Distributed Learning and Inference with Compressed Images

Sudeep Katakol, Basem Elbarahashy, Luis Herranz, Joost van de Weijer, Antonio M. López

Computer Vision Center, Universitat Autònoma de Barcelona

IEEE Transactions on Image Processing, 2021

CVC

UAB Universitat Autònoma de Barcelona

Data collection for onboard perception

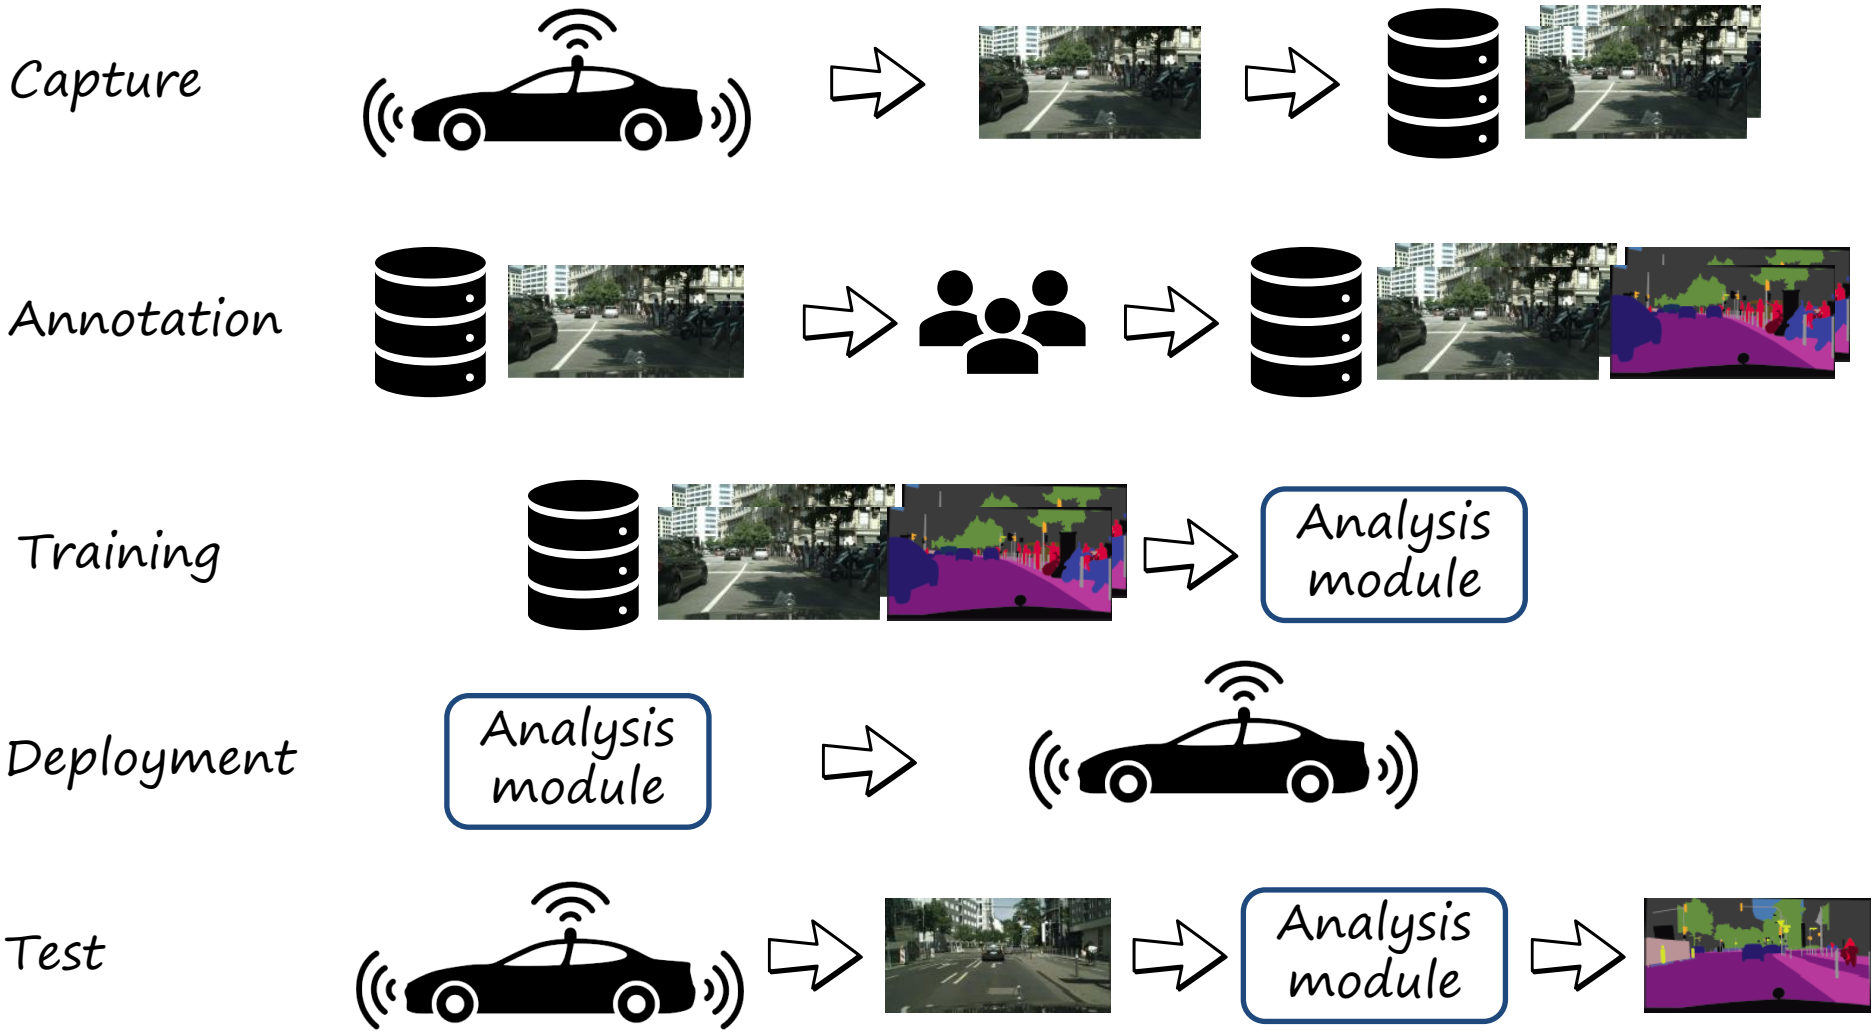

Data collection for onboard perception

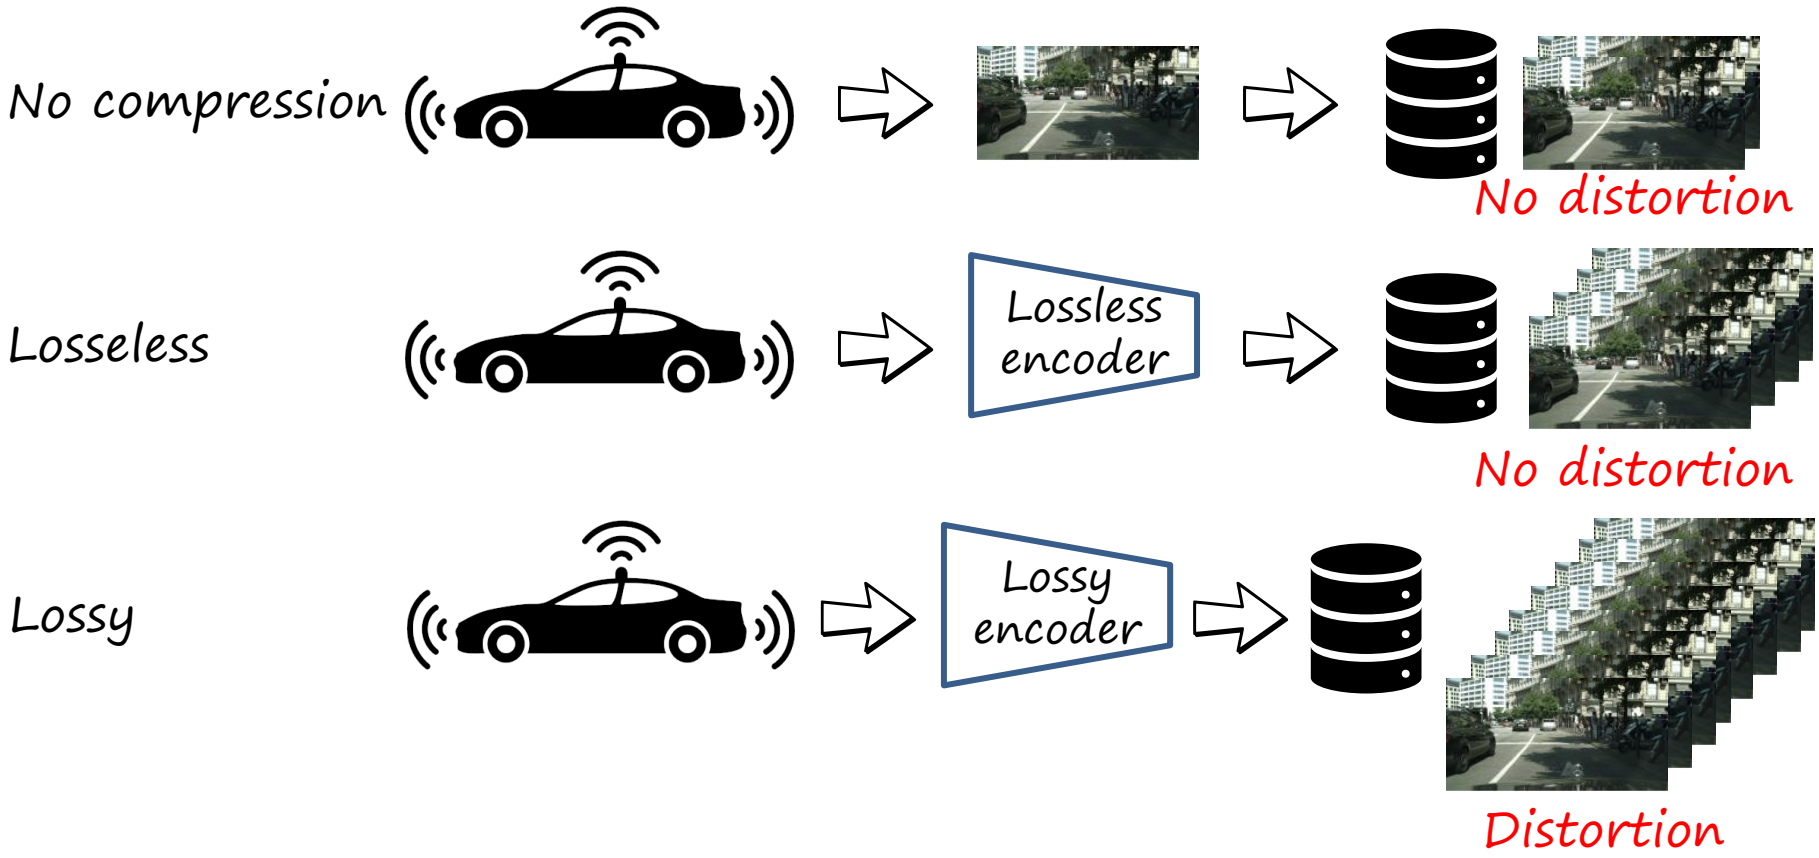

Distributed data collection

