

Learning to Drive from Simulation without Real World Labels

Alex Bewley, Jessica Rigley, Yuxuan Liu, Jeffrey Hawke, Richard Shen, Vinh-Dieu Lam, Alex Kendall

Abstract—Simulation can be a powerful tool for understanding machine learning systems and designing methods to solve real-world problems. Training and evaluating methods purely in simulation is often “doomed to succeed” at the desired task in a simulated environment, but the resulting models are incapable of operation in the real world. Here we present and evaluate a method for transferring a vision-based lane following driving policy from simulation to operation on a rural road without any real-world labels. Our approach leverages recent advances in image-to-image translation to achieve domain transfer while jointly learning a single-camera control policy from simulation control labels. We assess the driving performance of this method using both open-loop regression metrics, and closed-loop performance operating an autonomous vehicle on rural and urban roads.

I. INTRODUCTION

This paper demonstrates how to use machine translation techniques for unsupervised transfer of an end-to-end driving model from a simulated environment to a real vehicle. We trained a deep learning model to drive in a simulated environment (where complete knowledge of the environment is possible) and adapted it for the visual variation experienced in the real world (completely unsupervised and without real-world labels). This work goes beyond simple image-to-image translation by making the desired task of driving a differentiable component within a deep learning architecture. We show this shapes the learning process to ensure the driving policy is invariant to the visual variation across both simulated and real domains.

Learning to drive in simulation offers a number of benefits, including the ability to vary appearance, lighting and weather conditions along with more structural variations, such as curvature of the road and complexity of the surrounding obstructions. It is also possible to construct scenarios in simulation which are difficult (or dangerous) to replicate in the real world. Furthermore, simulation provides ground-truth representations like semantics and even privileged information such as the relative angle and position of the vehicle with respect to the road [25]. However, applying this knowledge to real-world applications has been limited due to the reality gap, often expressed as the difference in appearance from what can be rendered in simulation to how the physical world is viewed from a camera.

In contrast to the many efforts to improve the photometric realism of rendered images [5], [7], [25], [26], this work learns an invariant mapping between features observed in simulation and the real environment. This representation is jointly optimised for steering a vehicle along a road in

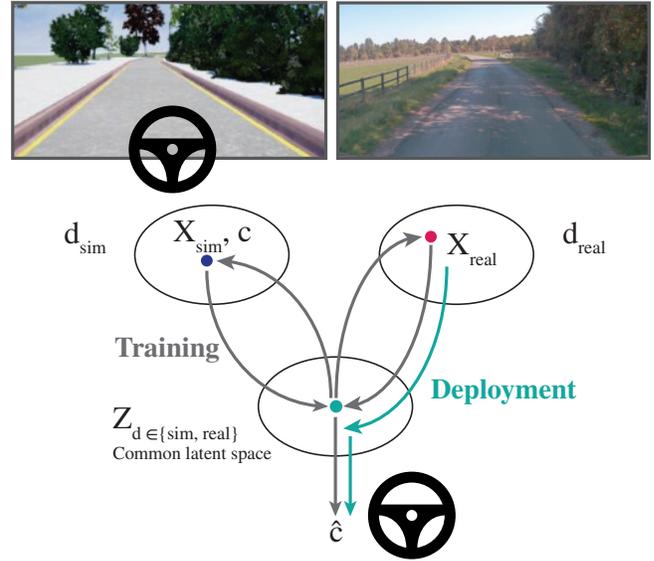


Fig. 1: We constructed a model for end-to-end driving (vision to action) for lane following by learning to translate between simulated and real-world imagery ($X_{sim,real}$ from domains $d_{sim,real}$), jointly learning a control policy from this common latent space Z using labels c from an expert driver in simulation. This method does not require real-world control labels (which are more difficult to acquire), yet learns a policy predicting a control signal \hat{c} which can be applied successfully to a real-world driving task.

addition to image translation. While enjoying the benefits of simulation, this method is also adept at inferring the common structure between the target environment and what has been seen in simulation. This is analogous to transferring skills learned though playing virtual racing games to a person’s first experience behind the wheel of a car [14].

In summary, the main contributions of this work are:

- We present the first example of an end-to-end driving policy transferred from a simulation domain with control labels to an unlabelled real-world domain.
- By exploiting the freedom to control the simulated environment we were able to learn manoeuvres in states beyond the common driving distribution in real-world imitation learning, removing the need for multiple camera data collection rigs or data augmentation.
- We evaluated this method against a number of baselines using open-loop (offline) and closed-loop metrics, driving a real vehicle over 3km without intervention on a rural road, trained without any real-world labels.
- Finally, we demonstrate this method successfully driving in closed-loop on public UK urban roads.

A supplementary video demonstrating the system behaviour is available here: <https://wayve.ai/sim2real>

II. RELATED WORK

There has been significant development in building simulators for the training and evaluation of robotic control policies. While recent simulated environments, such as [24], [5], have benefited from continued progress in photo realistic gaming engines, the reality gap remains challenging. [28] address this problem by exploiting a simulator’s ability to render significant amounts of random colours and textures with the assumption that visual variation observed in the real world is covered.

