

Simulation-Based Reinforcement Learning for Real-World Autonomous Driving

Błażej Osiński^{1,3}, Adam Jakubowski¹, Paweł Zięcina¹, Piotr Miłoś^{1,5},
Christopher Galias^{1,4}, Silviu Homoceanu² and Henryk Michalewski³

Abstract—We use reinforcement learning in simulation to obtain a driving system controlling a full-size real-world vehicle. The driving policy takes RGB images from a single camera and their semantic segmentation as input. We use mostly synthetic data, with labelled real-world data appearing only in the training of the segmentation network.

Using reinforcement learning in simulation and synthetic data is motivated by lowering costs and engineering effort.

In real-world experiments we confirm that we achieved successful sim-to-real policy transfer. Based on the extensive evaluation, we analyze how design decisions about perception, control, and training impact the real-world performance.

I. Introduction

This work focuses on verifying whether it is possible to obtain a driving system using data from a simulator, which can be deployed on a real car. Using synthetic data, in comparison to collecting it in the real world, reduces the cost of developing such a system. Our policies were trained using reinforcement learning (RL) and confirmed to be useful in driving a real, full-sized passenger vehicle with state-of-the-art equipment required for Level 4 autonomy. The real-world tests consisted of 9 driving scenarios of total length of about 2.5 km.

The driving policy is evaluated by its real-world performance on multiple scenarios outlined in Section III–a. To complete a scenario, the driving agent needs to execute from 250 to 700 actions at 10 Hz at speeds varying from 15 to 30 km/h. In some of our experiments, the learned controller outputs the steering command directly. In other, the controller outputs waypoints which are transformed to steering commands using a proprietary control system. In this work, we decided to limit intermediate human-designed or learned representations of the real world only to semantic segmentation. The semantic segmentator used in our system is the only component trained using the real-world data – its training process mixes real-world and synthetic images. Our driving policies are trained only in simulation and directly on visual inputs, understood as RGB images along with their segmentation masks. The input contains also selected car metrics and a high-level navigation command inspired by [8].

Using reinforcement learning and RGB inputs was a deliberate design decision. The goal behind this choice was to answer the following research questions: Is it possible to

train a driving policy in an end-to-end fashion? Is such a policy obtained in simulation capable of real-world driving?

Using simulation and reinforcement learning is instrumental in generating rich experience. The former enables avoiding hard safety constraints of real-world scenarios and lets the latter to explore beyond the imagination of hand-designed situations. Furthermore, RL enables end-to-end training with none or little human intervention. This property is highly desirable as it greatly reduces the engineering effort. It may also eliminate errors arising when gluing a heterogeneous system consisting of separate perception and control modules.

The quality of a simulator is crucial to obtain transferable policies. We use CARLA [13], which offers reasonable fidelity of physical and visual layers. This fidelity comes at a high computation cost, however. In our case, it was alleviated by implementing a parallelized training architecture inspired by IMPALA [14]. With our current infrastructure, we gathered as much as 100 years of simulated driving experience. This enabled us to conduct several lines of experiments testing various design choices. These experiments constitute the main contribution of this work:

1. In simulation: we verify the influence of visual randomizations on transfer between different scenarios in simulation; results are summarized in Section IV-A.

2. In the real world: we test 10 models listed in Table I on 9 driving scenarios. In total we report results gathered over more than 400 test drives. See Section IV-B for a detailed description.

model	description
CONTINUOUS-PLAIN	base experiment
CONTINUOUS-LOW-RAND	experiment using smaller number of randomizations
DISCRETE-PLAIN	model using discrete actions
CONTINUOUS-REG	experiment with additional l_2 regularization
DISCRETE-REG	analog of DISCRETE-PLAIN with additional l_2 regularization
SEMSEG-ONLY	model with semantic segmentation as only visual input
WAYPOINTS-DISCRETE	model driving with waypoints
AUXILIARY-DEPTH	model predicting depth as auxiliary task
DYNAMICS-RAND-FFW	feed-forward model trained with dynamics randomizations
DYNAMICS-RAND-RNN	model with memory trained with dynamics randomizations

TABLE I: Summary of models evaluated in this work.

In Section IV-C we describe two failure cases and in Section IV-D we assess a proxy metric potentially useful for offline evaluation of models. We provide recordings from 9 autonomous test drives at <https://bit.ly/2k8syvh>. The test drives correspond to the scenarios listed in Figures 1.