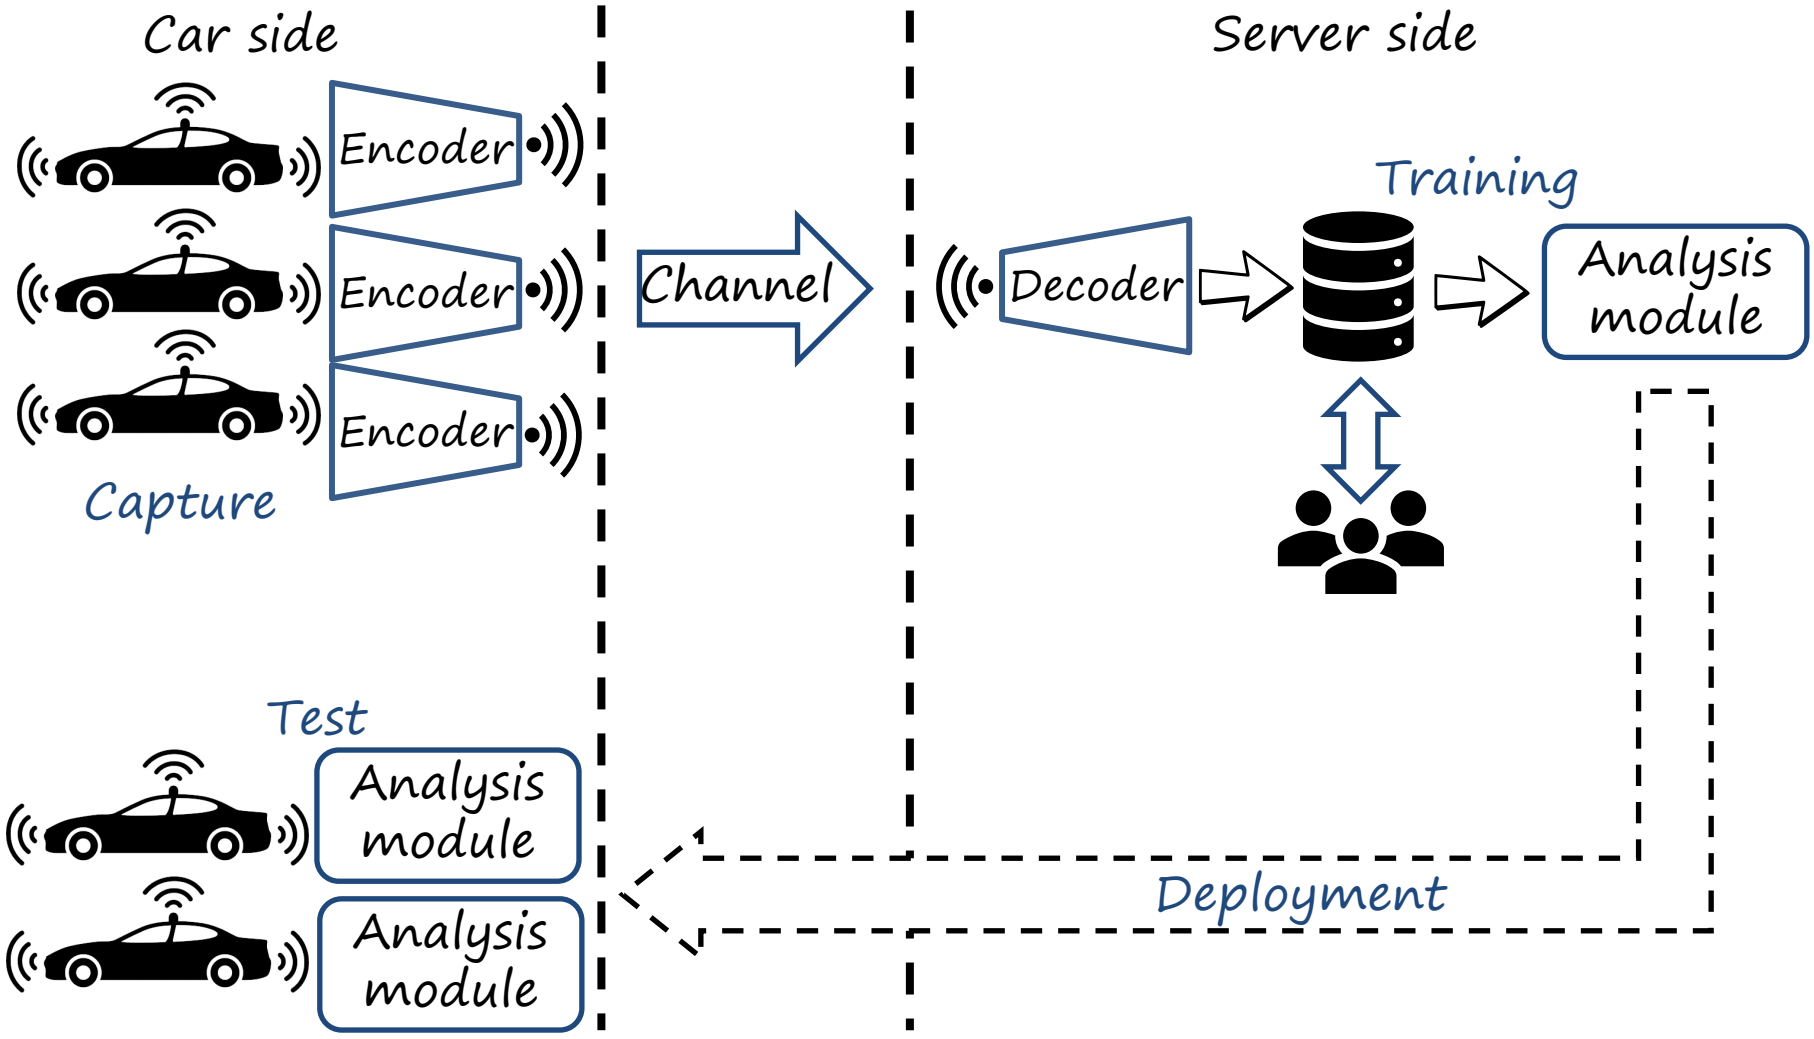

Distributed data collection

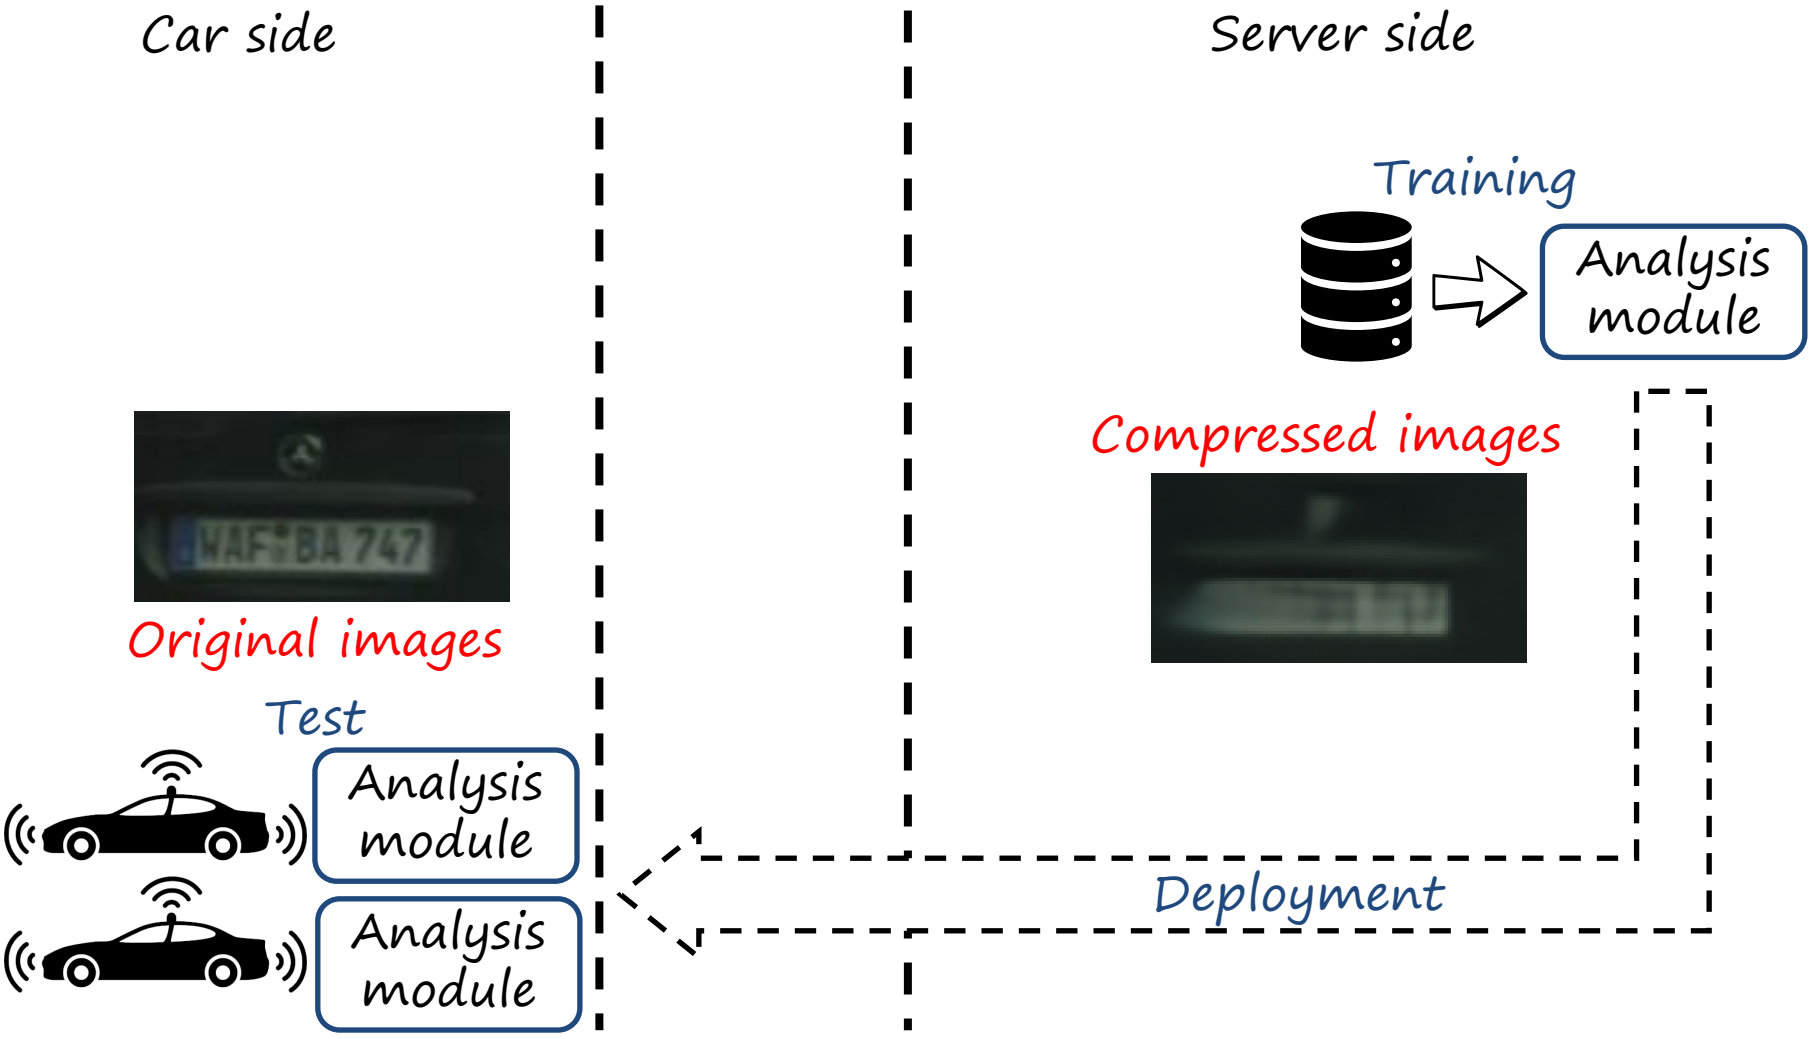

Training images vs test images

Training (original)

Test (compressed)

Training images vs test images

Training (compressed)

Test (original)

