

# Learning to Drive a Real Car in 20 Minutes

Martin Riedmiller

Neuroinformatics Group, Univ. of Osnabrueck

Email: martin.riedmiller@uos.de

Mike Montemerlo, Hendrik Dahlkamp

AI Lab, Stanford University

Email: {montemerlo, dahlkamp}@stanford.edu

## Abstract

*The paper describes our first experiments on Reinforcement Learning to steer a real robot car. The applied method, Neural Fitted Q Iteration (NFQ) is purely data-driven based on data directly collected from real-life experiments, i.e. no transition model and no simulation is used. The RL approach is based on learning a neural Q value function, which means that no prior selection of the structure of the control law is required. We demonstrate, that the controller is able to learn a steering task in less than 20 minutes directly on the real car. We consider this as an important step towards the competitive application of neural Q function based RL methods in real-life environments.*

## 1 Introduction

The interest in applying Reinforcement Learning (RL) methods to real life control applications is growing rapidly e.g. [5], [11], [6], [3]. In this paper we focus on situations, where the controller should learn by interacting with the real system only. In particular, for the design of the controller we will not assume, that a system model is available; neither in form of system equations nor in form of a simulator (the latter approach was successfully applied in a number of applications, see e.g. [2], [5]). In contrast to that, here, we only assume that the controller is able to collect state action transitions by observing the real system behaviour while controlling it.

Learning by interacting with the real system directly has an important advantage: the controller is tailored exactly to the behaviour of the real system at hand instead of a more or less exact model of it. The big challenge in learning with real systems lies in the fact that learning must occur in a reasonable amount of time with a reasonable effort: in a real application, one typically can not wait for hundreded thousands of episodes, until a controller is learned.

Another important aspect of this paper, is, that we do not need prior knowledge about the policy to be learned. This has the advantage, that we do not constrain the control law

a priori to a certain class of controllers. Additionally, the proposed approach is applicable even in situations, where a priori no idea about a working control law is available.

In this paper, NFQ [8] is used as the core learning algorithm for a learning controller that learns to steer a real car from scratch. The state space for this problem is 6 dimensional and continuous. We demonstrate, that our approach is able to learn a successful steering in less than 20 minutes of driving the real car. The approach is fully data-driven, which means that no simulation model of the car and no pre-defined control law is used for learning.

## 2 Steering a Robot Car

### 2.1 Task Description

The task considered is to follow a given track by controlling the steering angle of the steering wheel. The deviation of the track is measured in terms of the cross-track-error (cte), which should be kept small. A side goal of the controller is to show a 'smooth' steering behavior, that gently controls the car.

Steering is an intermediate step to a fully autonomous car. In its full version, the robot car shall also activate brakes and throttle. In the current state of the project, throttle and brake are activated by a human driver.

### 2.2 Hardware and Software

The robot car is equipped with several sensors that allow to determine its position both with respect to global coordinates (GPS), and to relative coordinates with respect to the current shape of the road (e.g. visual lane detection). This information delivers the pose estimation of the car. From this pose estimation, essential inputs to the controller are determined, like e.g. the deviation of the car from the given track.

On the output side, the car allows to set the steering angle of the wheel by activating a motor coupled to the steering wheel. Currently, the steering controller outputs the wheel

angle, which is then transformed into control signals to the motor controller.

## 2.3 Classical Controller Design

A good controller must not only reduce the error as quickly as possible, but must also consider the dynamic behavior of the vehicle - e.g. in order to avoid break-outs by too harsh steerings. Also, some kind of pro-activeness is expected: a good steering controller should also consider to some extent the future curvature of the road [1]. An analytical controller design based on these considerations can be found in [4]. This controller was developed based on careful analytical considerations, numerical parameter search methods and handtuning. In contrast, the learning controller proposed in the following, shall learn the control law from interaction with the real car only.

## 3 Reinforcement Learning and NFQ

### 3.1 RL Basics

The control problems considered can be described as Markovian Decision Processes (MDPs). An MDP is described by a set  $S$  of states, a set  $A$  of actions, a stochastic transition function  $p(s, a, s')$  describing the (stochastic) system behavior and an immediate reward or cost function  $c : S \times A \rightarrow \mathbf{R}$ . The goal is to find an optimal policy  $\pi^* : S \rightarrow A$ , that minimizes the expected cumulated costs

$$J^\pi(s) = E \sum_{t=0}^{\infty} c(s_t, \pi(s_t)), s_0 = s \quad (1)$$

for each state. In particular, we allow  $S$  to be continuous and assume  $A$  to be finite for our learning system. The transition model  $p$  is assumed to be unknown to our learning system (model-free approach). Decisions are taken in regular time steps with a constant cycle time.

