

Learn-to-Race: A Multimodal Control Environment for Autonomous Racing

James Herman^{1*} Jonathan Francis^{1,2*} Siddha Ganju³ Bingqing Chen¹ Anirudh Koul⁴
Abhinav Gupta¹ Alexey Skabelkin⁵ Ivan Zhukov⁵ Max Kumskey⁵ Eric Nyberg¹

¹School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

²Human-Machine Collaboration, Bosch Research, Pittsburgh, PA, USA

³NVIDIA, Santa Clara, CA, USA

⁴Pinterest, San Francisco, CA, USA

⁵Autonomous Driving, Arrival, London, UK

{jamesher, jmf1, bingqinc, agupta6, ehnl}@cs.cmu.edu, {sganju1, akoul}@alumni.cmu.edu,
{skabelkin, zhukov, kumskey}@arrival.com

Abstract

Existing research on autonomous driving primarily focuses on urban driving, which is insufficient for characterising the complex driving behaviour underlying high-speed racing. At the same time, existing racing simulation frameworks struggle in capturing realism, with respect to visual rendering, vehicular dynamics, and task objectives, inhibiting the transfer of learning agents to real-world contexts. We introduce a new environment, where agents Learn-to-Race (L2R) in simulated competition-style racing, using multimodal information—from virtual cameras to a comprehensive array of inertial measurement sensors. Our environment, which includes a simulator and an interfacing training framework, accurately models vehicle dynamics and racing conditions. In this paper, we release the Arrival simulator for autonomous racing. Next, we propose the L2R task with challenging metrics, inspired by learning-to-drive challenges, Formula-style racing, and multimodal trajectory prediction for autonomous driving. Additionally, we provide the L2R framework suite, facilitating simulated racing on high-precision models of real-world tracks. Finally, we provide an official L2R task dataset of expert demonstrations, as well as a series of baseline experiments and reference implementations. We make all code available: <https://github.com/learn-to-race/l2r>.

1. Introduction

Progress in the field of autonomous driving relies on the existence of challenging tasks and well-defined evalua-




Figure 1: Learn-to-Race interfaces with a racing simulator, which features numerous real-world racetracks such as the Thruxton Circuit (top-left) and Las Vegas Motor Speedway (top-right). Simulated race cars (bottom) are empowered with learning agents, tasked with the challenge of learning to race for the fastest lap-times and best metrics.

tion metrics, which enable researchers to effectively assess and improve algorithms. Models developed in learning-to-drive settings continue to struggle with issues in sample-complexity, safety, and unseen generalisation, calling for

*Equal contribution.

more suitable benchmarks [9, 16, 28]. We hypothesise that high-fidelity simulation environments, together with well-defined metrics and evaluation procedures, are conducive to developing more sophisticated agents; and, in turn, such agents will be better-suited to real-world deployment.

Simulated autonomous racing exhibits task complexity on several factors: (i) agents must perform real-time decision making, requiring computationally-efficient policy updates as well as robustness to latency; (ii) agents must be able to deal with realistic vehicle and environmental dynamics (whereas agents in less-realistic environments have been able to achieve super-human performance); (iii) agents must leverage more informative intrinsic reward schemes that enable replication of human-like driving behaviour, e.g., trading off safety and performance; and (iv) agents must use offline demonstrations effectively, without overfitting, and must leverage interactions with the environment sample-efficiently. We highlight simulated racing (Figure 1) as an opportunity for developing learning strategies that are capable of meeting these stringent requirements.

In this work, we release the Arrival Autonomous Racing Simulator, which includes numerous interfaces for both simulated and real vehicle instrumentation. Furthermore, we introduce `Learn-to-Race (L2R)`, a multimodal and continuous control environment for training and evaluating autonomous racing agents. Through the `L2R` environment, we simulate competition-style racetracks that are based off real-world counterparts, we provide mechanisms for fully-characterising realistic racing agents (e.g., flexible sensor placements, multimodal cameras, and various vehicle dynamics profiles), and we provide numerous tools for fine-grained agent evaluation (e.g., random and fixed spawn locations, custom racing map construction, and injection of external disturbances). Using these facilities, we enable research in problems that require agents to make safety-critical, sub-second decisions in dynamically unstable contexts, such as autonomous racing, real-time uncertainty analysis in highway driving, and trajectory forecasting. In this paper, we exemplify algorithm development and benchmarking of methods under learning from demonstrations, reinforcement learning, and model-predictive control.

Concretely, our contributions include: (i) the Arrival Autonomous Racing Simulator, which simulates high-fidelity competition-style tracks, vehicles, and various sensor signals; (ii) the `Learn-to-Race (L2R)` framework, a plug-and-play environment, which defines interfaces for various sensor modalities and provides an OpenAI-gym compliant training and testing environment for learning-based agents; (iii) an official `L2R` task and dataset with expert demonstrations, metrics, and reference evaluation procedures; and (iv) an academic release of the simulator, code for the `L2R` framework, and implementations of baseline agents to facilitate full reproducibility and extension.

