

LEARNING VISUAL SERVOING WITH DEEP FEATURES AND FITTED Q-ITERATION

Alex X. Lee[†], Sergey Levine[†], Pieter Abbeel^{†‡§}

[†] UC Berkeley, Department of Electrical Engineering and Computer Sciences

[‡] OpenAI

[§] International Computer Science Institute

{alexlee_gk, svlevine, pabbeel}@cs.berkeley.edu

ABSTRACT

Visual servoing involves choosing actions that move a robot in response to observations from a camera, in order to reach a goal configuration in the world. Standard visual servoing approaches typically rely on manually designed features and analytical dynamics models, which limits their generalization capability and often requires extensive application-specific feature and model engineering. In this work, we study how learned visual features, learned predictive dynamics models, and reinforcement learning can be combined to learn visual servoing mechanisms. We focus on target following, with the goal of designing algorithms that can learn a visual servo using low amounts of data of the target in question, to enable quick adaptation to new targets. Our approach is based on servoing the camera in the space of learned visual features, rather than image pixels or manually-designed keypoints. We demonstrate that standard deep features, in our case taken from a model trained for object classification, can be used together with a bilinear predictive model to learn an effective visual servo that is robust to visual variation, changes in viewing angle and appearance, and occlusions. A key component of our approach is to use a sample-efficient fitted Q-iteration algorithm to learn which features are best suited for the task at hand. We show that we can learn an effective visual servo on a complex synthetic car following benchmark using just 20 training trajectory samples for reinforcement learning. We demonstrate substantial improvement over a conventional approach based on image pixels or hand-designed keypoints, and we show an improvement in sample-efficiency of more than two orders of magnitude over standard model-free deep reinforcement learning algorithms. Videos are available at http://rll.berkeley.edu/visual_servoing.

1 INTRODUCTION

Visual servoing is a classic problem in robotics that requires moving a camera or robot to match a target configuration of visual features or image intensities. Many robot control tasks that combine perception and action can be posed as visual servoing, including navigation (DeSouza & Kak, 2002; Chen et al., 2006), where a robot must follow a desired path; manipulation, where the robot must servo an end-effector or a camera to a target object to grasp or manipulate it (Malis et al., 1999; Corke, 1993; Hashimoto, 1993; Hosoda & Asada, 1994; Kragic & Christensen, 2002); and various other problems, as surveyed in Hutchinson et al. (1996). Most visual servoing methods assume access to good geometric image features (Chaumette & Hutchinson, 2006; Collewet et al., 2008; Caron et al., 2013) and require knowledge of their dynamics, which are typically obtained from domain knowledge about the system. Using such hand-designed features and models prevents exploitation of statistical regularities in the world, and requires manual engineering for each new system.

In this work, we study how learned visual features, learned predictive dynamics models, and reinforcement learning can be combined to learn visual servoing mechanisms. We focus on target following, with the goal of designing algorithms that can learn a visual servo using low amounts of

data of the target in question, so as to be easy and quick to adapt to new targets. Successful target following requires the visual servo to tolerate moderate variation in the appearance of the target, including changes in viewpoint and lighting, as well as occlusions. Learning invariances to all such distractors typically requires a considerable amount of data. However, since a visual servo is typically specific to a particular task, it is desirable to be able to learn the servoing mechanism very quickly, using a minimum amount of data. Prior work has shown that the features learned by large convolutional neural networks on large image datasets, such as ImageNet classification (Deng et al., 2009), tend to be useful for a wide range of other visual tasks (Donahue et al., 2014). We explore whether the usefulness of such features extends to visual servoing.

To answer this question, we propose a visual servoing method that uses pre-trained features, in our case obtained from the VGG network (Simonyan & Zisserman, 2014) trained for ImageNet classification. Besides the visual features, our method uses an estimate of the feature dynamics in visual space by means of a bilinear model. This allows the visual servo to predict how motion of the robot’s camera will affect the perceived feature values. Unfortunately, servoing directly on the high-dimensional features of a pre-trained network is insufficient by itself to impart robustness on the servo: the visual servo must not only be robust to moderate visual variation, but it must also be able to pick out the target of interest (such as a car that the robot is tasked with following) from irrelevant distractor objects. To that end, we propose a sample-efficient fitted Q-iteration procedure that automatically chooses weights for the most relevant visual features. Crucially, the actual servoing mechanism in our approach is extremely simple, and simply seeks to minimize the Euclidean distance between the weighted feature values at the next time step and the target. The form of the servoing policy in our approach leads to an analytic and tractable linear approximator for the Q-function, which leads to a computationally efficient fitted Q-iteration algorithm. We show that we can learn an effective visual servo on a complex synthetic car following benchmark using just 20 training trajectory samples for reinforcement learning. We demonstrate substantial improvement over a conventional approach based on image pixels or hand-designed keypoints, and we show an improvement in sample-efficiency of more than two orders of magnitude over standard model-free deep reinforcement learning algorithms.

The environment for the synthetic car following benchmark is available online as the package CitySim3D¹, and the code to reproduce our method and experiments is also available online². Supplementary videos of all the test executions are available on the project’s website³.

2 RELATED WORK

Visual servoing is typically (but not always) performed with calibrated cameras and carefully designed visual features. Ideal features for servoing should be stable and discriminative, and much of the work on visual servoing focuses on designing stable and convergent controllers under the assumption that such features are available (Espiau et al., 2002; Mohta et al., 2014; Wilson et al., 1996). Some visual servoing methods do not require camera calibration (Jagersand et al., 1997; Yoshimi & Allen, 1994), and some recent methods operate directly on image intensities (Caron et al., 2013), but generally do not use learning to exploit statistical regularities in the world and improve robustness to distractors.

Learning is a relatively recent addition to the repertoire of visual servoing tools. Several methods have been proposed that apply ideas from reinforcement learning to directly acquire visual servoing controllers (Lampe & Riedmiller, 2013; Sadeghzadeh et al., 2015). However, such methods have not been demonstrated under extensive visual variation, and do not make use of state-of-the-art convolutional neural network visual features. Though more standard deep reinforcement learning methods (Lange et al., 2012; Mnih et al., 2013; Levine et al., 2016; Lillicrap et al., 2015) could in principle be applied to directly learn visual servoing policies, such methods tend to require large numbers of samples to learn task-specific behaviors, making them poorly suited for a flexible visual servoing algorithm that can be quickly repurposed to new tasks (e.g. to following a different object).

¹<https://github.com/alexlee-gk/citysim3d>

²https://github.com/alexlee-gk/visual_dynamics

³http://rll.berkeley.edu/visual_servoing

Instead, we propose an approach that combines learning of predictive models with pre-trained visual features. We use visual features trained for ImageNet (Deng et al., 2009) classification, though any pre-trained features could in principle be applicable for our method, so long as they provide a suitable degree of invariance to visual distractors such as lighting, occlusion, and changes in viewpoint. Using pre-trained features allows us to avoid the need for large amounts of experience, but we must still learn the policy itself. To further accelerate this process, we first acquire a predictive model that allows the visual servo to determine how the visual features will change in response to an action. General video prediction is an active research area, with a number of complex but data-hungry models proposed in recent years (Oh et al., 2015; Watter et al., 2015; Mathieu et al., 2015; Xue et al., 2016; Lotter et al., 2016; Jia et al., 2016; Walker et al., 2016; Vondrick et al., 2016).

However, we observe that convolutional response maps can be interpreted as images and, under mild assumptions, the dynamics of image pixels during camera motion can be well approximated by means of a bilinear model (Censi & Murray, 2015). We therefore train a relatively simple bilinear model for short-term prediction of visual feature dynamics, which we can use inside a very simple visual servo that seeks to minimize the error between the next predicted feature values and a target image.

Unfortunately, simply training predictive models on top of pre-trained features is insufficient to produce an effective visual servo, since it weights the errors of distractor objects the same amount as the object of interest. We address this challenge by using an efficient Q-iteration algorithm to train the weights on the features to maximize the servo’s long-horizon reward. This method draws on ideas from regularized fitted Q-iteration (Gordon, 1995; Ernst et al., 2005; Farahmand et al., 2009) and neural fitted Q-iteration (Riedmiller, 2005) to develop a sample-efficient algorithm that can directly estimate the expected return of the visual servo without the use of any additional function approximator.