II. Related work

a) *Synthetic data and real-world robotics:* Synthetic images were used in the ALVINN experiment [30]. [35]

¹ deepsense.ai, Warsaw, Poland

² Volkswagen AG, Wolfsburg, Germany

³ University of Warsaw, Warsaw, Poland

⁴ Jagiellonian University, Cracow, Poland

⁵ Institute of Mathematics of the Polish Academy of Sciences, Warsaw, Poland

proposed a training procedure for drones and [18], [29], [28], [44], [27] proposed experiments with robotic manipulators where training was performed using only synthetic data. Progressive nets and data generated using simulator were used in [33] to learn policies in the domain of real-world robot manipulation. A driving policy for a one-person vehicle was trained in [2]. The policy in [2] is reported to show good performance on a rural road and the training used mostly synthetic data generated by Unreal Engine 4. Our inclusion of segmentation as described in Section III–d is inspired by sim-to-real experiments presented in [19] and [25]. Visual steering systems inspired by [35] and trained using synthetic data were presented in [36], [34].

b) Synthetic data and simulated robotics: Emergence of high-quality general purpose physics engines such as MuJoCo [45], along with game engines such as Unreal Engine 4 and Unity, and their specialized extensions such as CARLA [13] or AirSim [40], allowed for creation of sophisticated photo-realistic environments which can be used for training and testing of autonomous vehicles. A deep RL framework for autonomous driving was proposed in [37] and tested using the racing car simulator TORCS. Reinforcement learning methods led to very good performance in simulated robotics – see, for example, solutions to complicated walking tasks in [16], [23]. In the context of CARLA, impressive driving policies were trained using imitation learning [8], [31], affordance learning [38], reinforcement learning [4], and a combination of model-based and imitation learning methods proposed in [31]. However, as stated in [2]: “training and evaluating methods purely in simulation is often ‘doomed to succeed’ at the desired task in a simulated environment” and indeed, in our suite of experiments described in Section III most of the simulated tasks can be relatively easily solved, in particular when a given environment is deterministic and simulated observations are not perturbed.

c) Reinforcement learning and real-world robotics: A survey of various applications of RL in robotics can be found in [9, Section 2.5]. The role of simulators and RL in robotics is discussed in [41, Section IV]. In [35], [29], [28], [44], [27], [19], [2], [33] policies are deployed on real-world robots and training is performed mostly using data generated by simulators. [21] proposes a system with dynamics trained using real-world data and perception trained using synthetic data. Training of an RL policy in the TORCS engine with a real-world deployment is presented in [42].

III. Environment and learning algorithm

We use CARLA [13], an open-source simulator for autonomous driving research based on Unreal Engine 4. CARLA features open assets, including seven built-in maps, 14 predefined weather settings, and multiple vehicles with different physical parameters. In our experiments we use input from simulated cameras; their settings, including position, orientation, and field of view, can be customized. Two visual quality levels (LOW and EPIC) are supported; the latter implements visual features including shadows, water reflections, sun flare effects, and antialiasing.

Below we describe our experimental setup as used in the basic CONTINUOUS–PLAIN experiment. We varied its various elements in other experiments. Details are provided in Section IV-B.

a) Simulated and real-world scenarios: The models were tested on 9 real-world scenarios presented in Figure 1. These scenarios contain diverse driving situations, including turns and the entry and exit of an overpass. For training, we developed new CARLA-compatible maps resembling the real-world testing area (approximately 50% of the testing scenarios were covered). We used these maps along with maps provided in CARLA for training, with some scenarios reserved for validation. In all scenarios the agent’s goal is to follow a route from start to finish. These routes are lists of checkpoints: they are generated procedurally in CARLA maps and predefined in maps developed by us.

In training, agents are expected to drive in their own lanes, but other traffic rules are ignored. Moreover, we assume that the simulated environment is static, without any moving cars or pedestrians, hence a number of a safety driver interventions during test deployments in real traffic is unavoidable.

b) Rewards in simulation and metrics of real-world performance: In simulation, the agent is rewarded for following a reference trajectory, which provides a dense training signal. The episode fails if the agent diverges from the trajectory more than 5 meters or collides with an obstacle. In the real world, for each scenario, we measure the percentage of distance driven autonomously (i.e. without human intervention); results are presented in Figure 3. Since tests were made in an uncontrolled environment with other vehicles and pedestrians, the safety driver was instructed to take over in all situations which were potentially risky. We also measure divergence from expert trajectories (see Figure 4).

c) Actions: Vehicles are controlled by two values: throttle and steering. The throttle is controlled by a PID controller with speed set to a constant, and thus our neural network policies only control the steering. We explore various possibilities for action spaces. Unless stated otherwise, in training the policy is modeled as a Gaussian distribution over the angle of the steering wheel. In evaluations we use the mean of the distribution.

