kinematic and camera models are subject to large errors. As has been said in the literature, an important issue in docking problems is robustness of terminal positioning, and in each simulated case above remarkable robustness was observed.

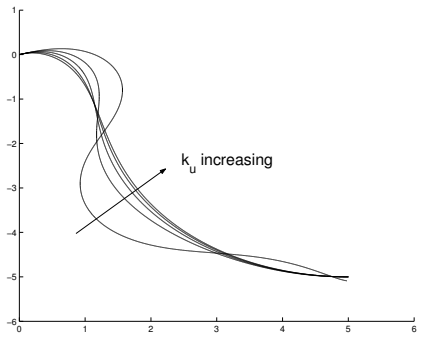

k_u	0.2	0.5	1	2	5
r(T) ($\times 10^{-2}$)	9.22	0.04	0.06	0.07	0.02
Angle Error (deg)	12.89	1.09	0.12	0.10	0.14

Fig. 4. The effect of incorrect $u \rightarrow \dot{\phi}$ gain

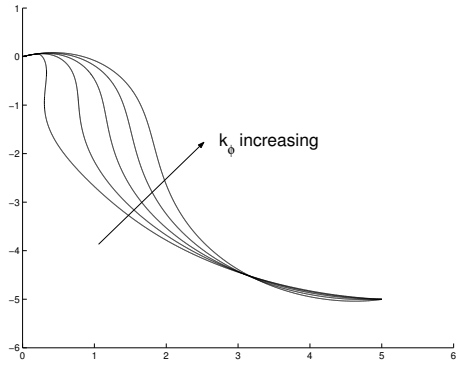

k_ϕ	0.6	0.8	1	1.2	1.4
r(T) ($\times 10^{-2}$)	0.03	0.08	0.06	0.06	0. 33
Angle Error (deg)	4.33	2.49	0.12	4.37	13.12

Fig. 5. The effect of erroneous measurement of ϕ

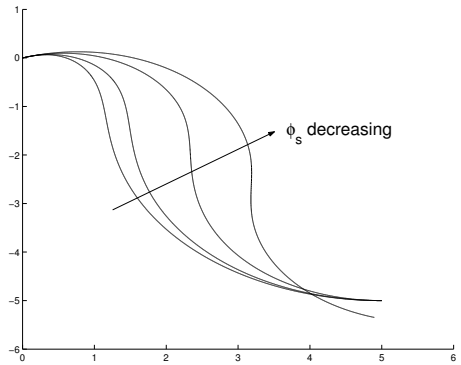

ϕ_s	∞	$\pi/4$	$\pi/6$	$\pi/8$
r(T) ($\times 10^{-2}$)	0.06	0.02	0.07	36.20
Angle Error (deg)	0.12	0.12	0.15	16.90

Fig. 6. The effect of saturation of the steering wheels

REFERENCES

R. C. Arkin. *Behavior-Based Robotics*. The MIT Press, Cambridge, MA, 1998.

Y. Bar-Cohen and C. Breazeal. *Biologically Inspired Intelligent Robots*. SPIE Press, Bellingham, WA, 2003.

G. L. Barrows, J. S. Chahl, and M. V. Srinivasan. Biologically inspired visual sensing and flight control. *Aeronautical Journal*, 107(1069):159–168, Mar 2003.

A. Cárdenas, B. Goodwine, S. Skaar, and M. Seelinger. Vision-based control of a mobile base and on-board arm. *International Journal*

of Robotics Research, 22(9):677–698, Sep 2003.

F. Conticelli, B. Allotta, and P. K. Khosla. Image-based visual servoing of nonholonomic mobile robots. In *38th IEEE Conference on Decision and Control*, pages 3496–3501, Dec 1999.

C. C. de Wit and O. J. Sørdalen. Exponential stabilization of mobile robots with nonholonomic constraints. *IEEE Transactions on Automatic Control*, 37(11):1791–1797, Nov 1992.

M. O. Franz and H. A. Mallot. Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153, 2000.

K. Hashimoto and T. Noritsugu. Visual servoing of nonholonomic cart. In *IEEE International Conference on Robotics and Automation*, pages 1719–1724, Apr 1997.

S. Hutchinson, G. D. Hager, and P. I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, Oct 1996.

A. Kelly and B.Nagy. Reactive nonholonomic trajectory generation via parametric optimal control. *International Journal of Robotics Research*, 22(7):583–601, Jul 2003.

H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River NJ, third edition, 1993.

J. P. Laumond, editor. *Robot Motion Planning and Control*. Springer-Verlag, New York, NY, 1998.

S. Lee, M. Kim, Y. Youm, and W. Chung. Control of a car-like mobile robot for parking problem. In *IEEE International Conference on Robotics and Automation*, pages 1–6, Sep 1999.

I. R. Manchester and A. V. Savkin. Circular navigation guidance law for precision missile target engagements. In *41st IEEE Conference on Decision and Control*, 2002.

I. R. Manchester and A. V. Savkin. Circular navigation missile guidance with incomplete information and uncertain autopilot model. *AIAA Journal of Guidance, Control and Dynamics*, Accepted for Publication. Tentatively scheduled for Nov-Dec Issue. 2004.

I. R. Manchester, A. V. Savkin, and F. A. Faruqi. Optical-flow based precision missile guidance inspired by honeybee navigation. In *42nd IEEE Conference on Decision and Control*, Dec 2003.

I. R. Petersen and A. V. Savkin. *Robust Kalman Filtering for Signals and Systems with Large Uncertainties*. Birkhauser, Boston, MA, 1999.

I. R. Petersen, V. A. Ugrinovskii, and A. V. Savkin. *Robust Control Design using H^∞ Methods*. Springer-Verlag, London, 2000.

J. Santos-Victor and G. Sandini. Visual behaviors for docking. *Computer Vision and Image Understanding*, 67(3):223–238, Sep 1997.

A. V. Savkin, P. Pathirana, and F. A. Faruqi. The problem of precision missile guidance: LQR and $\mathcal{H}_\infty$ frameworks. *IEEE Transactions on Aerospace and Electronic Systems*, 39(3):901–910, Jul 2003.

P. Souères and J. P. Laumond. Shortest paths synthesis for a car-like robot. *IEEE Transactions on Automatic Control*, 41(5):672–688, May 1996.

M. V. Srinivasan. An image-interpolation technique for the computation of optic flow and egomotion. *Biological Cybernetics*, 71:401–415, 1994.

M. V. Srinivasan, S.W. Zhang, J.S. Chahl, E. Barth, and S. Venkatesh. How honeybees make grazing landings on flat surfaces. *Biological Cybernetics*, 83:171–183, 2000.

H. Zhang and J. P. Ostrowski. Visual motion planning for mobile robots. *IEEE Transactions on Robotics and Automation*, 18(2):199–208, Apr 2002.