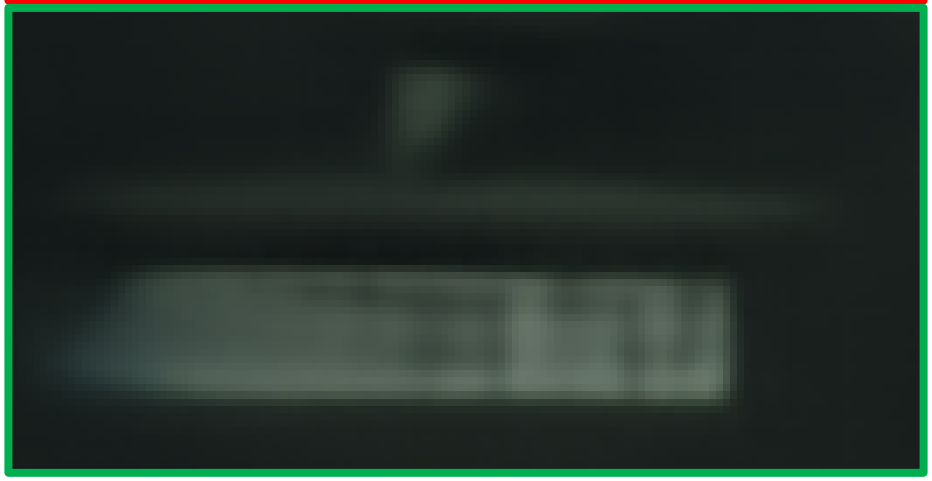

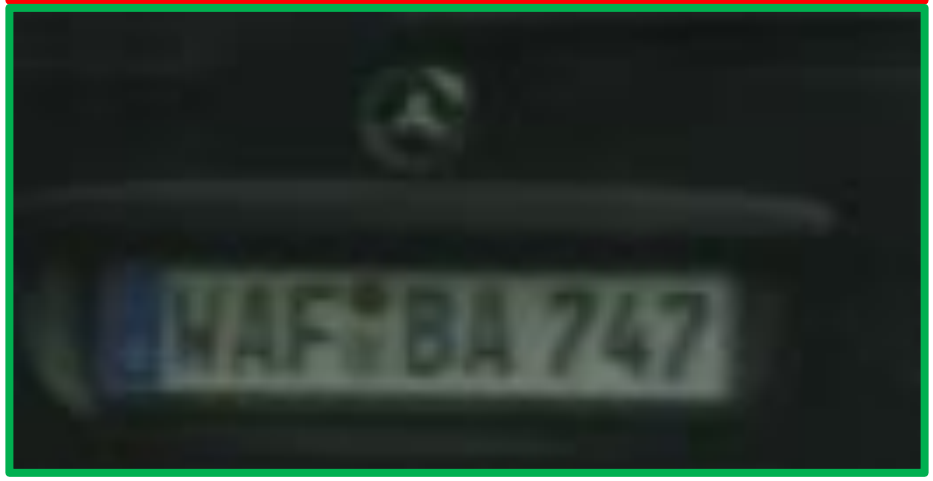

Training images vs test images

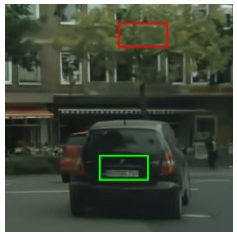

Training (compressed) Test (original)

Configuration CO:
compressed/original

Observation 1: training and test distributions are different (**covariate shift**)

Observation 2: training images have less information than test images (**loss of information**)

Training/test configurations

	Training	Test	Covariate shift	Information loss (training/test)
OO (ideal) original/original			No	No/No
CO compressed/ original			Yes	Yes/No
CC compressed/compressed			No	Yes/Yes
OC original/compressed			Yes	No/Yes

Effect on downstream task

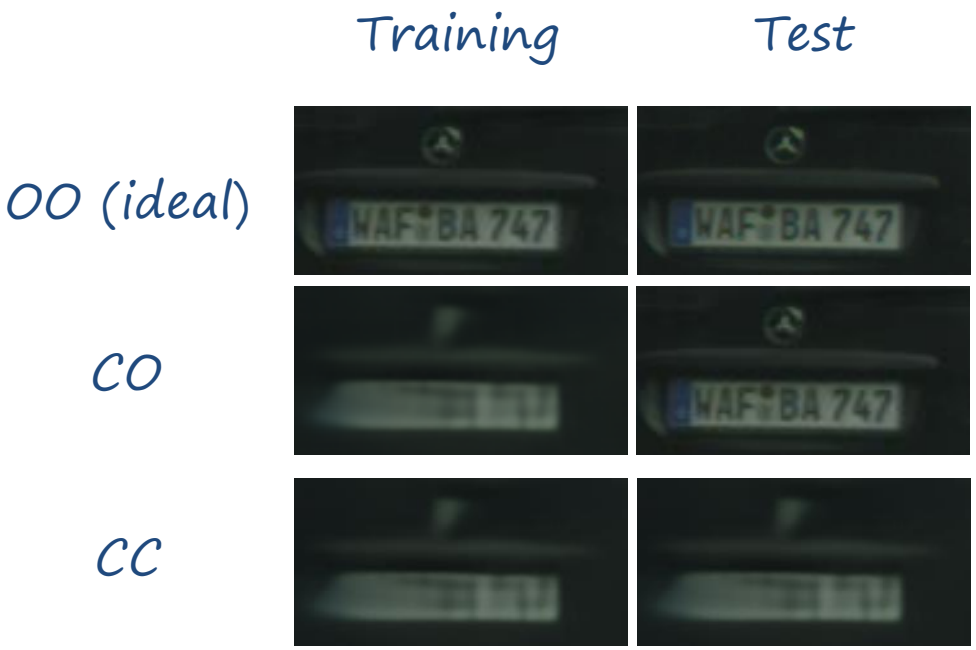

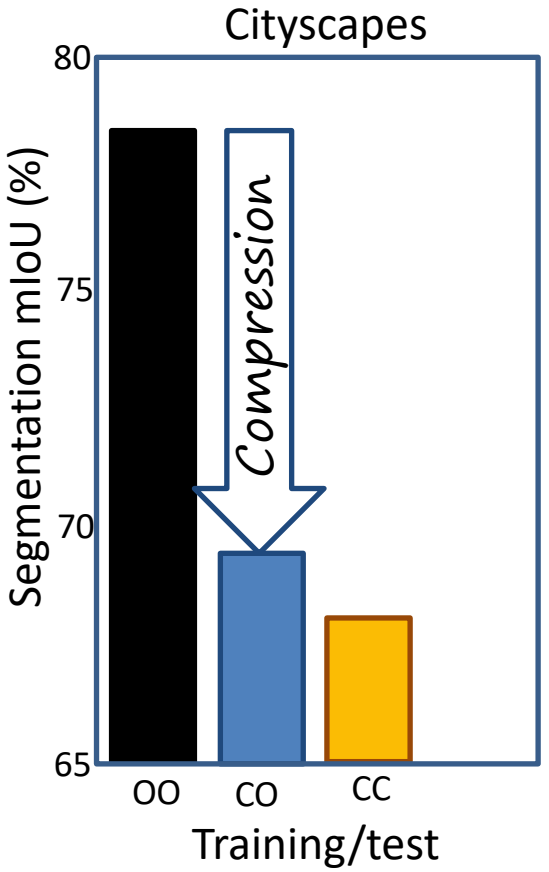

Conclusion (this dataset): better to keep more information in test than reduce the covariate shift