2. Related Work

2.1. Reinforcement Learning Environments

Research progress in the fields of Reinforcement Learning (RL), Planning, and Control has relied on various simulation environments, for benchmarking agent performance on game-playing and robot control tasks [32, 8, 33, 18, 30]. These tasks require sequential decision-making in order to complete objectives and are generally characterised by their state dimensionality, the nature of their action space (e.g., discrete or continuous), agent cardinality (i.e., single- or multi-agent), and by the capability of the underlying simulator in capturing real-world physical dynamics [18]. Whereas the vast majority of tasks offered by, e.g., the DeepMind Control Suite [32], OpenAI Gym [8], and the MuJoCo physics engine [33] have been solved—with agents often achieving superhuman performance—no existing environments focus on high-fidelity simulation of high-speed driving, in dynamically unstable contexts.

2.2. Simulation of Autonomous Driving

Urban driving. CARLA [14] is an open-source simulator for autonomous driving, wherein various tasks have been defined to challenge agents’ street-legal urban driving abilities. Duckietown [29, 10] provides a customisable platform for urban autonomous driving, as well as hardware support for miniature vehicles controlled via RaspberryPi’s. In this paper, we focus primarily on autonomous racing environments, which present challenges outside the conventional scope of urban and highway driving.

Track racing. In autonomous racing, agents must make sub-second decisions in regimes of unstable physical dynamics, wherein the ramifications of control actions can be amplified or suppressed, non-linearly, depending on vehicular and environmental state. `CarRacing-v0`, an OpenAI-gym environment [8], is a simple racing environment, which uses only bird’s-eye-view (BEV) observations. In [17], researchers trained agents to race in the video game `Gran Turismo Sport`, but have not yet released their environment. Moreover, instead of using sensory perception, agents were directly provided with privileged information, e.g., distance to obstacles and road boundaries. `TORCS` [3] is an open-source simulator and is used by the Simulated Car Racing Championship [27], despite its game-like qualities. As the goal of simulators should be to accurately model the dynamics of the real-world, we assert that the potential for model transfer from these frameworks remains limited.

2.3. Learning Paradigms

We discuss various learning paradigms that are enabled by the simulation of autonomous driving.

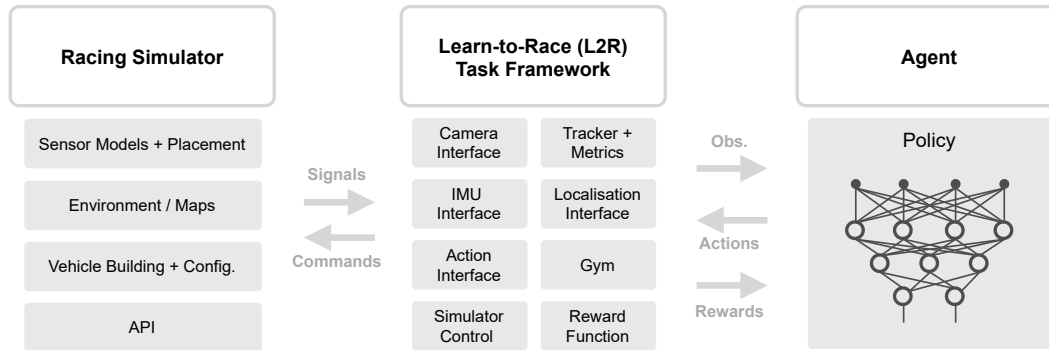


Figure 2: Learn-to-Race allows agents to interact with the racing simulator through a series of interfaces for observations, actions, and simulator control.

Simulation-to-real transfer. DeepRacer [5], developed by Amazon Web Services, provides an end-to-end framework for training and deploying 1/18th-scale autonomous racing cars. The Indy Autonomous Challenge [1] encourages institutions to create autonomous vehicle technology; participants are given the proprietary VRXPERIENCE driving simulator, which focuses more on optimising human-machine interactions within the vehicle, in the context of situational highway driving, which contrasts with our focus in this work on autonomous racing. Roborace [2] is the first international championship for full-size, real-world autonomous racing. Here, teams develop self-driving software and compete in challenges, using Roborace-owned vehicles. Roborace provides teams with proprietary software-in-the-loop (SIL) and hardware-in-the-loop (HIL) simulators, with a base driving stack. These simulators are predominately used for developing classical control methods, however, and do not include facilities for training learning-based agents [7, 22, 23, 31]. The authors obtained a podium finish at the Thruxton Circuit (UK), in the Season Beta Roborace competitions (2020-2021), and we now wish to enable new technologies through open-sourcing our autonomous racing research: to our knowledge, we publicly release the first environment that is specifically intended for simulating autonomous competition-style track racing and for transferring learning-based agents to the real world.

Safe and efficient learning. Imposing safety constraints in, e.g., RL algorithms, has become popular for the potential of reducing failures in simulation-to-real transfer settings and for enabling agent robustness to environmental stochasticity [20]. The goal is to embed safety guarantees in policies, without compromising their performance or sample-efficiency. While a few works consider detection and avoidance of unsafe states, in urban driving [13] and human-assistive robotics [19], no existing works focus on safe learning and control, for autonomous racing in dynamically unstable contexts. Popular Safe-RL benchmarks (e.g.,

OpenAI Safety Gym) lack realistic dynamics and they evaluate agents at much lower speeds; thus, the numerous limitations of existing methods cannot be studied comprehensively. We assert that the physical realism that L2R provides facilitates improvement of those underlying approaches.

