Unsupervised Magnification of Posture Deviations across Subjects

RUPRECHT-KARLS-UNIVERSITÄT HEIDELBERG · IWR Interdisciplinary Center for Scientific Computing · Computer Vision group heidelberg · Heidelberg Collaboratory for Image Processing

Michael Dorkenwald* Uta Büchler* Björn Ommer

CVPR SEATTLE WASHINGTON JUNE 16-18 2020

1 Introduction

Can you see a difference in posture?

Our Goal: Making subtle posture deviations across subjects visible

Generating new images showing *magnified* posture deviations

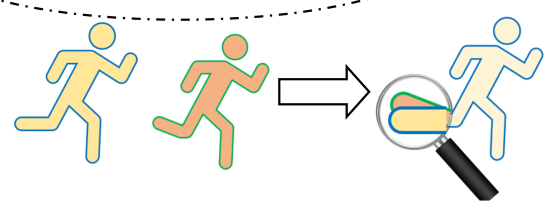

Applications

- Medicine — Increase visibility of motor function impairment for better treatments

Comparison with healthy movement (left): Impaired subject (right) walks with its legs apart

- Sport — Highlight incorrect postures to improve performance

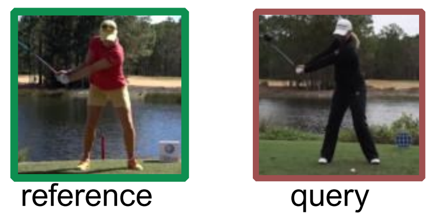

Comparison with successful golf swing (left): Knee is twisted inside (right)

Related Work

e.g. Facilitating perception of 'wobbling' effect of a pupil

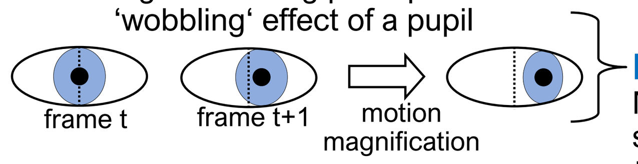

Image Generation
Synthesize images with specific appearance in different postures
Requirement: Disentanglement of appearance and posture

Motion Magnification
Magnify small motion within a specific video
Requirement: Sensitiv to *small* changes