An active area of research is in image-to-image translation which aims to learn a mapping function between two image distributions [16], [15], [11], [38]. Isola *et al.* [11] proposed a conditional adversarial network where the discriminator compares a pair of images corresponding to two modalities of the same scene. Zhu *et al.* [38] relaxed the need for corresponding images by proposing a cyclic consistency constraint for back translation of images. Similarly, Liu *et al.* [15] combine cross domain discriminators and a cycle consistent reconstruction loss to produce high quality translations. We extend this framework to accommodate learning control as an additional supervised decoder output.

A problem closer to this work is that of unsupervised domain adaptation combining supervision from the source domain with pairwise losses or adversarial techniques to align intermediate features for a supervised task [29], [34], [8]. However, feature based alignment typically requires some additional pairwise constraints, [29] or intermediate domains [34], [35], to bridge significant shifts in domain appearance. Other related work align both feature and pixel level distributions for the perception only task of semantic segmentation [8]. The very recent work of Wenzel *et al.* [33] learns transfer control knowledge across simulated weather conditions via a modular pipeline with both image translation and feature-wise alignment.

While the large majority of work in autonomous driving focuses on rule-based and traditional robotic control approaches, end-to-end learning approaches have long held the promise of learning driving policies directly from data with minimal use of brittle assumptions and expensive hand tuning. Pomerleau [23] first demonstrated the potential feasibility of this approach by successfully applying imitation learning (IL) to a simple real-world lane following task. Muller *et al.* [20] further successfully learned to imitate functional driving policies in more unconstrained off-road environments with toy remote control cars.

More recently, Bojarski *et al.* [1] demonstrated that simple behaviour cloning could be scaled to a larger and more complex set of lane following scenarios when using a multi-camera system and modern neural net architectures. Codevilla *et al.* [4] took a similar multi-camera approach, and learned to navigate successfully in simulated urban driving scenarios by conditioning on higher level driving decisions. Building on this, Müller *et al.* [19] transferred policies from simulation to reality by using a modular architecture predicting control decisions learned from semantic

image segmentation.

Mehta *et al.* [18] take a learning from demonstration approach, and augment the control prediction task to depend on intermediate predictions of both high-level visual affordances and low-level action primitives. Kendall *et al.* [12] demonstrated that pure reinforcement learning approaches with a single camera system could be successfully applied for learning real-world driving policies, though single camera systems have typically not been implemented to date, attributed largely to the robustness gained from training data with multiple views and synthetic control labels. In this work, robustness is achieved through noisy experiences captured in simulation.

In this work, we use an intermediate representation where an implicit latent structure is formed, as opposed to more explicit representations, such as semantic segmentation [9]. The work presented in this paper is closest to [36] who recently proposed to exploit image translation for learning to steer while also sharing a similar philosophy in reducing the complexity of the simulated environment to minimise the sufficient statistics for the desired task. Our work differs in that we use a bi-directional image translation model where we are not dependent on having real-world driving commands. We also demonstrate the efficacy of a driving policy on a real system, showing closed-loop performance metrics rather than solely open-loop metrics against a validation set.

III. METHOD

The problem of translating knowledge from expert demonstrations in simulation to improved real-world behaviour can be cast into the framework of unsupervised domain adaptation. Here, the training set consisted of image-label pairs from a simulated environment (X_{sim}, C_{sim}). An independent set of unlabelled images (X_{real}) from the real target environment was used to capture the distribution of real-world observations. Crucially, it is important to note there were no pairwise correspondences between images in the simulated training set and the unlabelled real-world image set.

A. Architecture

Our framework consisted of various modules, which can broadly be broken down into an image-translator, a set of domain discriminator networks and the controller network for imitation learning. Fig. 2 provides a high-level overview of our model, showing the flow of information between the different modules.

1) *Translator*: The image-translation module followed the general framework proposed in [15]. This consisted of two convolutional variational autoencoder-like networks where the latent embedding was swapped for translating between domains. More formally:

$$Z_d = E_d(X_d) + \epsilon \quad (1)$$

$$\hat{X}_d = G_d(Z_d) \quad (2)$$

where $d \in [sim, real]$ represents the domain and $\epsilon \sim N(0, 1)$ is noise added during training, but set to zero for

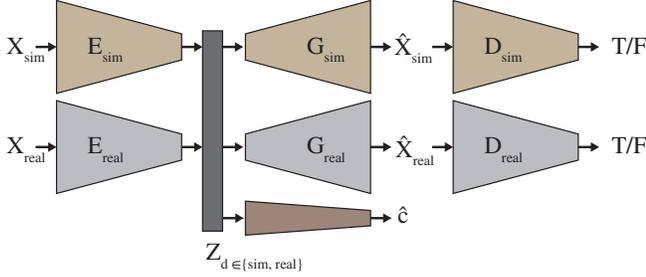


Fig. 2: Model architecture for domain-transfer from a simulated domain to real-world imagery, jointly learning control and domain translation. The encoders $E_{sim,real}$ map input images from their respective domains to a latent space Z which is used for predicting vehicle controls \hat{c} . This common latent space is learned through direct and cyclic losses as part of learning image-to-image translation, indicated conceptually in Figure 3 and in Section III-B.

inference. The process of translation consisted of computing Z_d using one domain encoder and then predicting \hat{X} with the other.