3. Simulation Environment

3.1. Arrival Autonomous Racing Simulator

The Arrival simulator is a powerful tool for the development and testing of autonomous vehicles. It is based on Unreal Engine 4 and includes such features as: (i) a vehicle prototyping framework; (ii) full software-in-the-loop (SIL) simulation, to model all vehicle control devices; (iii) controller area network (CAN) bus interface; (iv) camera, inertial measurement unit (IMU), light detection and ranging (LiDAR), ultrasonic, and radar sensor models; (v) semantic segmentation; (vi) sensor placement and configuration facilities; (vii) V2V/V2I interface subsystem; (viii) dynamic racing scenario creation; (ix) race track generation from scanned datasets; (x) support for full integration with the CARLA simulator [14]; and (xi) an application programming interface (API), which is automatically generated based on C++ code analysis. Details in supplementary.

3.2. Learn-to-Race Environment

Learn-to-Race (L2R) is a multimodal control environment that provides a series of interfaces for an agent to interact with the racing simulator, including the capabilities to send control commands and make observations of the environment and vehicle state via different sensors (Figure 2). L2R is implemented as a Gym environment [8], enabling quick prototyping of control policies. While we release the L2R environment and task (Section 4) alongside the Arrival simulator, we note that other simulators may be used with L2R, including those provided by [2].

Agent-Simulator Interaction. At each step t , an agent selects an action a_t , based on its current observation s_t , using its policy π_θ : $a_t \sim \pi_\theta(\cdot|s_t)$. The control action from the agent is forwarded to the simulator as a UDP message. L2R receives updates from the simulator, i.e., images from the virtual camera and/or measurements from other vehicle sensors, through TCP and UDP socket connections. As in reality, update frequencies across the various sensor modalities are not equal, so L2R synchronizes observations by providing agents with the most recent data from each (Algorithm 1). The `step` method of the environment returns the new observation s_{t+1} , along with a calculated reward to the agent, $r_t = R(s_t, a_t, s_{t+1})$, and a Boolean terminal state flag. The reward function and evaluation metrics are defined in Section 4.3.

Algorithm 1 Agent-Simulator Interaction

```

1: function SENSOR THREAD
2:    $data \leftarrow$  Initial value
3:   function GET DATA
4:     return  $data$ 
5:   while not terminated do
6:      $data \leftarrow$  Receive Data
7:
8:   function STEP( $a_t$ )
9:     Send  $a_t$  as UDP message
10:     $s_{t+1} \leftarrow$  Get Data  $\forall$  Sensor Threads
11:     $r_t \leftarrow R(s_t, a_t, s_{t+1})$ 
12:     $done \leftarrow$  IsTerminal( $s_t, s_{t+1}$ )
13:    return  $s_{t+1}, r_t, done$ 

```

Episodic control. The control interface communicates with the simulator to automatically setup and execute simulations in an episodic manner. L2R conveniently allows for training to be launched in one command, as all aspects of the racing simulator and learning environment are parameterised. A state is considered terminal if all laps are successfully completed, if at least 2 of the vehicle’s wheels go outside of the drivable area, or if progress is minimally insufficient. The episode begins by resetting the vehicle to a standing start position, at a parameterised location along with configured sensor interfaces and initialised reward function. Discrete steps are taken by the agent until one of the aforementioned episode termination criteria is met.

4. Task: Learn-to-Race

The Learn-to-Race (L2R) task tests an agent’s ability to execute the requisite behaviours for competition-style track racing, through multimodal perceptual input. In this section, we provide a task overview and describe task properties, dataset characteristics, and metrics.

4.1. Task Overview

L2R provides an OpenAI Gym [8] compliant learning environment, where researchers could flexibly select among the available sensor modalities. This early version of the environment enables single-agent racing on three racetracks (with custom track construction facility), modeled after their real-world counterparts. Included tracks are the Thruxton Circuit (`Track01:Thruxton`) and Anglesey National Circuit (`Track02:Anglesey`) from the United Kingdom, and the North Road Track at Las Vegas Motor Speedway (`Track03:Vegas`), located in the United States. Analogous to having separate town maps for training and testing in other simulation environments, e.g., CARLA [14], we use `Track01` and `Track02` for training and `Track03` for testing. Consequently, we generate expert traces from the training tracks, for inclusion in our initial dataset release (see Section 4.2). Many avenues for research can be explored within L2R, including various learning paradigms, such as: (constrained) reinforcement learning, learning from demonstrations, multitask learning, transfer learning and domain adaptation, simulation-to-real transfer, fast decision-making, classical/neural hybrid modeling, etc. Regardless of the method chosen, agents’ multimodal perception capabilities—i.e., their ability to fuse and align sensory information—are of critical importance.

