

Use-Case in Delta Learning

We address the domain shift for simulation-to-real and cross-sensor adaptation in terms of a top-down human pose estimation.

Technical Problem

Top-down human pose estimation involves two steps:

1. Object detection: Detect persons and their bounding boxes
2. Human pose estimation: Perform pose estimation on the predicted bounding boxes

Both steps are affected by a domain shift. We use a dataset consisting of synchronized

synthetic and real data to investigate this domain shift and to develop and evaluate our unsupervised domain adaptation method.

Technical Solution

Our method is an extension of CycleGAN [1] as shown in Figure 1. Domain transfer is guided by additional task specific losses.

Evaluation

In Figure 2, we report the mean average precision (mAP) for the object detection and the percentage of correct key points with an allowed deviation of 0.5 (PCKh@0.5) and 0.1 (PCKh@0.1) times the size of the persons head for human pose estimation.

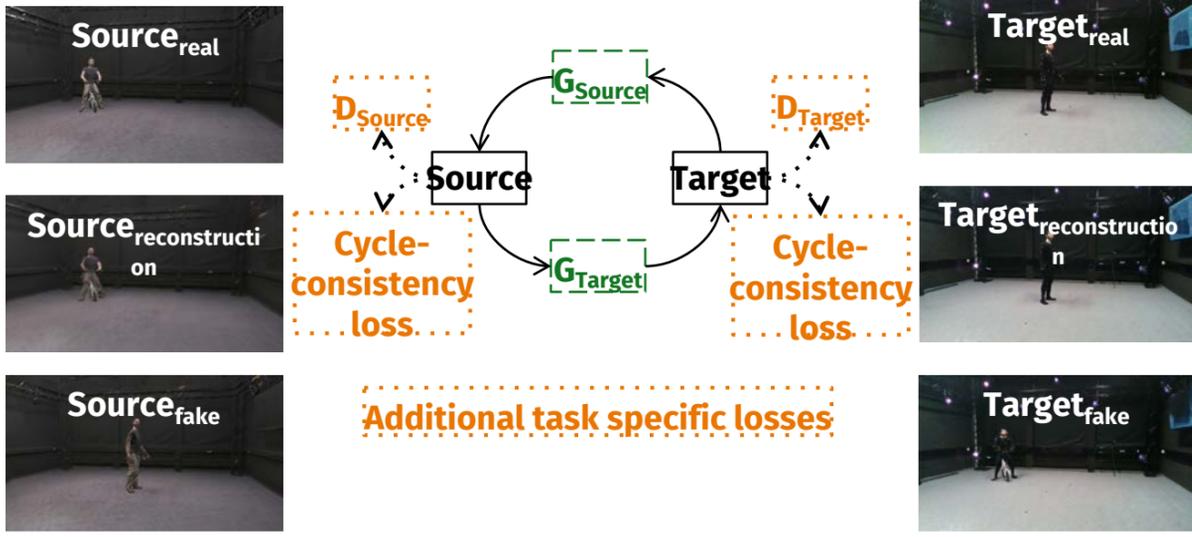


Figure 1: CycleGAN [1] learns to translate between the source and target domain. We extend CycleGAN with additional task specific losses to further guide the domain transfer.

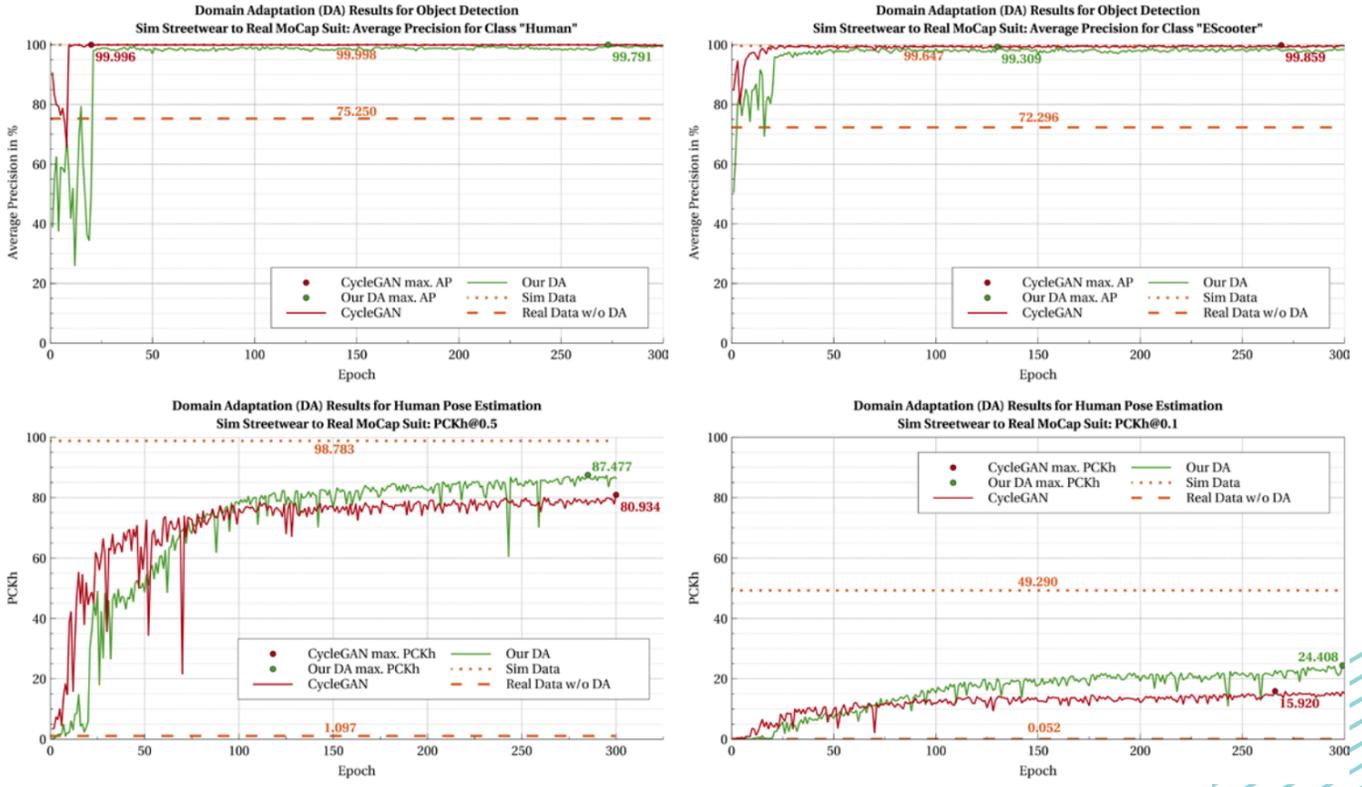


Figure 2: Preliminary results comparing CycleGAN and our domain adaptation method on the task of object detection as well as human pose estimation.

References:

- [1] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," Oct. 2017. doi: 10.1109/iccv.2017.244.

Partners



External partners