Learning Resolution-Invariant Deep Representations for Person Re-Identification

Yun-Chun Chen[1*]   Yu-Jhe Li[1,2*]   Xiaofei Du[3]   Yu-Chiang Frank Wang[1,2]

[1]National Taiwan University   [2]MOST Joint Research Center for AI Technology and All Vista Healthcare   [3]Umbo Computer Vision

## Cross-Resolution Person Re-ID

- **Goal**: Identify images of same person across camera views
- **Challenges**:
  - Viewpoint variations
  - Resolution issues due to person-camera distance

Long distance → **Low** resolution   **Short** Distance → **High** resolution
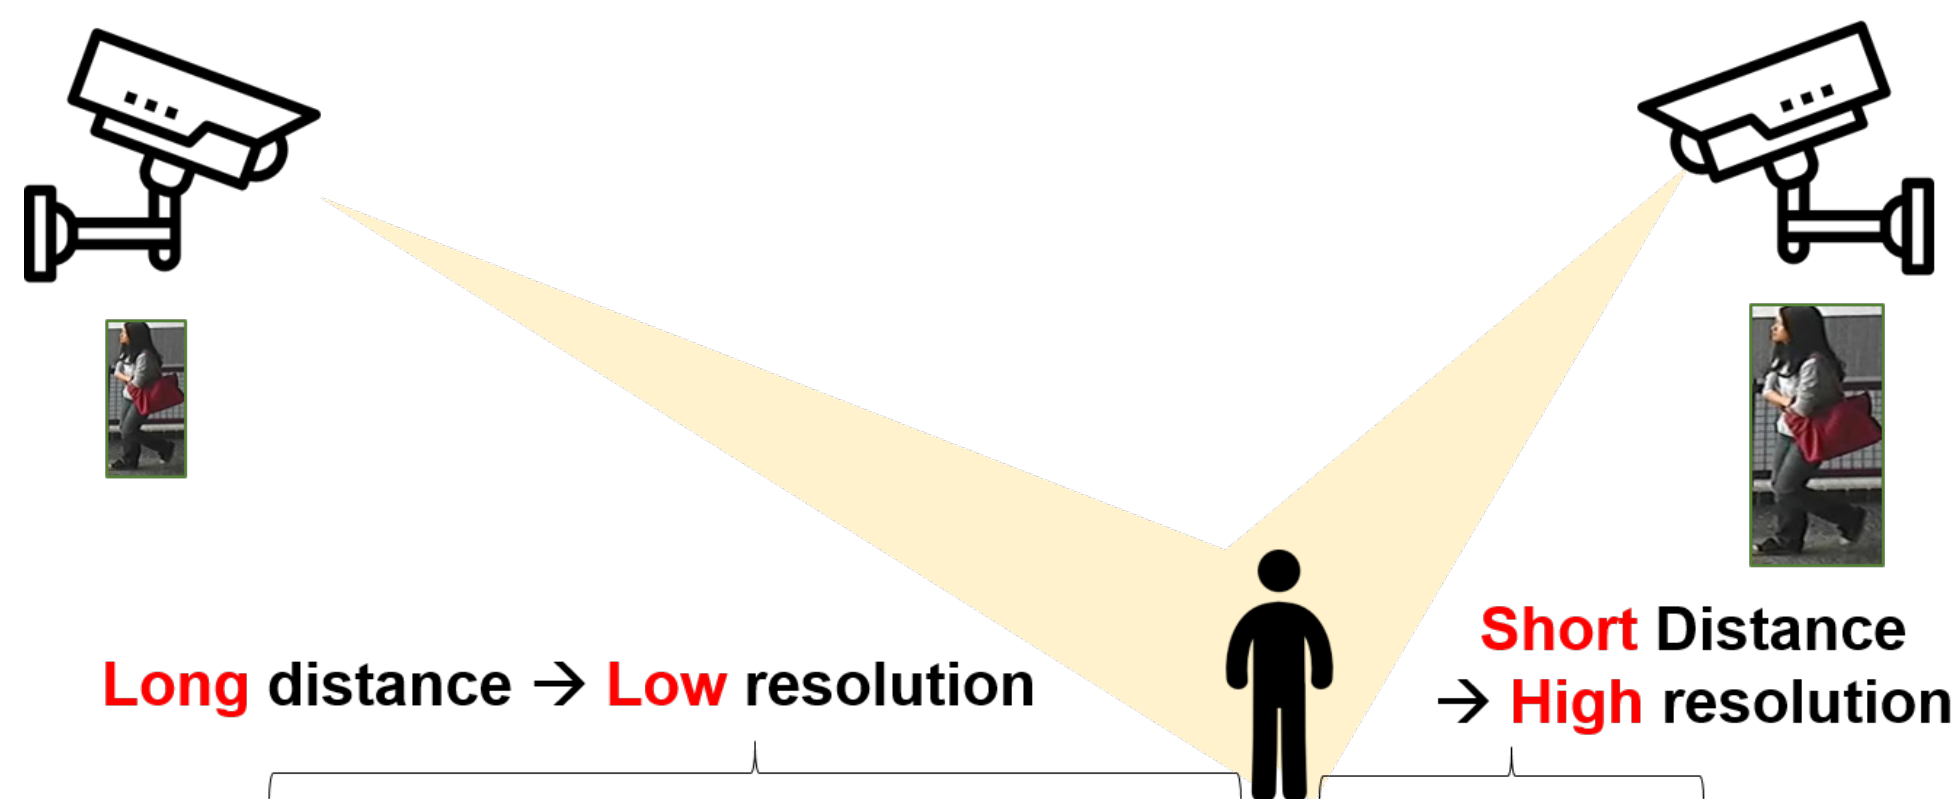
Contributions

