
Incorporating Voice Instructions in Model-Based Reinforcement Learning for Self-Driving Cars

Mingze Wang, Ziyang Zhang, Grace Hui Yang
InfoSense, Department of Computer Science
Georgetown University, United States
{mw1222, zz249, grace.yang}@georgetown.edu

Abstract

This paper presents a novel approach that supports natural language voice instructions to guide deep reinforcement learning (DRL) algorithms when training self-driving cars. DRL methods are popular approaches for autonomous vehicle (AV) agents. However, most existing methods are sample- and time-inefficient and lack a natural way to communicate with a human expert. In this paper, how new human drivers learn from human coaches motivates us to study new ways of human-in-the-loop learning and a more natural and approachable training interface. We propose incorporating natural language voice instructions (NLI) in model-based deep reinforcement learning for self-driving cars. We evaluate the proposed method together with a few state-of-the-art DRL methods in the CARLA simulator. The results show that natural language instructions can help ease the training process and significantly boost the agents' learning speed.

1 Introduction

Research in self-driving cars, also known as autonomous vehicles (AV), has made rapid progress due to significant deep learning advancements. Deep reinforcement learning (DRL) methods are popular approaches for self-driving cars. A DRL agent learns by interacting with the driving environment and gradually forms a good policy that satisfies constraints in the environment. An effective policy can take up many trial-and-error steps. In each learning episode, the agent improves its policy slightly towards an optimization objective, which in most cases is to maximize the long-term expected cumulative rewards [29]. While these algorithms show promising results, most of them are sample inefficient and require substantial training episodes. Model-based reinforcement learning has been proposed to improve sample efficiency by learning an environmental model and using the model to better plan the actions.

However, learning the policy and learning the environmental model are still time- and sample-consuming. When watching an agent's learning curve develop, it is not uncommon for an engineer to think and almost verbalize to the agent that –“Hey, that move is stupid”, “Good job!” or “Just go there.” When seeing good moves, she might want to praise the agent; she might like to remind when she sees terrible actions. However, such a “direct”, natural, and quick communication channel does not exist.

On the other hand, we notice that human novice drivers can learn driving with much less training. Human drivers have a human coach sit beside them, giving instructions and evaluative feedback. The student driver does the driving during a lesson, and the coach verbally delivers commands, praises, and scolds. Sometimes, the coach explains traffic rules and driving conventions, paired with stories and reasoning. Their way of communication is through natural language (NL) utterances. The student may not speak much when focusing on driving and learning; the coach does most of the talking.




Figure 1: Model-Based RL with Voice Instructions.




Figure 2: Example town in CARLA



Figure 3: Front camera inputs town in CARLA

Compared to an artificial learner, only a small amount of instructions from the coach is needed by the human student, and she learns pretty efficiently. On average, a new human driver takes three to five behind-the-wheel classes, each of which lasts around two hours, to pass the road test for a class-D license.¹

This motivates us to work on a new way to communicate and train autonomous car agents. In this paper, we propose to give natural language voice instructions to train AV agents. Compared to keyboard strokes or other control commands, natural language is very convenient, natural, and fast in communicating. Recent progress in automated speech recognition (ASR), natural language processing (NLP), and personal assistants encourages people to interact with the agents with their voices. Our goal is to provide the agents with the most natural and direct interference and improve and accelerate their learning.

The voice training interface that we present can potentially facilitate two scenarios. First, it can help machine learning practitioners interact with their algorithms more naturally, directly impacting the agent’s learning process. Second, the voice inputs allow non-technical end-users to continue the training after purchasing the agent home. It is like online learning after they have owned the product for personal use.

Particularly, we propose incorporating voice instructions in model-based deep reinforcement learning to train self-driving cars. First, we collect voice instructions offline and build a training dataset and a taxonomy of NL instructions. Our data analysis shows that most instructions are indeed about actions and rewards. Second, we create a human-in-the-loop framework that incorporates live voice instructions. Two places in a model-based RL algorithm can include natural language instructions. The first is policy learning, where the natural language instructions can improve or overwrite the agent’s actions. The second is the environmental reward function. Ideally, a reward function for a self-driving car would handle multiple objectives, such as safety, conforming to traffic rules, speed under control, and whether passengers are feeling comfortable. It is often pre-set before the learning starts. Allowing human interference during training can diversify the reward function from its pre-set formulation and better suit real-world situations.

We evaluate our method in the CARLA simulator [9, 6, 27]. It is an open-source simulator that supports flexible sensor suites, environmental conditions, complete control of static and dynamic actors, and map generation. However, we realize, compared to actual road tests, the simulator’s image resolution is lower and has fewer traffic signs (it only has traffic lights, stop signs, and speed limits). The simulated town is also small, and other vehicles are assumed to drive ideally. No aggressive drivers are present. Regardless of the idealization, we think CARLA is sufficient and use it to test our idea.