Proposed approach: dataset restoration

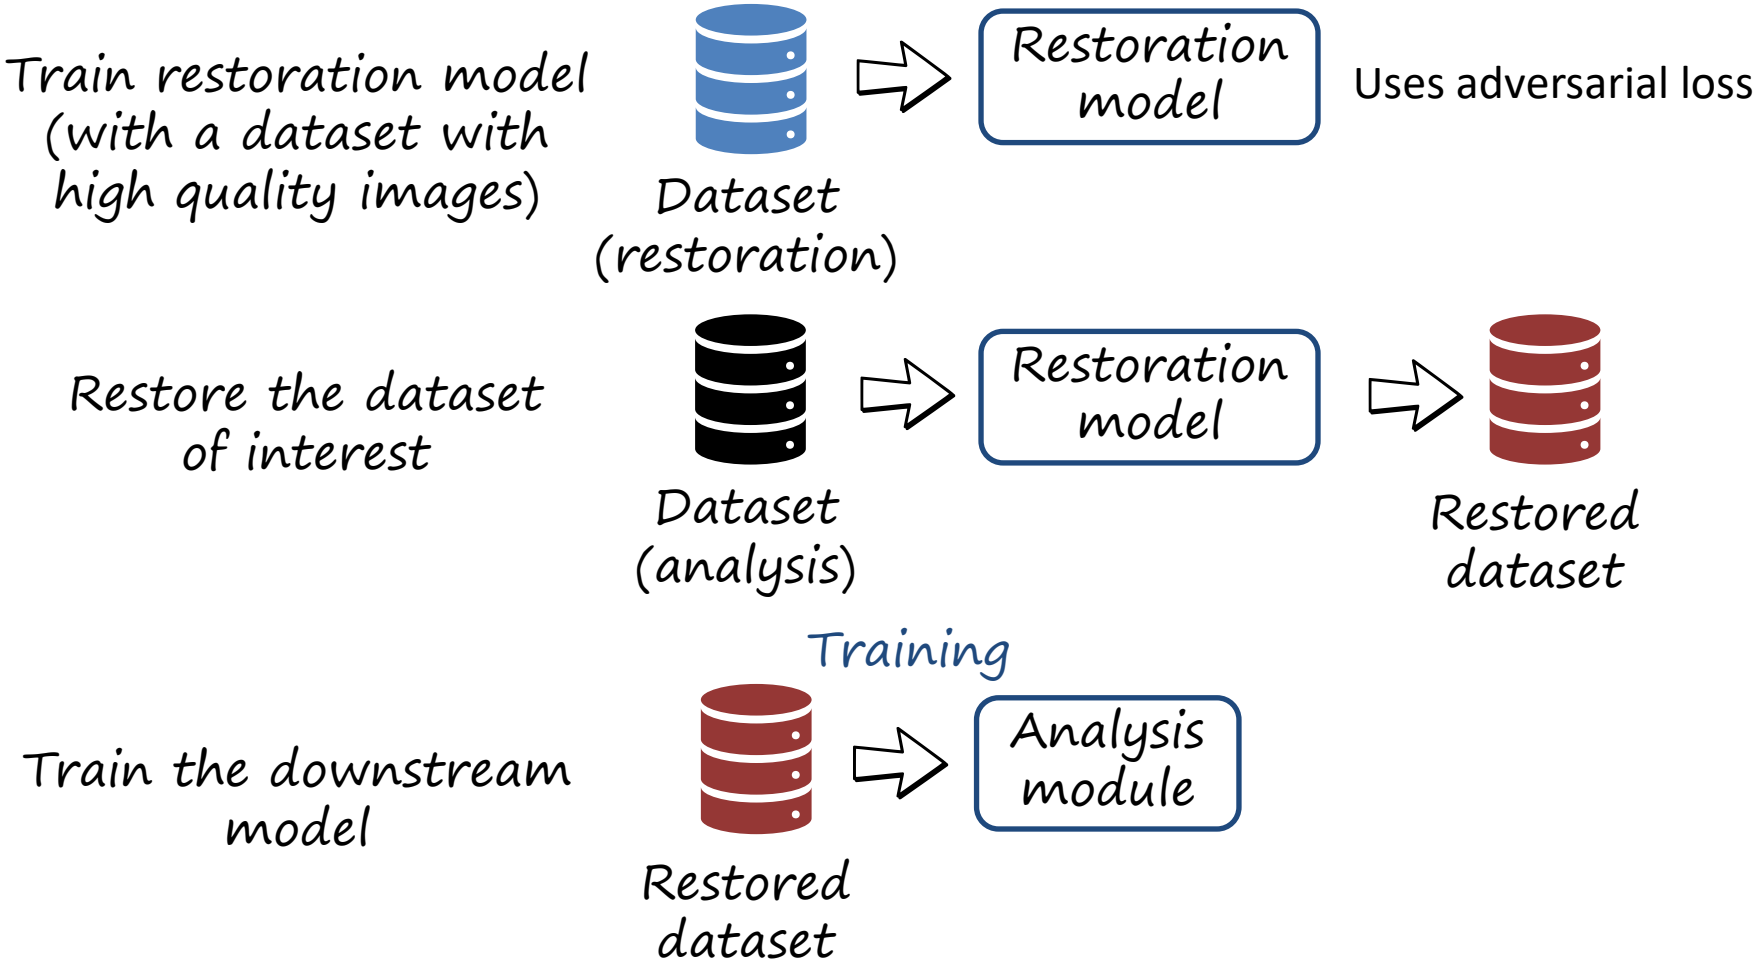

Training images vs test images

Original

Compressed

Restored

Training images vs test images

Original *Compressed* *Restored*





Effect on downstream task

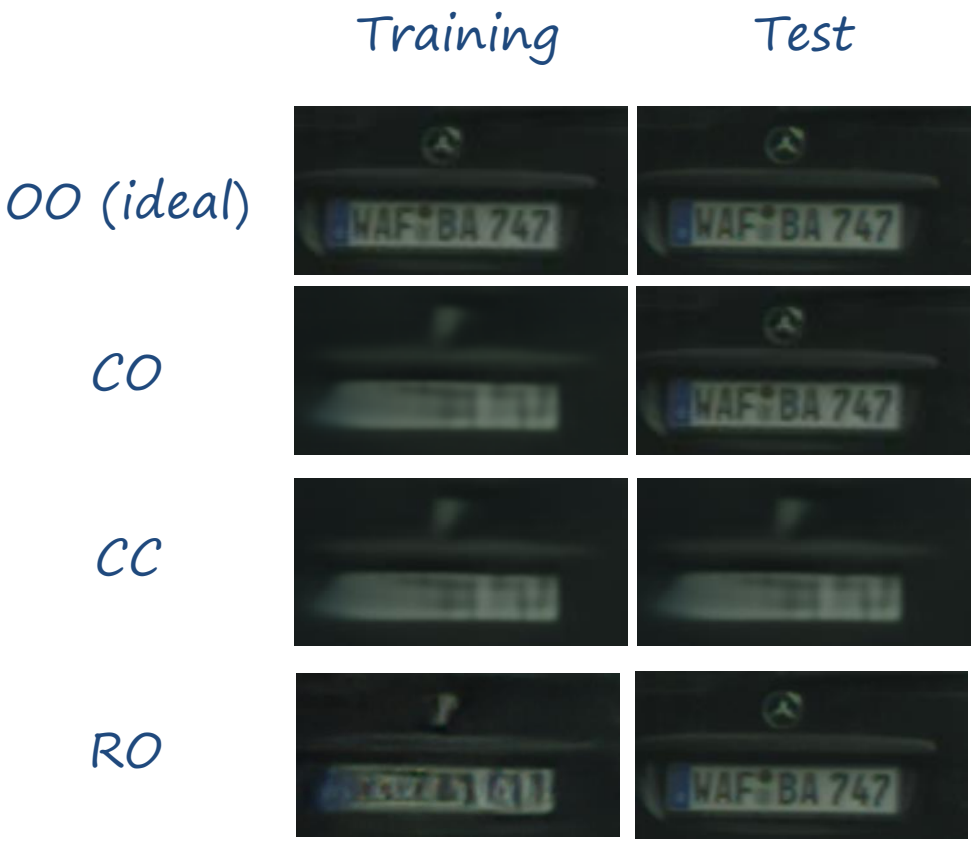

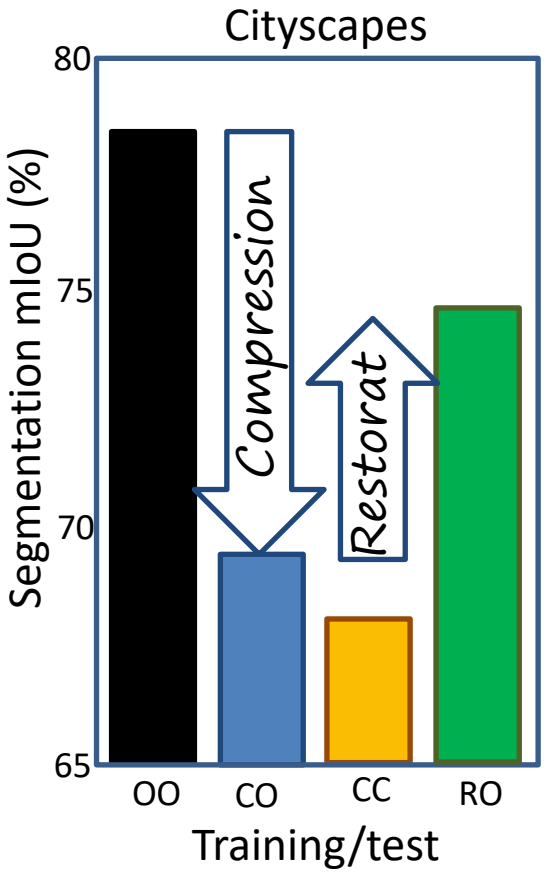

Why does it work?

- Alleviates the covariate shift