d) Semantic segmentation: The semantic segmentation model is trained in a supervised way separately from the reinforcement learning loop. We used the U-Net [32] architecture and synthetic data from CARLA (which can render both RGB images and their ground-truth semantic labels), the Mapillary dataset [26], as well as real-world labeled data from an environment similar to the one used in test drives. The output of the model is further simplified to include only the classes most relevant to our problem: road, road marking, and everything else (i.e. obstacles).

e) Observations: The observation provided to the agent consists of visual input and two car metrics (speed and acceleration). In order to disambiguate certain road situations (e.g. intersections), a high-level navigation command is also given: lane follow, turn right/left or go straight. The commands resemble what could be provided



Fig. 1: All real-world scenarios used in our experiments. Left map: (a) autouni-arc, (b) autouni-straight. Center map: (c) factory_city-overpass*, (d) factory_city-overpass_exit. Right map: (e) factory_city-tunnel-bt10*, (f) factory_city-bt10-u_turn, (g) factory_city-u_turn-sud_strasse, (h) factory_city-sud_strasse_u_turn*, (i) factory_city-u_turn-bt10*. Scenarios marked with * were used for training in simulation.

by a typical GPS-based navigation system.

The visual input is based on an RGB image from a single front camera which is downscaled to 134×84 pixels. The camera position and orientation in simulation was configured to reflect the real-world setup. The RGB observation is concatenated with its semantic segmentation as described in the previous subsection. Including this component was motivated by [25], which claims that it “contains sufficient information for following the road and taking turns, but it is abstract enough to support transfer”. We have further evaluated this claim in experiment SEMSEG-ONLY.

f) Domain randomizations: Randomizations are considered to be pivotal to achieve sim-to-real transfer (and robust policies in general; see e.g. [27]). In our experiments we used the following list of visual randomizations: 10 weather settings (we used CARLA weather presets, which affect only the visual features of the environment), the simulation quality (we used both LOW and EPIC), camera input randomizations (we used a set of visual augmentations, such as adding gaussian noise, varying brightness, and applying blur or cutout [10]). We recall that our policies are trained on multiple scenarios and different maps, which is also aimed to increase robustness. Unless stated otherwise these randomizations are used in all experiments. In a separate experiment we also evaluated using randomization of dynamics, see description in Section IV-B-e.

g) Network architecture: The RL policy is implemented using a neural network; see its simplified architecture in Figure 2. As the feature extractor for visual input (RGB and semantic segmentation) we use the network from [14]. Our choice was influenced by [6], where this network was shown to generalize well to different RL environments. Note that policy transfer between simulation and reality can be seen as a generalization challenge.

h) Learning algorithm: We used OpenAI Baselines [11] ppo2 (see website <http://bit.ly/34xh7z4> for training hyperparameters). We typically use 4 simultaneous PPO trainers, which share gradients via Horovod [39]. Each

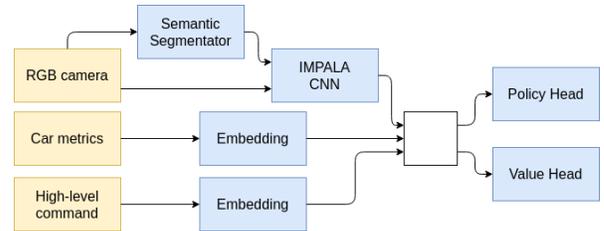


Fig. 2: Network architecture.

of them uses 2 Tesla K40 GPUs; one for running 10 Carla instances and the other for the optimisation of the policy network. In practice, this setup is able to gather $1.2 \cdot 10^7$ frames a day (equivalent to 13 days of driving).

Thanks to dense rewards the training in simulation was quite stable across models and hyperparameters. For deployments we have decided to use 1-4 models per experiment type, each trained using roughly 10^8 frames (equivalent to about 115 days of driving).

i) System identification: Inspired by the importance of system identification for sim-to-real transfer demonstrated in [43], we configured the CARLA simulator to mimic some values measured in the car used for deployment. These were the maximal steering angle and the time for a steering command to take effect (i.e. its delay).

IV. Experiments

Our models have been evaluated both in simulation and in reality, with much more focus on the latter. Below we provide detailed description of experiments conducted in both domains.

