Learning Visual Servoing for Nonholonomic Mobile Robots with Uncalibrated Cameras

Wang, Jen-Wei; Nikovski, Daniel N.

TR2025-042 April 01, 2025

Abstract

The paper proposes a method for visual servo control of nonholonomic robots with unknown dynamics using images captured by uncalibrated cameras. The method learns the transition dynamics of the robot directly in visual feature space and linearizes it successively in order to compute controls. Experiments both in simulation and on a real testbed using a unicycle-type mobile robot demonstrate that the use of planning and trajectory stabilization algorithms based on differential dynamic programming is much more effective in handling nonholonomic constraints in executing difficult maneuvers, such as parallel parking, than more traditional visual servoing schemes that linearize the learned dynamics around a single point. The ability of the proposed method to control nonholonomic robots without manually calibrating cameras and identifying robot dynamics could potentially significantly lower the cost of deployment of autonomous mobile robots at scale.

The 40th ACM/SIGAPP Symposium On Applied Computing 2025

Learning Visual Servoing for Nonholonomic Mobile Robots with Uncalibrated Cameras

Jen-Wei Wang
Mitsubishi Electric Research Laboratories
Cambridge, Massachusetts, USA
University of California, Berkeley
Berkeley, California, USA
jwwang@berkeley.edu

Daniel Nikovski
Mitsubishi Electric Research Laboratories
Cambridge, Massachusetts, USA
nikovski@merl.com

Abstract

The paper proposes a method for visual servo control of nonholonomic robots with unknown dynamics using images captured by uncalibrated cameras. The method learns the transition dynamics of the robot directly in visual feature space and linearizes it successively in order to compute controls. Experiments both in simulation and on a real testbed using a unicycle-type mobile robot demonstrate that the use of planning and trajectory stabilization algorithms based on differential dynamic programming is much more effective in handling nonholonomic constraints in executing difficult maneuvers, such as parallel parking, than more traditional visual servoing schemes that linearize the learned dynamics around a single point. The ability of the proposed method to control nonholonomic robots without manually calibrating cameras and identifying robot dynamics could potentially significantly lower the cost of deployment of autonomous mobile robots at scale.

CCS Concepts

• **Computer systems organization** → **Robotic control**; *Robotic autonomy*; • **Applied computing** → Engineering.

Keywords

Robotics, machine learning, mobile robots, visual servoing

1 Introduction

Visual servoing (VS) is an important class of robot control methods often used for executing common robotic tasks such as manipulation with articulated robot arms and navigation of mobile robots [1]. The general principle of operation of these methods is to move a robot observed by a camera so that a collection of features or image intensities extracted from the camera image will assume a desired configuration corresponding to a goal state for the robot. In this paper, we are considering one of the most difficult versions of the problem: visual servoing of underactuated planar mobile robots with nonholonomic constraints, without prior knowledge of their dynamics and using only uncalibrated cameras of fixed position. We assume that the pose of the robot is uniquely determined by two or more visual features that are always observable by a fixed camera (eye-to-hand VS). The assumption for planar movement is common for most mobile robot applications where the robot must navigate around a floor or a flat work surface, and the need for two or more visual features that are always observable can be satisfied easily by placing distinct markers on top of the robot and observing them by a single overhead camera.

Eliminating the need for camera calibration when deploying robots is very appealing, but controlling robots in this way requires the combined solution of the problem of perception from an uncalibrated camera with that of control of underactuated nonholonomic robots, each of which is difficult in its own right even when solved in isolation. We propose to use learning sequential control methods for the solution of the VS problem in this setting.

We are considering the visual servo control of a differentially-driven planar mobile robot whose configuration $q \in C$ is described by the vector $q = [x, y, \theta]$ in an inertial frame, where the configuration space C coincides with $\mathbb{R}^2 \times \mathbf{S}^1$. The robot is steered by control inputs $u = [v, \omega]$ representing its commanded linear and angular velocities and is subject to the nonholonomic differential constraint $\dot{x} \sin\theta - \dot{y} \cos\theta = 0$ that represents the robot's inability to move in a direction perpendicular to its current heading.