2) *Discriminators*: For each translated image a two-scale discriminator was used to align the appearance distribution of the translator output with the images in the corresponding domain.

3) *Controllers*: The control architecture is broadly based on the end-to-end control architecture in [4], where the image translator forms the primary convolutional encoder. The latent tensor Z_d is passed through CoordConv [21] layers to learn a spatially aware embedding suited for control from this latent translation representation. Finally, the spatially aware tensor is reduced with global average pooling and then followed by fully connected layers.

B. Losses

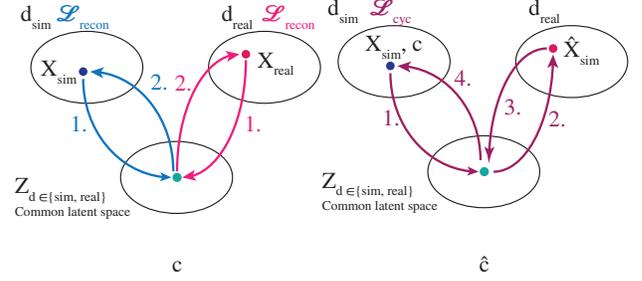
Figure 3 gives an overview of the main losses.

1) *Image Reconstruction Loss*: For a given domain d , E_d and G_d constitute a VAE. To improve image translation we used an L1 Loss \mathcal{L}_{recon} between an image X_d and the reconstructed image after passing it through the corresponding VAE, $X_d^{recon} = G_d(E_d(X_d))$, as shown in Figure 3a.

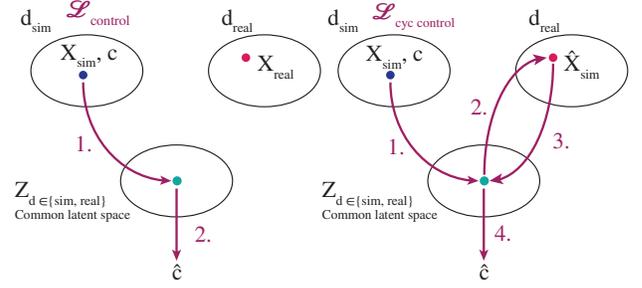
2) *Cyclic Reconstruction Loss*: Assuming a shared latent space implies the cycle-consistency constraint, which says that if an image is translated to the other domain and then translated back, the original image should be recovered [38], [15]. We applied a cyclic consistency loss \mathcal{L}_{cyc} to the VAEs, given by an L1 loss between an image X_d and the image after translating to the other domain, d' , and back, $X_d^{cyc} = G_d(E_{d'}(G_{d'}(E_d(X_d))))$, see Figure 3b.

3) *Control Loss*: To guide our model to learn features that are useful for driving, we also used a control loss $\mathcal{L}_{control}$, which is an L1 loss between the controller's predicted steering $\hat{c} = C(E_d(X_d))$ and the ground truth given by the autopilot, c , shown in Figure 3c.

Control should be based on the semantic content of an image and independent of which domain the original image came from. We therefore introduced a cyclic control loss $\mathcal{L}_{cyc\ control}$, as shown in Figure 3d, an L1 loss between



(a) Reconstruction loss \mathcal{L}_{recon} (b) Cyclic reconstruction loss \mathcal{L}_{cyc}



(c) Control loss $\mathcal{L}_{control}$ (d) Cyclic control loss $\mathcal{L}_{cyc\ control}$

Fig. 3: Learning the latent space Z for translation and control requires a combination of direct and cyclic losses: learning to reconstruct images intra-domain, translating them inter-domain, while jointly learning a supervised control signal from the simulation domain. This figure does not illustrate the *adversarial* or *perceptual* loss. The *latent reconstruction* loss is conceptually analogous to \mathcal{L}_{cyc} in 3b, but applied to the latent space. Section III-B describes these loss terms in further detail.

the predicted steering \hat{c} and the steering predicted by the controller when the image was translated to the other domain and then encoded, $\hat{c}^{cyc} = C(E_{d'}(G_{d'}(E_d(X_d))))$.

