Cycle-Consistent World Models for Domain Independent Latent Imagination

Tim Joseph

KI DELTA LEARNING
Scalable AI for Automated Driving

Semi- and Unsupervised Learning / Environment

Problem

For many tasks, such as autonomous driving, running an agent in the real world is prohibitively expensive and dangerous. For this reason, training the agent in a simulated environment before releasing it to the real world is necessary. However, many agents fail in the real world because they are not able to handle the domain gap [1] between the real and the simulated environment.

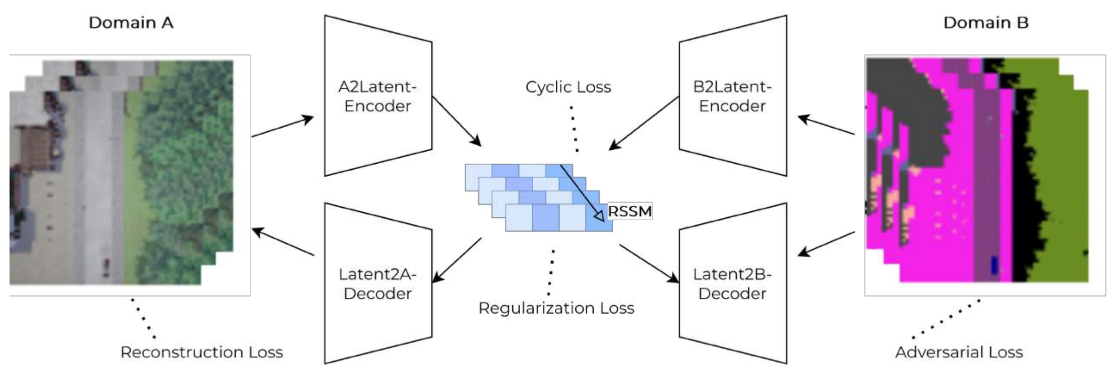

Figure 1: The general architecture of our approach. We train multiple encoder and decoder that project either into or from a common latent space on both domains.

Approach

In this work, we present an approach to create a world model [2] that projects observations from two different domains into a single common latent space. Such a world model can then be rolled out to generate data to train a policy on latent states which are similar for simulation and reality. Thus, it is not necessary to adapt policy training to either of the domains. This approach can be used for any sensor input in the form of 2D maps, such as RGB-camera input or BEV-projected LIDAR points. For performance reasons we evaluate our approach in an RGB to Semantic Segmentation set

Cycle-Consistent World Models (CCWM)

Our model consists of the following modules:

- Domain A/B encoder and decoder
- Latent transition model prior and posterior
- Reward decoder

The loss function consists of the reconstruction losses to each domain, the cycle-consistency losses and regularization losses that aligns the prior and posterior. An observation from either domain is projected in the shared latent space and latent transition model enforces that the latent space must be learned such that future latent states can be predicted independently from the domain.

Approach	Reward RSE	Reward RSE cross-modality	PSNR
Single Modality	0.25	3.86	10.21
RC	0.31	0.49	11.39
CycleGAN	0.28	0.57	12.28
Ours	0.23	**0.48**	**13.91**

Table 1: Our approach outperforms common related approaches.

Results

We compared our approach to related domain adaption approaches, such as Random Convolutions [3] and Cycle-GANs [4]. Our results show that CCWM yields better reward prediction than related approaches in the original domain and in the cross-modality setting CCWM can reconstruct the cross-modality observation from latent states Additionally, an advantage of our approach is that by translation on the other domain, we can make the model interpretable, e.g. we can analyze qualitative differences between the reconstruction in both domains.

References:

[1] Y. Ganin and V. Lempitsky. Unsupervised Domain Adaptation by Backpropagation, ICML 2015

[2] D. Ha and J. Schmidhuber, Recurrent World Models Facilitate Policy Evolution, NIPS 2015

[3] J.-Y. Zhu et. al., Unpaired image-to-image translation using cycle-consistent adversarial networks, ICCV 2017

[4] Z. Xu et. al., Robust and generalizable visual representation learning via random convolutions, ICLR 2021

Figure 2: Examples for translation from/to a semantic birds-eye-view domain from/to a rgb domain.

Partners

External partners

For more information contact:

joseph@fzi.de

KI Delta Learning is a project of the KI Familie. It was initiated and developed by the VDA Leitinitiative autonomous and connected driving and is funded by the Federal Ministry for Economic Affairs and Climate Action.

KI FAMILIE

www.ki-deltalearning.de @KI_Familie KI Familie

Supported by:

Federal Ministry for Economic Affairs and Climate Action

on the basis of a decision by the German Bundestag