4.2. Learn-to-Race Dataset

We generate a rich, multimodal dataset of expert demonstrations from training tracks, `Track01:Thruxton` and `Track02:Anglesey`, in order to facilitate pre-training of agents via, e.g., imitation learning (IL). The L2R dataset contains multi-sensory input at a 100-millisecond resolution, in both the observation and action spaces. Depending on the selected simulator perception mode, agents have access to one (*vision-only mode*) or all modalities (*multimodal mode*). See Table 1 for a complete list of available modalities. The action space is defined by continuous values for acceleration and steering, in the ranges $[-1.0, 1.0]$, where negative acceleration values will decelerate the vehicle to a halted position. Note that *Gear* is a controllable action, but fixed to *drive* in all our experiments.

The expert demonstrations were collected using a model predictive controller (MPC) (Section 5), which follows the centerline of the race track at a pre-specified reference speed. This first version of the dataset contains 10,600 samples of each sensor and action dimension, for 9 complete laps around each track. Future version releases of L2R will include access to new simulated tracks (modeled after other real tracks, around the world) as well as expert traces generated from these additional tracks—across various weather conditions, in challenging multi-agent settings, and within dangerous obstacle-avoidance scenarios.

Table 1: Summary of the observation and continuous action spaces, for the `Learn-to-Race` task. When the simulator is initialised in *vision-only* mode, the observation space consists of just the images from the ego-vehicle’s front-facing camera. The additional observation data, all of which is realistically accessible on a real racing car, is available in *multimodal* mode. *Whereas gear is permitted as a controllable parameter, we do not use it in our experiments.

	Signal	Description	Dimension
Action	Acceleration	Command in [-1.0, 1.0]	\mathbb{R}^1
	Steering	Command in [-1.0, 1.0]	\mathbb{R}^1
	Gear	{park, drive, neutral, reverse}	—
Observation	Image	RGB image	$\mathbb{R}^{W \times H \times 3}$
	Steering	Observed steering direction	\mathbb{R}^1
	Gear	{park, drive, neutral, reverse}	—
	Mode	Vehicle mode	\mathbb{R}^1
	Velocity	In ENU coordinate (m/s)	\mathbb{R}^3
	Acceleration	In ENU coordinate (m/s^2)	\mathbb{R}^3
	Yaw, Pitch, Roll	Orientation of the car (rad)	\mathbb{R}^3
	Angular Velocity	Rate of change of the orientation (rad/s)	\mathbb{R}^3
	Location	Location of the vehicle center in ENU (m)	\mathbb{R}^3
	Wheel Rotational Speed	<i>per wheel</i> (RPM)	\mathbb{R}^4
	Braking	Brake pressure <i>per wheel</i> (Pa)	\mathbb{R}^4
	Torque	<i>per wheel</i> (Nm)	\mathbb{R}^4

4.3. Task Metrics

The primary objective of the L2R task is to minimise the time taken for an agent to successfully complete racing laps, with additional requirements on the agent’s driving quality. We do not restrict the agent’s learning paradigms to, e.g., IL or RL; on the contrary, we can envision a wealth of combination strategies and other methods that are applicable to the task. While we do not include *planning-only* approaches as those that are consistent with the official L2R task, (i) we do encourage hybrid or model-based learning approaches; furthermore, (ii) we do encourage the simulator and the L2R interface to be used to further research in these areas, more generally. Agnostic to the learning paradigm used, and inspired by concepts from high-speed driving, robot navigation [18], and trajectory forecasting [28], we define the core modalities, metrics, and objectives that shall be used to train L2R agents and assess their performance. We summarise agents’ action and observation spaces in Table 1 and the official L2R task metrics in Table 2.

We define the successful completion of an episode in the L2R task to be 3 completed laps, from a standing start; *Episode Completion Percentage* (ECP) measures the amount of the episode completed, and *Episode Duration* (ED) measures the minimum amount of time that the agent took to progress to its furthest extent, through the episode. We define *Average Adjusted Track Speed* (AATS) as a metric that measures the average speed of the agent, across all three laps of the episode. Metrics may also include adjust-

ments for environmental factors, such as wheel slippage and weather effects as the task matures. *Average Displacement Error* (ADE), a common metric in trajectory forecasting [28], measures the agent’s average deviation from a reference path—in this case, the centerline of the track. *Trajectory Admissibility* (TrA) is the dimensionless metric α , defined in Equation 1, where t_e is the duration of the episode and t_u is the cumulative time spent driving unsafely with exactly one wheel outside of the drivable area.

$$\alpha = 1 - \sqrt{\frac{t_u}{t_e}} \quad (1)$$

We also utilise metrics that measure the smoothness of agent behaviour: *Trajectory Efficiency* (TrE) measures the ratio of track curvature to agent trajectory curvature, i.e., in terms of agent heading deviations; *Movement Smoothness* (MS) quantifies the smoothness of the agent’s acceleration profile, adjusted for gravity, using the negated log dimensionless jerk, η_{dj} in Equation 2, inspired by [6]:

$$\eta_{dj} = \ln \left(\frac{(t_2 - t_1)^3}{v_{peak}^2} \int_{t_1}^{t_2} \left| \frac{d^2v}{dt^2} \right|^2 dt \right) \quad (2)$$

Rather than restricting agents to predefined incentive policies, input dimensions, or even input modalities, L2R allows and encourages flexibility so that agents can learn to race effectively. The default reward function for L2R is inspired by [17]: this policy provides dense rewards for progression, consistent with the goal of minimising lap times, and negative rewards for going out-of-bounds.