3 PROBLEM STATEMENT

Let \mathbf{y}_t be a featurization of the camera’s observations \mathbf{x}_t and let \mathbf{y}_* be some given goal feature map. For the purposes of this work, we define *visual servoing* as the problem of choosing controls \mathbf{u}_t for a fixed number of discrete time steps t as to minimize the error $\|\mathbf{y}_* - \mathbf{y}_t\|$.

We use a relatively simple gradient-based servoing policy that uses one-step feature dynamics, $f : \{\mathbf{y}_t, \mathbf{u}_t\} \rightarrow \mathbf{y}_{t+1}$. The policy chooses the control that minimizes the distance between the goal feature map and the one-step prediction:

$$\pi(\mathbf{x}_t, \mathbf{x}_*) = \arg \min_{\mathbf{u}} \|\mathbf{y}_* - f(\mathbf{y}_t, \mathbf{u})\|^2. \quad (1)$$

Learning this policy amounts to learning the robot dynamics and the distance metric $\|\cdot\|$.

To learn the robot dynamics, we assume that we have access to a dataset of paired observations and controls $\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}$. This data is relatively easy to obtain as it involves collecting a stream of the robot’s observations and controls. We use this dataset to learn a general visual dynamics model that can be used for any task.

To learn the distance metric, we assume that the robot interacts with the world and collects tuples of the form $\mathbf{x}_t, \mathbf{u}_t, c_t, \mathbf{x}_{t+1}, \mathbf{x}_*$. At every time step during learning, the robot observes \mathbf{x}_t and takes action \mathbf{u}_t . After the transition, the robot observes \mathbf{x}_{t+1} and receives an immediate cost c_t . This cost is task-specific and it quantifies how good that transition was in order to achieve the goal. At the beginning of each trajectory, the robot is given a goal observation \mathbf{x}_* , and it is the same throughout the trajectory. We define the goal feature map to be the featurization of the goal observation. We learn the distance metric using reinforcement learning and we model the environment as a Markov Decision Process (MDP). The state of the MDP is the tuple of the current observation and the episode’s target observation, $\mathbf{s}_t = (\mathbf{x}_t, \mathbf{x}_*)$, the action \mathbf{u}_t is the discrete-time continuous control of the robot, and the cost function maps the states and action $(\mathbf{s}_t, \mathbf{u}_t, \mathbf{s}_{t+1})$ to a scalar cost c_t .




Figure 1: Multiscale bilinear model. The function h maps images \mathbf{x} to feature maps $\mathbf{y}^{(0)}$, the operator d downsamples the feature maps $\mathbf{y}^{(l-1)}$ to $\mathbf{y}^{(l)}$, and the bilinear function $f^{(l)}$ predicts the next feature $\hat{\mathbf{y}}^{(l)}$. The number of channels for each feature map is n_c , regardless of the scale l .




Figure 2: Dilated VGG-16 network. The intermediate feature maps drawn in a lighter shade are outputs of max-pooling layers. The feature maps in the conv4 and conv5 blocks are outputs of dilated convolutions with dilation factors of 2 and 4, respectively.

4 VISUAL FEATURES DYNAMICS

We learn a multiscale bilinear model to predict the visual features of the next frame given the current image from the robot’s camera and the action of the robot. An overview of the model is shown in Figure 1. The learned dynamics can then be used for visual servoing as described in Section 5.

4.1 VISUAL FEATURES

We consider both pixels and semantic features for the visual representation. We define the function h to relate the image \mathbf{x} and its feature $\mathbf{y} = h(\mathbf{x})$. Our choice of semantic features are derived from the VGG-16 network (Simonyan & Zisserman, 2014), which is a convolutional neural network trained for large-scale image recognition on the ImageNet dataset (Deng et al., 2009). Since spatial invariance is undesirable for servoing, we remove some of the max-pooling layers and replace the convolutions that followed them with dilated convolutions, as done by Yu & Koltun (2015). The modified VGG network is shown in Figure 2. We use the model weights of the original VGG-16 network, which are publicly available as a Caffe model (Jia et al., 2014). The features that we use are the outputs of some of the intermediate convolutional layers, that have been downsampled to a 32×32 resolution (if necessary) and standardized with respect to our training set.

We use multiple resolutions of these features for servoing. The idea is that the high-resolution representations have detailed local information about the scene, while the low-resolution representations have more global information available through the image-space gradients. The features at level l of the multiscale pyramid are denoted as $\mathbf{y}^{(l)}$. The features at each level are obtained from the features below through a downsampling operator $d(\mathbf{y}^{(l-1)}) = \mathbf{y}^{(l)}$ that cuts the resolution in half.

4.2 BILINEAR DYNAMICS

The features $\mathbf{y}_t^{(l)}$ are used to predict the corresponding level’s features $\mathbf{y}_{t+1}^{(l)}$ at the next time step, conditioned on the action \mathbf{u}_t , according to a prediction function $f^{(l)}(\mathbf{y}_t^{(l)}, \mathbf{u}_t) = \hat{\mathbf{y}}_{t+1}^{(l)}$. We use a bilinear model to represent these dynamics, motivated by prior work (Censi & Murray, 2015). In order to servo at different scales, we learn a bilinear dynamics model at *each* scale. We consider two variants of the bilinear model in previous work in order to reduce the number of model parameters.

The first variant uses *fully connected* dynamics as in previous work but models the dynamics of each channel independently. When semantic features are used, this model interprets the feature maps as

being abstract images with spatial information within a channel and different entities or factors of variation across different channels. This could potentially allow the model to handle moving objects, occlusions, and other complex phenomena.

The fully connected bilinear model is quite large, so we propose a bilinear dynamics that enforces sparsity in the parameters. In particular, we constrain the prediction to depend only on the features that are in its local spatial neighborhood, leading to the following *locally connected* bilinear model:

$$\hat{\mathbf{y}}_{t+1,c}^{(l)} = \mathbf{y}_{t,c}^{(l)} + \sum_j \left(\mathbf{W}_{c,j}^{(l)} * \mathbf{y}_{t,c}^{(l)} + \mathbf{B}_{c,j}^{(l)} \right) \mathbf{u}_{t,j} + \left(\mathbf{W}_{c,0}^{(l)} * \mathbf{y}_{t,c}^{(l)} + \mathbf{B}_{c,0}^{(l)} \right). \quad (2)$$

The parameters are the 4-dimensional tensor $\mathbf{W}_{c,j}^{(l)}$ and the matrix $\mathbf{B}_{c,j}^{(l)}$ for each channel c , scale l , and control coordinate j . The last two terms are biases that allow to model action-independent visual changes, such as moving objects. The $*$ is the locally connected operator, which is like a convolution but with untied filter weights⁴.

4.3 TRAINING VISUAL FEATURE DYNAMICS MODELS

The loss that we use for training the bilinear dynamics is the sum of the losses of the predicted features at each level, $\sum_{l=0}^L \ell^{(l)}$, where the loss for each level l is the squared ℓ -2 norm between the predicted features and the actual features of that level, $\ell^{(l)} = \|\mathbf{y}_{t+1}^{(l)} - \hat{\mathbf{y}}_{t+1}^{(l)}\|^2$.

We optimize for the dynamics while keeping the feature representation fixed. This is a supervised learning problem, which we solve with ADAM (Kingma & Ba, 2014). The training set, consisting of triplets $\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}$, was obtained by executing a hand-coded policy that moves the robot around the target with some Gaussian noise.

5 LEARNING VISUAL SERVOING WITH REINFORCEMENT LEARNING

We propose to use a multiscale representation of semantic features for servoing. The challenge when introducing multiple scales and multi-channel feature maps for servoing is that the features do not necessarily agree on the optimal action when the goal is unattainable or the robot is far away from the goal. To do well, it's important to use a good weighing of each of the terms in the objective. Since there are many weights, it would be impractically time-consuming to set them by hand, so we resort to learning. We want the weighted one-step lookahead objective to encourage good long-term behavior, so we want this objective to correspond to the state-action value function Q . So we propose a method for learning the weights based on fitted Q-iteration.