The robot is observed by a camera at a fixed position in the inertial frame, such that its image plane is generally not parallel to the xy plane in which the robot moves. The intrinsic and extrinsic parameters of the camera are unknown, so a correspondence between a point in the image plane and one in the robot's plane of motion cannot be established. However, we assume that at all control times t_k, a vector of m visual features $s[k] \in \mathbb{R}^m$ related to and uniquely determining the robot's pose is available. The objective of visual servoing is to bring the features $s[k]$ to some desired goal state s^* described directly in feature space, by manipulating the control variable $u[k]$ at discrete times t_k [1]. How the control action $u[k]$ affects the change in features $s[k]$ is expressed by means of an interaction matrix L_s whose inverse can be used to compute a control that reduces the feedback error $e[k] = s[k] - s^*$.

The choice of image features s and the estimation of the interaction matrix L_s are central to the success of VS control schemes. In image-based VS (IBVS), the features are usually the image-plane normalized coordinates of trackable points in the image. When the interaction matrix cannot be computed analytically, it can often be estimated by means of least-squares regression, and neural networks and other machine learning methods can be used to learn position-dependent interaction matrices [1].

Nonholonomic constraints and/or underactuation introduce yet another major level of difficulty. Most VS algorithms in this setting estimate the pose of the robot explicitly and employ path-planning algorithms to find and follow a trajectory in the true state space of the robot. However, this is impossible when the camera is uncalibrated and/or the plane of the robot's motion is unknown. A relatively small proportion of research on VS technology has addressed this very difficult version of the problem. Zhang et al. [7]

describe a method for visual servoing on a mobile robot equipped with an uncalibrated camera, but only in the eye-in-hand setting (camera moving with robot). Other algorithms [2, 6] solve the problem assuming that the image plane is parallel to the plane of the robot's motion, which is not always true or feasible to enforce.

2 Proposed Method for Visual Servoing for Planar Robots with Nonholonomic Constraints

We propose a novel VS method that learns the dynamics of the robot directly in terms of the image coordinates of identifiable feature points on the robot and then uses differential dynamic programming (DDP) algorithms on the learned dynamics to derive sequential feedback controllers that can steer the robot to the desired goal state. The residual dynamics function on the feature space is expressed as $\Delta[k] = f(s[k], u[k])$, where $\Delta[k]$ is a residual vector defined as $s[k+1] - s[k]$ and f is an unknown, generally nonlinear residual dynamics function describing how the robot moves in the feature space in response to a given control command $u[k]$. Because the camera is uncalibrated and f is unknown, we rely on robot exploration data and machine learning algorithms to approximate f. The desired optimality of the computed control law is expressed by means of a cumulative cost J_0 that is the sum of running costs l and a final cost l_f, where the summation is computed over a sequence of H control steps:

$$J_0(s_0, U) = \sum_{k=0}^{H-1} l(s[k], u[k]) + l_f(s[H]) \tag{1}$$

where the states $s[k]$, $k > 0$ follow the dynamics defined above starting from $s[0] = s_0$, and $U = \{u[0], u[1], \ldots, u[H-1]\}$ is the control sequence applied over a finite horizon of length H time steps. The differential dynamics programming (DDP) and iterative linear quadratic regulator (iLQR) algorithms solve this trajectory optimization problem very efficiently when the dynamics f and stage costs l are differentiable ([3–5]). The solution of the optimal control problem defines a set of linear feedback controllers

$$u[k] = K_{ilqr}[k](s[k] - \bar{s}[k]) + \bar{u}[k], \tag{2}$$

where K_{ilqr} is the closed-loop control gain and $\bar{s}$ and $\bar{u}$ are the open-loop reference state and control trajectories.

3 Experimental Verification

Simulated and real-world experiments were performed to evaluate different learning-based control methods for nonholonomic unicycle robots under the setting of uncalibrated cameras. In simulation, we use an analytical unicycle model as the ground-truth dynamics to simulate the rollout of a differential-drive robot under control commands. As shown in Figure 1, the simulated robot is operated on a virtual table, where its boundary (green line in the figure) limits the exploration space of the robot. To simulate perspective transformation, we attach a virtual camera above the table having a top-down view on the table. The camera is tilted in order to evaluate whether controllers can still perform well under perspective transformation and without camera calibration. The simulation environment in Figure 1 shows the view from the top-down camera, and accordingly the green lines do not form a perfect rectangle.