A. Experiment in simulation

In this experiment, we measure in simulation how randomizations affect performance. To this end, we apply fewer randomizations than the set used throughout all other experiments. Precisely, we used only one weather setting, trained only on the LOW quality settings and did not augment

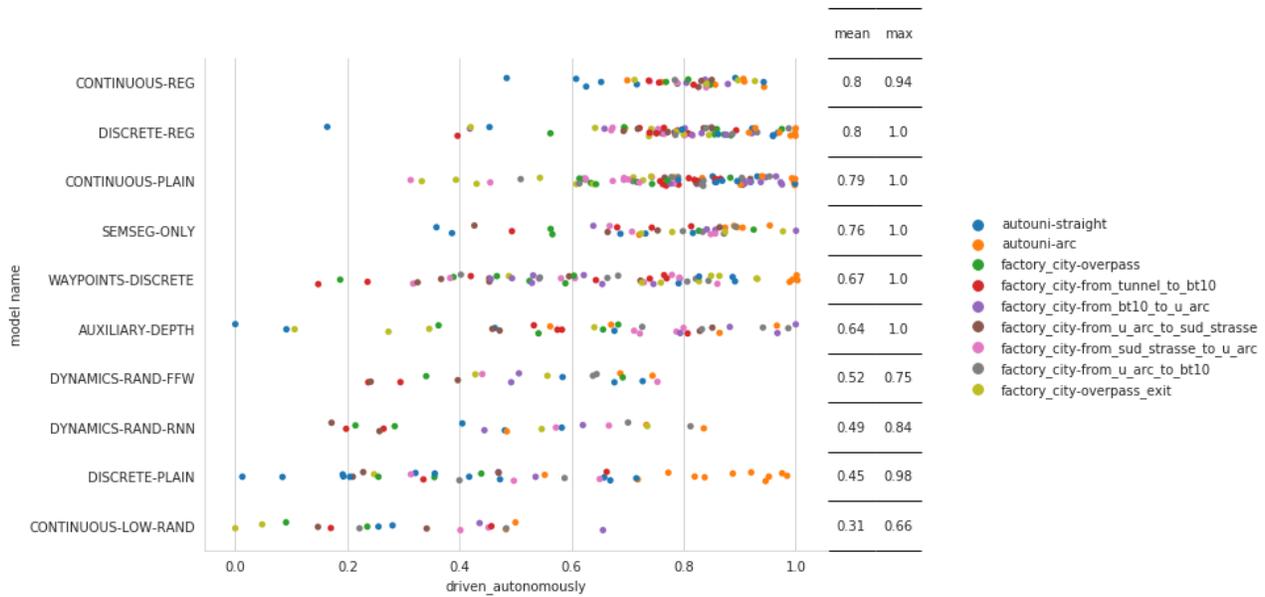


Fig. 3: Summary of experiments with baselines across nine real-world scenarios. The columns to the right show the mean and max of autonomy (the percentage of distance driven autonomously). Models are sorted according to their mean performance. Print in color for better readability.

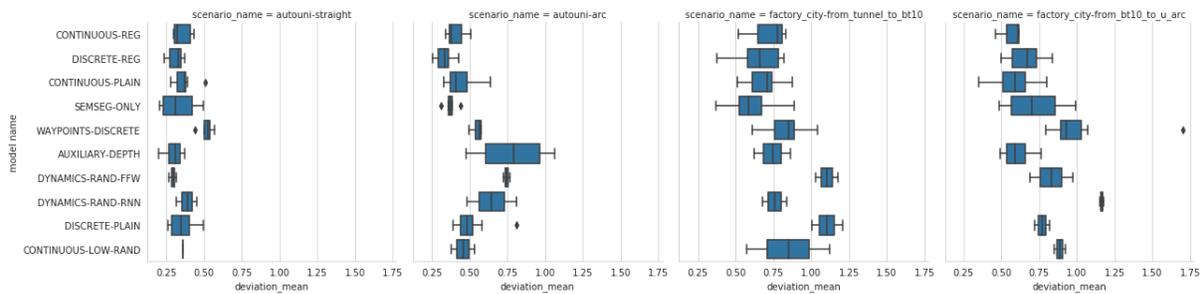


Fig. 4: Average deviation of models from expert trajectories. Measurements based on GPS. The graphs for all scenarios can be found on the website <http://bit.ly/34xh7z4>

the camera inputs. We conclude that the model trained with the standard set of randomizations generalizes better to the holdout town and with the holdout weather setting (see Figure 5).

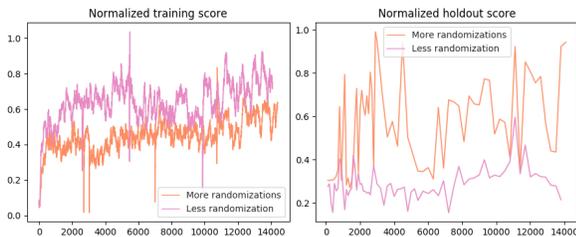


Fig. 5: Left: Episode scores obtained during training. The variant with less randomization is easier and faster to train. Right: On a holdout town with holdout weather better results are achieved by a model trained with more randomization.