- Propose an end-to-end trainable network that learns **resolution-invariant** representations

- Develop a **multi-level** adversarial network that effectively aligns feature representations across resolutions.

- Can handle images of a wide range of (and even **unseen**) low resolutions.

- Perform favorably against the state-of-the-art methods.

- Extensible to practical cross-resolution re-ID tasks under **semi-supervised** settings.

Resolution Adaptation and re-Identification Network (RAIN)
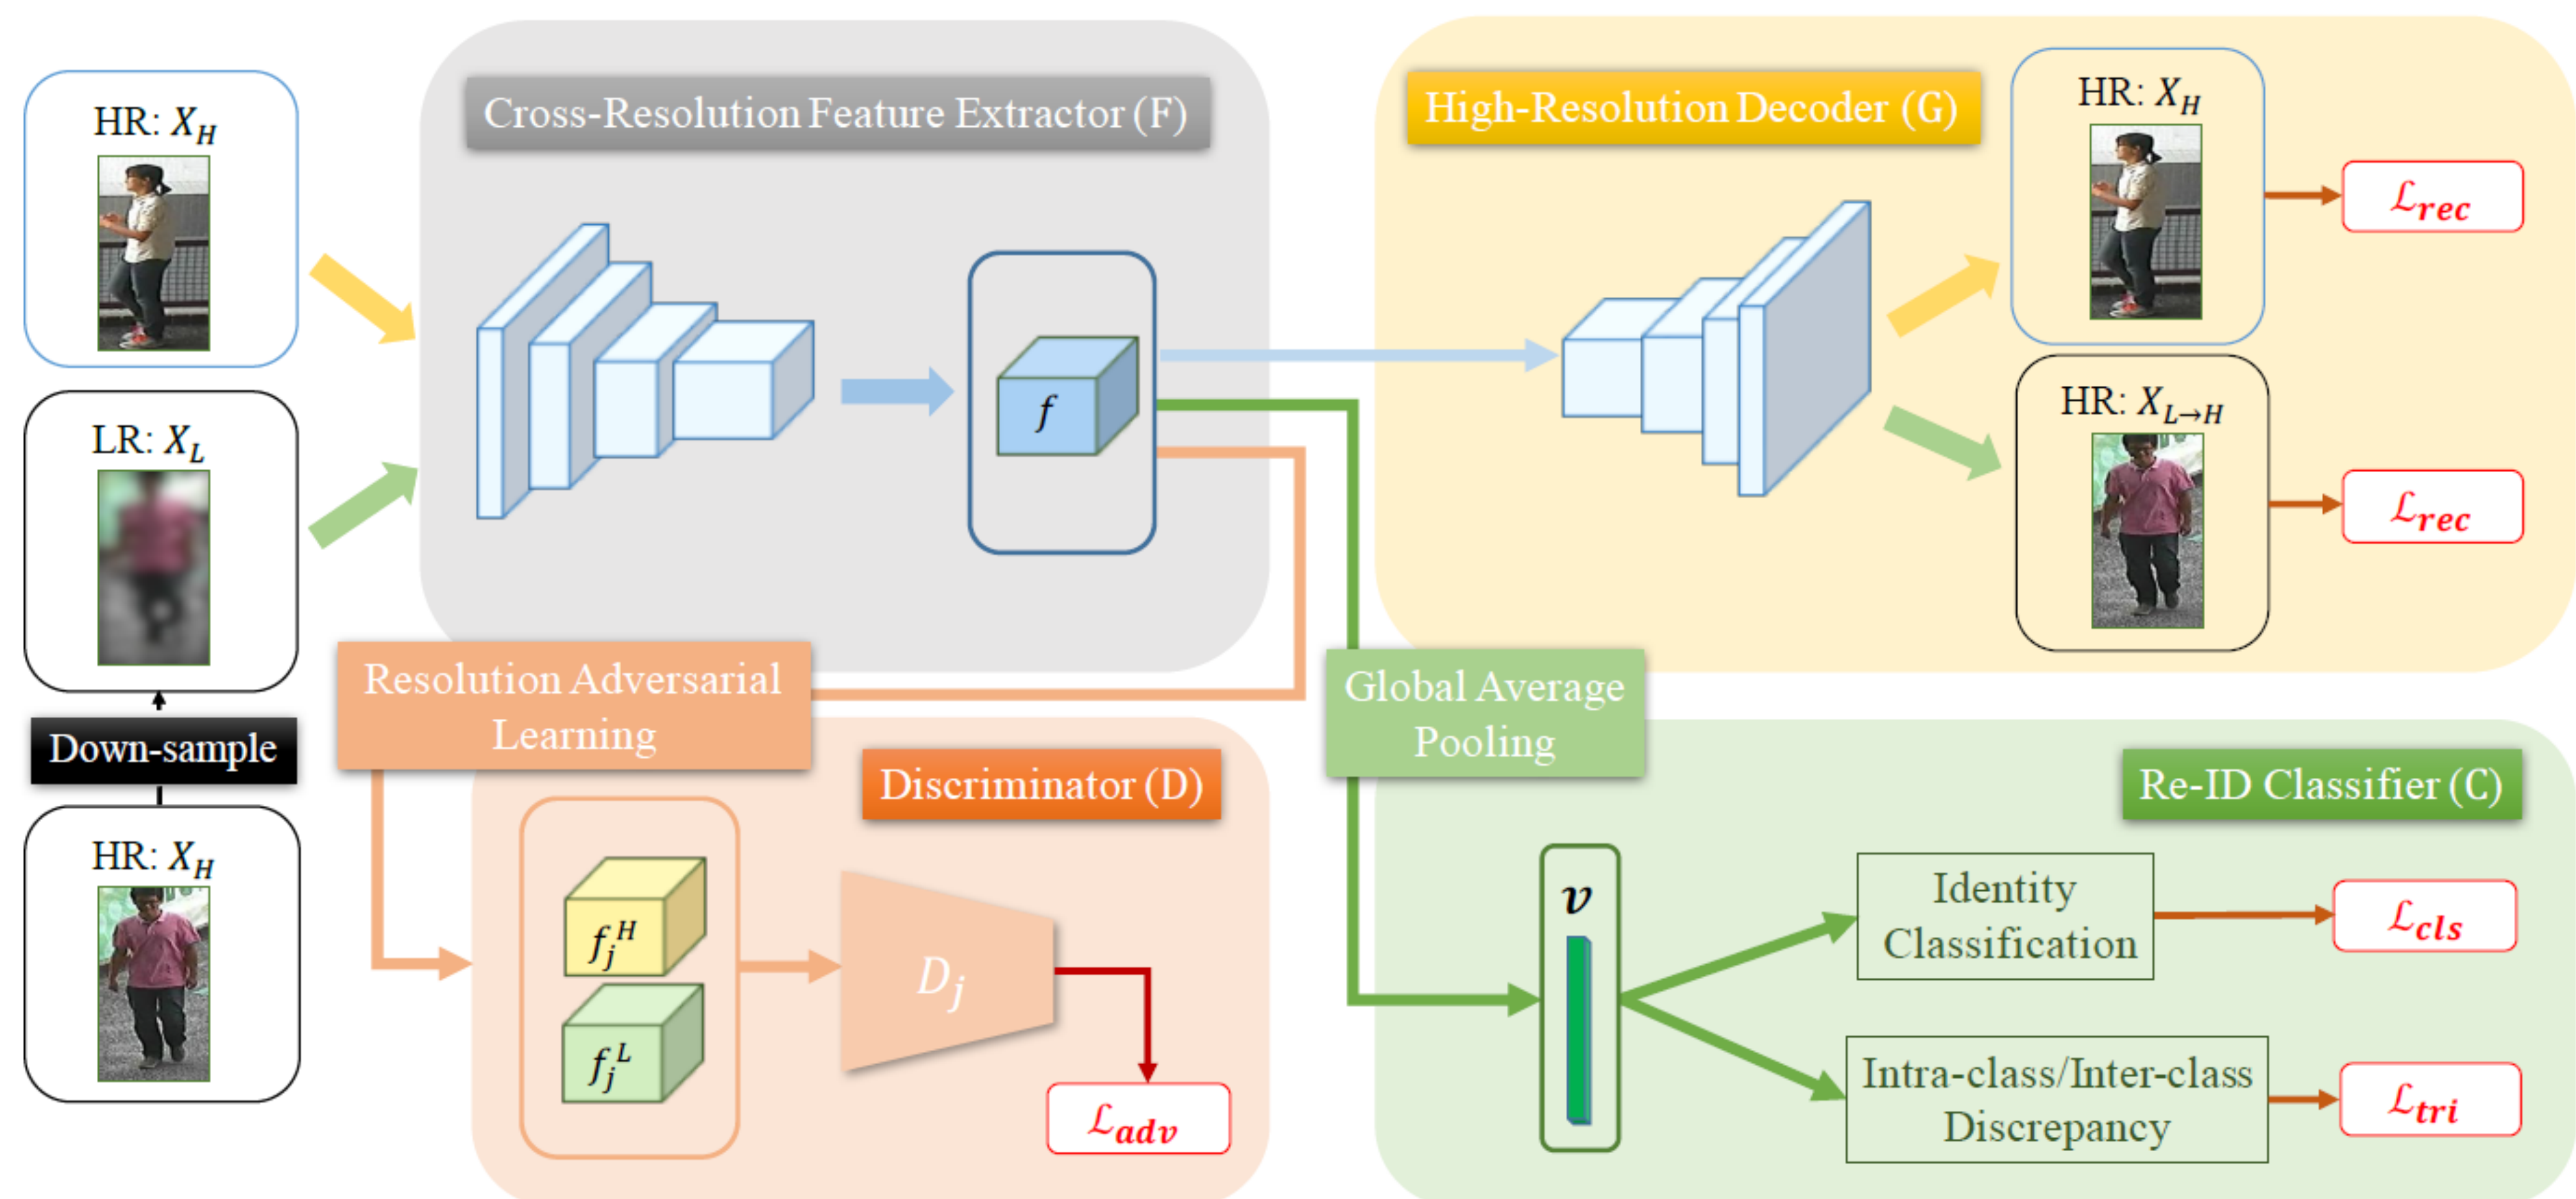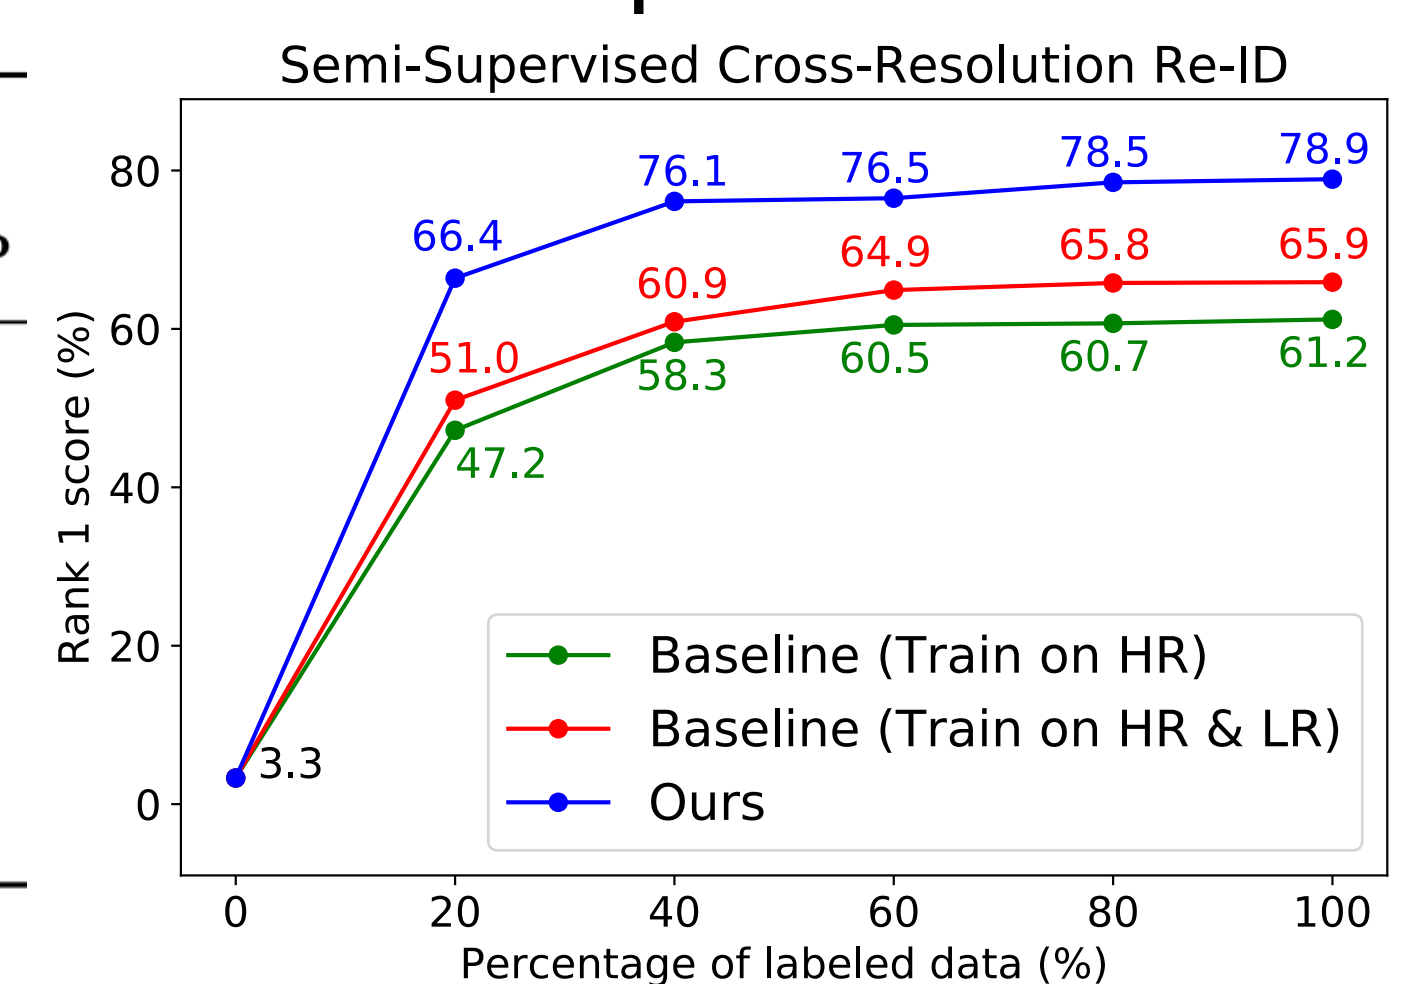