5.1 SERVOING WITH WEIGHTED MULTISCALE FEATURES

Instead of attempting to build an accurate predictive model for multi-step planning, we use the simple greedy servoing method in Equation (1), where we minimize the error between the target and predicted features for all the scales. Typically, only a few objects in the scene are relevant, so the errors of some channels should be penalized more than others. Similarly, features at different scales might need to be weighted differently. Thus, we use a weighting $\mathbf{w}_c^{(l)} \geq 0$ per channel c and scale l :

$$\pi(\mathbf{x}_t, \mathbf{x}_*) = \arg \min_{\mathbf{u}} \sum_c \sum_{l=0}^L \frac{\mathbf{w}_c^{(l)}}{|\mathbf{y}_{\cdot,c}^{(l)}|} \left\| \mathbf{y}_{*,c}^{(l)} - f_c^{(l)}(\mathbf{y}_{t,c}^{(l)}, \mathbf{u}) \right\|_2^2 + \sum_j \lambda_j \mathbf{u}_j^2, \quad (3)$$

where $|\cdot|$ denotes the cardinality operator and the constant $1/|\mathbf{y}_{\cdot,c}^{(l)}|$ normalizes the feature errors by its spatial resolution. We also use a separate weight λ_j for each control coordinate j . This optimization can be solved efficiently since the dynamics is linear in the controls (see Appendix A).

⁴ The locally connected operator, with a local neighborhood of $n_f \times n_f$ (analogous to the filter size in convolutions), is defined as:

$$(\mathbf{W} * \mathbf{y})_{k_h, k_w} = \sum_{i_h=k_h-\lfloor n_f/2 \rfloor}^{k_h+\lfloor n_f/2 \rfloor} \sum_{i_w=k_w-\lfloor n_f/2 \rfloor}^{k_w+\lfloor n_f/2 \rfloor} \mathbf{W}_{k_h, k_w, i_h-k_h, i_w-k_w} \mathbf{y}_{i_h, i_w}.$$

5.2 Q-FUNCTION APPROXIMATION FOR THE WEIGHTED SERVOING POLICY

We choose a Q-value function approximator that can represent the servoing objective such that the greedy policy with respect to the Q-values results in the policy of Equation (3). In particular, we use a function approximator that is linear in the weight parameters $\theta^\top = [\mathbf{w}^\top \ \lambda^\top]^\top$:

$$Q_{\theta,b}(\mathbf{s}_t, \mathbf{u}) = \phi(\mathbf{s}_t, \mathbf{u})^\top \theta + b, \quad \phi(\mathbf{s}_t, \mathbf{u})^\top = \left[\left[\frac{1}{|\mathbf{y}_{t,c}^{(l)}|} \left\| \mathbf{y}_{*,c}^{(l)} - f_c^{(l)}(\mathbf{y}_{t,c}^{(l)}, \mathbf{u}) \right\|_2^2 \right]_{c,l}^\top \quad [\mathbf{u}_j^2]_j^\top \right].$$

We denote the state of the MDP as $\mathbf{s}_t = (\mathbf{x}_t, \mathbf{x}_*)$ and add a bias b to the Q-function. The servoing policy is then simply $\pi_\theta(\mathbf{s}_t) = \arg \min_{\mathbf{u}} Q_{\theta,b}(\mathbf{s}_t, \mathbf{u})$. For reinforcement learning, we optimized for the weights θ but kept the feature representation and its dynamics fixed.

5.3 LEARNING THE Q-FUNCTION WITH FITTED Q-ITERATION

Reinforcement learning methods that learn a Q-function do so by minimizing the Bellman error:

$$\left\| Q(\mathbf{s}_t, \mathbf{u}_t) - \left(c_t + \gamma \min_{\mathbf{u}} Q(\mathbf{s}_{t+1}, \mathbf{u}) \right) \right\|_2^2. \quad (4)$$

In fitted Q-iteration, the agent iteratively gathers a dataset $\{\mathbf{s}_t^{(i)}, \mathbf{u}_t^{(i)}, c_t^{(i)}, \mathbf{s}_{t+1}^{(i)}\}_i^N$ of N samples according to an exploration policy, and then minimizes the Bellman error using this dataset. We use the term *sampling iteration* to refer to each iteration j of this procedure. At the beginning of each sampling iteration, the current policy with added Gaussian noise is used as the exploration policy.

It is typically hard or unstable to optimize for both Q-functions that appear in the Bellman error of Equation (4), so it is usually optimized by iteratively optimizing the current Q-function while keeping the target Q-function constant. However, we notice that for a given state, the action that minimizes its Q-values is the same for any non-negative scaling α of θ and for any bias b . Thus, to speed up the optimization of the Q-function, we first set $\alpha^{(k-\frac{1}{2})}$ and $b^{(k-\frac{1}{2})}$ by jointly solving for α and b of *both* the current and target Q-function:

$$\min_{\alpha \geq 0, b} \frac{1}{N} \sum_{i=1}^N \left\| Q_{\alpha\theta^{(k-1)}, b}(\mathbf{s}_t^{(i)}, \mathbf{u}_t^{(i)}) - \left(c_t^{(i)} + \gamma \min_{\mathbf{u}} Q_{\alpha\theta^{(k-1)}, b}(\mathbf{s}_{t+1}^{(i)}, \mathbf{u}) \right) \right\|_2^2 + \nu \|\theta\|_2^2. \quad (5)$$

This is similar to how, in policy evaluation, state values can be computed by solving a linear system. We regularize the parameters with an ℓ_2 penalty, weighted by $\nu \geq 0$. We use the term *FQI iteration* to refer to each iteration k of optimizing the Bellman error, and we use the notation $(k-\frac{1}{2})$ to denote an intermediate step between iterations $(k-1)$ and (k) . The parameters θ can then be updated with $\theta^{(k-\frac{1}{2})} = \alpha^{(k-\frac{1}{2})} \theta^{(k-1)}$. Then, we update $\theta^{(k)}$ and $b^{(k)}$ by optimizing for θ and b of the current Q-function while keeping the parameters of the target Q-function fixed:

$$\min_{\theta \geq 0, b} \frac{1}{N} \sum_{i=1}^N \left\| Q_{\theta, b}(\mathbf{s}_t^{(i)}, \mathbf{u}_t^{(i)}) - \left(c_t^{(i)} + \gamma \min_{\mathbf{u}} Q_{\theta^{(k-\frac{1}{2})}, b^{(k-\frac{1}{2})}}(\mathbf{s}_{t+1}^{(i)}, \mathbf{u}) \right) \right\|_2^2 + \nu \|\theta\|_2^2. \quad (6)$$

A summary of the algorithm used to learn the feature weights is shown in Algorithm 1.

Algorithm 1 FQI with initialization of policy-independent parameters

- 1: **procedure** FQI($\theta^{(0)}, \sigma_{\text{exploration}}^2, \nu$)
 - 2: **for** $s = 1, \dots, S$ **do** ▷ sampling iterations
 - 3: Gather dataset $\{\mathbf{s}_t^{(i)}, \mathbf{u}_t^{(i)}, c_t^{(i)}, \mathbf{s}_{t+1}^{(i)}\}_i^N$ using exploration policy $\mathcal{N}(\pi_{\theta^{(0)}}, \sigma_{\text{exploration}}^2)$
 - 4: **for** $k = 1, \dots, K$ **do** ▷ FQI iterations
 - 5: Fit $\alpha^{(k-\frac{1}{2})}$ and $b^{(k-\frac{1}{2})}$ using (5)
 - 6: $\theta^{(k-\frac{1}{2})} \leftarrow \alpha^{(k-\frac{1}{2})} \theta^{(k-1)}$
 - 7: Fit $\theta^{(k)}$ and $b^{(k)}$ using (6)
 - 8: $\theta^{(0)} \leftarrow \theta^{(K)}$
-



Figure 3: Cars used to learn the dynamics and the feature weights. They were also used in some of the test experiments.



Figure 4: Novel cars used only in the test experiments. They were never seen during training or validation.



Figure 5: Costs of test executions using various feature dynamics models, where the feature weights are optimized with FQI. We test on cars that were used during learning (left plot) and on novel cars that were only used at test time (right plot). The reported values are the mean and standard error across 100 trajectories, of up to 100 time steps each. The policies based on pixel intensities use either fully connected or locally connected dynamics, whereas all the policies based on VGG features use locally connected dynamics. The policies based on deeper VGG features generally achieve better performance, except for the deepest feature representation, VGG conv5_3, which is not as suitable for approximating Q-values. The policies based on pixel intensities and VGG conv5_3 features perform worse on the novel cars. However, VGG features conv1_2 through conv4_3 achieve some degree of generalization on the novel cars.