B. Experiments in the real world

Perhaps unsurprisingly, in our real-world experiments we observed a high level of noise. Moreover, simulation results

are a poor predictor of real-world performance. More precisely, in the set of policies above some threshold (of “decent driving”) the correlation of simulation and real-world scores is poor. This is perhaps the most evident instantiation of the sim-to-real gap (see also Section IV-D for a positive example of an offline metric correlating with real-world performance). These factors make it hard to put forward definitive conclusions. However, we were able to observe some trends and formulate recommendations for future research. Generally, we observe the positive influence of regularization and augmentation. It also seems beneficial to have an intermediate representation layer (semantic segmentation) and branching architectures (see Section IV-B-d).

Figure 3 summarizes the performance of our models in terms of the percentage of distance *driven autonomously* in all scenarios. See the project website <http://bit.ly/34xh7z4> for a more fine-grained presentation. Below we present a detailed per-experiment analysis.

a) *Base experiments:* The model CONTINUOUS-PLAIN exhibited very good performance and serves as a strong baseline for comparisons with other

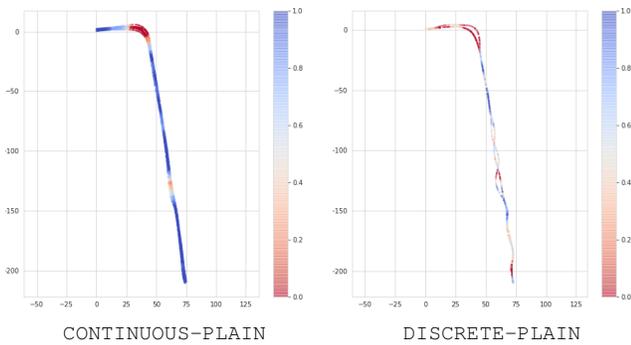


Fig. 6: Qualitative comparison of autonomy between a good model and a lagging one. The graphs show aggregation of few trials (test drives) by these models on the route `factory_city-tunnel-bt10`. The color depicts scale between full autonomy in all trials (blue) and human takeover in all trials (red). Graphs for other models and other routes are available on the project webpage <http://bit.ly/34xh7z4>.

variants discussed below. We aimed at creating a relatively simple model and training procedure. In other experiments we show that further simplifications deteriorate performance. Analogously to the experiment in simulation described above, we have verified the impact of training with fewer randomizations on real-world performance. As expected, the resulting model `CONTINUOUS-LOW-RAND` performs significantly worse, being in fact the worst model tested.

b) Discrete action space: This experiment aimed at measuring the impact of using a discrete distribution for the action space. The steering angles were discretized into unevenly distributed atoms. More of them were placed around 0 to improve smoothness of driving without increasing the action space (viz. $[0., \pm 0.01, \pm 0.02, \pm 0.03, \pm 0.05, \pm 0.08, \pm 0.12, \pm 0.15, \pm 0.2, \pm 0.25, \pm 0.3, \pm 0.4]$; values are in radians). During training the action was sampled, while during evaluation a deterministic policy output the expected action (i.e. the sum of the atom values multiplied by their probabilities).

The resulting model `DISCRETE-PLAIN` performed badly in real-world evaluations, mostly due to severe side-to-side wobbling. We performed more experiments with discrete actions, with results being mostly weak (see <http://bit.ly/34xh7z4>).

c) Regularization: Improved performance in RL generalization when using regularization was reported in [6]. We evaluated using regularization in two experiments. In the first experiment, `DISCRETE-REG`, we fine-tuned `DISCRETE-PLAIN` by further training the model in a slightly altered setup: including l_2 regularization and reducing the policy entropy coefficient from 0.01 to 0.001. The resulting model behaved significantly better (for example the wobbling observed before almost disappeared). In the second experiment, the performance of the continuous model trained with regularization – `CONTINUOUS-REG` – was only slightly improved over `CONTINUOUS-PLAIN`.

d) Control via waypoints: Following the approach presented in [25], in experiment `WAYPOINTS-DISCRETE` we trained a model to predict the next waypoint instead of

steering. Given a waypoint, low-level steering of the driving wheel is executed in order to reach this point. In simulation, this is realized by a PID controller, while in the case of the real car, we use a proprietary control system. To ensure similar performance in simulation and reality, we limit the action space of the RL agent to waypoints reachable by both of the controllers (this consists of points within a radius of 5 meters from the car). The action space is discrete – potential waypoints are located every 5 degrees between -30 and 30 , where 0 is the current orientation of the vehicle.