**Q-Learning and Neural Networks** The idea of classical Q-learning is to allow model-free Reinforcement Learning by iteratively learning an optimal value function over state-action pairs [12]. Typically, it is applied on-line, which means, that after each observation of a transition of the system, the corresponding value of the Q-function is updated by the following rule:

$$Q_{k+1}(s, a) := (1-\alpha)Q(s, a) + \alpha(c(s, a) + \gamma \min_b Q_k(s', b))$$

where  $s$  denotes the state where the transition starts,  $a$  is the action that is applied, and  $s'$  is the resulting state.  $\alpha$  is a learning rate that has to be decreased in the course of learning in order to fulfill the conditions of stochastic

approximation and  $\gamma$  is a discounting factor (see e.g. [10]). It can be shown, that under mild assumptions Q-learning converges for finite state and action spaces, where a table-based representation of the Q-function can be used. If every state action pair is updated infinitely often, in the limit, the optimal Q-function is reached. The greedy exploitation of this optimal Q-function finally yields the optimal policy.

To deal with continuous state spaces, the above Q-learning rule can be adapted to be realized in a function approximator. In the following, we will use a neural network of type multi-layer perceptron to store the Q-value function. Although the theoretical convergence guarantees from above do not hold any more in this case, empirically a lot of successful applications of neural Q functions have been reported. One particular problem when directly implementing the Q-learning rule in a multi-layer perceptron, is that each update for one state-action pair might induce unforeseeable changes at the Q-values for other state-action pairs - disturbing or even destroying the effort done so far. In practice, this might lead to very long training times, often requiring several hundred thousands of trials until a successful policy is learned. The following will report about an idea to better deal with that problem and as a consequence to reduce training effort drastically.

**Neural Fitted Q Iteration (NFQ)** The crucial idea underlying NFQ is the following: Instead of updating the neural value function on-line after each sample (which leads to the above problems), the update is performed off-line considering the entire set of transition experiences done so far [8]. Experiences therefore are collected in triples of the form  $(s, a, s')$  by interacting with the (real) system. Here,  $s$  is the original state,  $a$  is the action applied and  $s'$  is the resulting state. The set of experiences is called the sample set  $\mathcal{D}$ .

The NFQ algorithm [8] consists of two major steps: The generation of the training set  $P$  and the training of these patterns within a multi-layer perceptron. The input part of each training pattern consists of the state  $s^l$  and action  $a^l$  of training experience  $l$ . The target value is computed by the sum of the transition costs  $c(s^l, a^l, s^{l+1})$  and the expected minimal path costs for the successor state  $s^l$ , computed on the basis of the current estimate of the Q-function,  $Q_k$ .

The consideration of the entire training information instead of updating on-line after each sample, has an important further consequence: It allows the application of advanced supervised learning methods, that converge faster and more reliably than online gradient descent methods. In our implementation of NFQ, we use the Rprop algorithm for fast supervised learning [9]. The training of the pattern set is repeated for several epochs (=complete sweeps through the pattern set), until the pattern set is learned successfully.

**Sampling of transitions** In principle, the transitions used in NFQ can be sampled arbitrarily. However, when learning with a real system, collecting transitions that lie on actual trajectories will be probably the most practical implementation. A straight-forward way to do this is to greedily exploit the current neural Q-function to determine the current policy. We call this procedure 'greedy-sampling'. It is implemented by sampling transitions following the current greedy policy until an episode is terminated. Then the data collected on the trajectory is added to the transition set and one or more NFQ iterations are performed. The new policy is determined by greedily exploiting the resulting new Q function. In principle, exploration can be added; however in the following experiments, the policy was determined without exploration.

## 4 An RL Controller for Steering

The control task is to autonomously steer a robot car close to a given track by controlling the angle of the steering wheel. Sensory inputs are the pose of the car (provided by a GPS system), and the speed of the car. The controller output is the desired angle of the steering wheel, in the range of  $\pm 520$  degrees. The controller acts with 20 Hz, corresponding to a control interval of  $50ms$ . The deviation of the car from the given track is measured by the 'cross-track-error (cte)'. As a strict constraint, we request the cte to be always less than  $0.5m$ .

The learning controller shall learn by directly interacting with the real car itself. Observations of sensory values of car behavior recorded during driving are the only source of information, meaning that the approach is purely data-driven. No model, neither analytical nor in terms of a simulator, are considered for learning.