6 EXPERIMENTS

We evaluate the performance of the model for visual servoing in a simulated environment. The simulated quadcopter is governed by rigid body dynamics. The robot has 4 degrees of freedom, corresponding to translation along three axis and yaw angle. This simulation is inspired by tasks in which an autonomous quadcopter flies above a city, with the goal of following some target object (e.g., a car).

6.1 LEARNING FEATURE DYNAMICS AND WEIGHTS WITH FQI

The dynamics for each of the features were trained using a dataset of 10000 samples (corresponding to 100 trajectories) with ADAM (Kingma & Ba, 2014). A single dynamics model was learned for each feature representation for all the training cars (Figure 3). This training set was generated by executing a hand-coded policy that navigates the quadcopter around a car for 100 time steps per trajectory, while the car moves around the city.

We used the proposed FQI algorithm to learn the weightings of the features and control regularizer. At every sampling iteration, the current policy was executed with Gaussian noise to gather data from 10 trajectories. All the trajectories in our experiments were up to 100 time steps long. The immediate cost received by the agent encodes the error of the target in image coordinates (details in Appendix B). Then, the parameters were iteratively updated by running $K = 10$ iterations of FQI. We ran the overall algorithm for only $S = 2$ sampling iterations and chose the parameters that achieved the best performance on 10 validation trajectories. These validation trajectories were obtained by randomly choosing 10 cars from the set of training cars and randomly sampling initial states, and executing the policy with the parameters of the current iteration. All the experiments share the same set of validation trajectories.







Feature Dynamics	Observations from Test Executions	Cost
pixel, locally connected		0.95
		6.26
		14.49
VGG conv4_3		0.38
		0.48
		1.02

Table 1: Sample observations from test executions in our experiments with the novel cars, and the costs for each trajectory, for different feature dynamics. We use the weights learned by our FQI algorithm. In each row, we show the observations of every 10 steps and the last one. The first observation of each trajectory is used as the target observation. The trajectories shown here were chosen to reflect different types of behaviors. The serving policy based on pixel feature dynamics can generally follow cars that can be discriminated based on RGB pixel intensities (e.g., a yellow car with a relatively uniform background). However, it performs poorly when distractor objects appear throughout the execution (e.g., a lamp) or when they appear in the target image (e.g., the crosswalk markings on the road). On the other hand, VGG conv4_3 features are able to discriminate the car from distractor objects and the background, and the feature weights learned by the FQI algorithm are able to leverage this. Additional sample executions with other feature dynamics can be found in Table 3 in the Appendix.

6.2 COMPARISON OF FEATURE REPRESENTATIONS FOR SERVOING

We compare the servoing performance for various feature dynamics models, where the weights are optimized with FQI. We execute the learned policies on 100 test trajectories and report the average cost of the trajectory rollouts on Figure 5. The cost of a single trajectory is the (undiscounted) sum of costs c_t . We test the policies with cars that were seen during training as well as with a set of novel cars (Figure 4), to evaluate the generalization of the learned dynamics and optimized policies.

The test trajectories were obtained by randomly sampling 100 cars (with replacement) from one of the two sets of cars, and randomly sampling initial states (which are different from the ones used for validation). For consistency and reproducibility, the same sampled cars and initial states were used across all the test experiments, and the same initial states were used for both sets of cars. These test trajectories were never used during the development of the algorithm or for choosing hyperparameters.

From these results, we notice that policies based on deeper VGG features, up to VGG conv4_3, generally achieve better performance. However, the deepest feature representation, VGG conv5_3, is not as suitable for approximating Q-values. We hypothesize that this feature might be too spatially invariant and it might lack the necessary spatial information to differentiate among different car positions. The policies based on pixel intensities and VGG conv5_3 features perform worse on the novel cars. However, VGG features conv1_2 through conv4_3 achieve some degree of generalization on the novel cars.

We show sample trajectories in Table 1. The policy based on pixel-intensities is susceptible to occlusions and distractor objects that appear in the target image or during executions. This is because distinguishing these occlusions and distractors from the cars cannot be done using just RGB features.

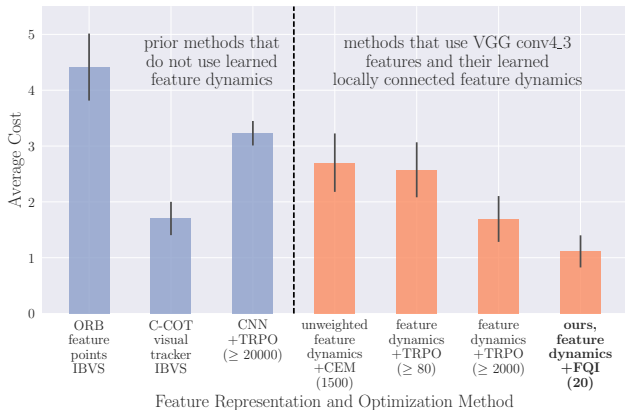


Figure 6: Comparison of costs on test executions of prior methods against our method based on VGG conv4.3 feature dynamics. These costs are from executions with the training cars; the costs are comparable when testing with the novel cars (Table 2). The first two methods use classical image-based visual servoing (IBVS) with feature points from an off-the-shelf keypoint detector and descriptor extractor (ORB features), and with feature points extracted from bounding boxes predicted by a state-of-the-art visual tracker (C-COT tracker), respectively. The third method trains a convolutional neural network (CNN) policy end-to-end with Trust Region Policy Optimization (TRPO). The other methods use the servoing policy based on VGG conv4.3 feature dynamics, either with unweighted features or weights trained with TRPO for either 2 or 50 iterations. In the case of unweighted features, we learned the weights λ and a single weight w with the cross entropy method (CEM). We report the number of training trajectories in parenthesis for the methods that require learning. For TRPO, we use a fixed number of training *samples* per iteration, whereas for CEM and FQI, we use a fixed number of training *trajectories* per iteration. We use a batch size of 4000 samples for TRPO, which means that at least 40 trajectories were used per iteration (since trajectories can terminate early, i.e. in less than 100 time steps).

6.3 COMPARISON OF WEIGHTINGS FROM OTHER OPTIMIZATION METHODS

We compare our policy using conv4.3 feature dynamics, with weights optimized by FQI, against policies that use these dynamics but with either no feature weighting or weights optimized by other algorithms.

For the case of no weighting, we use a single feature weight w but optimize the relative weighting of the controls λ with the cross entropy method (CEM) (De Boer et al., 2005). For the other cases, we learn the weights with Trust Region Policy Optimization (TRPO) (Schulman et al., 2015). Since the servoing policy is the minimizer of a quadratic objective (Equation (3)), we represent the policy as a neural network that has a matrix inverse operation at the output. We train this network for 2 and 50 sampling iterations, and use a batch size of 4000 samples per iteration. All of these methods use the same feature representation as ours, the only difference being how the weights w and λ are chosen.

We report the average costs of these methods on the right of Figure 6. In 2 sampling iterations, the policy learned with TRPO does not improve by much, whereas our policy learned with FQI significantly outperforms the other policies. The policy learned with TRPO improves further in 50 iterations; however, the cost incurred by this policy is still about one and a half times the cost of our policy, despite using more than 100 times as many trajectories.

6.4 COMPARISON TO PRIOR METHODS

We also consider other methods that do not use the dynamics-based servoing policy that we propose. We report their average performance on the left of Figure 6.

For one of the prior methods, we train a convolutional neural network (CNN) policy end-to-end with TRPO. The policy is parametrized as a 5-layer CNN, consisting of 2 convolutional and 3 fully-connected layers, with ReLU activations except for the output layer; the convolutional layers use

16 filters (4×4 , stride 2) each and the first 2 fully-connected layers use 32 hidden units each. The policy takes in raw pixel-intensities and outputs controls.

This policy achieves a modest performance (although still worse than the policies based on conv4_3 feature dynamics) but it requires significantly more training samples than any of the other learning-based methods. We also trained CNN policies that take in extracted VGG features (without any dynamics) as inputs, but they perform worse (see Table 4 in the Appendix). This suggests that given a policy parametrization that is expressive enough and given a large number of training samples, it is better to directly provide the raw pixel-intensity images to the policy instead of extracted VGG features. This is because VGG features are not optimized for this task and their representation loses some information that is useful for servoing.