Our Method: Combining both worlds

Magnification of *subtle* posture deviations across subjects

2 Disentanglement for Magnification

Goal: Magnifying *only* posture differences

Unsupervised disentanglement of posture π and appearance α

Model: Autoencoder with two encoders E_π and E_α

Problematic Shortcut: Ignoring encoding of E_α

Our Solution: Discourage correct reconstuction if one encoding is ignored

Input: Perturbed video frames

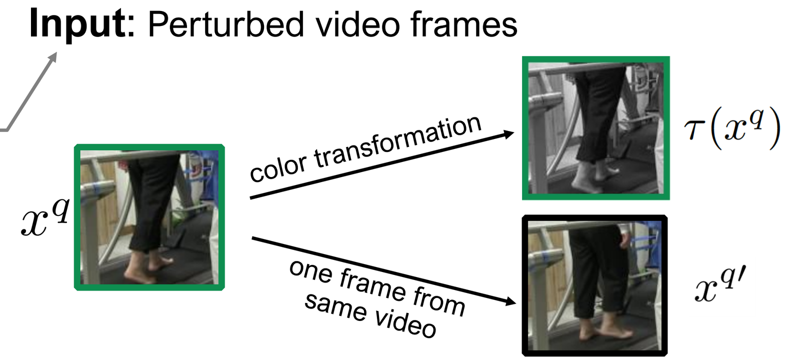

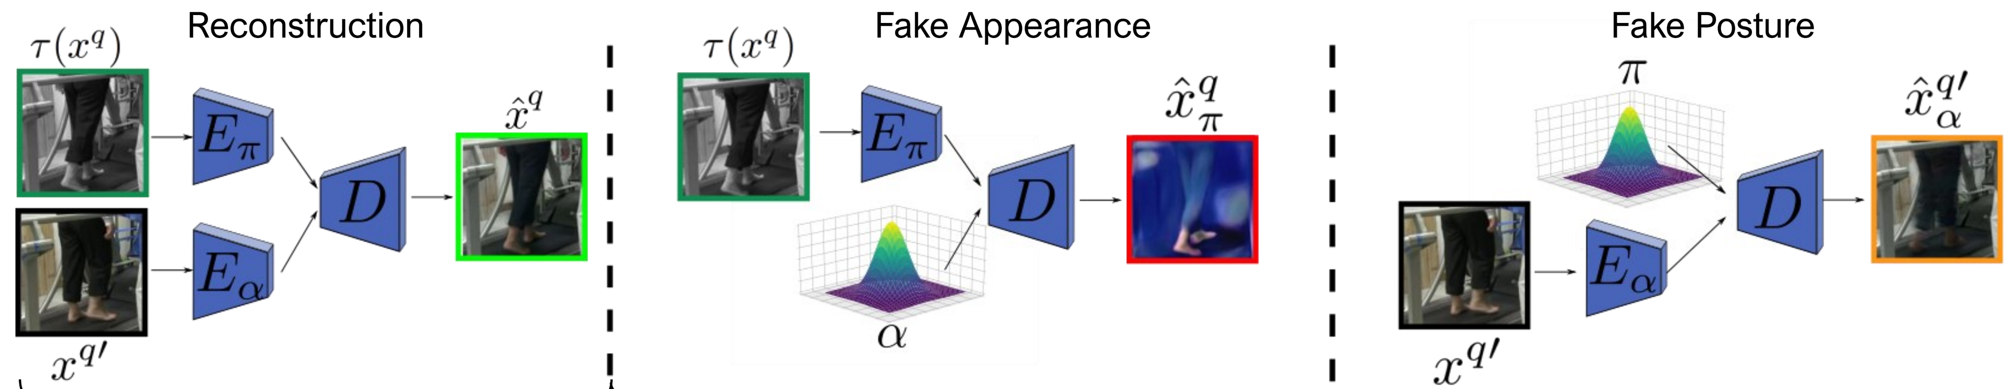

$$\mathcal{L}_{\text{rec}} = d(\hat{x}^q, x^q)$$

$$\mathcal{L}_{\text{dis}} = \|d(x^q, \hat{x}^q_\pi) - t_\pi\|_1 + \|d(x^q, \hat{x}^{q'}_\alpha) - t_\alpha\|_1$$

target values of distance between x^q and fake reconstructions

3 Learning to Magnify

Goal: Produce realistic and valid magnifications

Challenges:

1) How do we magnify? → Extrapolation
2) Magnification produces new postures
   - No ground-truth → Fixpoint loss $\mathcal{L}_{\text{fix}}$
   - No guarantee for realistic postures → Adverserial loss $\mathcal{L}_{\text{A}}$

Our Solution: Extrapolation; Fixpoint loss $\mathcal{L}_{\text{fix}}$; Adverserial loss $\mathcal{L}_{\text{A}}$

magnification intensity

$$m_\pi := E_\pi(x^r) + \lambda\,(E_\pi(x^q) - E_\pi(x^r))$$

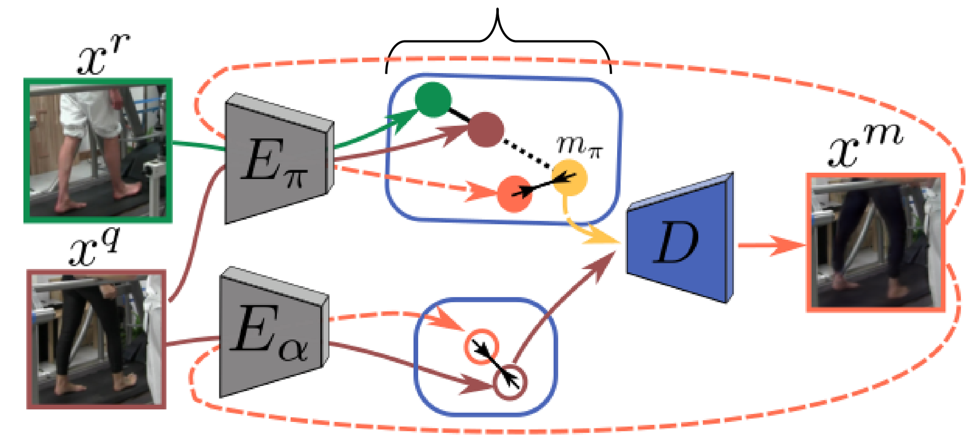

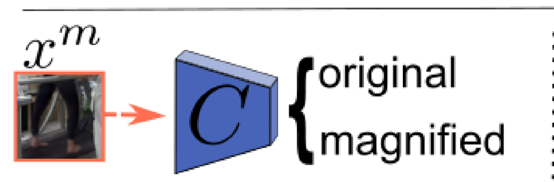

$$\min \mathbf{d}(\bullet, \bullet) \qquad \min \mathbf{d}(\circ, \circ)$$

$$\mathcal{L}_A \qquad \mathcal{L}_{\text{fix}} = \|E_\pi(x^m) - m_\pi\|_2^2 + \|E_\alpha(x^m) - E_\alpha(x^q)\|_2^2$$