### 4.1 Setup of the learning task

**Immediate costs** The RL task is formulated as the minimisation of expected cumulated costs over time. The cost function  $c(s, u) : S \times U \rightarrow \mathbb{R}$  is chosen to be:

$$c(s, u) = \begin{cases} 0 & , \text{ if } |cte| < 0.05m \text{ (success)} \\ +1 & , \text{ if } |cte| > 0.5m \text{ (failure)} \\ 0.01 & , \text{ else} \end{cases}$$

The cost function reflects the desire for a fast reduction to a low error with some tolerance ( $\pm 0.05m$ ) in accuracy. A cte-value of more than  $\pm 0.5m$  is considered a failure and must be avoided under all circumstances. This is reflected by punishing failure states with a final cost of +1, which is the upper bound of the sigmoidal output neuron of the neural Q function.

**Controller inputs** The controller input consists of 5 continuous variables. The aim is to represent kinematic and dynamic information about the state of the car, that is relevant for the steering task. The most important information is the cross-track-error (*cte*) and its first-order time derivative (*cte*). Since the dynamic behavior of the car crucially depends on its speed, the speed of the rear wheels,  $v_{rear}$  is the third input. The fourth input is the 'heading-error', which measures the difference between the heading angle of the car and the curvature of the track. The fifth input variable is the so called 'yaw-rate-matching', which measures the difference of the speed of change of the yaw angle with respect to the speed of change of the track curvature. The latter two variables shall provide the controller with the possibility of a 'pro-active' steering, by taking the development of the track into account.

The considerations that lead to the selection of the above variables were based on the discussions of relevant inputs for an analytical controller design [4]. We found this to be a reasonable starting point; testing other combinations of input variables is an area of future research.


**Q function and policy** To represent the Q-function, a neural network of type multi-layer perceptron is used. The control action is determined by selecting the action that minimizes the Q-value of the current state, i.e.

$$\pi(s) = \arg \min_u Q(s, u).$$

To do this, for each action respective input vectors have to be propagated through the net and the minimum output value has to be determined. Since the action set typically consists of only few actions, and the neural net computations are fast, this operation is not time-critical.

**Actions** The output of the controller delivers the angle of the steering wheel. The RL approach as we use it here, requires a discretisation of actions. However, when too few actions are provided this might result in a very harsh and uncomfortable policy (e.g. 'bang-bang' control).

The idea of 'dynamic output elements' (DOEs) [7] is to enhance the core RL controller by an additional dynamic operator, that takes the RL action as input and computes the final control signal as its output. For the task at hand, we use a DOE with a simple integrator property (I-DOE), i.e. the DOE sums the RL actions over time. The state of the I-DOE accordingly evolves by  $s_{DOE}(t) = s_{DOE}(t-1) + u'(t)$ , where  $u'(t)$  is the action selected by the RL core controller. The DOE output  $u(t) = s_{DOE}(t)$  is the control signal (i.e. the steering wheel angle) applied to the car. To describe the complete state of the original plant and DOE, the input of the RL core controller is extended by the DOE state information  $s_{DOE}()$ .



**Figure 1. Preliminary experiments on a simulated car: behavior of an RL controller enhanced with an I-DOE. Upper figure: The cross-track-error (cte) is perfectly kept close to zero, never leaving its valid range. Lower figure: Thanks to the I-DOE method, the steering actions are considerably smooth and cover a wide range of steering angles, although only 5 basic actions are used.**

The action set for the RL DOE-controller is accordingly selected to be  $U' = \{\pm 60^\circ, \pm 10^\circ, 0\}$ . For example, to reach a steering wheel angle of  $-360^\circ$ , the RL core controller therefore has to select the action  $-60^\circ$  six times in a row. The resulting I-DOE controller shows a very smooth behaviour of the final control signal, although only 5 basic actions are provided (see figure 1).

**The final controller structure** Driving the robot car with people sitting inside requires that the car is controlled safely even when learning. Since at the beginning of learning, the RL controller will not be able to stay on the track, precautions have to be taken in order to assure safety. There are several possibilities to deal with that issue:

A. Stop the car automatically, whenever the RL controller failed. Ask the human driver to bring the car back

to a non-failure state and start a new RL training episode (a refinement of this without a human driver would be to virtually reduce the cross-track-error to a valid state by virtually moving the target track).

B. Use an automatic 'recovery' control policy, that takes over, whenever the RL controller fails. This has the advantage, that driving can be continued without stopping and with only limited human interaction. Of course, this solution requires a recovery policy that brings the car back closer to the track, but as we argue in the sequel, this recovery policy can be arbitrary simple.

