28-01-2022

Unpaired day-night domain adaptation using CycleGANs combined with CIConv

1. Definitions

- Generative Adversarial Networks (GAN)

Neural network (NN) capable of generating realistic fake images. It is trained using two opposing models. The Generator aims to generate fake images of sufficient quality to fool the Discriminator into classifying them as real. They learn by becoming better at beating the other. See figure 1 for examples of GAN generated images.

Figure 1: GAN generated faces.

- Cycle-Consistent GAN (CycleGAN)

NN capable of performing domain adaptation consisting of two GANs setup in a cyclic fashion in order to maintain structural image properties. See figure 2 for examples for CycleGAN converted images.

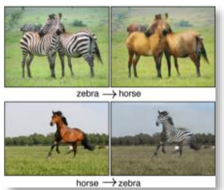

Figure 2: CycleGAN converted images [2].

- Color Invariant Convolution (CIConv)

Learnable layer in a neural network for making images more robust to illumination changes, such as the difference between day and night [1]. See figure 3 for an example of CIConv applied to an image of a bike.

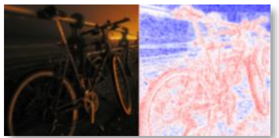

Figure 3: CIConv applied to image of bike [1].

2. Introduction

- Lack of night-time training data for self-driving cars' NNs (usually CNNs).
- Daytime training data readily available.
- CycleGANs for day-night domain adaptation, see figure 4.
- Combined with CIConv for potentially improved results.

Figure 4: Example of day-night domain adaptation.

3. Research question

What is the influence of CIConv in combination with CycleGANs on the generation of labelled training data in a domain for which only unlabelled data is available?

4. Method

- Three CycleGAN architectures:

CG => No CIConv
CG_{disc} => CIConv in first layer Discriminators
CG_{gen} => CIConv in first layer Generators

- Two loss functions:

$\mathcal{L}_{CycleGAN}$ => Adversarial loss and cycle-consistency loss [2]
$\mathcal{L}_{CycleWGAN_{GP}}$ => Wasserstein distance, gradient penalty and cycle-consistency loss [3]

- Train on daytime images and night-time images:

Daytime images => CityScapes
Night-time images => Dark Zurich

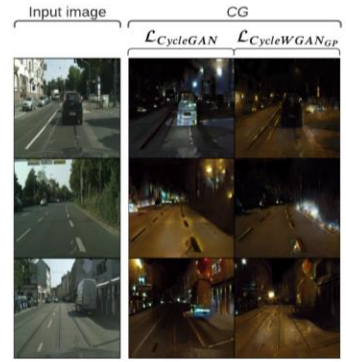

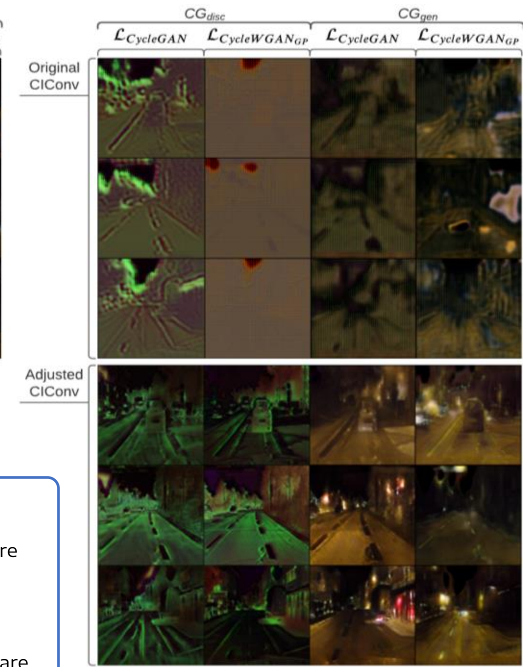

Figure 5: Examples of day-night domain adaptation performed by the three architectures trained for only 10 epochs with the two different loss functions and CIConv versions.

5. Experiments

- CG_{disc} and CG_{gen} with original CIConv are incapable of training stably.
- We propose an adjusted CIConv layer where its output is capped before normalization.
- CG_{disc} and CG_{gen} with adjusted CIConv are capable of training stably.
- However, only CG_{gen} seems capable of domain adaptation.
- See figure 5 for examples of day-night domain adaptation by the three architectures, the two loss functions and both CIConv versions.

6. Conclusion

- Original CIConv layer causes exploding gradients and is therefore not suitable in a CycleGAN performing day-night domain adaptation.
- Adjusted CIConv layer allows for stable training and is therefore potentially a useful addition to a CycleGAN performing day-night domain adaptation. However, this seems to only hold for the architecture with CIConv in the Generators.

7. Future work

- Test different CIConv cap values.
- Implement different CIConv adjustments.
- Compare quality of generated images of architectures with adjusted CIConv and without CIConv.
- Compare effectiveness in training of CNNs with images generated by the architectures with adjusted CIConv and without CIConv.

Contact

- By: Thomas Streefkerk (t.s.streefkerk@student.tudelft.nl)
- Supervisor: Atilla Lengyel (a.lengyel@tudelft.nl)
- Responsible professor: Jan van Gemert (j.c.vangemert@tudelft.nl)

[1] A. Lengyel, S. Garg, M. Milford, and J. C. van Gemert, "Zero-shot day-night domain adaptation with a physics prior,"
[2] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks,"
[3] W. Hu, M. Li, and X. Ju, "Improved CycleGAN for image-to-image translation,"