Although our method is general and can be used in combination with various learning methods for self-driving cars, we investigate the Dyna algorithm [29] with PPO in this paper[28]. Experiments show that using voice instructions to interfere with self-driving cars’ learning can significantly improve the learning efficiency and make the agents converge much faster to a good policy.

2 Related Work

Research in Autonomous Vehicles. Research on autonomous driving agents can be categorized as modular models and end-to-end models [11]. Traditionally, self-driving cars use a pipeline of perception-planning-action modular models to compute their driving decisions [15]. More recently,

¹<https://dmv.dc.gov/page/mandatory-driver-education>.

end-to-end learning methods, which directly map from raw sensor inputs to control commands [11], have become popular for autonomous driving. Many are deep reinforcement learning (DRL) methods that directly optimize the whole system’s effectiveness from raw inputs and environmental rewards. For instance, [14] asynchronous actor-critic agents (A3C) [24] were used to map image inputs from a forward-facing camera to control a car agent in a game simulator. The algorithm uses millions of interactions with the environment to converge to a good racing policy. [15] employed deep deterministic policy gradient [19] for their self-driving car agents. Their agent learns to follow a 250m rural lane within only half an hour of training. [31] applied proximal policy optimization (PPO) [28] to a lane following task in the simulator CARLA, with an average of eight hours of training. All three pieces of work can solve the simple lane following problem, where the car drives on a route with no intersections, no traffic lights, and no other vehicles. However, even with such a simple task setting, the training process is still largely inefficient, especially at the beginning of the training.

Simulated training. It can be an extremely high cost to train self-driving cars on real roads. Therefore, many algorithms use simulators to train the agents. CARLA [9] is a popular simulator used in recent research [6, 27] for autonomous driving research. It can support the development, training, and validation of autonomous driving systems. The simulator provides a flexible specification of sensor suites, environmental conditions, full control of all static and dynamic actors, map generation, and many other functionalities. As a simulator, CARLA’s limitations are mostly its simplification to real-world situations. For example, it has limited types of traffic signs, low texture resolution, and small towns. But for the essential tasks of a self-driving car, such as following the traffic lanes and watching out the pedestrians, CARLA is sufficient. This work uses the “Town02” environment with random weather settings and a limited refresh rate of 10 per second.

Human-in-the-loop learning. Our work belongs to human-in-the-loop learning (HILL), a recent machine learning paradigm that has generated significant practical and theoretical interests [20]. We group HILL research into the following categories by the *usages of the human interventions* and by the *types of human feedback*.

Based on where an RL algorithm uses human intervention, we can classify the works into direct policy learning, reward learning, and policy and reward learning. The first and most common usage of human intervention is providing evaluative feedback directly to the action generated by a policy [4, 21]. These methods are simple and effective. They help the policy learning quickly converge. The second type is to incorporate human feedback in reward learning, aka inverse reinforcement learning. The idea is to use human trainers to control the robots or agents to generate a demonstration [3, 18] and the inverse reinforcement learning can benefit from the generated demonstrations to learn the reward functions [1].

Various types of human interventions have been explored in HILL. For example, the simplest one is the hardware delivered feedback, mainly including mouse clicks, keyboard keys, and other sensors [16, 21]. Although this method is precise, it needs reaction time for human trainers to express the feedback by using a keyboard or mouse. Instead, we can use mores natural ways to provide feedback such as facial feedback [2, 5], natural language feedback [17, 22] and gesture feedback [7, 32]. VNLA [25] use natural language in a navigation task. If the agent is stuck somewhere or gets lost, the human trainers give the agent detailed instructions on reaching the next sub-goal. The instructions contain the direction and the features of the sub-goal, which is easy for the agent to detect and reach. However, the natural language feedback in VNLA is not in real-time. For AV tasks, real-time instructions are more effective than non-real-time ones. Our approach uses real-time instructions to coach the agents on driving and achieve higher AV tasks.

The most similar piece of work to ours is perhaps the Deep COACH (CONvergent Actor-Critic by Humans) method [4]. Deep COACH works with a human trainer as part of the deep actor-critic reinforcement learning algorithm. It considers human feedback as a more accurate evaluation of the agent’s choice than a reward function. The human trainer acts as the critic, and the objective policy acts as the actor. Thus the relationship between human trainers and the objective policy is suitable for the actor-critic algorithm. The types of human interventions used in Deep COACH are simple positive or negative values. However, an actual human trainer can indeed provide more informative information to the agent beyond simple positive or negative feedback. In this paper, we propose to use natural language as the communication channel between the human trainer and the robotic car agent, aiming to provide the agents with richer and more precise instructions that make full use of human trainers.