Experimental Results

- ### Results on three standard benchmarks

Method	MLR-CUHK03			MLR-VIPeR			CAVIAR		
	Rank 1	Rank 5	Rank 10	Rank 1	Rank 5	Rank 10	Rank 1	Rank 5	Rank 10
JUDEA [1]	26.2	58.0	73.4	26.0	55.1	69.2	22.0	60.1	80.8
SLD²L [2]	-	-	-	20.3	44.0	62.0	18.4	44.8	61.2
SDF [3]	22.2	48.0	64.0	9.25	38.1	52.4	14.3	37.5	62.5
SING [4]	67.7	90.7	94.7	33.5	57.0	66.5	33.5	72.7	89.0
CSR-GAN [5]	71.3	92.1	97.4	37.2	62.3	71.6	34.7	72.5	87.4
Baseline (HR)	60.6	89.4	95.0	32.5	59.2	69.0	27.5	63.2	79.3
Baseline (HR & LR)	65.9	92.1	97.4	36.6	62.3	70.9	31.7	68.4	84.2
Ours (single-level)	77.6	96.2	98.5	41.2	66.3	75.6	41.5	75.3	85.6
Ours (multi-level)	**78.9**	**97.3**	**98.7**	**42.5**	**68.3**	**79.6**	**42.0**	**77.3**	**89.6**

- ### Ablation study

Method	MLR-CUHK03		
	Rank 1	Rank 10	mAP
Ours	**78.9**	**98.7**	**74.5**
Ours w/o $\mathcal{L}_{cls}$	70.8	97.7	68.0
Ours w/o $\mathcal{L}_{tri}$	69.1	96.6	64.1
Ours w/o $\mathcal{L}_{rec}$	67.3	94.5	64.2
Ours w/o $\mathcal{L}_{adv}$	65.9	97.4	62.3