4) *Adversarial Loss*: Both the image-translator and the discriminators were optimised with the Least-Squares Generative Adversarial Network (LSGAN) objective proposed by [17]. The discriminator (3) and generator (4) adversarial losses ensured that translated images resembled those from the chosen domain.

$$\mathcal{L}_{LSGAN}(D) = \mathbb{E}_{X \sim p_{data}} [(D(X) - 1)^2] + \mathbb{E}_{Z \sim p_Z(Z)} [(D(G(Z)) - 0)^2] \quad (3)$$

$$\mathcal{L}_{LSGAN}(G) = \mathbb{E}_{Z \sim p_Z(Z)} [(D(G(Z)) - 1)^2] \quad (4)$$

5) *Perceptual Loss*: To encourage consistent semantic content across the two domains, we employed the use of a pre-trained VGG [27] model, which was applied to both the original and translated images. The perceptual loss $\mathcal{L}_{perceptual}$ was expressed as the difference between the features extracted from the last convolutional layer in the VGG16 model for a given input image and its translated counterpart. Extracted features were normalised via Instance Norm (IN) [31] following the result from Huang *et al.* [10] demonstrating that applying IN before computing the feature distance makes the perceptual loss more domain-invariant.

6) *Latent Reconstruction Loss*: Ideally we wanted to encode the semantic content of the images within the latent space such that Z is independent of the domain from which an image came. We therefore applied a latent reconstruction loss \mathcal{L}_{Zrecon} , an L1 loss between the latent representation of an image Z_d and the reconstruction of the latent representation after it was decoded to the other domain and then encoded once more, $Z_d^{recon} = E_{d'}(G_{d'}(Z_d))$.

The total discriminator loss for a given domain was $\mathcal{L}_{LSGAN}(D_d)$ (3). The VAEs and the controller were trained jointly using a weighted sum of losses with weights λ_i , given by

$$\begin{aligned} \mathcal{L}_{tot} = & \lambda_0 \mathcal{L}_{recon} + \lambda_1 \mathcal{L}_{cyc} + \lambda_2 \mathcal{L}_{control} \\ & + \lambda_3 \mathcal{L}_{cyc\ control} + \lambda_4 \mathcal{L}_{LSGAN}(G) \\ & + \lambda_5 \mathcal{L}_{perceptual} + \lambda_6 \mathcal{L}_{Zrecon}. \end{aligned} \quad (5)$$

The image translator was optimised using ADAM [13], with a learning rate of 0.0001 and momentum terms of 0.5 and 0.999. The controller was trained using stochastic gradient descent with a learning rate of 0.01.

IV. EVALUATION

To evaluate our approach in Section III, we considered data gathered from a simulated domain (with a procedurally generated environment) and a real domain (gathered from driving on real roads). We trained models using the methods in IV-D first in a rural setting, to compare our method to a number of baseline approaches. We then demonstrate our method scaling to the complexity of public urban roads in the United Kingdom.

Table II presents in-domain and cross domain open-loop (offline) metrics for the learned controller, and Table III presents closed loop imitation control results evaluated on a physical vehicle in the rural environment. Qualitative results of the image-to-image translation are shown in Figure 7 for both the rural and urban setting. The following subsections detail the performance metrics used and test scenarios considered, followed by an analysis of the obtained results.

A. Simulation Domain

1) *Procedural Environment Generation*: For the purpose of this work we created a simulated environment where we could build up a significant training set of image- and control-label pairs on procedurally generated virtual roads. Figure 4 illustrates the rural simulated world, while Figure 6 shows the urban simulated world.

For each data collection run in the rural world, a single road was created using a piece-wise Bézier curve generated by sampling 1-dimensional Simplex noise [22]. The curvature of the road could be coarsely controlled using the frequency parameter of the noise. The road surface was then assigned a random texture from a finite set of example road textures. Once the rural road was constructed, trees and foliage were placed using Poisson Disk Sampling [2] according to a foliage density parameter. In simulation we were able to vary environmental factors, such as cloud

cover, rainfall, surface water accumulation, and time of day. However, we fixed these as we wished to learn an image-to-action control policy in a very narrow source domain: we did not require high variance in the source distribution, rather to learn how the higher variance target distribution translates to the source.

For the urban world, we use a similar approach to procedurally generate a road network. We procedurally add buildings, trees and parked cars. Care was made to approximate the layout and topology of the real-world urban environment, but not the visual complexity or photo-realism. Instead, we rely on the image translation model to transfer the policy from the cartoon-like simulated world to the visually rich real-world.

2) *Simulated Expert Agent*: The labelled training set was generated from an expert autopilot agent. The expert driver has a simple proportional controller empirically tuned to track the lane based on ground-truth distance from the centre, maintaining a constant vehicle speed. The curve from Section IV-A.1 used to generate the road was used to generate a set of lane paths which were offset from the central curve. One end of the road was chosen at random and the simulated vehicle was placed at the end and told to follow the corresponding lane path.

To perturb the vehicle state, we used additive Ornstein-Uhlenbeck (OU) process noise [30] to the expert driver’s actions. This had the effect of generating a more robust training dataset by moving the vehicle throughout the drivable lane space, observing the expert driver’s response from each perturbed pose. The OU noise parameters θ and σ were selected to maximise the magnitude of perturbation while still allowing the expert driver to largely stay within the lane. This simulation environment ran asynchronously, with image-label pairs captured at 10Hz.