4.4. Task Evaluation Procedure

Agent assessment is conducted through a leaderboard competition, with two distinct stages: (1) pre-evaluation and (2) evaluation. Predicated on industry standards, we adopt a racing-centric pre-evaluation step, for assessing agent performance, giving agents a warm start on the test track before formal evaluation. Much like how human racing drivers are permitted to acquaint themselves with a new racing track, before competition, we run a pre-evaluation on models, with unfrozen weights, allowing for some initial (albeit constrained) exploration. In this pre-evaluation period, agents may explore the environment for a fixed time of 60 minutes, defined in the number of time-steps of discrete observation from the L2R framework. In the pre-evaluation, we further define a “competency check” that agents must pass, in order to successfully proceed through to the main evaluation phase. For the North Road track at Las Vegas Motor Speedway, the only competency check is that agents are able to successfully complete a lap during the pre-evaluation period with acceleration capped at 50% of maximum allowed in the action space. A successful episode is defined the completion of 3 laps from a standing start and the agent not going out of the driveable areas of the track. If the agent is unsuccessful in the pre-evaluation phase, it is disqualified and not evaluated further. As we continue to provide support for new tracks (necessitating more novel driving maneuvers), we will also continue to add and permute the driving competency checks, to maintain fairness of evaluation on those tracks.

Post a successful pre-evaluation stage, the final test stage occurs: agents are provided all the various input modalities and have to compete on the metrics defined Section 4.3. When the agent successfully passes through the pre-evaluation stage, the user is not provided with the results of the competency checks and instead is able to view the results of the complete evaluation directly on the leaderboard.

5. Baseline Agents

We define a series of learning-free (e.g., RANDOM, MPC) and learning-based (e.g., reinforcement learning, imitation learning) baselines, to illustrate the performance of various algorithmic classes on the L2R task. We also benchmark human performance, through a series of expert trials.

Random. The RANDOM agent is mainly intended as a simple demonstration of how to interface with the L2R environment. The RANDOM agent is spawned at the start of the track, and uniformly samples actions, i.e., steering and acceleration, from the action space. The agent then proceeds to execute these random actions.

MPC. The MPC was used to generate expert demonstrations (Section 4.2) and is intended as a reference solution

of L2R via classical control approaches. The MPC minimizes the tracking error with respect to the centerline of the race track at a pre-specified reference speed. We use the iterative linear quadratic regulator (iLQR) proposed in [26], which iteratively linearizes the non-linear dynamics along the current estimate of trajectory, solves a linear quadratic regulator problem based on the linearized dynamics, and repeats the process until convergence. Specifically, we used the implementation for iLQR from [4]. We adopt the kinematic bike model [25] to characterize the vehicle dynamics. Further MPC details are provided in the supplementary.

While MPC implies optimal control performance, we want to point out the limitations of our current implementation. Firstly, the ground truth vehicle parameters were not known to us and we used estimated values. Secondly, we asked the MPC to follow the centerline of the track, which is not the trajectory expert drivers would have taken, especially when cornering. Finally, we pre-specified the MPC to drive at a conservative speed (12.5m/s), which makes the expert demonstrations easier to learn from.

Conditional Imitation Learning. We adopted the same neural architecture from [11], except that we do not have different commands in our case, e.g., turn left, turn right, go straight, and stop. Thus, we used a single branch for decoding actions. We assume both front view images and sensor measurements are available for the IL agent. In each sample, the input consists of a 512×384 image and 30 sensor measurements, and output is 2 actions (as listed in Table 1). The implementation of CIL automatically adjusts the neural network architecture based on specified input-output dimensions. The imitation loss (Equation 3) is the mean squared error between the predicted action, \hat{a}_t , and the action taken by the expert, a_t .

$$\mathcal{L} = \sum_{i=1}^n \|\hat{a}_i - a_i\|_2^2 \quad (3)$$

Soft Actor-Critic. We provide a reference implementation of Soft-Actor Critic (SAC) [12, 21], which is generally performant and known to be robust [15]. SAC belongs to the family of maximum entropy reinforcement learning (RL) algorithms, wherein an agent maximizes expected return, subject to an entropy regularization term (Equation 4), as a principled way to trade-off exploration and exploitation.

$$\mathcal{J}(\theta) = \sum_{t=1}^T \mathbb{E}_{\pi_{\theta}} [R(s_t, a_t) - \mathcal{H}(\pi_{\theta}(a_t|s_t))] \quad (4)$$

Our RL-SAC agent demonstrates several of features in the environment: it operates in vision-only mode, but rather than learning directly from pixels, we pre-trained a convolutional, variational auto-encoder [24] made on sample camera images. Therefore, our agent only need to learn to decode actions from image embeddings using a multi-layer

Table 2: `Learn-to-Race` defines multiple metrics for the assessment of agent performance. These metrics measure overall success—e.g., whether and how fast the task is completed—along with more specific properties, such as trajectory admissibility and smoothness.