In contrast to the experiments with direct steering, the continuous version of this experiment – `WAYPOINTS-CONTINUOUS` – was weaker and exhibited strong wobbling, even in simulation.

In this experiment we used a branched neural network architecture, again inspired by [25]. Namely, we use separate heads for each of the four high-level navigation commands (see Section III-e). Such architectures are considered to learn semantically different behaviors for the commands (e.g. *go straight vs turn left*) more easily and better handle command frequency imbalance. Our experimental results are in line with this interpretation – our models performed turns better than ones with the standard architecture. We expect that they will offer better general performance, which we plan to investigate in further research.

e) Dynamics randomizations: In [28], [27] dynamics randomization is pointed out as an important ingredient for successful sim-to-real transfer. In order to verify this in our context we introduced randomization to the following aspects of the environment: target speed, steering response (including a random multiplicative factor and bias), latency (the delay between observation and applying the policy’s response to it), and noise in car metric observation (speed, acceleration, and wheel angle). Dynamics randomization parameters were sampled once at the beginning of each training episode.

For both experiments with dynamics randomization – `DYNAMICS-RAND-RNN` and `DYNAMICS-RAND-FFW` – the performance during evaluation on the real car was substandard. Somewhat surprisingly, the feed-forward model `DYNAMICS-RAND-FFW` performed slightly better than `DYNAMICS-RAND-RNN` using a GRU memory cell [5]. This is in contrast to literature, e.g. [28] which highlights importance of using memory cells: intuitively, an agent with memory should infer the dynamics parameters at the beginning of the episode and utilize them for better driving. We intend to further evaluate the possibility of using dynamics randomization for sim-to-real autonomous driving in future work. As a first step we will look for an explanation of the described mediocre performance. We speculate that this might be due to poor alignment of our randomizations with real-world requirements or overfitting when using high-capacity models with memory.

f) Auxiliary depth prediction: Auxiliary tasks are an established method of improving RL training (see e.g. [17]). Following that, in experiment `AUXILIARY-DEPTH` the model also predict the depth. The depth prediction is learned in a supervised way, along with RL training. This auxiliary

task slightly speeds up the training in simulation. However, in real-world evaluations it does not improve over the baseline experiments.

g) *Segmentation only*: Similarly to [25] we test the hypothesis that semantic segmentation is a useful common representation space for sim-to-real transfer. In experiment SEMSEG-ONLY the model takes in only segmentation as input and performs only slightly weaker than the baseline.

C. Selected failure cases

Besides major design decisions (as described in previous sections) there are a number of small tweaks and potential pitfalls. Below we show two examples, which we find illustrative and hopefully useful for other researches and practitioners.

1) *Single-line versus double-line road markings*: In initial experiments we have used CARLA’s TOWN1 and TOWN2 maps, which feature only double-line road markings. When evaluated on real-world footage, policies trained only on the above were not sensitive to single-line road markings. This problem was fixed after introducing our custom maps, which feature single-line road markings.

2) *Bug in reward function resulting in driving over the curb*: Our reward function includes a term which incentivizes the agent to follow scenarios routes defined by checkpoints connected with straight lines. In one of our scenarios, based on the real-world testing area, the checkpoints were too sparse, resulting in one of the connecting lines going over a curb on a bend of the road. We noticed this only after doing a real-world test, where the car also exhibited similar behavior. This illustrates the well-known tendency of RL methods to overfit to idiosyncrasies of the reward function’s design. Somewhat ironically, in this case our system overcame the sim-to-real gap and transferred the unexpected behaviour precisely. This suggests that methods with stronger generalization are required – we would hope they would generalize from the other bends the car could drive on without touching the curb.

D. Offline models evaluation

A fundamental issue in sim-to-real experiments is that good performance in simulation does not necessarily transfer to the real-world. This is aggravated by the fact that real-world testing is costly both in time and other resources. Inspired by [7] we introduced a proxy metric, which can be calculated offline and correlates with real-world evaluations. Namely, for the seven scenarios with prefix `factory_city` we obtained a human reference drive. Frame by frame, we compared the reference steering with the one given by our models by calculating the mean absolute error. We observe a clear trend (see Figure 7). While this result is still statistically rather weak, we consider it to be a promising future research direction. We present an additional offline evaluation metric (F_1) on the project website <http://bit.ly/34xh7z4>.