Leveraging this expert agent to learn to drive with imitation learning is only possible in simulation. It requires privileged information, such as the distance to the centre of the lane, which is not available in the real world. Furthermore, we perturb the state with noise to generate richer training data; this would be dangerous in the real world as it would require swerving on the road.

B. Real-World Domain

1) *Rural road*: We considered a single driving environment for initial real-world testing, consisting of a 250m private one-lane rural road, shown in Figure 4. For safety, the vehicle was only driven at 10 kmh^{-1} in the absence of any other vehicles wishing to use the road, both in data collection and under test.

2) *Urban road*: In addition, we consider an urban road environment in Cambridge, UK. We select minor public roads in dense suburban areas and opportunistically test when they are void of other traffic. Figure 6 illustrates typical scenes. We test across sun, overcast and raining weather conditions.

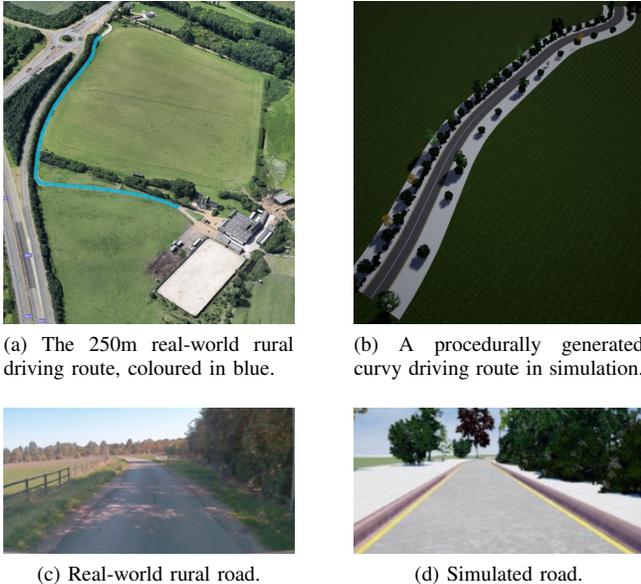


Fig. 4: Rural setting: aerial views of the real-world (a) and simulated driving routes (b), along with example images from each domain showing the perspective from the vehicle (c, d).

C. Data

Table I outlines the training, and test, lane following datasets used. For each domain, we collected approximately 60k frames of training data and 20k frames of test data.

TABLE I: Datasets used for driving policy domain adaptation. Each simulation frame has associated expert control labels.

| | Training Frames | Test Frames |
|------------|-----------------|-------------|
| Simulation | 60014 | 17741 |
| Real | 57916 | 19141 |

1) *Data Balancing*: The maximum steering range of each vehicle was parameterised to ± 1.0 . Due to the nature of driving, the labelled training set was heavily dominated by data with near-zero steering angles (i.e., driving straight), as shown in Figure 5. If the controller is trained with a naive mean absolute error (MAE) loss, it could feasibly predict zero steering for all outputs and fail to learn the correct image-steering mapping.

To address this problem, we split the data into eight bins according to the steering angle, and uniformly sampled from each bin, upsampling the data such that every bin contained the same number of samples as the bin with the maximum number of data. The bins edges used in our experiments were $[-1, -0.075, -0.05, -0.025, 0.0, 0.025, 0.05, 0.075, 1]$. The steering limits differ between vehicles, hence we applied a linear calibration to the real vehicle steering output.

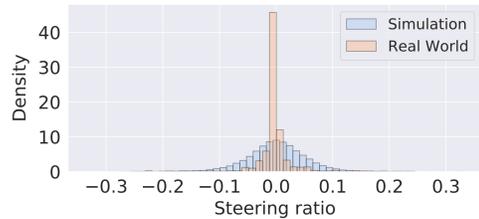


Fig. 5: The steering data distribution (in the range ± 1.0) differs dramatically between the simulation and real world, where simulation has much greater variance. The real-world data has shorter tails and is heavily concentrated around the zero position, due to the profile of the rural driving route.

D. Transfer Learning Results

1) *Baselines*: We compared our method to the following baselines:

- **Simple Transfer** this takes a model pre-trained in simulation and directly applies it to the real-world data. Note: compared to the following this model only sees examples from simulation.
- **Real-to-Sim Translation** uses the unsupervised image-to-image translation to convert a real-world image to the simulation domain for direct application of the controller pre-trained in simulation.
- **Sim-to-Real Translation** uses the unsupervised image-to-image translation for training a controller on translated images (sim-to-real) with corresponding simulated actions. At test time no translation was performed and only the controller was used.
- **Latent Feature ADA** applies Adversarial Domain Adaptation (ADA) [34] to the feature space to align encoders from simulation and real data.

For evaluation, we also compared a Drive-Straight policy as a proxy to assess road curvature, and to quantitatively assess the efficacy of offline metrics.