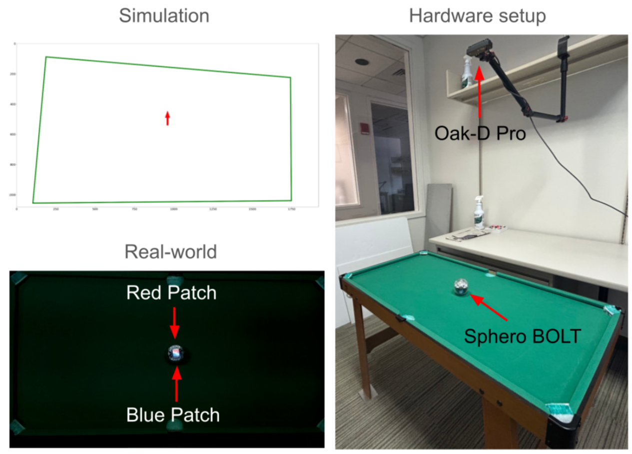

Figure 1: Experimental testbeds in simulation and real world. The simulation environment is the view from a top-down camera. Green lines show the boundary of the pool table and a red arrow represents robot's position and heading. The hardware setup consists of a pool table, a Sphero BOLT robot, and an Oak-D Pro camera. We display a red and blue patches on top of the robot in order to track two representative points.

In the hardware setup shown in Figure 1, we use a Sphero BOLT robot as our target robot and operate it on a pool table. Although the Sphero BOLT looks like an omnidirectional robot, its internal mechanism makes it close to a unicycle robot, because it cannot move sideways without turning its heading first. A Sphero BOLT can be controlled via a commanded rotation angle and translation velocity. In our experiments, we used a control rate of 7.5 Hz. We mounted a Luxonis Oak-D Pro camera above the table as shown in the hardware setup in Figure 1. We designed a tracking algorithm and a corresponding pattern to identify two points on the robot, making use of the 8X8 LED array on top of the robot.

We collect data for dynamics learning both in simulation and from the real robot. We decomposed the dynamics into a translation and rotation part and separately collected translation data and rotation data for learning the two components. The first evaluation metric is the one-step prediction error on a testing set. We compared three different ML algorithms that are often used for learning dynamics functions: Locally Weighted Regression (LWR), Gaussian Process Regression (GPR), and Gaussian Mixture Models. Based on Tables 1 and 2, the GPR model has the best one-step prediction performance on average for both the simulation and real-world datasets.

We compared greedy feedback VS control (using a single interaction matrix obtained by linearizing the learned dynamics around the goal state [1]) against sequential control (using iLQR) for several goal-reaching tasks on a differential drive robot. We set the goal state as the center of the table with the heading pointing at the positive x axis (red arrow in Figure 2) and placed the robot at a

Table 1: One-step prediction error on the simulation dataset with several ML methods.

Models	Dyn. component	RMSE (pixel)	NRMSE	Dyn. component	RMSE (pixel)	NRMSE
LWR		0.0714	0.0319		0.0484	**0.0052**
GPR	Rot	0.0035	**0.0016**	Trans	0.0849	0.0091
GMM		2.7140	1.2077		5.4671	0.5858

Table 2: One-step prediction error on the real-world dataset.

Models	Dyn. component	RMSE (pixel)	NRMSE	Dyn. component	RMSE (pixel)	NRMSE
LWR		2.0535	1.0518		2.3017	0.4968
GPR	Rot	0.8669	**0.4440**	Trans	1.9516	**0.4212**
GMM		1.6733	0.7446		5.3474	1.1123

random position on the table with a random heading (green arrow in Figure 2). Since the robot needs to turn and move at the same time in order to reach the goal, this general task can evaluate the ability of a controller for jointly determining two control inputs: v and ω, under the nonholonomic constraint.

Across different goal-reaching tasks, we prepared tasks with two difficulty levels. The simpler among them is to place the robot at the negative x axis (on the left-hand side of the goal state) pointing at −45 degrees. The more difficult task is to place the robot at the negative y axis (the lower side of the goal state) with the same heading as the goal state. We differentiate the level by the magnitude of the projection of the initial feedback error $e[0] = s[0] - s^*$ onto the initial heading vector of the robot. The simpler task has a non-zero projection and the more difficult one has a nearly zero projection. The more difficult one is commonly known as the parallel parking scenario, where a nonholonmic robot cannot achieve the goal state by simply moving sideways.