Table 1: Taxonomy of Voice Instructions.

Type	Sub-type	Example	Implementation
Action	Go Straight	Just keep your car straight. Hold steering wheel steady. That’s too much oscillation.	$steer = 0$
	Turn Left	To the left. Take left, take left, left.	$steer = -0.2$
	Turn Right	Take a right. Move, right.	$steer = 0.2$
	Speed up	You can probably speed up. Go go, a bit faster.	$throttle+ = 0.15$
	Slow down	oming up slow, slow down. You need to slow down.	$throttle- = 0.15$
	Stop	Stop sign. Stop! There’s a stop sign so we have to stop.	$brake = 1$
Reward	Great	Cool! Yeah, you are doing well.	$r_h = 30$
	Good	All right. It’s ok.	$r_h = 10$
	Bad	Oh, no. That’s that’s not good.	$r_h = -10$
	Terrible	What are you doing? You can’t do that.	$r_h = -30$
	Mistake	You blew another red light. You’re on the sidewalk.	$r_h = -30$
	Accident	You ran into a wall. You blew a stop sign.	$r_h = -30$
Reasoning	Causal	There’s a stop sign so we have to stop. You’re approaching a yield. So you may want to look for cars. No need to stop. You. Didn’t need to stop there. There’s no sign.	
State	Describe	There’s a stop sign. You’re at a red light. Watch out the grass!	

3 Constructing a Voice Instruction Taxonomy

Our research aims to give natural language voice instructions directly to self-driving car agents for more effective and efficient training. The challenge of this project is to train the agent to understand natural language coaching instructions. To serve this goal, we collected a training dataset and built an NL instruction taxonomy from it. The dataset has been annotated and will be made available for research use upon publication.

Collecting Natural Language Instructions. To build the dataset, we gather audio recordings from machine learning practitioners for their voice instructions when they are watching the agents’ driving. Our goal is to create a representative, natural, free-form voice instruction dataset.

The participants were instructed to imagine self-driving car coaches teaching a new driver (the self-driving car) in a driving school. Their instructions to the agent should be natural conversations instead of from a controlled vocabulary. For instance, the instructions can be direct commands, praises, curses, or rules that they wish the agent to learn. The participants were encouraged to be creative and talk freely, as long as the instructions fit the role of a driving coach. Six participants who are familiar with machine learning and driving participated in the data collection. The participants’ names or user ids are fully anonymized and the data are password-protected and stored in a secure place. The participants were also instructed to only talk in the context of teaching an agent to drive and not to talk about personal, medical, or other sensitive information.

Four types of sample driving videos are provided based on their training levels: (1) “Non-trained” are videos from a random self-driving car agent who receives no training at all. It can act quite weird and full of crashes. (2) “Mid-trained” are videos from self-driving car agents whose learning curves are not yet converged and only half-trained. They can follow the lane twisty. (3) “Trained” are the videos from the self-driving car agents who have been trained by a model-free RL method and can follow the lane correctly. (4) “Autopilot” are videos produced by the autopilot control in the simulator. It

can be considered as the gold standard that an artificial agent can achieve. When watching the videos, the participants were instructed to record their voice instructions using an audio recording device. They were also asked to try their best to synchronize the voice to the video. In the end, we collected 20 audio clips.

Speech-to-Text Transcription. We used Google’s speech-to-text (STT) [10] to convert these audio clips into text transcripts. The text transcripts are then segmented into utterances and aligned with the videos. The utterances range from short phrase-like commands to long sentences. The content includes direct commands such as “go straight”, positive recognition, expressions of surprises and disappointment, and reasoning. These utterances may not be 100% grammatically correct. However, they show us a representative group of expressions that a human coach would use to teach a car agent.

Note that the speech-to-text tool can also perform real-time speech recognition. Each time the microphone receives an instruction, it sends it to the server of STT and returns the corresponding text scripts quickly. We present our use of it for real-time transcription in later sections. In addition, we are also aware that audios contain more information than their text transcripts. For instance, intonations, volume level, exclamations, and signs would be important signals to express evaluative feedback to the agents or infer the human’s real intention. However, we leave these to future work and concentrate on textual analysis of the dataset.

Building Taxonomy of Instructions. According to the real-world driving experience, we analyze and semi-automatically organize the voice instruction dataset into a taxonomy. Table 1 shows our two-level taxonomy, examples and how to implement them into the learning algorithm.