2) *Performance Metrics*: Open-loop evaluation of control policies is an open question: [3] and [36] present a range of metrics, demonstrating weak correlation between offline metrics and online performance. We adopted two metrics to compare driving policies: mean average error (MAE), and Balanced MAE. These metrics for all approaches tested here are outlined in Table II. The MAE between the predicted and ground-truth steering is a useful loss function, but poor reflection of actual performance due to the data imbalance. We propose a Balanced-MAE metric: discretising the validation data as per the bins in Section IV-C.1, computing the MAE of the data per bin, then taking an equally weighted average across all bins (similar in principle to mean AP in object detection or recognition tasks). Qualitatively, we found that this metric correlates with closed-loop driving performance in simulation.

For closed-loop testing on our rural driving route, we used a simple measure of distance travelled per safety driver intervention in metres travelled, averaged over 3km of driving. During test, we ran the models at close to the camera’s 30Hz



Fig. 6: Urban setting: illustrations of the simulated and real-world domains for urban road scenes. Despite quite a large appearance change from the cartoon-like simulated world to the real-world, our model is able to successfully transfer the learned control policy to drive in the real world, with no real world labels.

frame rate. Table III outlines the performance of our system, as well as a number of baseline models. We found that the open-loop metrics in Table II hold only a weak correlation with real-world performance, demonstrated by the difference in closed- and open-loop performance between our model and the baseline approaches. Clearly, further work is needed to develop open-loop metrics which support real-world performance.

TABLE II: Open-loop control metrics on the simulation and real test datasets from Table I.

| | Simulation | | Real | |
|-------------------------|--------------|--------------|--------------|--------------|
| | MAE | Bal-MAE | MAE | Bal-MAE |
| Drive-Straight | 0.043 | 0.087 | 0.019 | 0.093 |
| Simple Transfer | 0.05 | 0.055 | 0.265 | 0.272 |
| Real-to-Sim Translation | - | - | 0.261 | 0.234 |
| Sim-to-Real Translation | - | - | 0.059 | 0.045 |
| Latent Feature ADA | 0.040 | 0.047 | 0.032 | 0.071 |
| Ours | 0.017 | 0.018 | 0.081 | 0.087 |

TABLE III: On-vehicle performance, driving 3km on the rural driving route in 4a. For policies unable to drive a lap with ≤ 1 intervention, we terminated after one 250m lap (\dagger).

| | Mean distance / intervention (metres) |
|-------------------------|---------------------------------------|
| Drive-straight | 23 \dagger |
| Simple Transfer | 9 \dagger |
| Real-to-Sim Translation | 10 \dagger |
| Sim-to-Real Translation | 28 \dagger |
| Latent Feature ADA | 15 \dagger |
| Ours | No intervention over 3km |

3) *Loss Ablation*: The various terms in Eqn. 5 are designed to work in concert to maintain a structured latent space to facilitate zero-shot transfer from simulated label to the real domain. The following ablation study investigates the influence of each term by setting the corresponding weighting λ to zero, effectively removing it from the training procedure. Table IV shows that all terms play a significant role in reducing the error with the exception of the cyclic control factor. This indicated that the other cyclic

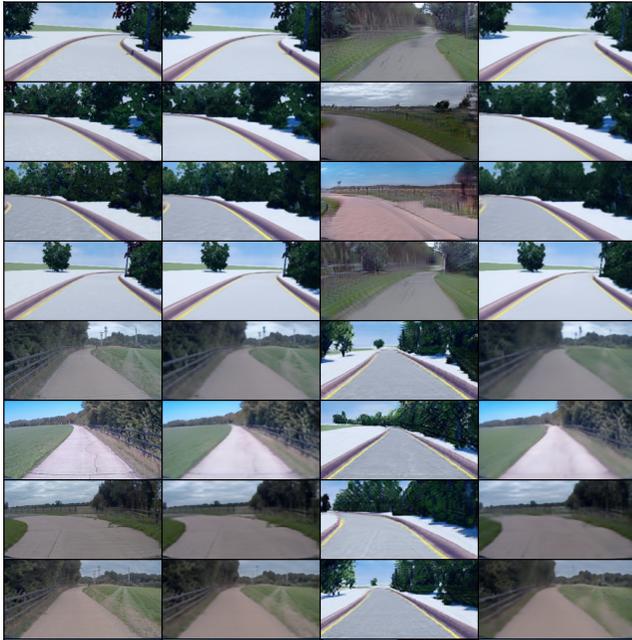
losses are sufficient to maintain a shared latent structure. Not surprisingly, the L1 cyclic reconstruction and GAN loss from [15] are most critical for control transfer. Finally, as removing the control term prevented the network from learning anything sensible, we instead evaluated the effect of the gradients of the controller on the image translator. Here we zeroed the gradients from the controller before flowing back into the translator, effectively turning the controller into a passive neural stethoscope, as described in [6]. These gradients would normally help inject structure for the task into the translator in an auxiliary fashion leading to better real-world performance. This increase in error by not optimising the encoders for both tasks could also be a contributing factor to the poor performance of the translation baselines (real-to-sim and sim-to-real).