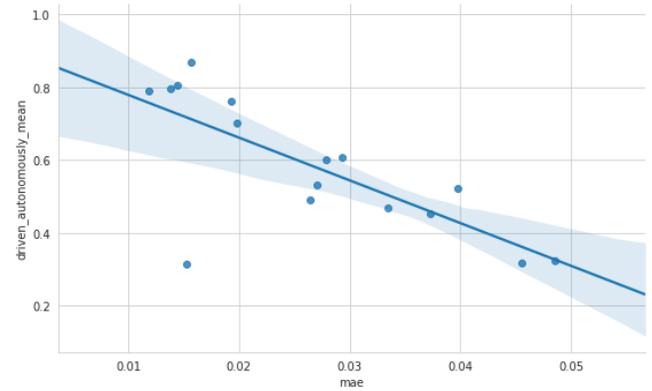


Fig. 7: Dependence of the mean of the *driven autonomously* metric on the mean absolute error with reference drives. Models utilizing waypoints are not included due to a different action space.

V. Conclusions and future work

We presented an overview of a series of experiments intended to train an end-to-end driving policy using the CARLA simulator. Our policies were deployed and tested on a full-size car exhibiting a substantial level of autonomy in a number of restricted driving scenarios.

The current results let us to speculate about the following promising directions: using more regularization, control via waypoints, and using offline proxy metrics. While we obtained poor results with memory-augmented architectures, we plan to investigate the topic further.

We also consider other training algorithms which use a replay buffer such as V-trace [14] and SAC [15]. The asymmetric actor-critic architecture presented in [29] and a generator-discriminator pair similar to the one in [3] can be also beneficial for training of driving policies. Another interesting and challenging direction is integration of an intermediate representation layer — for example a 2D-map or a bird’s-eye view, as proposed in [4], [31], [12], [1]. Focusing RL training on fragments of scenarios with the highest uncertainty, see, e.g., [22] might improve driving stability. Integration of model-based methods similar to [24], [20] would be a desirable step towards better sample efficiency.

VI. Acknowledgements

We would like to thank Simon Barthel, Frederik Kanning, and Roman Vaclavik for their help and dedication during deployment on the physical car. We would also like to thank Tomasz Grel and Przemysław Podczasi for their contribution to the initial stage of the project. The work of Piotr Miłoś was supported by the Polish National Science Center grant UMO2017/26/E/ST6/00622. The work of Henryk Michalewski was supported by the Polish National Science Center grant UMO-2018/29/B/ST6/02959. This research was supported by the PL-Grid Infrastructure. We extensively used the Prometheus supercomputer, located in the Academic Computer Center Cyfronet in the AGH University of Science and Technology in Kraków, Poland. We managed our experiments using <https://neptune.ai>. We would like to thank the Neptune team for providing us access to the team version and technical support.