- ### Semi-supervised results

Semi-Supervised Cross-Resolution Re-ID

- Baseline (Train on HR)
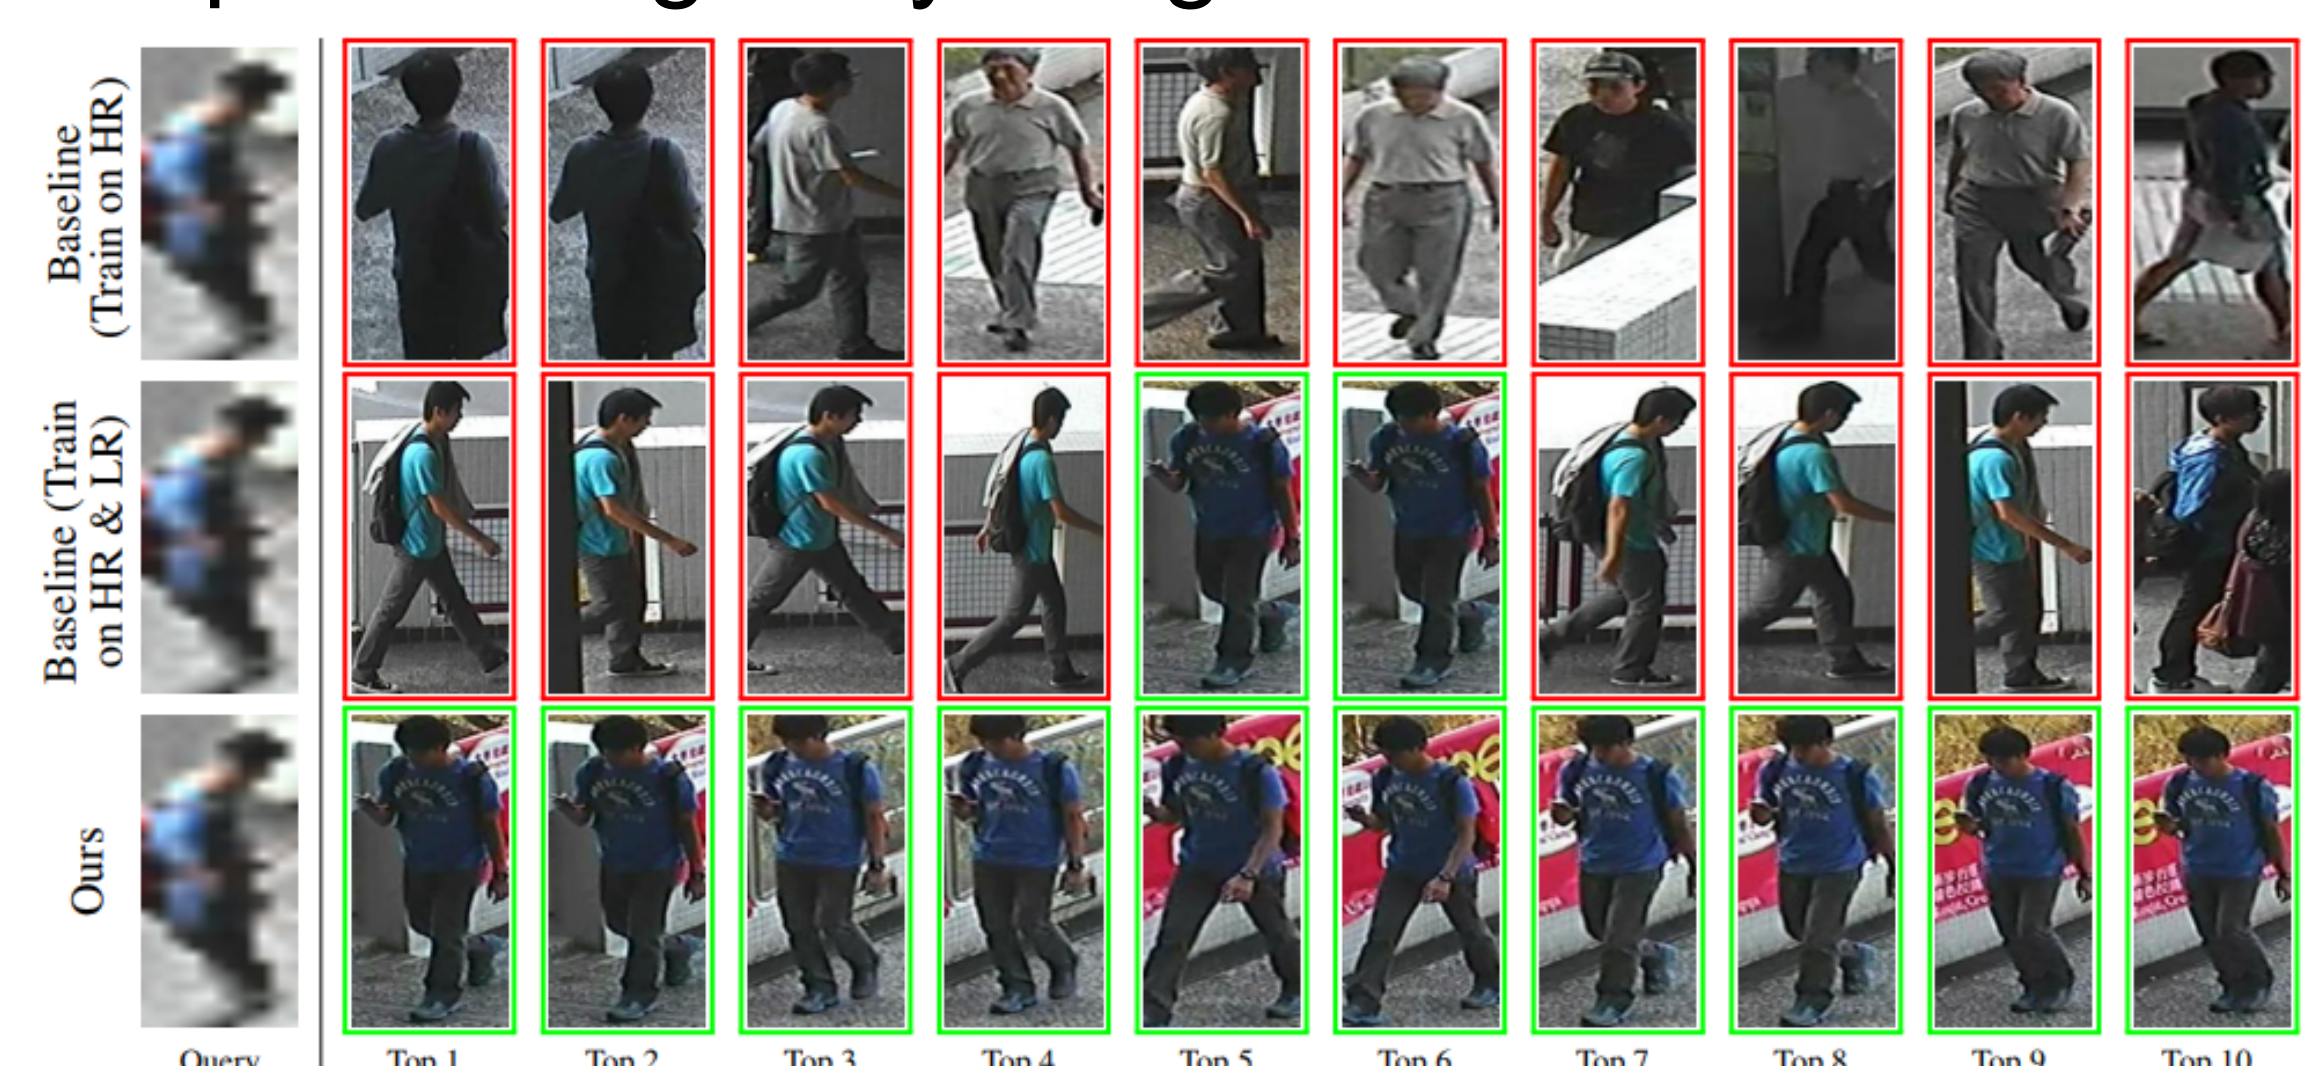- Baseline (Train on HR & LR)
- Ours