TABLE IV: Ablation by removing single terms in eqn. 5. Values expressed as the average error multiplier across three runs for the open-loop metrics when transferred to the real domain test dataset. For the Control term the gradients applied to the image translator for this loss are removed (\ddagger).

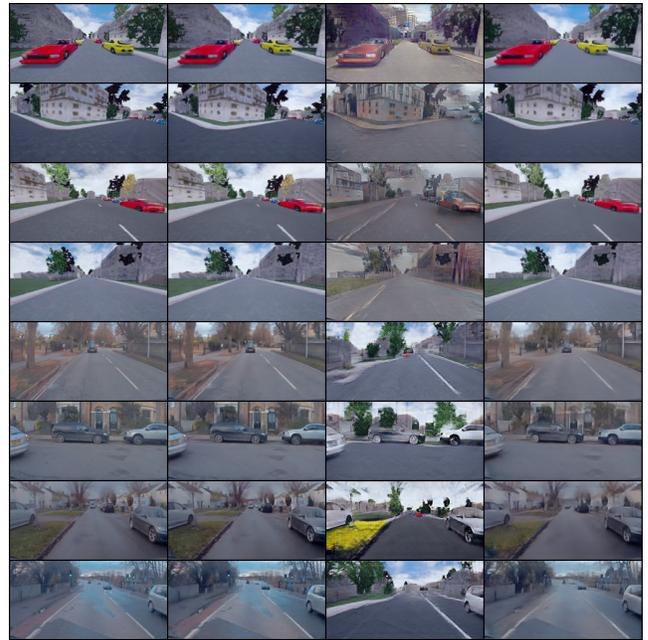
| | MAE | Bal-MAE |
|----------------------------|-------|---------|
| Recon. L1 | 1.602 | 1.259 |
| Recon. Cyclic L1 | 2.980 | 2.395 |
| Z Recon. | 1.766 | 1.270 |
| Perceptual | 1.693 | 1.435 |
| Cyc Control | 0.923 | 1.068 |
| GAN | 1.893 | 1.507 |
| Control-grad(\ddagger) | 1.516 | 1.433 |

4) *Urban Road Results*: To further demonstrate the efficacy of our method, we performed closed-loop testing on public UK urban roads. Using our model, we were able to successfully demonstrate the vehicle lane-following on dense urban side streets in Cambridge, UK. Examples are shown in Figure 7b.

We observed one main conclusion: the model is typically able to drive if, and only if, the visual translation is successful. For example, when the topology of the road and all cars are translated correctly between domains, the car's control policy gives the desired lane-following closed-loop behaviour. When a car is mis-translated into a footpath, the



(a) Rural environment.



(b) Urban environment.

Fig. 7: Qualitative results of the image translation model for (a) rural and (b) urban environments. In each example, the top four rows show images originating in our simulated environment, and the bottom half showing translation from real world images. The columns left to right are as follows: original, image reconstructed from latent back to original domain, translated to other domain, and taking the translated image and translating it back to the original domain.

control policy requires intervention. Therefore, we conclude that this method may be able to scale to more complex domains, beyond simple lane following tasks, if the simulator is capable of simulating these scenarios and we can learn a successful domain translation model.

5) *Visualisation*: Understanding the performance of a control policy coupled with image domain transfer can be difficult. The bi-directional image translator provides the ability to inspect the model’s interpretation of the road surface as shown in Fig. 7. Here we can observe that the curvature and offset of the road is appropriately translated across domains facilitating a consistent transfer of the steering signal. Interestingly, in the third row of Figure 7a (where the vehicle is nearly driving off the virtual road, which is far from the distribution of collected real images) we notice the model will generate an image closer to the distribution of real-world observations. In such scenarios, the controller is trained with a high magnitude steering command to correct its course in simulation, which results in robust behaviour in the real world as it controlled the vehicle back into the distribution of real observations.

V. CONCLUSIONS

Learning a driving policy from simulation has many advantages: training data is cheap, auxiliary ground-truth information can be provided with ease, and the vehicle can be put in situations that are difficult or dangerous to undertake in reality. Previously, with the substantial gap in complexity between the two domains, it was considered

infeasible to transfer driving policies from simulation to the real world.

In this work, we present the first system that is capable of leveraging simulation to learn an end-to-end driving policy to directly transfer to real-world scenarios without any additional human demonstrations. This model jointly learns to translate images between the real-world operating domain and a procedurally generated simulation environment, while also learning to predict control decisions off of the latent space between the two domains given only the ground-truth labels from simulation. We empirically validated our proposed model in closed-loop against several baselines, successfully driving 3km between interventions on a real-world lane following task. We further evaluated the model using several standard open-loop metrics, observing that these metrics ultimately did not prove predictive of driving performance. Finally, we demonstrated this system driving in closed-loop on public urban roads in the United Kingdom.