The two most common types of instructions in the utterances that we collected are actions and rewards. Actions mean the instructions tell the agent how to drive in the next step. In this type of instructions, we have steer commands *turn right*, *turn left* and *go straight* and speed commands *speed up*, *speed down* and *stop*. The reward-related instructions give evaluative feedback to the agent. Some of them indicate severity of an issue, so we use them as different reward levels: *great*, *good*, *bad* and *terrible*. Others directly mention the event type, for instance, *mistake* or *accident*.

4 Model-Based RL with Voice Instructions

This paper proposes incorporating natural language voice instructions in model-based reinforcement learning to train self-driving car agents. Our goal is to support more natural and direct interference from human engineers or end-users to the agents and improve and accelerate the agents’ learning.

Model-Based Deep RL. Similar to most reinforcement learning methods, model-based reinforcement learning is built on top of the Markov Decision Process (MDP). It consists of a tuple $\langle \mathbb{S}, \mathbb{A}, \mathbb{R}, \mathbb{T} \rangle$ to represent state, action, reward, and state transition function.

- State: $s \in \mathbb{S}$, consisting two parts $s = (i, v)$, $i \in \mathbb{I}$, where \mathbb{I} is the set of all possible images captured from cameras, and $v \in \mathbb{R}^+$, the current car running speed. The camera image set \mathbb{I} can include front, side, and rearview. But in this work, we only use the front view, which is adequate for most tasks in CARLA.
- Action: $a \in \mathbb{A}$, consisting of steer $a_{steer} \in [-1, 1]$, throttle $a_{throttle} \in [0, 1]$ and brake $a_{brake} \in [0, 1]$, represented by a vector of length three $a = [a_{steer}, a_{throttle}, a_{brake}]^T$.
- Reward function \mathbb{R} : part of the environmental model in a model-based RL algorithm and is learned via model learning.
- Transition function \mathbb{T} : also part of the environmental model and is learned via model learning. However, it is not in the interest of this paper and we will skip it.

Model-based reinforcement learning has two concurrent learning components. One is similar to model-free reinforcement learning, known as “policy learning” or “direct RL.” Policy learning’s goal is to find the best policy π^* directly via environmental rewards, where assuming the transition model and the reward function are deterministic. Policy learning aims to maximize the objective function $J(\pi)$, which is the expected value of the cumulative future reward $E_{\pi} \left[\sum_{k=t}^T \gamma^k R(s_k, \pi(s_k)) \right]$. Model learning is the other learning component. In our work, we focus more on reward function learning because the CARLA simulator decides the environmental state transition.

Algorithm 1 Teaching AV Agents with Voice Instructions.

```
1: Input: training environment  $E$ ; voice action  $a_h$ , voice reward  $r_h$ ; number of trajectories each epoch  $N$ ;
   replay buffer  $B$ ; reward coefficients  $\beta$ , policy  $\pi_\theta$  with initial parameters  $\theta$ 
2: repeat
3:   for  $t = 0$  to  $N$  do
4:     initialize  $done, traj, E$ 
5:     while not  $done$  do
6:       Get current state  $s$  from  $E$ .
7:       Sample action  $a$  from  $\hat{\pi}_n$ 
8:       if voice action  $a_h \neq \emptyset$  then
9:          $a = a_h$ 
10:      end if
11:      Get reward  $r$  and next state  $s'$  after executing  $a$ 
12:      if voice reward  $r_h \neq \emptyset$  then
13:        Update  $\beta$  for an earlier reward  $r^{prev}$  with  $r_h$ 
14:        for  $i = prev$  to  $t - 1$  do
15:          Update  $(s, a, r)$  at timestep  $i$  in  $traj$  based on  $\beta$ 
16:        end for
17:      end if
18:      Append  $(s, a, r)$  to  $traj$ 
19:    end while
20:     $B$  stores  $traj$ 
21:  end for
22:  Update  $\hat{\pi}_n$  with  $(s, a, r)$  in buffer  $B$ 
23: until Algorithm converges.
```

In this work, we propose incorporating voice instructions into both the model-based RL’s policy learning and model learning. Next, we first present how we obtain and process the voice inputs in real-time, followed by details of incorporating the processed instructions in policy learning and model learning, respectively. The overall algorithm is presented in Algorithm 1.

Mapping Voice Instructions to Actions and Rewards. Our project aims to use human instructions in real-time to teach a learning car agent. One option is to train an end-to-end neural network classifier to map raw audio clips to our instruction categories (Table 1). However, this could be impractical due to the live interaction constraint. We thus propose a two-stage process that first performs speech-to-text transcription over the speech data and then classifies the textual utterances into action- or reward-related commands that the agent can easily operate.