Note that in both cases, recovery policies are only used to re-establish a valid state for the RL controller and *not* to produce examples of good steering. No data collection is done during the regime of the recovery control policy; it therefore can be arbitrary simple without influencing the behavior of the RL learning process. For our experiments, we pursued the second approach, since it allows for a fully-automated and convenient operation of the learning process during driving. For the recovery policy, we used an analytically derived steering controller developed in [4].

**Implementation details** Whenever the RL controller reports a failure ( $cte > 0.5m$ ), the recovery controller immediately takes over. It controls the car, until the state is back within the valid region of the RL controller. Then the RL controller takes over again. To avoid the system to permanently switch between the two controllers, a hysteresis mechanism was applied: After failure, the RL controller is allowed to take over again, only if the error is below a certain threshold, i.e.  $cte < 0.1m$ .

During RL control, the controller continuously stores the observed transitions triples  $(s, u, s')$  to its transition database. When a failure occurred, the following things happen: 1. the episode is terminated, 2. the controller triggers a background process to do one iteration of NFQ to produce a new Q-net based on the complete transition database collected so far, 3. the recovery controller takes over (for safety reasons). In the cycles following a failure, the RL controller actively tests, if the cte fulfills the re-entry criterion ( $cte < 0.1m$ ) and if the background learning process has produced a new neural Q-net. If both pre-conditions are fulfilled, the RL controller takes over again and starts its next episode.

This procedure allows the human driver to only activate throttle and brake - steering is done completely autonomously and continuously either by the RL controller or by the recovery controller. The learning process is controlled completely autonomously. In a successful experiment, we expect the RL controller to incrementally increase its active time in course of learning.

## 5 Learning on the Real Car


The control architecture applied was the I-DOE RL controller enhanced with a recovery controller that takes over in case of failure as described in section 4. Control interval is 50ms, control signal is the steering wheel angle of the car. The action set for the RL controller is  $U' = \{\pm 60^\circ, \pm 10^\circ, 0\}$ . The controller is running on a laptop, that communicates with the car steering software via inter process communication.

The neural RL controller starts completely from scratch, with the weights of the neural Q function randomly initialized within the range  $\pm 0.5$ . As the neural Q function, a multi-layer perceptron was applied consisting of one input layer, two hidden layers with 10 neurons each and an output layer with one output neuron. For all neurons, sigmoidal activation functions with range  $(0, 1)$  were used. The inputs consist of 5 continuous variables describing the state of the car, one additional input for the state of the I-DOE, and another input for the candidate action. Transition data was collected in a greedy-sampling manner, i.e. in each episode, the current Q function was greedily exploited to determine the new control (and sampling) policy. After each episode, one iteration of NFQ is performed. The neural network training method applied is Rprop, running for 300 epochs. Rprop was applied with its standard parametrization [9]. For the training pattern set, 100 extra-patterns according to the 'hint-to-goal' heuristic were generated [8]. The immediate cost function used was the one described in section 4.1.

Activating throttle and brake was done by a human driver. Speeds of the car therefore vary with the mood of the driver. For learning, speeds were typically between 4m/s to 7.5m/s. The learned controller was also tested for speeds up to 9m/s.

### 5.1 Learning results

As described in section 4 learning is triggered completely autonomously. The only thing the human driver has to do is to operate throttle and brake of the real car. Whenever the pre-conditions for the RL controller are fulfilled ( $cte < 0.1$  and new Q-net available), RL takes over and steers the car. Figure 2 shows the behavior of the of the cross-track-error of RL steering the real car in the first 350 seconds of driving and learning. In the (short) pause between two RL episodes, the recovery controller has control while the RL controller is doing NFQ learning. In the first 90s, whenever the RL controller is in charge, it nearly immediately fails by provoking a cte of more than 0.5m. However, after only 90s of gross driving time, the RL controller shows its first successes by driving about 30s until failure. After 250s gross driving time, RL managed 80s of



**Figure 2. Behavior of the cross-track-error for the RL controller in the first 350 seconds of driving. In the first couple of trials, the controller always quickly fails. After only 90s driving, RL successfully drives for about 30s, after about 250s, RL drives for more than 80s.**


autonomous steering.

The first controller, that manages a complete round (which was a round trip of about 800m length, corresponding to about 130 seconds driving) was obtained after 56 episodes, after about 11 minutes of driving with the car. That first successful controller still acted pretty 'shaky' putting the goal of a small error higher than the desire (of the passengers!) for a smooth steering.