Metric	Definition
<i>Episode Completion Percentage</i>	Percentage of the 3-lap episode completed
<i>Episode Duration</i>	Duration of the episode, in seconds
<i>Average Adjusted Track Speed</i>	Average speed, across all three laps, adjusted for environmental conditions, in km/h
<i>Average Displacement Error</i>	Euclidean displacement from (unobserved) track centerline, in meters
<i>Trajectory Admissibility</i>	Complement of the square root of the proportion of cumulative time spent unsafe
<i>Trajectory Efficiency</i>	Ratio of track curvature to trajectory curvature (i.e., in agent heading)
<i>Movement Smoothness</i>	Log dimensionless jerk based on accelerometer data, adjusted for gravity

Table 3: Baseline agent results on `Learn-to-Race` task while training on Thruxton track, with respect to the task metrics in Table 2: Episode Completion Percentage (**ECP**), Episode Duration (**ED**), Average Adjusted Track Speed (**AATS**), Average Displacement Error (**ADE**), Trajectory Admissibility (**TrA**), Trajectory Efficiency (**TrE**), and Movement Smoothness (**MS**). Arrows (\uparrow / \downarrow) indicate directions of better performance. Asterisks (*) in Tables 3 and 4 indicate metrics which may be misleading, for incomplete racing episodes.

Agent	ECP (\uparrow)	ED (\downarrow)	AATS (\uparrow)	ADE (\downarrow)	TrA (\uparrow)	TrE (\uparrow)	MS (\uparrow)
HUMAN	100.0(\pm 0.0)	235.8(\pm 1.7)	171.2(\pm 3.4)	2.4(\pm 0.1)	0.93(\pm 0.01)	1.00(\pm 0.02)	11.7(\pm 0.1)
RANDOM	0.5(\pm 0.3)	14.0(\pm 5.5)	11.9(\pm 3.8)	1.5(\pm 0.6)	0.81(\pm 0.04)	0.33(\pm 0.38)*	6.7(\pm 1.1)
MPC	100.0(\pm 0.0)	904.2(\pm 0.7)	45.1(\pm 0.0)	0.9(\pm 0.1)	0.98(\pm 0.01)	0.85(\pm 0.03)	10.4(\pm 0.6)
RL-SAC	31.1(\pm 0.0)	251.2(\pm 1.4)	50.5(\pm 0.3)	0.5(\pm 0.0)	0.97(\pm 0.0)	0.48(\pm 0.0)*	11.1(\pm 0.4)

Table 4: Baseline agent results on `Learn-to-Race` task while testing on Las Vegas track.

Agent	ECP (\uparrow)	ED (\downarrow)	AATS (\uparrow)	ADE (\downarrow)	TrA (\uparrow)	TrE (\uparrow)	MS (\uparrow)
HUMAN	100.0(\pm 0.0)	176.2(\pm 3.4)	114.2(\pm 2.3)	1.7(\pm 0.1)	0.88(\pm 0.01)	1.09(\pm 0.02)	10.1(\pm 0.3)
RANDOM	1.0(\pm 0.6)	21.9(\pm 9.6)	9.2(\pm 1.5)	1.4(\pm 0.3)	0.74(\pm 0.01)	0.18(\pm 0.05)*	8.4(\pm 1.0)
MPC	69.5(\pm 10.7)	353.2(\pm 54.8)	40.5(\pm 0.1)	0.8(\pm 0.1)	0.91(\pm 0.02)	1.07(\pm 0.01)*	10.4(\pm 0.2)
RL-SAC	11.8(\pm 0.1)	109.9(\pm 7.5)	22.1(\pm 1.5)	1.3(\pm 0.1)	0.95(\pm 0.01)	0.58(\pm 0.01)*	9.9(\pm 0.2)

perceptron with two hidden layers of 64 hidden units each. Our agent’s reward function was the environment’s default with the inclusion of a bonus if the agent remained near the center of the track.

Human. We additionally establish a HUMAN performance baseline, by collecting simulated racing results from human expert players. The collection procedure involved a private crowd-sourcing event, which was split into two separate phases—practice/training and recording/testing. Expert players were already familiar with the simulator, task, and objective, prior to engaging in the event. In the training phase, players were instructed to engage in the race, until the variance in finished lap-times, for three consecutive runs, fell below a certain threshold. After this training phase, players were allowed to proceed to the testing phase, for which their top-3 laps were recorded. We averaged the top-3 results in the testing phase, from all experts, for each track; the training results were discarded.

6. Experiments and Results

We evaluate each of the baseline agents—HUMAN, RANDOM, MPC, and RL-SAC—on the L2R task, with the objective of finishing 3 consecutive laps in minimal time. For all approaches, agents complete model training and tuning on `Track01:Thruxton`. We present the average of each metric across 3 consecutive episodes, in Table 3. Afterwards, agents are evaluated based on their performance on `Track03:Vegas`, following the 1-hour pre-evaluation period described in Section 4.4. Learning-free agents, RANDOM and MPC, simply perform inference in the testing environment. The RL-SAC agent, a learning-based approach, operates in vision-only mode and utilizes the pre-evaluation stage to perform simple transfer learning to the new racetrack. The agent’s image encoder does not have access to the test track prior to pre-evaluation and is not updated during this phase, but the model weights of the agent do update as new experience becomes available. Following

the pre-evaluation phase, agents completed 3 consecutive episodes, and we present metric averages in Table 4.