- Keeps useful information for segmentation (e.g. texture)

Experiments. Rate-distortion

Dataset: Cityscapes. Codecs: BPG (traditional), MSH (neural)

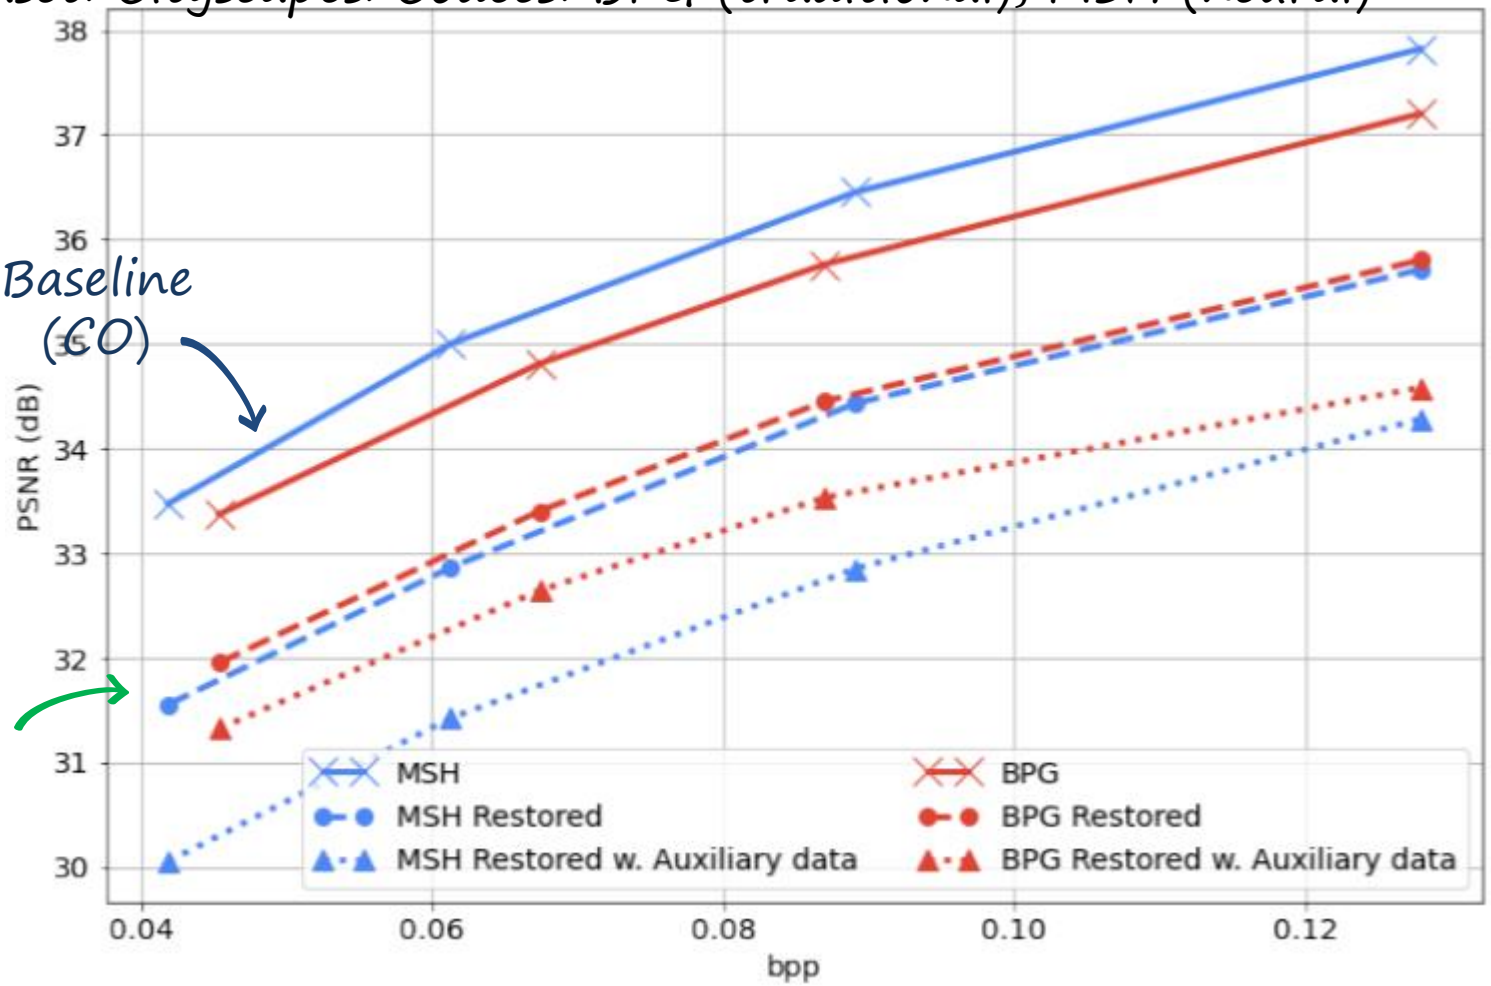

Baseline (CO)

Restoration (RO)

Restoration harms R-D performance

Experiments. Rate-distortion

Dataset: Cityscapes. Codecs: BPG (traditional), MSH (neural)