We combined the learned dynamics and the controller together and evaluated them in the real-world environment. Since closed-loop controllers are known for regulating dynamics well and are generally robust to uncertainties in dynamics, the focus here is to examine whether a closed-loop controller can still achieve the goal when using imperfect dynamics. Figure 2 demonstrates that the greedy controller will fail in most of the general tasks, whereas the sequential controller will successfully reach the goal.

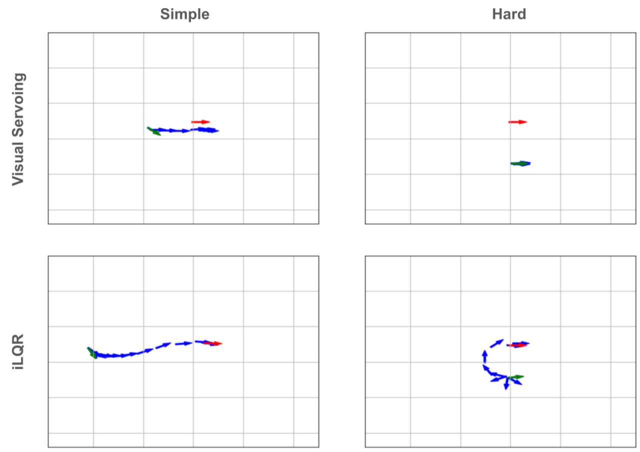

Figure 2: Comparison between VS and iLQR for the general task in the real-world. VS fails to complete either scenario, because it cannot balance between translation and rotation regulation. In contrast, iLQR can complete both scenarios successfully and in a largely time-optimal manner.

4 Conclusion and Future Work

The paper proposes a novel method for visual servo control of nonholonomic robots with unknown dynamics using images captured by uncalibrated cameras. The method learns the transition dynamics of the robot directly in visual feature space and linearizes it as needed to compute controls, either by means of non-sequential visual servoing control as well as sequential planning methods based on DDP. Experiments both in simulation and on a real test bed using a unicycle-type mobile robot demonstrate that the iLQR algorithm is much more effective in handling nonholonomic constraints and executing difficult maneuvers than more traditional visual servoing schemes. Learning of the transition dynamics of the robot can be performed in a self-supervised manner and might be easy to parallelize over a swarm of multiple robots, as long as they can be tracked individually. The ability of the proposed method to control nonholonomic robots without manually calibrating cameras and identifying robot dynamics could potentially significantly lower the cost of deployment of autonomous mobile robots at scale.

References

[1] François Chaumette, Seth Hutchinson, and Peter Corke. 2016. Visual servoing. *Handbook of Robotics, 2nd ed.* (2016), 841–866.

[2] Hua Chen, Shihong Ding, Xi Chen, Lihua Wang, Changping Zhu, and Wen Chen. 2014. Global finite-time stabilization for nonholonomic mobile robots based on visual servoing. *International Journal of Advanced Robotic Systems* 11, 11 (2014), 180.

[3] David H Jacobson and David Q Mayne. 1970. *Differential dynamic programming*. Elsevier.

[4] Weiwei Li and Emanuel Todorov. 2004. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *First International Conference on Informatics in Control, Automation and Robotics*, Vol. 2. SciTePress, 222–229.

[5] Yuval Tassa, Tom Erez, and Emanuel Todorov. 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4906–4913.

[6] Chaoli Wang, Yingchun Mei, Zhenying Liang, and Qingwei Jia. 2010. Dynamic feedback tracking control of non-holonomic mobile robots with unknown camera parameters. *Transactions of the Institute of Measurement and Control* 32, 2 (2010), 155–169.

[7] Xuebo Zhang, Yongchun Fang, Baoquan Li, and Jin Wang. 2016. Visual servoing of nonholonomic mobile robots with uncalibrated camera-to-robot parameters. *IEEE Transactions on Industrial Electronics* 64, 1 (2016), 390–400.