This work provides evidence that end-to-end policy learning and simulation-to-reality transfer are highly promising directions for the development of autonomous driving systems. We note that standard open-loop metrics for this problem need to be improved, and leave this question to future work. Furthermore, the addition of orthogonal, but relevant, approaches on temporal motion consistency [32], [37] could further advance this field. We hope this work inspires further investigation of both learning driving policies directly from data, and exploiting simulation for removing the constraints of the real world.

REFERENCES

- [1] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [2] Robert Bridson. Fast poisson disk sampling in arbitrary dimensions. In *ACM SIGGRAPH 2007 Sketches*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.
- [3] Felipe Codevilla, Antonio M. Lopez, Vladlen Koltun, and Alexey Dosovitskiy. On offline evaluation of vision-based driving models. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [4] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *International Conference on Robotics and Automation (ICRA)*, 2018.
- [5] 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.
- [6] Fabian B Fuchs, Oliver Groth, Adam R Kosorick, Alex Bewley, Markus Wulfmeier, Andrea Vedaldi, and Ingmar Posner. Neural stethoscopes: Unifying analytic, auxiliary and adversarial network probing. *arXiv preprint arXiv:1806.05502*, 2018.
- [7] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4340–4349, 2016.
- [8] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola Openai, Kate Saenko, Alexei A Efros, and Trevor Darrell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *International Conference on Machine Learning (ICML)*, 2018.
- [9] Zhang-Wei Hong, Chen Yu-Ming, Shih-Yang Su, Tzu-Yun Shann, Yi-Hsiang Chang, Hsuan-Kung Yang, Brian Hsi-Lin Ho, Chih-Chieh Tu, Yueh-Chuan Chang, Tsu-Ching Hsiao, et al. Virtual-to-real: Learning to control in visual semantic segmentation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [10] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal Unsupervised Image-to-Image Translation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [12] Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen, Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to drive in a day. *arXiv preprint arXiv:1807.00412*, 2018.
- [13] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [14] Li Li, Rongrong Chen, and Jing Chen. Playing action video games improves visuomotor control. *Psychological science*, 27(8):1092–1108, 2016.
- [15] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised Image-to-Image Translation Networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [16] Ming-Yu Liu and Oncel Tuzel. Coupled Generative Adversarial Networks. In *Advances in Neural Information Processing Systems*, 2016.
- [17] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2813–2821. IEEE, 2017.
- [18] Ashish Mehta, Adithya Subramanian, and Anbumani Subramanian. Learning end-to-end autonomous driving using guided auxiliary supervision. *arXiv preprint arXiv:1808.10393*, 2018.
- [19] Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladen Koltun. Driving policy transfer via modularity and abstraction. *arXiv preprint arXiv:1804.09364*, 2018.
- [20] Urs Muller, Jan Ben, Eric Cosatto, Beat Flepp, and Yann L Cun. Off-road obstacle avoidance through end-to-end learning. In *Advances in neural information processing systems*, pages 739–746, 2006.
- [21] David Novotny, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. Semi-convolutional operators for instance segmentation. *arXiv preprint arXiv:1807.10712*, 2018.
- [22] Ken Perlin. Improving noise. *ACM Trans. Graph.*, 21(3):681–682, July 2002.
- [23] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.
- [24] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *International conference on computer vision (ICCV)*, volume 2, 2017.
- [25] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision*, pages 102–118. Springer, 2016.
- [26] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3234–3243, 2016.
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [28] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [29] Eric Tzeng, Coline Devin, Judy Hoffman, Chelsea Finn, Pieter Abbeel, Sergey Levine, Kate Saenko, and Trevor Darrell. Adapting deep visuomotor representations with weak pairwise constraints. In *Workshop on Algorithmic Foundations in Robotics (WAFR)*, 2016.
- [30] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Phys. Rev.*, 36:823–841, Sep 1930.
- [31] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [33] Patrick Wenzel, Qadeer Khan, Daniel Cremers, and Laura Leal-Taixe. Modular vehicle control for transferring semantic information between weather conditions using gans. In *Conference on Robot Learning*, pages 253–269, 2018.
- [34] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. Addressing appearance change in outdoor robotics with adversarial domain adaptation. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [35] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. Incremental adversarial domain adaptation for continually changing environments. In *International Conference on Robotics and Automation (ICRA)*, 2018.
- [36] Luona Yang, Xiaodan Liang, Tairui Wang, and Eric Xing. Real-to-Virtual Domain Unification for End-to-End Autonomous Driving. Technical report.
- [37] Jingwei Zhang, Lei Tai, Yufeng Xiong, Ming Liu, Joschka Boedecker, and Wolfram Burgard. Vr goggles for robots: Real-to-sim domain adaptation for visual control. *arXiv preprint arXiv:1802.00265*, 2018.
- [38] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.