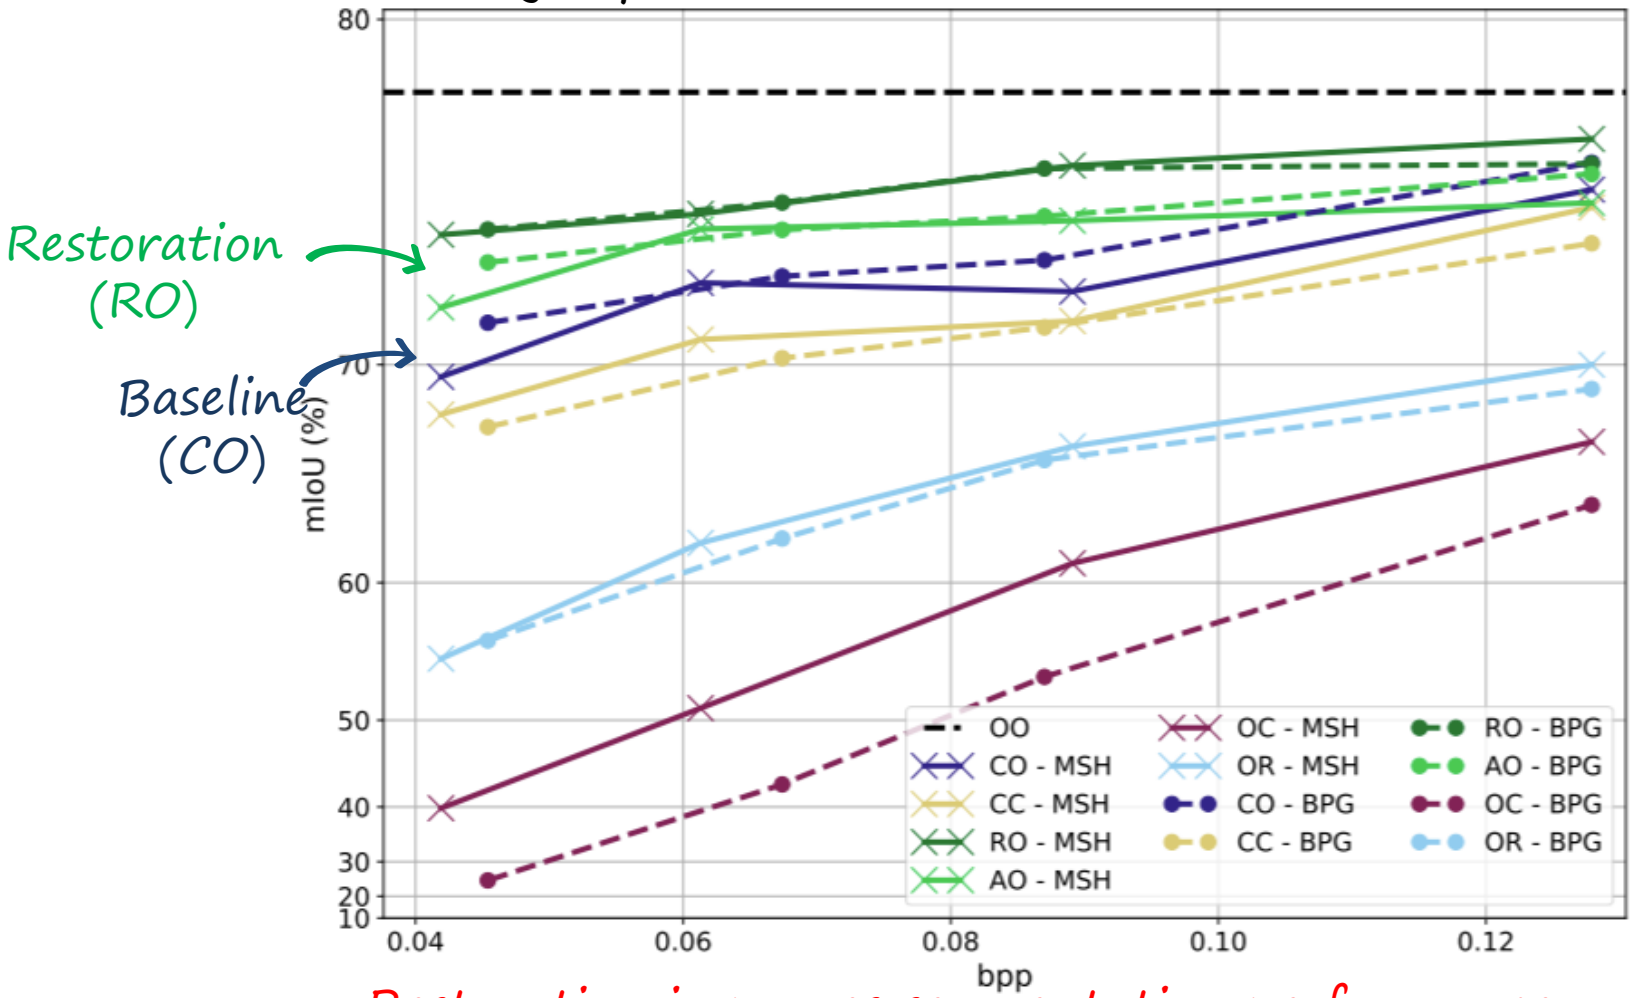

Restoration improves segmentation performance

Adversarial vs non-adversarial restoration

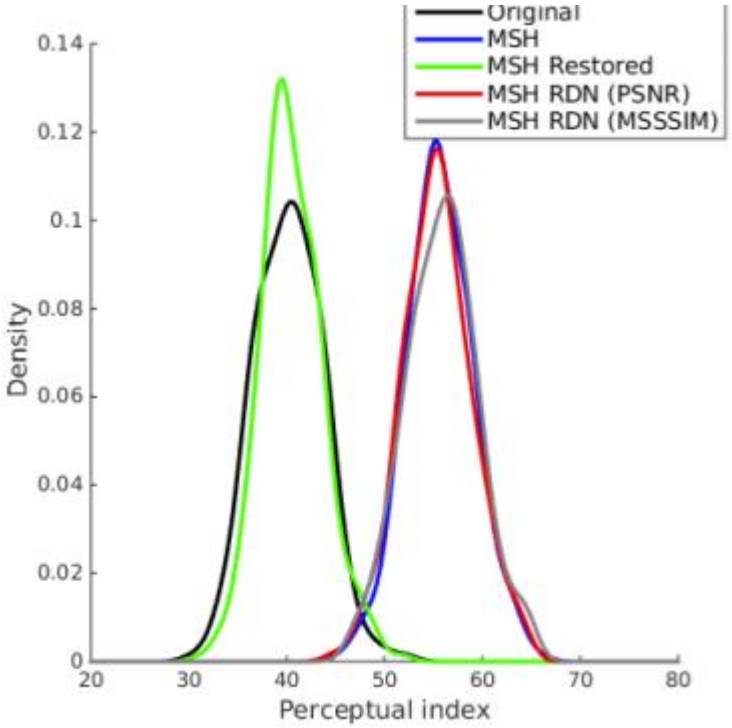

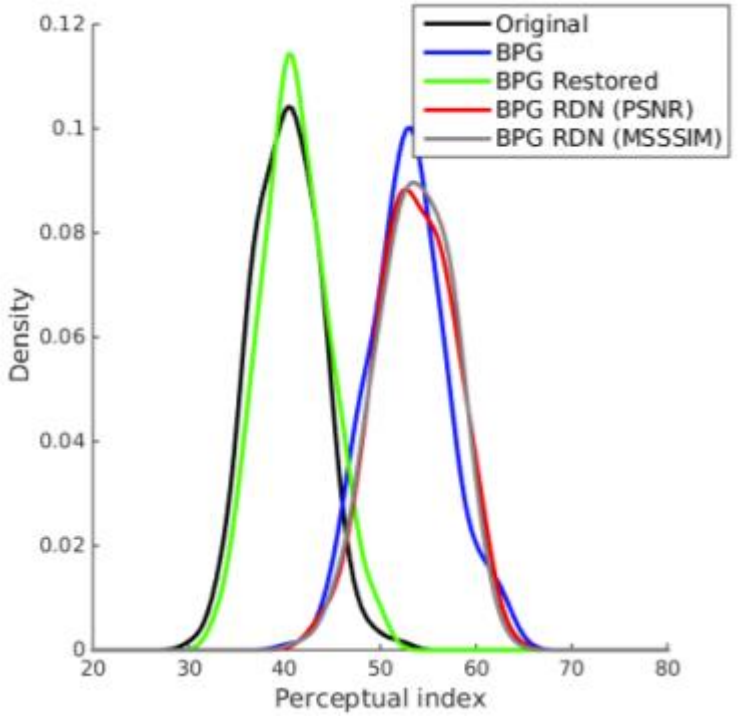

Restoration must be adversarial

Perception-distortion tradeoff

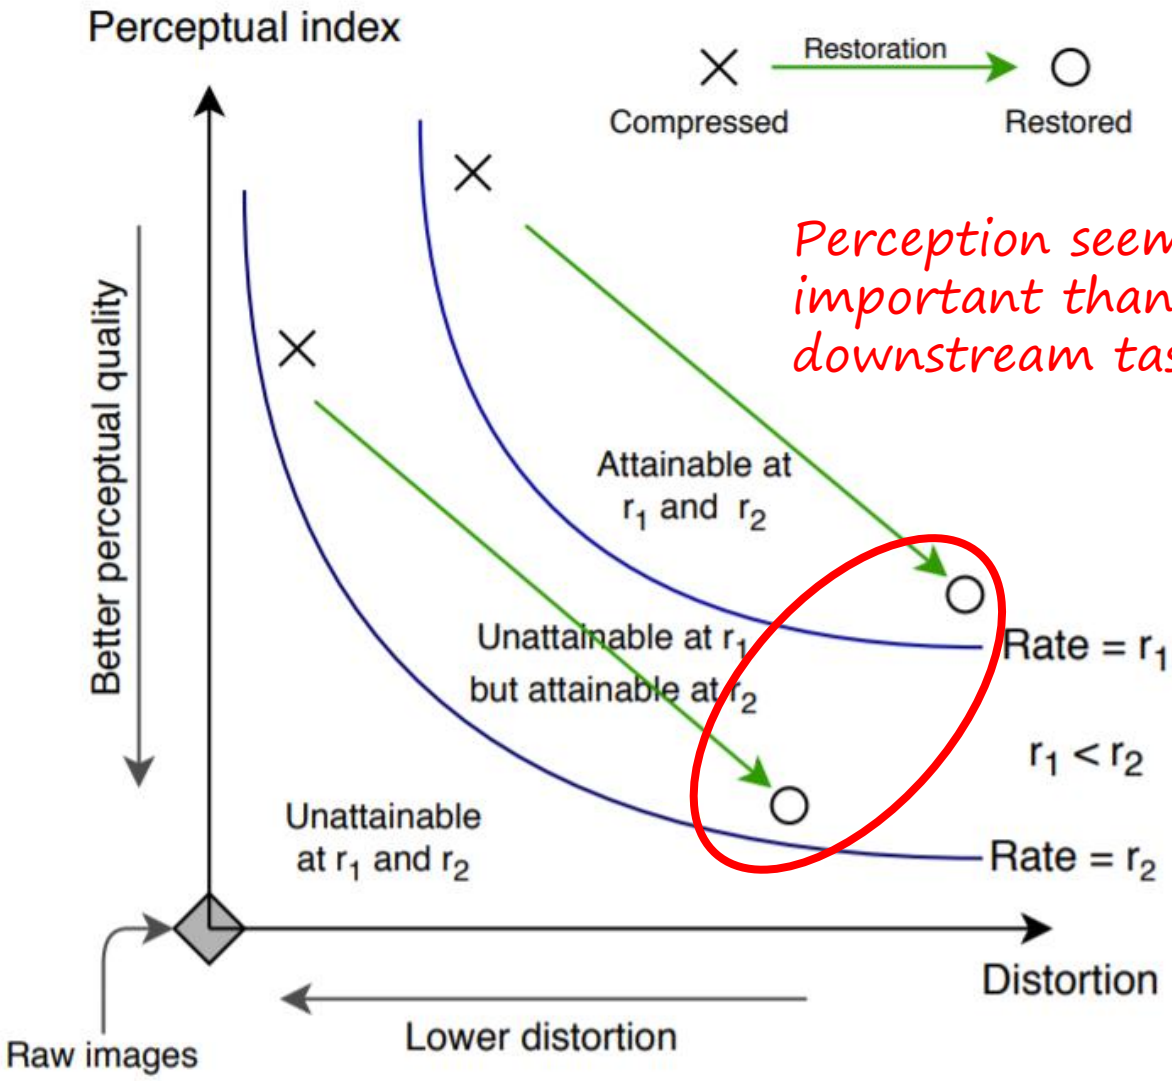

Perception seems to be more important than distortion for downstream tasks