First, we use Google’s speech-to-text (STT) tool [10], which can convert audio into text with real-time speech recognition. Each time the microphone receives voice instructions, it sends it to the server of STT and returns the corresponding text quickly. This service enables us to process a human coach’s live instructions like in a behind-the-wheel driving class. The accuracy of STT depends on many factors, such as speech tones, accent, and microphone quality. The accuracy of STT reported by Google Team is 93.1% (6.9% Word Error Rate) [26].

Second, we employ a BERT-based neural network text classifier [8] to classify each transcribed utterance, i.e., a verbal sentence that the human trainer spoke, to one of the actions or reward subtypes. BERT is a general-purpose and robust pre-trained language model widely used in natural language processing (NLP) tasks. We mainly use BERT-Base’s model [30] to translate the sentences into vectors. The vectors are then fed to a three-layer multi-layer perceptron (MLP) network to make the final classification decision. Our method can reach 94.4% accuracy in the experiments.

NL Instructions for Policy Learning. The neural network classifier maps a real-time NL instruction ins to action a_h in the agent’s action space. We consider the new action command from the human coach a_h has the highest priority so that it can overwrite the agent’s sampled action generated from the agent’s policy. In Algorithm 1, at each interaction loop, if a piece of instruction ins comes, the agent replaces the sampled action a with a_h for t_d seconds. The idea of replacing the agent’s actions with a_h intuitively works with RL because the human coach’s policy, once executed, would be reinforced if it leads to higher rewards.

When using the original PPO [28] as the direct policy learner, one issue arises that after the sampled actions are replaced with a_h , the ratio $\eta(\theta)$ becomes $\frac{\pi_\theta(a_h|s_t)}{\pi_{\theta_{old}}(a_h|s_t)}$. Note that NL instruction a_h is not sampled from the agent’s policy. Thus, it may drastically change the agent’s policy (if the human trainer’s policy deviates far from the agent’s policy) and the policy may not be able to converge. We propose to modify PPO’s regularization by changing the clipping function. Our solution is to clip the old probability by a fix number: $\eta(\theta) = \frac{\pi_\theta(a_h|s_t)}{\max\{\pi_{\theta_{old}}(a_h|s_t), 0.05\}}$. Empirically, we find it helpful to clip the old probability to be larger than 0.05.

NL Instructions for Model Learning. Our reward function r is designed similarly to [31]’s and based on three factors: speed, deviation, and orientation. For a non-terminal state, the reward function is defined as the multiplication of all three factors. For a terminal states, it uses a penalty score r_T :

$$r = \begin{cases} \log \exp \sum \beta_i r_i & \text{Non-terminal states} \\ r_T & \text{Terminal states} \end{cases}, \quad (1)$$

where r_i is one of r_s, r_o, r_d and β_i is the weights to each reward component. The *speed* factor r_s is a function of the car speed s : $r_s(s) = \mathbb{1}_{\{0 \leq s < s_{min}\}} \times \frac{s}{s_{min}} + \mathbb{1}_{\{s_{min} \leq s < s_{target}\}} + \mathbb{1}_{\{s_{target} \leq s < +\infty\}} \times (1 - \frac{s - s_{target}}{s_{max} - s_{target}})$, where $s_{min} = 10\text{km/h}$, $s_{target} = 15\text{km/h}$, $s_{max} = 20\text{km/h}$. The *deviation* factor r_d is a function of the distance d between the center of the car and the center of the road: $r_d(d) = \max\{1 - \frac{d}{d_{max}}, 0\}$, where $d_{max} = 2.5\text{m}$. The *orientation* factor r_o is a function of the angle a between the orientation of the car and the orientation of the road: $r_o(a) = \max\{1 - \frac{a}{a_{max}}, 0\}$, where $a_{max} = 20^\circ$. The self-driving car reaches a terminal state if it drives off track ($r_T = -50$) or collides with any other object ($r_T = -100$) or stops moving for a period of time ($r_T = 0$).

The NL voice reward r_h gives evaluative feedback to the agent. After receiving it, we use r_h to update a previous state’s reward function instead of the current state’s. This is because of the latency of proceeding with the human voice input in real-time. According the frame rate fps and pre-set latency lat , we find the previous state s^{t-g} based on the time gap $g = t - \frac{lat}{fps}$ between the current state s^t and s^{t-g} . We then use supervised regression to learn the reward parameters at $t - g$ that minimize the cost to r_h^t at t :

$$\beta^{*t-g} = \arg \min_{\beta} \sum \|r^{t-g}, r_h^t\|. \quad (2)$$

The newly learned reward function would then be used to update the time steps forward until updated by another human input.