The other two prior methods use classical image-based visual servoing (IBVS) (Chaumette & Hutchinson, 2006) with respect to Oriented FAST and Rotated BRIEF (ORB) feature points (Rublee et al., 2011), or feature points extracted from a visual tracker. For the former, the target features consist of only the ORB feature points that belong to the car, and this specifies that the car is relevant for the task. For the tracker-based method, we use the Continuous Convolution Operator Tracker (C-COT) (Danelljan et al., 2016) (the current state-of-the-art visual tracker) to get bounding boxes around the car and use the four corners of the box as the feature points for servoing. We provide the ground truth car’s bounding box of the first frame as an input to the C-COT tracker. For all of the IBVS methods, we provide the ground truth depth values of the feature points, which are used in the algorithm’s interaction matrix⁵.

The first method performs poorly, in part because ORB features are not discriminative enough for some of the cars, and the target feature points are sometimes matched to feature points that are not on the car. The tracker-based method achieves a relatively good performance. The gap in performance with respect to our method is in part due to the lack of car dynamics information in the IBVS model, whereas our method implicitly incorporates that in the learned feature dynamics. It is also worth noting that the tracker-based policy runs significantly slower than our method. The open-source implementation of the C-COT tracker⁶ runs at about 1 Hz whereas our policy based on conv4_3 features runs at about 16 Hz. Most of the computation time of our method is spent computing features from the VGG network, so there is room for speedups if we use a network that is less computationally demanding.

7 DISCUSSION

Manual design of visual features and dynamics models can limit the applicability of visual servoing approaches. We described an approach that combines learned visual features with learning predictive dynamics models and reinforcement learning to learn visual servoing mechanisms. Our experiments demonstrate that standard deep features, in our case taken from a model trained for object classification, can be used together with a bilinear predictive model to learn an effective visual servo that is robust to visual variation, changes in viewing angle and appearance, and occlusions. For control we propose to learn Q-values, building on fitted Q-iteration, which at execution time allows for one-step lookahead calculations that optimize long term objectives. Our method can learn an effective visual servo on a complex synthetic car following benchmark using just 20 training trajectory samples for reinforcement learning. We demonstrate substantial improvement over a conventional approach based on image pixels or hand-designed keypoints, and we show an improvement in sample-efficiency of more than two orders of magnitude over standard model-free deep reinforcement learning algorithms.

ACKNOWLEDGEMENTS

This research was funded in part by the Army Research Office through the MAST program and the Berkeley DeepDrive consortium. Alex Lee was also supported by the NSF GRFP.

⁵The term interaction matrix, or feature Jacobian, is used in the visual servo literature to denote the Jacobian of the features with respect to the control.

⁶<https://github.com/martin-danelljan/Continuous-ConvOp>

REFERENCES

- Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 404–417. Springer, 2006, 2006.
- Guillaume Caron, Eric Marchand, and El Mustapha Mouaddib. Photometric visual servoing for omnidirectional cameras. *Autonomous Robots*, 35(2-3):177–193, 2013, 2013.
- Andrea Censi and Richard M Murray. Bootstrapping bilinear models of simple vehicles. *The International Journal of Robotics Research*, 34(8):1087–1113, 2015, 2015.
- Francois Chaumette and Seth Hutchinson. Visual servo control. I. Basic approaches. *IEEE Robotics & Automation Magazine*, 13(4):82–90, 2006, 2006.
- Jian Chen, Warren E Dixon, M Dawson, and Michael McIntyre. Homography-based visual servo tracking control of a wheeled mobile robot. *IEEE Transactions on Robotics*, 22(2):406–415, 2006, 2006.
- Christophe Collewet and Eric Marchand. Photometric visual servoing. *IEEE Transactions on Robotics*, 27(4):828–834, 2011, 2011.
- Christophe Collewet, Eric Marchand, and Francois Chaumette. Visual servoing set free from image processing. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 81–86. IEEE, 2008, 2008.
- Peter I Corke. Visual control of robot manipulators – A review. *Visual servoing*, 7:1–31, 1993, 1993.
- Martin Danelljan, Andreas Robinson, Fahad Shahbaz Khan, and Michael Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 472–488. Springer, 2016, 2016.
- Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinfeld. A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67, 2005, 2005.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255. IEEE, 2009, 2009.
- Guilherme N DeSouza and Avinash C Kak. Vision for mobile robot navigation: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 24(2):237–267, 2002, 2002.
- Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. DeCAF: A deep convolutional activation feature for generic visual recognition. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 32, pp. 647–655, 2014, 2014.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005, 2005.
- Bernard Espiau, Francois Chaumette, and Patrick Rives. A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 8(3):313–326, 2002, 2002.
- Amir Massoud Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor. Regularized fitted Q-iteration for planning in continuous-space Markovian decision problems. In *American Control Conference, 2009. ACC'09.*, pp. 725–730. IEEE, 2009, 2009.
- John T Feddema and Owen Robert Mitchell. Vision-guided servoing with feature-based trajectory generation (for robots). *IEEE Transactions on Robotics and Automation*, 5(5):691–700, 1989, 1989.
- Geoffrey J Gordon. Stable function approximation in dynamic programming. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1995, 1995.
- Koichi Hashimoto. *Visual servoing*, volume 7. World scientific, 1993, 1993.

- Koh Hosoda and Minoru Asada. Versatile visual servoing without knowledge of true Jacobian. In *Intelligent Robots and Systems '94: Advanced Robotic Systems and the Real World', IROS'94. Proceedings of the IEEE/RSJ/GI International Conference on*, volume 1, pp. 186–193. IEEE, 1994, 1994.
- Seth Hutchinson, Gregory D Hager, and Peter I Corke. A tutorial on visual servo control. *IEEE transactions on robotics and automation*, 12(5):651–670, 1996, 1996.
- Martin Jagersand, Olac Fuentes, and Randal Nelson. Experimental evaluation of uncalibrated visual servoing for precision manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pp. 2874–2880. IEEE, 1997, 1997.
- Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic filter networks. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 667–675, 2016, 2016.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM International Conference on Multimedia*, pp. 675–678. ACM, 2014, 2014.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014, 2014.
- Danica Kragic and Henrik I Christensen. Survey on visual servoing for manipulation. *Computational Vision and Active Perception Laboratory, Fiskartorpsv*, 15, 2002, 2002.
- Thomas Lampe and Martin Riedmiller. Acquiring visual servoing reaching and grasping skills using neural reinforcement learning. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2013, 2013.
- Sascha Lange, Martin Riedmiller, and Arne Voigtlander. Autonomous reinforcement learning on raw visual input data in a real world application. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2012, 2012.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016, 2016.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *CoRR*, abs/1509.02971, 2015, 2015.
- William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. *CoRR*, abs/1605.08104, 2016, 2016.
- David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004, 2004.
- Ezio Malis, Francois Chaumette, and Sylvie Boudet. 2 1/2 D visual servoing. *IEEE Transactions on Robotics and Automation*, 15(2):238–250, 1999, 1999.
- Michaël Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *CoRR*, abs/1511.05440, 2015, 2015.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing Atari with deep reinforcement learning. *CoRR*, abs/1312.5602, 2013, 2013.
- Kartik Mohta, Vijay Kumar, and Kostas Daniilidis. Vision-based control of a quadrotor for perching on lines. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3130–3136. IEEE, 2014, 2014.
- Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in Atari games. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2863–2871, 2015, 2015.

- Martin Riedmiller. Neural fitted Q iteration – First experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pp. 317–328. Springer, 2005, 2005.
- Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 2564–2571. IEEE, 2011, 2011.
- Mehdi Sadeghzadeh, David Calvert, and Hussein A Abdullah. Self-learning visual servoing of robot manipulator using explanation-based fuzzy neural networks and Q-learning. *Journal of Intelligent & Robotic Systems*, 78(1):83–104, 2015, 2015.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael I Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1889–1897, 2015, 2015.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014, 2014.
- Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 613–621, 2016, 2016.
- Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 835–851. Springer, 2016, 2016.
- Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2746–2754, 2015, 2015.
- Lee E Weiss, Arthur C Sanderson, and Charles P Neuman. Dynamic sensor-based control of robots with visual feedback. *IEEE Journal on Robotics and Automation*, 3(5):404–417, 1987, 1987.
- William J Wilson, Carol C Williams Hulls, and Graham S Bell. Relative end-effector control using cartesian position based visual servoing. *IEEE Transactions on Robotics and Automation*, 12(5): 684–696, 1996, 1996.
- Tianfan Xue, Jiajun Wu, Katherine Bouman, and Bill Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 91–99, 2016, 2016.
- Billibon H Yoshimi and Peter K Allen. Active, uncalibrated visual servoing. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 156–161. IEEE, 1994, 1994.
- Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *CoRR*, abs/1511.07122, 2015, 2015.