Cost of collecting data

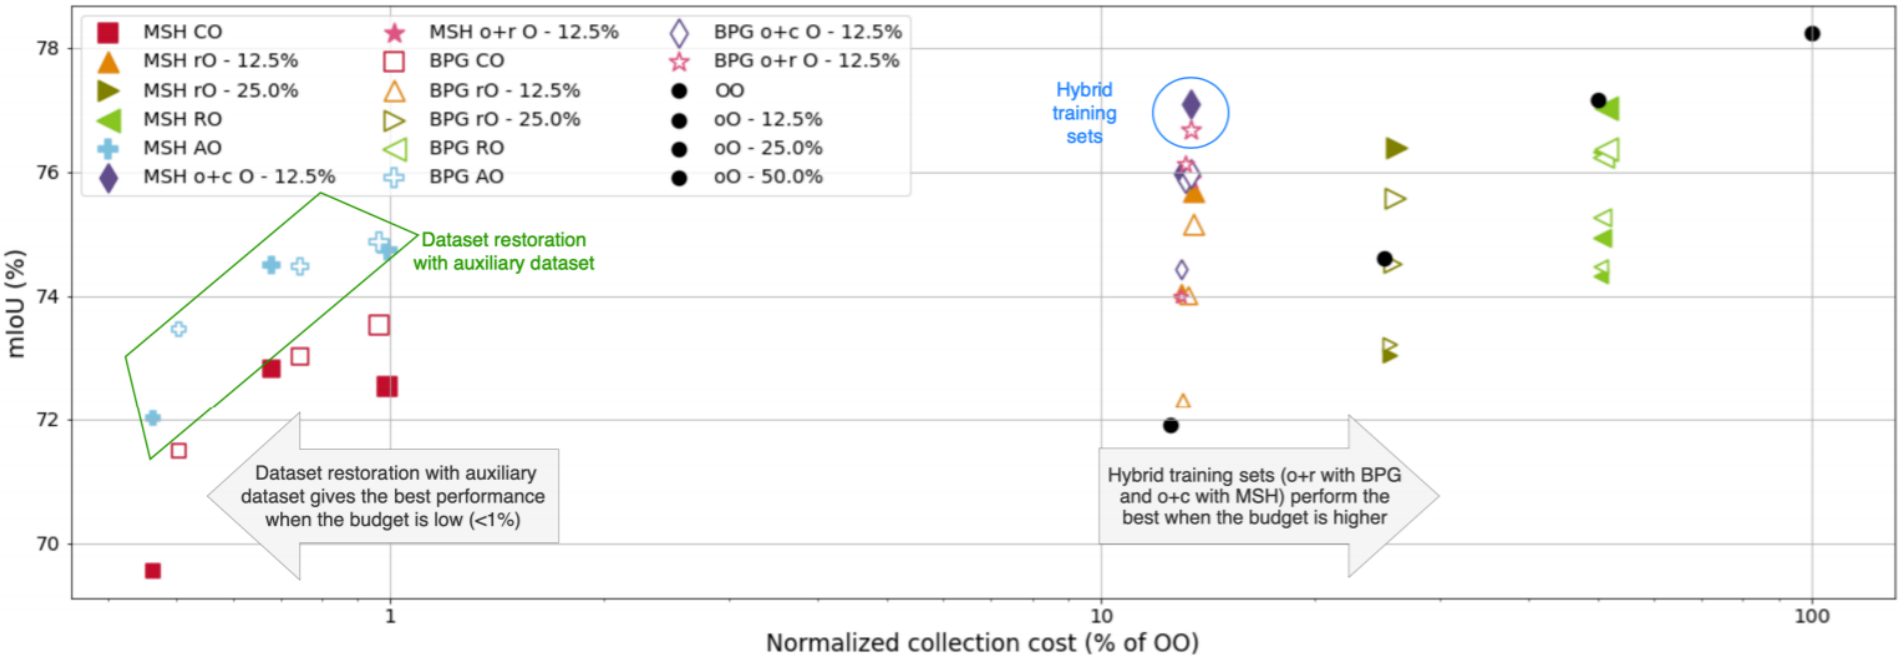

Thanks!

S. Katakol, B. Elbarashy, L. Herranz, J. van de Weijer, A. M. Lopez, “Distributed Learning and Inference with Compressed Images”, *IEEE Transactions on Image Processing*, 2021

https://arxiv.org/abs/2004.10497

CVC

UAB Universitat Autònoma de Barcelona