5 Experimental Results

We evaluate our method on the CARLA simulator [9], a popular open-source simulator for autonomous driving cars. The tasks we experiment on include *Lane following*, *red lights and stop signs*, *full lap*, and *avoidance*. We test in both situations, with or without other vehicles running and pedestrians walking in the street. The map of Town02 is used, with random weather conditions initiated. We limit the frame rate to 10 per second and use the asynchronous model.

5.1 Experimental Setup

Metrics. The following metrics are used in the evaluation. 1) The cumulative reward $G(\pi)$ of each episode: $G(\pi) = \sum_{k=1}^T \gamma^k R(s_k, \pi(s_k))$ 2) The success rate out of 100 times for each task. Finishing a whole episode without driving out of the lane is considered a success in the lane following task. Finishing a whole episode without crashing into other vehicles or pedestrians is considered a success in the avoidance task. The final success rate $SuccessRate = \frac{\#of\ successful\ episodes}{100}$. 3) The completion percentage: $P = \frac{Length_c}{Length} \times 100\%$, where $Length$ is the total length of the testing lap and $Length_c$ is the length of the lap a car agent goes through before the episode terminates. 4) Lane Deviation: Lane deviation is the distance in meters between the car and the center of the driving lane. It measures the self-driving car’s ability to stay in its lane.

Baselines. To show the effectiveness of our method, we compare with several best-performing model-free RL methods such as DQN [23], PPO [31], SAC [13] and model-based RL baseline

Table 2: Effectiveness in tasks without other vehicles and pedestrians. The full lap task tests if an agent can drive stably and complete a lap.

	Lane following	Full Lap		
	Success rate	Laps Completed	Cum. Rw	Avg Dev.
DQN [23]	17	13.61 (std 6.96)	81.19 (47.13)	1.01 (0.24)
SAC [13]	57	15.91 (8.05)	90.72 (84.10)	0.84 (0.17)
DA-RB+ [27]	45	1.98 (0.94)	6.32 (3.28)	13.65 (5.25)
DDQ [12]	55	26.29 (5.40)	181.61 (55.78)	1.1 (1.02)
Dyna-PPO [31, 29]	60	18.65 (6.89)	272.87 (131.68)	0.5 (0.07)
Dyna-PPO+NLI (this paper)	100	42.95 (10.24)	635.96 (170.89)	0.4 (0.03)
Deep COACH [4]	21	21.22 (5.30)	82.09 (14.32)	0.61 (0.03)

Table 3: Effectiveness in tasks with vehicles and pedestrians. The avoidance task tests whether the agent can correctly avoid crashing into other vehicles and pedestrians and simultaneously follow the lane.

	Avoidance	Full Lap (Avoidance)		
	Success rate	Laps Completed	Cum. Rw	Avg Dev.
DQN [23]	7	3.61 (std 5.77)	65.63 (107.20)	0.373 (0.158)
SAC [13]	40	12.71 (8.45)	142.31 (100.10)	0.8 (0.24)
DA-RB+ [27]	27	0.89 (2.67)	-7.94 (6.16)	32.66 (9.49)
DDQ [12]	40	17.49 (8.79)	121.22 (57.00)	1.01 (0.77)
Dyna-PPO [31, 29]	57	16.85 (8.52)	266.80 (163.10)	0.71 (0.20)
Dyna-PPO+NLI (this paper)	73	18.88 (6.40)	296.55 (113.65)	0.64 (0.14)
Deep COACH [4]	13	6.41 (4.94)	30.84 (61.33)	0.75 (0.18)

methods such as Dyna (with PPO) and DDQ [12]. DA-RB+ [27] is the SOTA for CARLA at the moment, and it is an imitation learning method. We report our implementation of it here. In addition, we compare with a human-in-the-loop learning method, Deep COACH [4]. These methods share the same training experiences and use similar convolutional neural network architectures for image feature extraction.

Implementation Details. We implement the above algorithms in Tensorflow and CARLA 0.9.10. We use a five-layer residual CNN to process the image input, and the input size is 160×80 RGB. After five convolution layers, the network receives the speed v and combines it with CNN’s output. After another two dense layers, it returns the final result. The action sample rate is $10HZ$, which means that each interaction step takes 0.1 seconds.

5.2 Effectiveness

We compare the proposed algorithms with state-of-the-art self-driving car systems in different fields such as model-based learning, model-free learning, and human-in-the-loop learning. The results are summarized in Table 2 and Table 3. Table 2 shows the performance of AV agents without the presence of other vehicles and pedestrians. Our method outperforms the baselines and reaches a 100% successful rate in lane following task, while the best result of other models is 57%. Our agent drives 63% further in the full lap task than others and consequently reaches a higher cumulative reward. Also, we have the lowest average deviation of 0.4, which means our agent drives more stable than the baselines. The baseline agents continuously change their steer angles; eventually, they cannot find that the best solution is to hold the steer and drive straight. With natural language guidance, we can inform the agent to go straight, which leads to a higher cumulative reward. The agent learns not to change steer too frequently and achieves lower deviation.