References

- [1] Mayank Bansal, Alex Krizhevsky, and Abhijit S. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019.*, 2019.
- [2] Alex Bewley, Jessica Rigley, Yuxuan Liu, Jeffrey Hawke, Richard Shen, Vinh-Dieu Lam, and Alex Kendall. Learning to drive from simulation without real world labels. In *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*, pages 4818–4824, 2019.
- [3] Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, Sergey Levine, and Vincent Vanhoucke. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 4243–4250, 2018.
- [4] Jianyu Chen, Bodi Yuan, and Masayoshi Tomizuka. Model-free deep reinforcement learning for urban autonomous driving. *CoRR*, abs/1904.09503, 2019.
- [5] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734, 2014.
- [6] Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, pages 1282–1289, 2019.
- [7] Felipe Codevilla, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. On offline evaluation of vision-based driving models. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*, pages 246–262, 2018.
- [8] Felipe Codevilla, Matthias Müller, Alexey Dosovitskiy, Antonio López, and Vladlen Koltun. End-to-end driving via conditional imitation learning. *CoRR*, abs/1710.02410, 2017.
- [9] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2(1-2):1–142, 2013.
- [10] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [11] Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. OpenAI Baselines. <https://github.com/openai/baselines>, 2017.
- [12] Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, and Jeff Schneider. Motion prediction of traffic actors for autonomous driving using deep convolutional networks. *CoRR*, abs/1808.05819, 2018.
- [13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. *arXiv olopreprint arXiv:1711.03938*, 2017.
- [14] Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, pages 1406–1415, 2018.
- [15] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, pages 1856–1865, 2018.
- [16] Nicolas Heess, Dhruva TB, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, S. M. Ali Eslami, Martin A. Riedmiller, and David Silver. Emergence of locomotion behaviours in rich environments. *CoRR*, abs/1707.02286, 2017.
- [17] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z. Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- [18] Stephen James, Andrew J. Davison, and Edward Johns. Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task. *CoRR*, abs/1707.02267, 2017.
- [19] Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irpan, Julian Ibarz, Sergey Levine, Raia Hadsell, and Konstantinos Bousmalis. Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12627–12637, 2019.
- [20] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H. Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Ryan Sepassi, George Tucker, and Henryk Michalewski. Model-based reinforcement learning for atari. *CoRR*, abs/1903.00374, 2019.
- [21] Katie Kang, Suneel Belkale, Gregory Kahn, Pieter Abbeel, and Sergey Levine. Generalization through simulation: Integrating simulated and real data into deep reinforcement learning for vision-based autonomous flight. In *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*, pages 6008–6014, 2019.
- [22] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In *British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017*, 2017.
- [23] Lukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F. Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmaszczyk, Piotr Jarosik, Mikhail Pavlov, Sergey Kolesnikov, Sergey Plis, Zhibo Chen, Zhizheng Zhang, Jiale Chen, Jun Shi, Zhuobin Zheng, Chun Yuan, Zhihui Lin, Henryk Michalewski, Piotr Milos, Blazej Osinski, Andrew Melnik, Malte Schilling, Helge Ritter, Sean F. Carroll, Jennifer Hicks, Sergey Levine, Marcel Salathé, and Scott Delp. Learning to run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments. In Sergio Escalera and Markus Weimer, editors, *The NIPS '17 Competition: Building Intelligent Systems*, pages 121–153, Cham, 2018. Springer International Publishing.
- [24] Kendall Lowrey, Aravind Rajeswaran, Sham M. Kakade, Emanuel Todorov, and Igor Mordatch. Plan online, learn offline: Efficient learning and exploration via model-based control. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- [25] Matthias Mueller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun. Driving policy transfer via modularity and abstraction. In *2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings*, pages 1–15, 2018.
- [26] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kontschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *International Conference on Computer Vision (ICCV)*, 2017.
- [27] OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *arXiv preprint arXiv:1808.00177*, 2018.
- [28] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 1–8, 2018.
- [29] Lrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor critic for image-based robot learning. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018.
- [30] Dean Pomerleau. ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems 1, [NIPS Conference, Denver, Colorado, USA, 1988]*, pages 305–313, 1988.
- [31] Nicholas Rhinehart, Rowan McAllister, and Sergey Levine. Deep imitative models for flexible inference, planning, and control. *CoRR*, abs/1810.06544, 2018.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015*

- 18th International Conference Munich, Germany, October 5 - 9, 2015, *Proceedings, Part III*, pages 234–241, 2015.

- [33] Andrei A. Rusu, Matej Vecerík, Thomas Rothörl, Nicolas Heess, Razvan Pascanu, and Raia Hadsell. Sim-to-real robot learning from pixels with progressive nets. In *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, pages 262–270, 2017.
- [34] Fereshteh Sadeghi. Divis: Domain invariant visual servoing for collision-free goal reaching. In *Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019.*, 2019.
- [35] Fereshteh Sadeghi and Sergey Levine. CAD2RL: real single-image flight without a single real image. In *Robotics: Science and Systems XIII, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, July 12-16, 2017*, 2017.
- [36] Fereshteh Sadeghi, Alexander Toshev, Eric Jang, and Sergey Levine. Sim2real viewpoint invariant visual servoing by recurrent control. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4691–4699, 2018.
- [37] Ahmad El Sallab, Mohammed Abdou, Etienne Perot, and Senthil Yogamani. Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, 2017(19):70–76, Jan 2017.
- [38] Axel Sauer, Nikolay Savinov, and Andreas Geiger. Conditional affordance learning for driving in urban environments. In *2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings*, pages 237–252, 2018.
- [39] Alexander Sergeev and Mike Del Balso. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.
- [40] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics, Results of the 11th International Conference, FSR 2017, Zurich, Switzerland, 12-15 September 2017*, pages 621–635, 2017.
- [41] Niko Sünderhauf, Oliver Brock, Walter J. Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, and Peter Corke. The limits and potentials of deep learning for robotics. *I. J. Robotics Res.*, 37(4-5):405–420, 2018.
- [42] Bowen Tan, Nayun Xu, and Bingyu Kong. Autonomous driving in reality with reinforcement learning and image translation. *CoRR*, abs/1801.05299, 2018.
- [43] Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. In Tom Howard Hadas Kress-Gazit, Siddhartha Srinivasa and Nikolay Atanasov, editors, *Robotics: Science and System XIV*, 2018.
- [44] Josh Tobin, Lukas Biewald, Rocky Duan, Marcin Andrychowicz, Ankur Handa, Vikash Kumar, Bob McGrew, Alex Ray, Jonas Schneider, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Domain randomization and generative models for robotic grasping. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2018, Madrid, Spain, October 1-5, 2018*, pages 3482–3489, 2018.
- [45] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pages 5026–5033, 2012.