$$\Rightarrow \mathcal{L}_{\text{mag}} = \mathcal{L}_{\text{A}} + \beta\mathcal{L}_{\text{fix}}$$

4 Results and Ablation Studies

Qualitative Results on a Human Gait

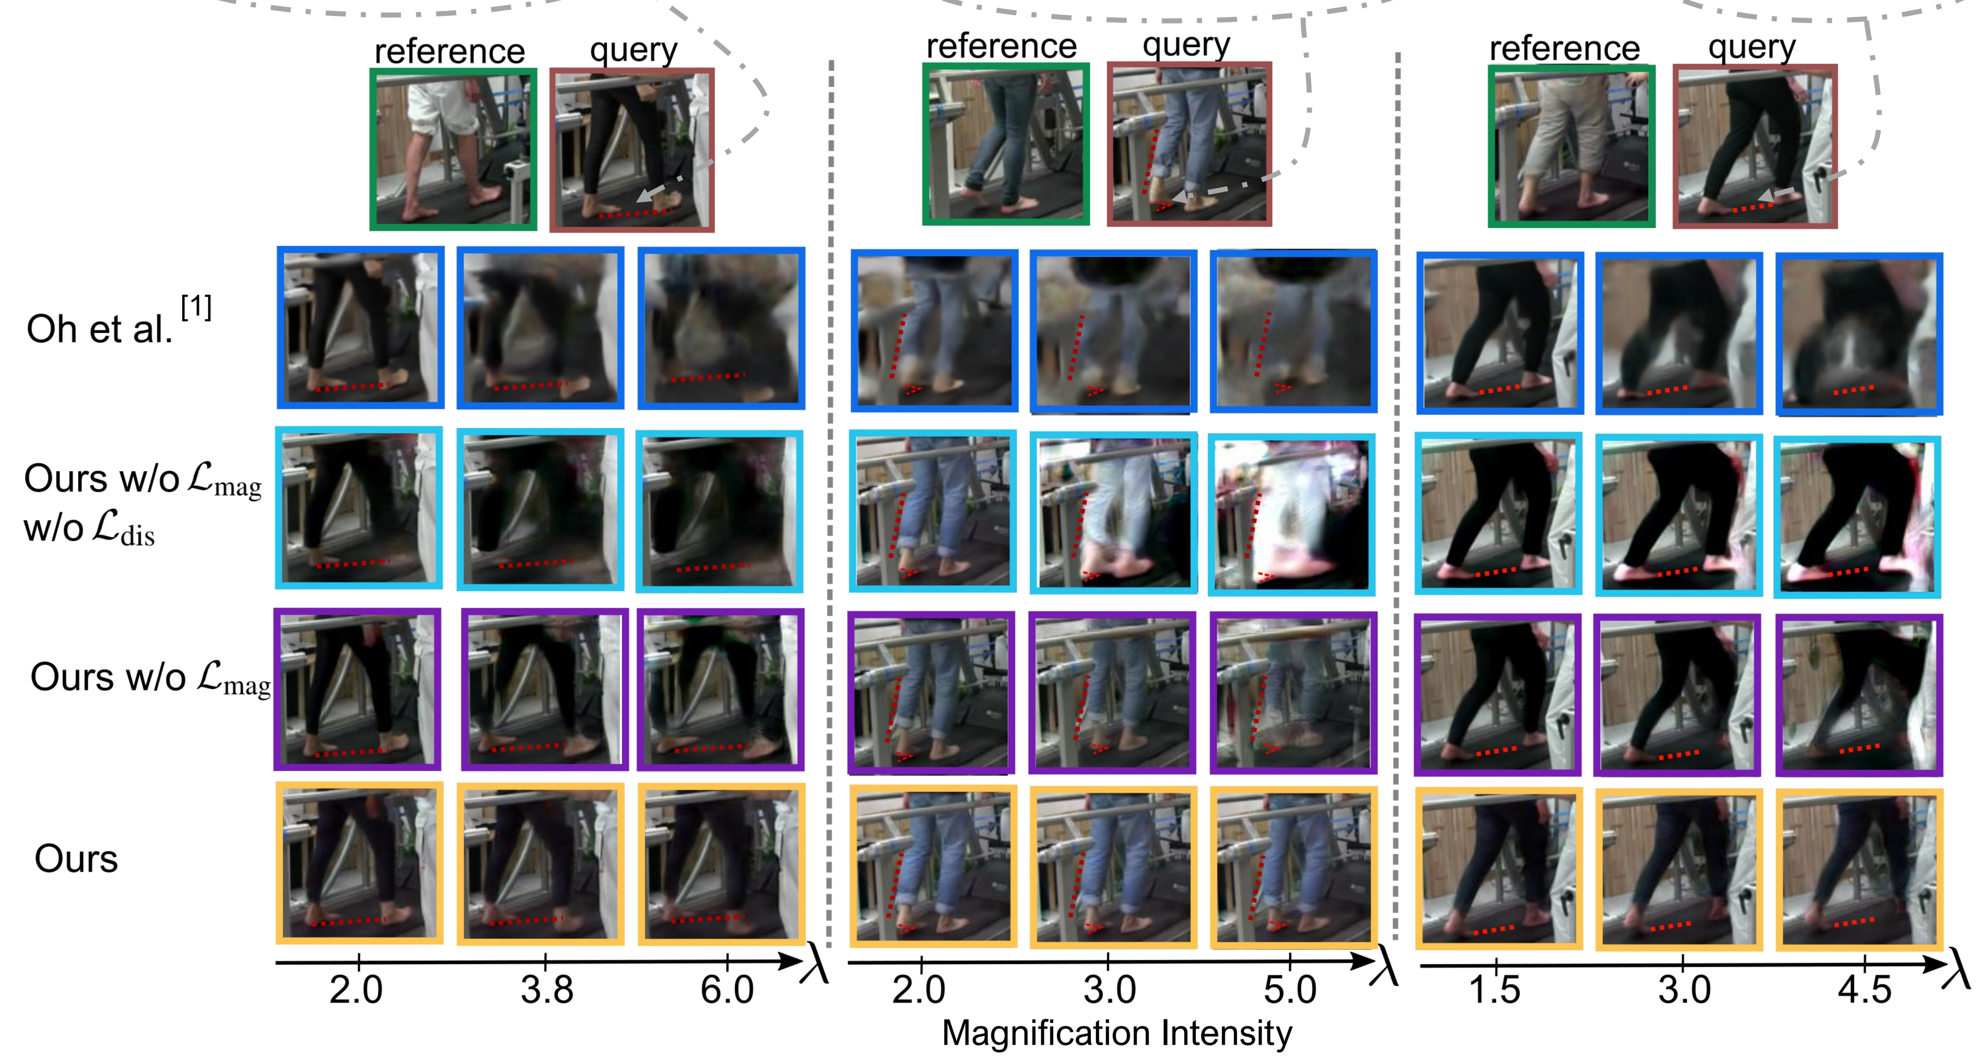

Qualitative Results on Golf Swing

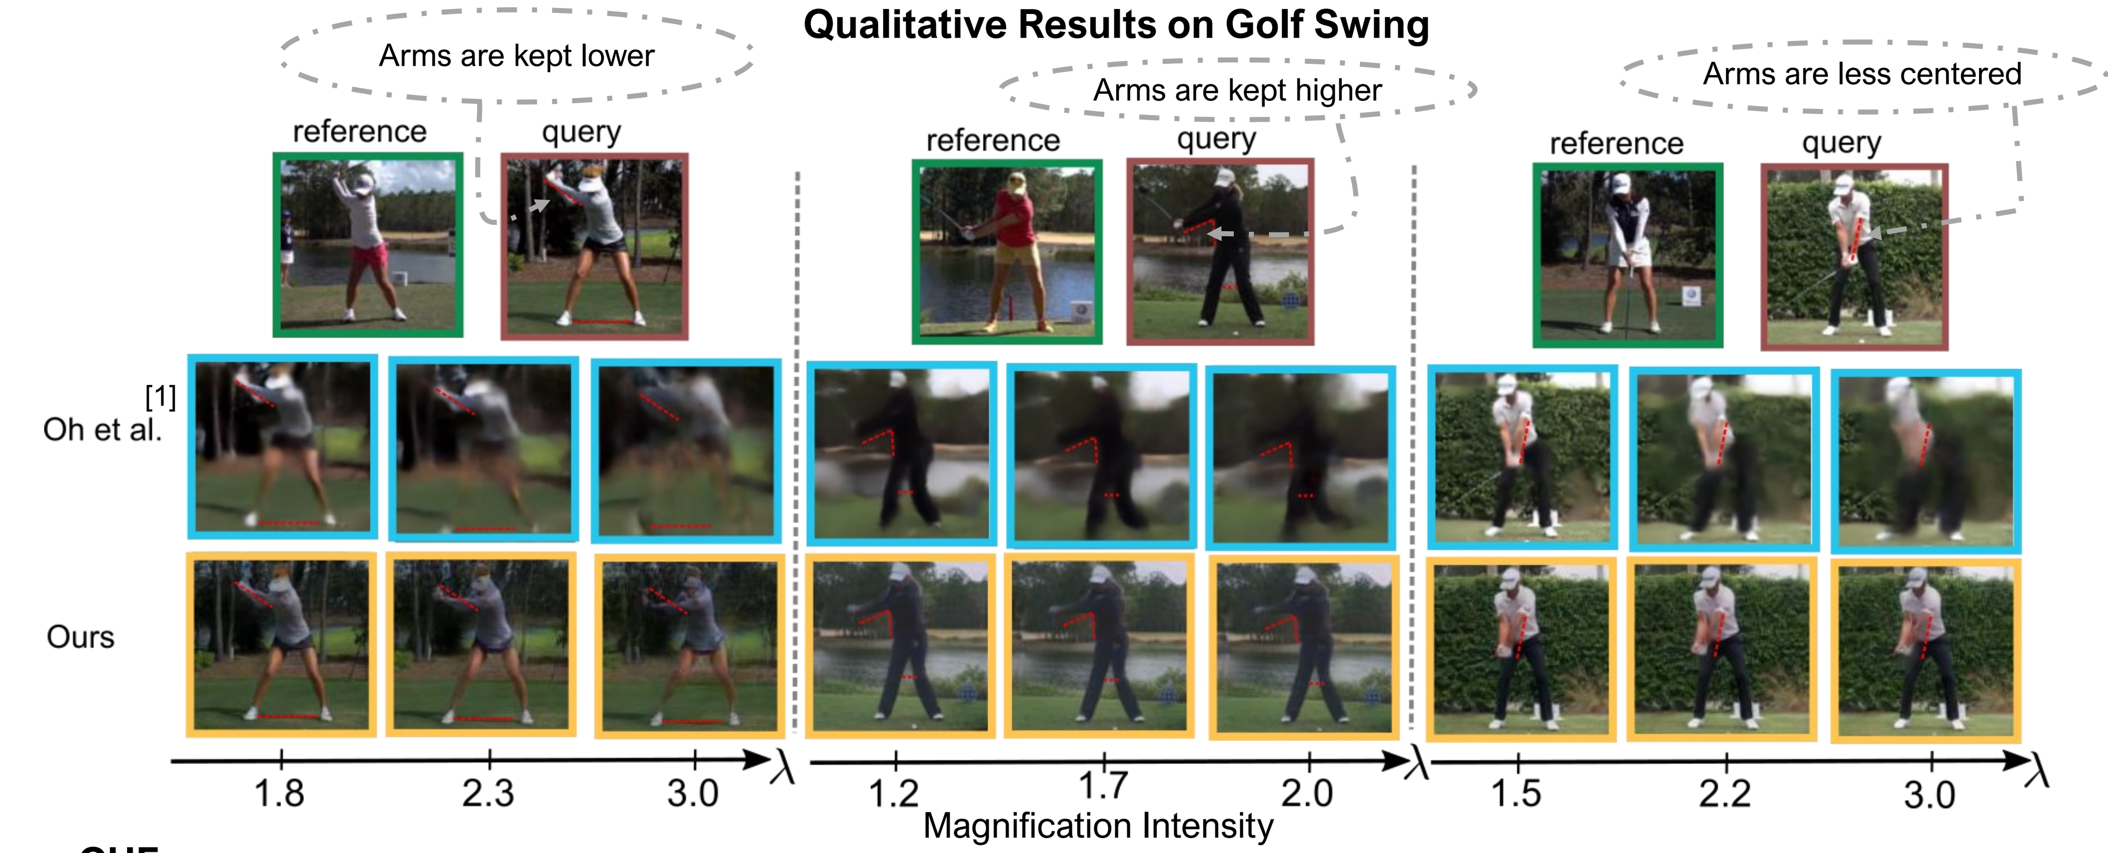

Qualitative Results on CUEye