A LINEARIZATION OF THE BILINEAR DYNAMICS

The optimization of Equation (3) can be solved efficiently by using a linearization of the dynamics,

$$f_c^{(l)}(\mathbf{y}_{t,c}^{(l)}, \mathbf{u}) = f_c^{(l)}(\mathbf{y}_{t,c}^{(l)}, \bar{\mathbf{u}}) + \mathbf{J}_{t,c}^{(l)}(\mathbf{u} - \bar{\mathbf{u}}) = f_c^{(l)}(\mathbf{y}_{t,c}^{(l)}, \mathbf{0}) + \mathbf{J}_{t,c}^{(l)}\mathbf{u}, \quad (7)$$

where $\mathbf{J}_{t,c}^{(l)}$ is the Jacobian matrix with partial derivatives $\frac{\partial f_c^{(l)}}{\partial \mathbf{u}}(\mathbf{y}_{t,c}^{(l)}, \bar{\mathbf{u}})$ and $\bar{\mathbf{u}}$ is the linearization point. Since the bilinear dynamics are linear with respect to the controls, this linearization is exact and the Jacobian matrix does not depend on $\bar{\mathbf{u}}$. Without loss of generality, we set $\bar{\mathbf{u}} = \mathbf{0}$.

Furthermore, the bilinear dynamics allows the Jacobian matrix to be computed efficiently by simply doing a forward pass through the model. For the locally bilinear dynamics of Equation (2), the j -th column of the Jacobian matrix is given by

$$\mathbf{J}_{t,c,j}^{(l)} = \frac{\partial f_c^{(l)}}{\partial \mathbf{u}_j}(\mathbf{y}_{t,c}^{(l)}, \mathbf{0}) = \mathbf{W}_{c,j}^{(l)} * \mathbf{y}_{t,c}^{(l)} + \mathbf{B}_{c,j}^{(l)}. \quad (8)$$

B SERVOING COST FUNCTION FOR REINFORCEMENT LEARNING

The goal of reinforcement learning is to find a policy that maximizes the expected sum of rewards, or equivalently, a policy that minimizes the expected sum of costs. The cost should be one that quantifies progress towards the goal. We define the cost function in terms of the position of the target object (in the camera’s local frame) after the action has been taken,

$$c(\mathbf{s}_t, \mathbf{u}_t, \mathbf{s}_{t+1}) = \begin{cases} \sqrt{\left(\frac{\mathbf{p}_{t+1}^x}{\mathbf{p}_{t+1}^z}\right)^2 + \left(\frac{\mathbf{p}_{t+1}^y}{\mathbf{p}_{t+1}^z}\right)^2 + \left(\frac{1}{\mathbf{p}_{t+1}^z} - \frac{1}{\mathbf{p}_*^z}\right)^2}, & \text{if } \|\mathbf{p}_{t+1}\|_2 \geq \tau \text{ and car in FOV} \\ (T - t + 1) c(\cdot, \cdot, \mathbf{s}_t), & \text{otherwise,} \end{cases} \quad (9)$$

where T is the maximum trajectory length. The episode terminates early if the camera is too close to the car (less than a distance τ) or the car’s origin is outside the camera’s field of view (FOV). The car’s position at time t is $\mathbf{p}_t = (\mathbf{p}_t^x, \mathbf{p}_t^y, \mathbf{p}_t^z)$ and the car’s target position is $\mathbf{p}_* = (0, 0, \mathbf{p}_*^z)$, both in the camera’s local frame (z-direction is forward). Our experiments use $T = 100$ and $\tau = 4$ m.

C EXPERIMENT DETAILS

C.1 TASK SETUP

The camera is attached to the vehicle slightly in front of the robot’s origin and facing down at an angle of $\pi/6$ rad, similar to a commercial quadcopter drone. The robot has 4 degrees of freedom, corresponding to translation and yaw angle. Pitch and roll are held fixed.

In our simulations, the quadcopter follows a car that drives at 1 m s^{-1} along city roads during training and testing. The quadcopter’s speed is limited to within 10 m s^{-1} for each translational degree of freedom, and its angular speed is limited to within $\pi/2 \text{ rad s}^{-1}$. The simulator runs at 10 Hz. For each trajectory, a car is chosen randomly from a set of cars, and placed randomly on one of the roads. The quadcopter is initialized right behind the car, in the desired relative position for following. The image observed at the beginning of the trajectory is used as the goal observation.

C.2 LEARNING FEATURE DYNAMICS

The dynamics of all the features were trained using a dataset of 10000 triplets $\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}$. The observations are 128×128 RGB images and the actions are 4-dimensional vectors of real numbers encoding the linear and angular (yaw) velocities. The actions are normalized to between -1 and 1 .

The training set was generated from 100 trajectories of a quadcopter following a car around the city with some randomness. Each trajectory was 100 steps long. Only 5 training cars were shown during learning. The generation process of each trajectory is as follows: First, a car is chosen at random from the set of available cars and it is randomly placed on one of the roads. Then, the quadcopter is placed at some random position relative to the car’s horizontal pose, which is the car’s pose that has been rotated so that the vertical axis of it and the world matches. This quadcopter position is uniformly sampled in cylindrical coordinates relative to the car’s horizontal pose, with heights in the interval 12 m to 18 m, and azimuthal angles in the interval $-\pi/2$ rad to $\pi/2$ rad (where the origin of the azimuthal angle is the back of the car). The radii and yaw angles are initialized so that the car is in the middle of the image. At every time step, the robot takes an action that moves it towards a target pose, with some additive Gaussian noise ($\sigma = 0.2$). The target pose is sampled according to the same procedure as the initial pose, and it is sampled once at the beginning of each trajectory.

We try the fully and locally connected dynamics for pixel intensities to better understand the performance trade-offs when assuming locally connected dynamics. We do not use the latter for the semantic features since they are too high-dimensional for the dynamics model to fit in memory. The dynamics models were trained with ADAM using 10000 iterations, a batch size of 32, a learning rate of 0.001, and momentums of 0.9 and 0.999, and a weight decay of 0.0005.

Feature Dynamics	Policy Optimization Algorithm				
	unweighted feature dynamics + CEM (1500)	feature dynamics + CEM (3250)	feature dynamics + TRPO (≥ 80)	feature dynamics + TRPO (≥ 2000)	ours, feature dynamics + FQI (20)
pixel, FC	8.20 ± 0.66	7.77 ± 0.66	9.56 ± 0.62	8.03 ± 0.66	7.92 ± 0.67
pixel, LC	8.07 ± 0.74	7.13 ± 0.74	10.11 ± 0.60	7.97 ± 0.72	7.98 ± 0.77
VGG conv1_2	2.22 ± 0.38		2.06 ± 0.35	1.66 ± 0.31	1.89 ± 0.32
VGG conv2_2	2.40 ± 0.47		2.42 ± 0.47	1.89 ± 0.40	1.40 ± 0.29
VGG conv3_3	2.91 ± 0.52		2.87 ± 0.53	1.59 ± 0.42	1.56 ± 0.40
VGG conv4_3	2.70 ± 0.52		2.57 ± 0.49	1.69 ± 0.41	1.11 ± 0.29
VGG conv5_3	3.68 ± 0.47		3.69 ± 0.48	3.16 ± 0.48	2.49 ± 0.35

(a) Costs when using the set of cars seen during learning.

Feature Dynamics	Policy Optimization Algorithm				
	unweighted feature dynamics + CEM (1500)	feature dynamics + CEM (3250)	feature dynamics + TRPO (≥ 80)	feature dynamics + TRPO (≥ 2000)	ours, feature dynamics + FQI (20)
pixel, FC	8.84 ± 0.68	8.66 ± 0.70	10.01 ± 0.62	8.75 ± 0.67	9.00 ± 0.70
pixel, LC	8.37 ± 0.75	7.17 ± 0.75	11.29 ± 0.57	8.25 ± 0.71	8.36 ± 0.79
VGG conv1_2	2.03 ± 0.43		1.79 ± 0.36	1.42 ± 0.33	1.78 ± 0.37
VGG conv2_2	2.01 ± 0.44		2.00 ± 0.45	1.26 ± 0.30	1.28 ± 0.30
VGG conv3_3	2.03 ± 0.47		2.08 ± 0.47	1.46 ± 0.37	1.04 ± 0.31
VGG conv4_3	2.40 ± 0.50		2.57 ± 0.53	1.48 ± 0.36	0.90 ± 0.26
VGG conv5_3	3.31 ± 0.45		3.55 ± 0.50	2.76 ± 0.42	2.56 ± 0.41

(b) Costs when using novel cars, none of which were seen during learning.