Human experts. Human experts strongly outperform during both training and testing, suggesting a general understanding of racing: they can quickly adapt to a new track, despite different features, including frequent and severe turns. Human experts fully complete 3 lap episodes at speeds near the vehicle’s physical limits and estimate their lap-time performance to be within 10% of optimal. We expect strong agents to execute trajectories which are of lower curvature than the racetrack’s centerline, or a **TrE** of at least 1.0, allowing the vehicle to maintain higher **AATS**. Only human experts were able to achieve this, considering that failure to complete an episode distorts the metric. However, such trajectories are aggressive and risky, because they often involve cutting corners with the wheels nearly outside of the driveable area; this is apparent by higher **ADE** values and relatively low **TrA**. Additionally, human experts performed well relative to other agents terms of **MS**, demonstrating the ability to anticipate the need for acceleration and to apply smooth control.

Baseline agents. There are several notable conclusions that we make based on the performance of our baseline agents which we do not claim to be state-of-the-art. The first is that the task is indeed challenging, as even the **MPC** agent with an approximate car model failed to consistently complete laps on the test track. Even after over 1 million steps environment steps on the training track, the **RL-SAC** agent only completes about 90% of a lap due to the challenging speed trap near the finish line at Thruxton. However, the **RL-SAC** agent demonstrates better control than the **MPC** in training in both **ADE** and **MS**. Second, we note the lack of generalization and poor sample efficiency of the **RL-SAC** agent whose performance dropped significantly in terms of ability to progress down the track, **ECP**, and stay near the centerline, **ADE**, despite being directly incentivised to do so. The agent learns to simply stop altogether to avoid going out-of-bounds about 1/3rd of the way around the test track. We note that imitation learning has potential for providing agents with strong priors. However, in our experiments, automatic network sizing based on input/output dimensions and step-wise supervision alone, suggested by [14], did not yield good performance. This demonstrates the challenge that **L2R** poses to this family of approaches, necessitating consideration of, e.g., joint **IL/RL** strategies.

7. Discussion

We are confident that agents can achieve superhuman performance for any given track given that (1) they are sufficiently complex and (2) that they have interacted with that environment enough times. What is not clear, is how well agents can generalize to new racetracks in a realistic simulation environment. We believe the **Learn-to-Race** task

will effectively assess models, based on their general understanding of vehicle dynamics, high-speed and high-risk control, racetrack perception, and intelligent racing tactics.

To challenge state-of-the-art learning approaches, which continue to demonstrate superhuman performance in simplistic environments, we believe that the direction of future tasks must be towards higher complexity and realism. Our racing simulator has been used as a primary modeling tool for autonomous agents which have demonstrated real-world racing speeds in excess of 200 km/h, an order of magnitude faster, and more complex, than comparable environments. Limitations of our simulation environment with respect to competing simulators include multi-agent racing and a (currently) limited supply of tracks—however, both multi-agent racing and additional tracks will follow. Future enhancements also include additional vehicle sensors, domain randomization, and support for distributed training in learning-based approaches. We believe these enhancements serve as a precursor to real-world transfer and safety learning.

8. Conclusion

We have presented: (i) a high-fidelity simulator for the development and testing of autonomous race cars, (ii) the **Learn-to-Race** environment which enables rapid prototyping, training, and testing in this simulated environment, and (iii) the **L2R** task which defined dataset characteristics and concrete driving-inspired metrics for evaluation. **L2R** addresses the lack of complex learning environments and introduces the challenging task of simulated, high-performance racing. While human experts have demonstrated strong results on this task, both using the **L2R** framework as well as in competition racing, learning agents have not. We have provided relevant racing metrics and baseline results for classical control, **RL**, and **IL** agents as well as human experts, and we are releasing reference implementations and model checkpoints to further advance the research. The **L2R** suite of tasks and metrics will continue to expand in the future including the introduction of multi-agent racing. We hope to someday see agents reach superhuman, real-world performance in autonomous racing.

Acknowledgements

There are many people that helped create this task environment. We are grateful to the **RoboRace** community, in particular, the **Hive** team which has developed the simulated racetrack maps. We thank developers Bohui Fang, Ignacio Maronna Musetti, Jikai Lu, Zihang Zhang, and Xinnan Du for their support. We also thank Maxim Integrated Products for supporting our efforts. This work was supported, in part, by a fellowship from Bosch Research Pittsburgh.