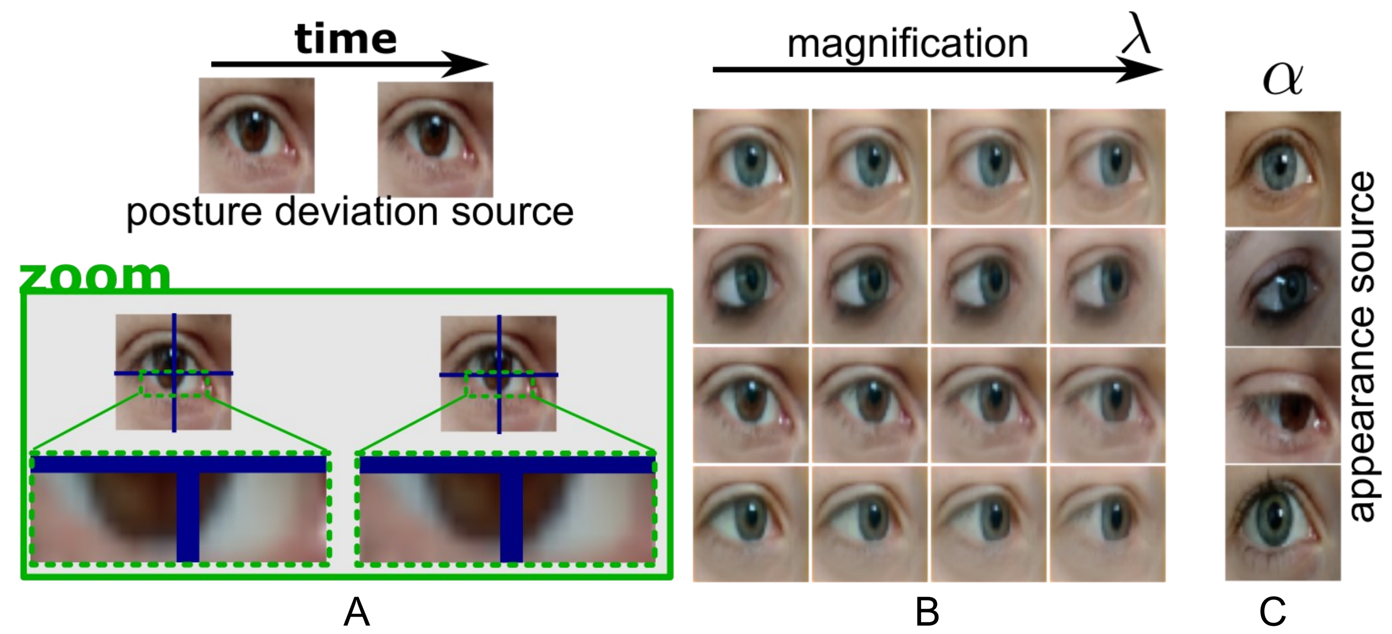

Transferring subtle movements of the query's pupil (A) to various target appearances (C). A: Tiny motion from left to right. **B:** Given one of the target appearances (C) our model can *transfer and magnify* the left-right movement from the query to the target appearance with different magnification factors λ.

[1] Tae-Hyun Oh, Ronnachai Jaroensri, Changil Kim, Mohamed Elgharib, Frédo Durand, William T Freeman, and Wojciech Matusik. Learning-based video motion magnification. *In Proceedings of the European Conference on Computer Vision (ECCV)*

5 Take-Aways

- Introduction of the new task of *magnifying posture deviations across subjects* with applications in different fields
- Proposal of an unsupervised disentanglement loss $\mathcal{L}_{\text{dis}}$ for separating posture from appearance
- Proposal of a magnification loss $\mathcal{L}_{\text{mag}}$ for learning to generate realistic and valid magnifications

Project Page