Table 2: Costs on test executions of the dynamics-based servoing policies for different feature dynamics and weighting of the features. The reported numbers are the mean and standard error across 100 test trajectories, of up to 100 time steps each. We test on executions with the training cars and the novel cars; for consistency, the novel cars follow the same route as the training cars. We compare the performance of policies with unweighted features or weights learned by other methods. For the case of unweighted feature dynamics, we use the cross entropy method (CEM) to learn the relative weights λ of the control and the single feature weight w . For the other cases, we learn the weights with CEM, Trust Region Policy Optimization (TRPO) for either 2 or 50 iterations, and our proposed FQI algorithm. CEM searches over the full space of policy parameters w and λ , but it was only ran for pixel features since it does not scale for high-dimensional problems. We report the number of training trajectories in parenthesis. For TRPO, we use a fixed number of training *samples* per iteration, whereas for CEM and FQI, we use a fixed number of training *trajectories* per iteration. We use a batch size of 4000 samples for TRPO, which means that at least 40 trajectories were used per iteration, since trajectories can terminate early, i.e. in less than 100 time steps.

C.3 LEARNING WEIGHTING OF FEATURE DYNAMICS WITH REINFORCEMENT LEARNING

We use CEM, TRPO and FQI to learn the feature weighting and report the performance of the learned policies in Table 2. We use the cost function described in Appendix B, a discount factor of $\gamma = 0.9$, and trajectories of up to 100 steps. All the algorithms used initial weights of $w = 1$ and $\lambda = 1$, and a Gaussian exploration policy with the current policy as the mean and a fixed standard deviation $\sigma_{\text{exploration}} = 0.2$.

For the case of unweighted features, we use CEM to optimize for a single weight w and for the weights λ . For the case of weighted features, we use CEM to optimize for the full space of parameters, but we only do that for the pixel feature dynamics since CEM does not scale for high-dimensional problems, which is the case for all the VGG features. Each iteration of CEM performs a certain number of noisy evaluations and selects the top 20% for the elite set. The number of noisy evaluations per iteration was 3 times the number of parameters being optimized. Each noisy evalua-















Feature Dynamics	Observations from Test Executions	Cost
pixel, fully connected		24.74
		16.69
pixel, locally connected		24.92
		16.47
VGG conv1_2		15.91
		1.57
VGG conv2_2		7.53
		2.56
VGG conv3_3		6.01
		3.76
VGG conv4_3		5.94
		4.31
VGG conv5_3		15.51
		17.39

Table 3: Sample observations from test executions in our experiments, and the costs for each trajectory, for different feature dynamics. We use the weights learned by our FQI algorithm. This table follows the same format as Table 1. Some of the trajectories were shorter than 100 steps because of the termination condition (e.g. the car is no longer in the image). The first observation of each trajectory is used as the target observation. The trajectories shown in here were chosen to reflect different types of behaviors. In the first trajectory, the blue car turns abruptly to the right, making the view significantly different from the target observation. In the second trajectory, a distractor object (i.e. the lamp) shows up in the target image and an occluder object (i.e. the traffic light) appears through the execution. The policies based on deeper VGG features, up to VGG conv4_3, are generally more robust to the appearance changes between the observations and the target observation, which are typically caused by movements of the car, distractor objects, and occlusions.

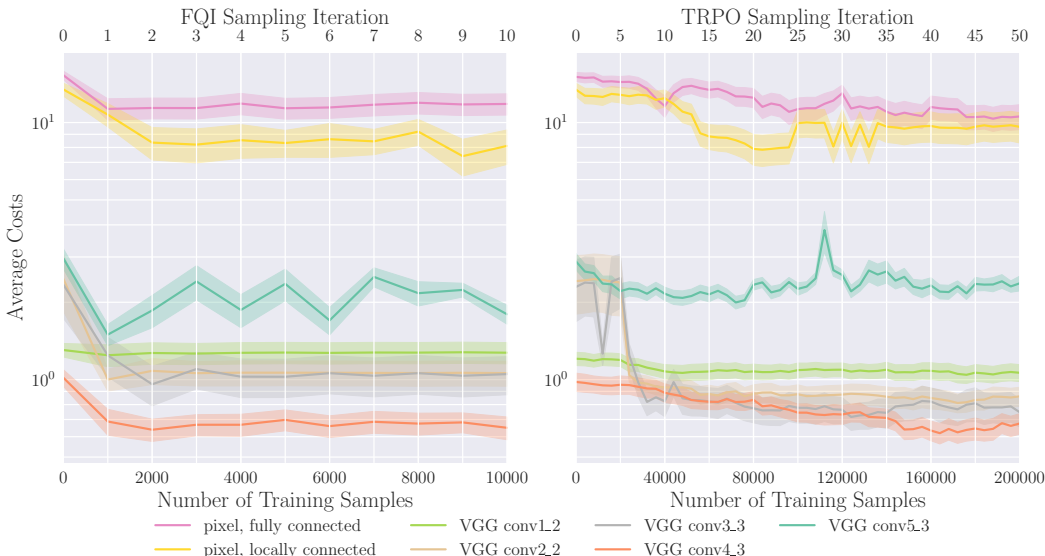


Figure 7: Costs of validation executions using various feature dynamics models, where the feature weights are optimized with FQI (left plot) or TRPO (right plot). The reported values are the mean and standard error across 10 validation trajectories, of up to 100 time steps each.

tion used the average sum of costs of 10 trajectory rollouts as its evaluation metric. The parameters of the last iteration were used for the final policy. The policies with unweighted features dynamics and the policies with pixel features dynamics were trained for 10 and 25 iterations, respectively.

We use TRPO to optimize for the full space of parameters for each of the feature dynamics we consider in this work. We use a Gaussian policy, where the mean is the servoing policy of Equation (3) and the standard deviation is fixed to $\sigma_{\text{exploration}} = 0.2$ (i.e. we do not learn the standard deviation). Since the parameters are constrained to be non-negative, we parametrize the TRPO policies with $\sqrt{\mathbf{w}}$ and $\sqrt{\boldsymbol{\lambda}}$. We use a Gaussian baseline, where the mean is a 5-layer CNN, consisting of 2 convolutional and 3 fully connected layers, and a standard deviation that is initialized to 1. The convolutional layers use 16 filters (4×4 , stride 2) each, the first 2 fully-connected layers use 32 hidden units each, and all the layers except for the last one use ReLU activations. The input of the baseline network are the features (either pixel intensities or VGG features) corresponding to the feature dynamics being used. The parameters of the last iteration were used for the final policy. The policies are trained with TRPO for 50 iterations, a batch size of 4000 samples per iteration, and a step size of 0.01.

We use our proposed FQI algorithm to optimize for the weights \mathbf{w} , $\boldsymbol{\lambda}$, and surpass the other methods in terms of performance on test executions, sample efficiency, and overall computation efficiency⁷. The updates of the inner iteration of our algorithm are computationally efficient; since the data is fixed for a given sampling iteration, we can precompute $\phi(\mathbf{s}_t, \mathbf{u}_t)$ and certain terms of $\phi(\mathbf{s}_{t+1}, \cdot)$. The parameters that achieved the best performance on 10 validation trajectories were used for the final policy. The policies are trained with FQI for $S = 2$ sampling iterations, a batch size of 10 trajectories per sampling iteration, $K = 10$ inner iterations per sampling iteration, and a regularization coefficient of $\nu = 0.1$. We found that regularization of the parameters was important for the algorithm to converge. We show sample trajectories of the resulting policies in Table 3.

The FQI algorithm often achieved most of its performance gain after the first iteration. We ran additional sampling iterations of FQI to see if the policies improved further. For each iteration, we evaluated the performance of the policies on 10 validation trajectories. We did the same for the policies trained with TRPO, and we compare the learning curves of both methods in Figure 7.