Table 3 shows the performance when pedestrians and other vehicles are present. All methods have reduced performance compared with Table 2. It is reasonable because the pedestrians can suddenly walk across the street, and other vehicles may overtake the agent and cut the lane. They add many unpredictable factors to the environment. As human drivers, we are taught to stop when facing these situations. However, the self-driving car agent cannot learn this insight quickly. We observe that most baseline methods take a sharp turn to avoid the pedestrians and then lose control. We can ask the agent to stop in front of pedestrians and vehicles during training with natural language instructions. Thus, our runs show better performance than the baselines.

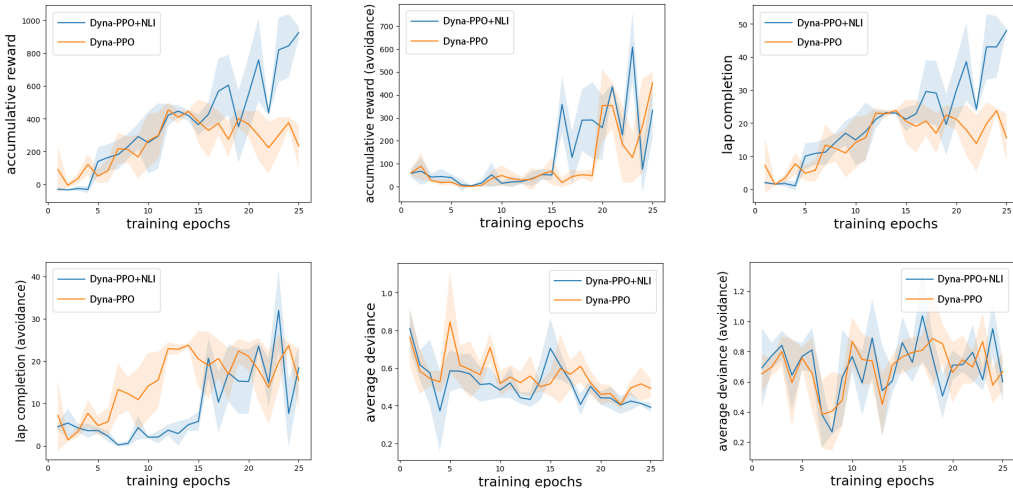


Figure 4: Learning curves comparison in different metrics.

5.3 Effects of using NL Voice Instructions

To show the effects of using NLI, we experiment on the full lap task. We plot the performance of the two models. One model is trained with NLI, and the other is not. The two models share the same hyper-parameter setting and training and test environments. Figure 4 shows the learning curves over six different metrics. In each training epoch, both models sample three trajectories for a round of gradient updates. The total training time is less than one hour for both models. We start coaching the model at the fifteenth epochs, after which we can see Dyna-PPO+NLI’s performance boosts, showing the effectiveness of NLI. The testing videos ² of the AV agent record the self-driving car smoothly turns with a handful of training episodes using NLI. The two models’ behaviors at the turning point lead to the performance difference. The lap completion rate of the Dyna-PPO agent does not increase after the fifteenth epoch because it never successfully turns. The whole turning process takes about 30 to 60 successive steps, making it difficult for a random exploration strategy to find the correct actions in all steps. However, this is not the case for Dyna-PPO+NLI. A human expert can talk to the agent about the correct steps. The agent then reinforces this skill because it leads to a higher expected return. This shows that the RL agent’s random exploration strategy can fail or take too long to learn an optimal policy, but utilizing human expertise for searching optimal policy and teaching the agent during the training may help.

6 Conclusion

This paper proposes to use natural language instructions to help model-based RL when training autonomous driving cars. First, we collect voice instructions offline and build a training dataset and a taxonomy of NLIs. The data analysis shows that most natural language instructions are indeed about actions and rewards for the agent. Second, we build a model-based DRL framework that can incorporate live human voice instructions. Our method can allow both machine learning engineers and everyday end-users to 1) provide feedback to shape the reward function for various driving behaviors and 2) directly instruct the car agent to act. Experiment results show the effectiveness of our method. In the future, we plan to further improve the processing speed of speech-to-text and text classification due to the live nature of our work. Moreover, for the actual deployment of the system, we must have a careful plan to handle real-time speech data in a privacy-preserving manner. Currently, speech recognition is done “in the cloud” and relies on an IT company’s service. This can be improved when we develop highly effective in-house speech recognition tools. This work presents a more natural way for AI agents to be trained through a human-friendly communication channel and contributes to the bigger picture of human-robot cooperation.