References

- [1] Indy autonomous challenge. <https://www.indyautonomouschallenge.com/>. Last accessed: 2021-01-30.
- [2] Roborace. <https://roborace.com/>. Last accessed: 2021-01-30.
- [3] Torcs, the open racing car simulator. <http://torcs.sourceforge.net/index.php?name=Sections&op=viewarticle&artid=19>. Last accessed: 2021-01-30.
- [4] Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J Zico Kolter. Differentiable mpc for end-to-end planning and control. *arXiv preprint arXiv:1810.13400*, 2018.
- [5] B. Balaji, S. Mallya, S. Genc, S. Gupta, L. Dirac, V. Khare, G. Roy, T. Sun, Y. Tao, B. Townsend, E. Calleja, S. Muralidhara, and D. Karuppasamy. Deep-racer: Autonomous racing platform for experimentation with sim2real reinforcement learning. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2746–2754, 2020.
- [6] S. Balasubramanian, A. Melendez-Calderon, and E. Burdet. A robust and sensitive metric for quantifying movement smoothness. *IEEE Transactions on Biomedical Engineering*, 59(8):2126–2136, 2012.
- [7] J. Betz, A. Wischnewski, A. Heilmeyer, F. Nobis, T. Stahl, L. Hermansdorfer, and M. Lienkamp. A software architecture for an autonomous racecar. In *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*, pages 1–6, 2019.
- [8] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- [9] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Conference on Robot Learning*, pages 66–75. PMLR, 2020.
- [10] M. Chevalier-Boisvert, F. Golemo, Y. Cao, B. Mehta, and L. Paull. Duckietown environments for openai gym. <https://github.com/duckietown/gym-duckietown>, 2018.
- [11] Felipe Codevilla, Matthias M’uller, Antonio Lopez, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4693–4700. IEEE, 2018.
- [12] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov. Openai baselines. <https://github.com/openai/baselines>, 2017.
- [13] Wenhao Ding, Baiming Chen, Bo Li, Kim Ji Eun, and Ding Zhao. Multimodal safety-critical scenarios generation for decision-making algorithms evaluation. *IEEE Robotics and Automation Letters*, 6(2):1551–1558, 2021.
- [14] 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.
- [15] Benjamin Eysenbach and Sergey Levine. Maximum entropy rl (provably) solves some robust rl problems. *arXiv preprint arXiv:2103.06257*, 2021.
- [16] Angelos Filos, Panagiotis Tigas, Rowan McAllister, Nicholas Rhinehart, Sergey Levine, and Yarin Gal. Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *International Conference on Machine Learning (ICML)*, 2020.
- [17] F. Florian, S. Yunlong, E. Kaufmann, D. Scaramuzza, and P. Duerr. Super-human performance in gran turismo sport using deep reinforcement learning, 2020.
- [18] Jonathan Francis, Nariaki Kitamura, Felix Labelle, Xiaopeng Lu, Ingrid Navarro, and Jean Oh. Core challenges in embodied vision-language planning. *arXiv preprint arXiv:2106.13948*, 2021.
- [19] David Fridovich-Keil, Andrea Bajcsy, Jaime F Fisac, Sylvia L Herbert, Steven Wang, Anca D Dragan, and Claire J Tomlin. Confidence-aware motion prediction for real-time collision avoidance1. *The International Journal of Robotics Research*, 39(2-3):250–265, 2020.
- [20] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [21] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1861–1870, Stockholm, Sweden, 10–15 Jul 2018. PMLR.
- [22] T. Herrmann, F. Passigato, J. Betz, and M. Lienkamp. Minimum race-time planning-strategy for an autonomous electric racecar. *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, Sep 2020.
- [23] T. Herrmann, A. Wischnewski, L. Hermansdorfer, J. Betz, and M. Lienkamp. Real-time adaptive velocity optimization for autonomous electric cars at the limits of handling. *IEEE Transactions on Intelligent Vehicles*, pages 1–1, 2020.

- [24] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [25] Jason Kong, Mark Pfeiffer, Georg Schildbach, and Francesco Borrelli. Kinematic and dynamic vehicle models for autonomous driving control design. In *2015 IEEE Intelligent Vehicles Symposium (IV)*, pages 1094–1099. IEEE, 2015.
- [26] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, pages 222–229. Citeseer, 2004.
- [27] D. Loiacono, L. Cardamone, and P. L. Lanzi. Simulated car racing championship: Competition software manual. *arXiv preprint arXiv:1304.1672*, 2013.
- [28] Seong Hyeon Park, Gyubok Lee, Manoj Bhat, Jimin Seo, Minseok Kang, Jonathan Francis, Ashwin R Jadhav, Paul Pu Liang, and Louis-Philippe Morency. Diverse and admissible trajectory forecasting through multimodal context understanding. In *European Conference on Computer Vision (ECCV)*, 2020.
- [29] L. Paull, J. Tani, H. Ahn, J. Alonso-Mora, L. Carlone, M. Cap, Y. F. Chen, C. Choi, J. Dusek, Y. Fang, D. Hoehener, S. Liu, M. Novitzky, I. F. Okuyama, J. Pazis, G. Rosman, V. Varricchio, H. Wang, D. Yershov, H. Zhao, M. Benjamin, C. Carr, M. Zuber, S. Karaman, E. Frazzoli, D. Del Vecchio, D. Rus, J. How, J. Leonard, and A. Censi. Duckietown: An open, inexpensive and flexible platform for autonomy education and research. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1497–1504, 2017.
- [30] K. Shao, Z. Tang, Y. Zhu, N. Li, and D. Zhao. A survey of deep reinforcement learning in video games, 2019.
- [31] T. Stahl, A. Wischnewski, J. Betz, and M. Lienkamp. Multilayer graph-based trajectory planning for race vehicles in dynamic scenarios. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 3149–3154, 2019.
- [32] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. de Las Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, T. Lillicrap, and M. Riedmiller. Deepmind control suite, 2018.
- [33] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.