⁷Our policy based on conv4.3 features takes around 650 s to run $K = 10$ iterations of FQI for a given batch size of 10 training trajectories.

Observation Modality	
ground truth car position	0.59 ± 0.24
raw pixel-intensity images	3.23 ± 0.22
VGG conv1_2 features	7.45 ± 0.40
VGG conv2_2 features	13.38 ± 0.53
VGG conv3_3 features	10.02 ± 0.49

(a) Costs when using the set of cars seen during learning.

Observation Modality	
ground truth car position	0.59 ± 0.24
raw pixel-intensity images	5.20 ± 0.40
VGG conv1_2 features	8.35 ± 0.44
VGG conv2_2 features	14.01 ± 0.47
VGG conv3_3 features	10.51 ± 0.65

(b) Costs when using a new set of cars, none of which were seen during learning.

Table 4: Costs on test executions of servoing policies that were trained end-to-end with TRPO. These policies take in different observation modalities: ground truth car position or image-based observations. This table follows the same format as Table 2. The mean of the first policy is parametrized as a 3-layer MLP, with tanh non-linearities except for the output layer; the first 2 fully connected layers use 32 hidden units each. For the other policies, each of their means is parametrized as a 5-layer CNN, consisting of 2 convolutional and 3 fully-connected layers, with ReLU non-linearities except for the output layer; the convolutional layers use 16 filters (4×4 , stride 2) each and the first 2 fully-connected layers use 32 hidden units each. All the policies are trained with TRPO, a batch size of 4000 samples, 500 iterations, and a step size of 0.01. The car position observations are not affected by the appearance of the cars, so the test performance for that modality is the same regardless of which set of cars are used.

C.4 LEARNING END-TO-END SERVOING POLICIES WITH TRPO

We use TRPO to train end-to-end servoing policies for various observation modalities and report the performance of the learned policies in Table 4. The policies are trained with the set of training cars, and tested on both this set and on the set of novel cars. The observation modalities that we consider are ground truth car positions (relative to the quadcopter), images of pixel intensities from the quadcopter’s camera, and VGG features extracted from those images. Unlike our method and the other experiments, no feature dynamics are explicitly learned for these experiments.

We use a Gaussian policy, where the mean is either a multi-layer perceptron (MLP) or a convolutional neural net (CNN), and the standard deviation is initialized to 1. We also use a Gaussian baseline, which is parametrized just as the corresponding Gaussian policy (but no parameters are shared between the policy and the baseline). For the policy that takes in car positions, the mean is parametrized as a 3-layer MLP, with tanh non-linearities except for the output layer; the first 2 fully connected layers use 32 hidden units each. For the other policies, each of their means is parametrized as a 5-layer CNN, consisting of 2 convolutional and 3 fully-connected layers, with ReLU non-linearities except for the output layer; the convolutional layers use 16 filters (4×4 , stride 2) each and the first 2 fully-connected layers use 32 hidden units each.

The CNN policies would often not converge for several randomly initialized parameters. Thus, at the beginning of training, we tried multiple random seeds until we got a policy that achieved a relatively low cost on validation trajectories, and used the best initialization for training. The MLP policy did not have this problem, so we did not have to try multiple random initializations for it. All the policies are trained with a batch size of 4000 samples, 500 iterations, and a step size of 0.01. The parameters of the last iteration were used for the final policy.

Observation Modality (Feature Points)		
corners of bounding box from C-COT tracker	(0.75)	1.70 ± 0.30
corners of ground truth bounding box	(0.75)	0.86 ± 0.25
corners of next frame’s bounding box from C-COT tracker	(0.65)	1.46 ± 0.22
corners of next frame’s ground truth bounding box	(0.65)	0.53 ± 0.05
SIFT feature points	(0.30)	14.47 ± 0.75
SURF feature points	(0.60)	16.37 ± 0.78
ORB feature points	(0.30)	4.41 ± 0.60

Table 5: Costs on test executions when using classical image-based visual servoing (IBVS) with respect to feature points derived from bounding boxes and keypoints derived from hand-engineered features. Since there is no learning involved in this method, we only test with one set of cars: the cars that were used for training in the other methods. This table follows the same format as Table 2. This method has one hyperparameter, which is the gain for the control law. For each feature type, we select the best hyperparameter (shown in parenthesis) by validating the policy on 10 validation trajectories for gains between 0.05 and 2, in increments of 0.05. The servoing policies based on bounding box features achieve low cost, and even lower ones if ground truth car dynamics is used. However, servoing with respect to hand-crafted feature points is significantly worse than the other methods.

C.5 CLASSICAL IMAGE-BASED VISUAL SERVOING

Traditional visual servoing techniques (Feddema & Mitchell, 1989; Weiss et al., 1987) use the image-plane coordinates of a set of points for control. For comparison to our method, we evaluate the servoing performance of feature points derived from bounding boxes and keypoints derived from hand-engineered features, and report the costs of test executions on Table 5.

We use bounding boxes from the C-COT tracker (Danelljan et al., 2016) (the current state-of-the-art visual tracker) and ground truth bounding boxes from the simulator. The latter is defined as the box that tightly fits around the visible portions of the car. We provide the ground truth bounding box of the first frame to the C-COT tracker to indicate that we want to track the car. We use the four corners of the box as the feature points for servoing to take into account the position and scale of the car in image coordinates.

We provide the ground truth depth values of the feature points for the interaction matrices. In classical image-based visual servoing, the control law involves the interaction matrix (also known as feature Jacobian), which is the Jacobian of the points in image space with respect to the camera’s control (see Chaumette & Hutchinson (2006) for details). The analytical feature Jacobian used in IBVS assumes that the target points are static in the world frame. This is not true for a moving car, so we consider a variant where the feature Jacobian incorporates the ground truth dynamics of the car. This amounts to adding a non-constant translation bias to the output of the dynamics function, where the translation is the displacement due to the car’s movement of the 3-dimensional point in the camera’s reference frame. Note that this is still not exactly equivalent to having the car being static since the roads have different slopes but the pitch and roll of the quadcopter is constrained to be fixed.

For the hand-crafted features, we consider SIFT (Lowe, 2004), SURF (Bay et al., 2006) and ORB (Rublee et al., 2011) keypoints. We filter out the keypoints of the first frame that does not belong to the car and use these as the target keypoints. However, we use all the keypoints for the subsequent observations.

The servoing policies based on bounding box features achieve low cost, and even lower ones if ground truth car dynamics is used. However, servoing with respect to hand-crafted feature points is significantly worse than the other methods. This is, in part, because the feature extraction and matching process introduces compounding errors. Similar results were found by Collewet & Marchand (2011), who proposed photometric visual servoing (i.e. servoing with respect to pixel intensities) and showed that it outperforms, by an order of magnitude, classical visual servoing that uses SURF features.

Observation Modality (Pose)	Policy Variant					
	Use Rotation			Ignore Rotation		
car pose	(1.55)	0.58	± 0.25	(1.90)	0.51	± 0.25
next frame’s car pose	(1.00)	0.0059 ± 0.0020		(1.00)	0.0025 ± 0.0017	

Table 6: Costs on test executions when using classical position-based visual servoing (PBVS). Since there is no learning involved in this method, we only test with one set of cars: the cars that were used for training in the other methods. This table follows the same format as Table 2. This method has one hyperparameter, which is the gain for the control law. For each condition, we select the best hyperparameter (shown in parenthesis) by validating the policy on 10 validation trajectories for gains between 0.05 and 2, in increments of 0.05. These servoing policies, which use ground truth car poses, outperforms all the other policies based on images. In addition, the performance is more than two orders of magnitude better if ground truth car dynamics is used.

C.6 CLASSICAL POSITION-BASED VISUAL SERVOING

Position-based visual servoing (PBVS) techniques use poses of a target object for control (see Chaumette & Hutchinson (2006) for details). We evaluate the servoing performance of a few variants, and report the costs of test executions on Table 6.

Similar to our IBVS experiments, we consider a variant that uses the car pose of the next time step as a way to incorporate the ground truth car dynamics into the interaction matrix. Since the cost function is invariant to the orientation of the car, we also consider a variant where the policy only minimizes the translational part of the pose error.

These servoing policies, which use ground truth car poses, outperforms all the other policies based on images. In addition, the performance is more than two orders of magnitude better if ground truth car dynamics is used.