The next controller that completed the round was obtained after 60 episodes in total, after about 16 minutes of driving. Figures 3 show the cross-track error which could be successfully held below 0.5m deviation for the entire round. The controller has successfully learned to use its integrating DOE to produce large steering angles if required. At that time, the database contained about 20 000 transition tuples. With that amount of data, one iteration of NFQ learning took about 30s training time on the laptop used.

Having shown that our approach is able to learn from scratch based on a very reasonable amount of real data, we stopped the experiment after 70 episodes in total. Up to that point, the total driving time was only 25 minutes. We consider that to be a very reasonable time for learning a controller completely from scratch.

**Discussion** Considering that our learning approach is not based neither on an initial guess of a parametrized policy nor on a (simulation) model, NFQ has successfully shown its ability of being very data-efficient. A good part of this performance is due to the good generalisation abilities of multi-layer perceptrons, which in combination with the explicit memorisation of transitions seems to be one of the key



**Figure 3. Behavior of RL controller on real car after 60 episodes (about 16 minutes driving). Above: The cross-track-error is successfully kept below 0.5m. Below: Control signal as produced by the RL-DOE controller. The controller learned to generate steering wheel angles of up to 400 degrees, in order to keep the car successfully on track.**

ingredients of successful value function learning.

The policy learned is successful, but not optimal yet. Further learning should decrease the cumulated cte further. Also, the steering is acceptable, but definitely could be smoother. Further learning means two things: First, more transitions might be collected to have a better representation of the car behavior - which means that learning is continued in the online-greedy manner as described above. In a comparable simulator experiment, we observed very good controllers to occur after about 120 episodes of online-greedy learning.

Another possibility for more learning would be to increase the number of NFQ iterations offline, i.e. without sampling data on the real car in every iteration. Again, in a corresponding simulator experiment, this procedure also produced better controllers when the number of NFQ iterations is increasing.

## 6 Conclusion

This paper shows the successful application of neural Q value function learning on a real car directly. Using NFQ as the core training method, learning to steer from scratch was possible in less than 20 minutes of driving a real car. No prior policy and no model are given. We describe an architecture, where the use of a recovery controller allows simultaneous learning during driving. However, the use of such a recovery controller is only optional for convenience; at no point it is used to collect data or to serve as an example policy for steering.

## References

- [1] T. Gillespie. *Fundamentals of Vehicle Dynamics*. SAE Publications, Warrendale, PA, 1992.
- [2] R. Hafner and M. Riedmiller. Reinforcement learning on an omnidirectional mobile robot. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003), Las Vegas, 2003*.
- [3] R. Hafner and M. Riedmiller. Neural Reinforcement Learning Controllers for a Real Robot Application. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 07)*, Rome, Italy, 2007.
- [4] G. M. Hoffmann, C. J. Tomlin, M. Montemerlo, and S. Thrun. Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing. In *To appear in the Proceedings of the 26th American Control Conference*, New York, NY, July 2007.
- [5] A. Y. Ng, A. Coates, M. Diel, V. Ganapathi, J. Schulte, B. Tse, E. Berger, and E. Liang. Inverted autonomous helicopter flight via reinforcement learning. In *International Symposium on Experimental Robotics*, 2004.
- [6] J. Peters and S. Schaal. Policy gradient methods for robotics. In *Proceedings of the IEEE international conference on intelligent robotics systems (iros 2006)*, 2006.
- [7] M. Riedmiller. Generating continuous control signals for reinforcement controllers using dynamic output elements. In *European Symposium on Artificial Neural Networks, ESANN'97*, Bruges, 1997.
- [8] M. Riedmiller. Neural Fitted Q Iteration - First experiences with a data efficient neural Reinforcement Learning Method. In *Proc. of the European Conference on Machine Learning, ECML 2005*, Porto, Portugal, October 2005.
- [9] M. Riedmiller and H. Braun. A direct adaptive method for faster backpropagation learning: The RPROP algorithm. In H. Ruspini, editor, *Proceedings of the IEEE International Conference on Neural Networks (ICNN)*, pages 586 – 591, San Francisco, 1993.
- [10] R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, Cambridge, MA, 1998.
- [11] R. Tedrake, T. W. Zhang, and H. S. Seung. Learning to walk in 20 minutes. In *Proceedings of the Fourteenth Yale Workshop on Adaptive and Learning Systems*, 2005.
- [12] C. J. Watkins. *Learning from Delayed Rewards*. Phd thesis, Cambridge University, 1989.