²<https://youtu.be/JXU4o7CkttM>.

References

- [1] Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-First International Conference on Machine Learning, ICML '04*, page 1, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138385. doi: 10.1145/1015330.1015430. URL <https://doi.org/10.1145/1015330.1015430>.
- [2] Riku Arakawa, Sosuke Kobayashi, Yuya Unno, Yuta Tsuboi, and Shin ichi Maeda. DQN-TAMER: Human-in-the-loop reinforcement learning with intractable feedback, 2018.
- [3] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [4] Dilip Arumugam, Jun Ki Lee, Sophie Saskin, and Michael L Littman. Deep reinforcement learning from policy-dependent human feedback. *arXiv preprint arXiv:1902.04257*, 2019.
- [5] Joost Broekens. Emotion and reinforcement: affective facial expressions facilitate robot learning. In *Artificial intelligence for human computing*, pages 113–132. Springer, 2007.
- [6] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating, 2019.
- [7] Francisco Cruz, German I Parisi, Johannes Twiefel, and Stefan Wermter. Multi-modal integration of dynamic audiovisual patterns for an interactive reinforcement learning scenario. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 759–766. IEEE, 2016.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019.
- [9] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.
- [10] Google. Google cloud speech-to-text, 2016. <https://cloud.google.com/speech-to-text/>.
- [11] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu. A survey of deep learning techniques for autonomous driving, 2020.
- [12] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2829–2838, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/gu16.html>.
- [13] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. ICML, 2018.
- [14] M. Jaritz, R. d. Charette, M. Toromanoff, E. Perot, and F. Nashashibi. End-to-end race driving with deep reinforcement learning. ICRA, 2018.
- [15] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J. Allen, V. Lam, A. Bewley, and A. Shah. Learning to drive in a day. Technical report, Wayve, 2018.
- [16] W Bradley Knox and Peter Stone. Interactively shaping agents via human reinforcement: The TAMER framework. In *Proceedings of the fifth international conference on Knowledge capture*, pages 9–16, 2009.
- [17] Gregory Kuhlmann, Peter Stone, Raymond Mooney, and Jude Shavlik. Guiding a reinforcement learner with natural language advice: Initial results in RoboCup soccer. *AAAI Workshop - Technical Report*, 06 2004.
- [18] H. Lieberman. Your wish is my command: Programming by example. 01 2001.

- [19] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *ICLR*, 2016.
- [20] J. Lin, Z. Ma, R. Gomez, K. Nakamura, B. He, and G. Li. A review on interactive reinforcement learning from human social feedback. *IEEE Access*, 8:120757–120765, 2020. doi: 10.1109/ACCESS.2020.3006254.
- [21] James MacGlashan, Mark K Ho, Robert Loftin, Bei Peng, David Roberts, Matthew E. Taylor, and Michael L. Littman. Interactive learning from policy-dependent human feedback, 2017.
- [22] Richard Maclin and Jude W Shavlik. Creating advice-taking reinforcement learners. *Machine Learning*, 22(1):251–281, 1996.
- [23] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *NIPS*, 2013.
- [24] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Harley, T. P. Lillicrap, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. *ICML*, 2016.
- [25] Khanh Nguyen, Debadeepta Dey, Chris Brockett, and Bill Dolan. Vision-based navigation with language-based assistance via imitation learning with indirect intervention, 2019.
- [26] Rohit Prabhavalkar, Tara N. Sainath, Yonghui Wu, Patrick Nguyen, Zhifeng Chen, Chung-Cheng Chiu, and Anjali Kannan. Minimum word error rate training for attention-based sequence-to-sequence models, 2017.
- [27] Aditya Prakash, Aseem Behl, Eshed Ohn-Bar, Kashyap Chitta, and Andreas Geiger. Exploring data aggregation in policy learning for vision-based urban autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11763–11773, 2020.
- [28] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. 2017.
- [29] R. Sutton and A. G. Barto. *Reinforcement Learning*. The MIT Press, Cambridge, Massachusetts, 2nd edition, 2018.
- [30] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: On the importance of pre-training compact models, 2019.
- [31] M. L. Vergara. Accelerating training of deep reinforcement learning-based autonomous driving agents through comparative study of agent and environment designs, 2019.
- [32] Paul M Yanik, Joe Manganelli, Jessica Merino, Anthony L Threath, Johnell O Brooks, Keith Evan Green, and Ian D Walker. A gesture learning interface for simulated robot path shaping with a human teacher. *IEEE Transactions on Human-Machine Systems*, 44(1):41–54, 2013.