- ### Resolution-invariant feature vector v
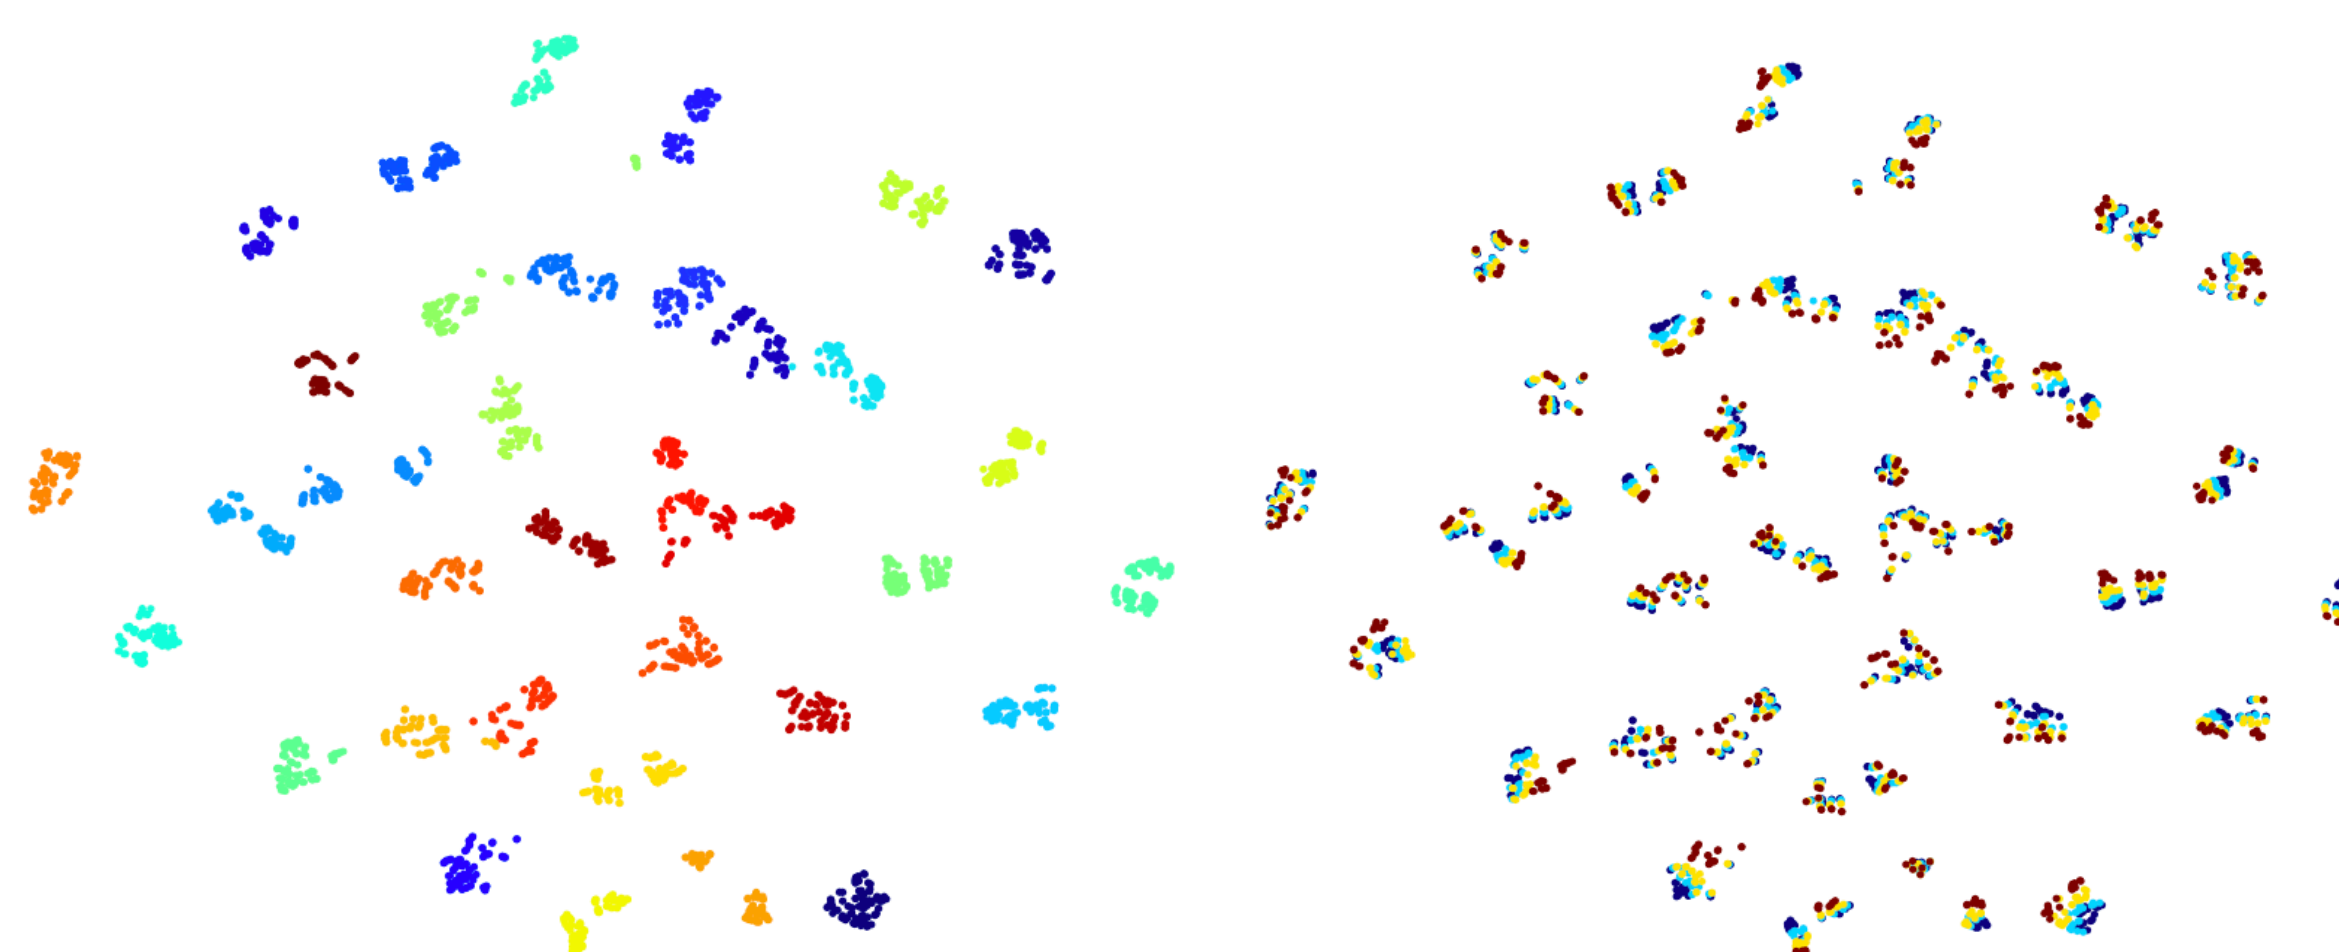
(a) Colorization with respect to **identity**.   (b) Colorization with respect to **resolution**.

- ### Top ranked gallery images

Reference

[1] Li et al. Multi-Scale Learning for Low-Resolution Person Re-Identification. In *ICCV*, 2015.
[2] Jing et al. Super-Resolution Person Re-Identification with Semi-Coupled Low-Rank Discriminant Dictionary Learning. In *CVPR*, 2015.
[3] Wang et al. Scale-Adaptive Low-Resolution Person Re-Identification via Learning a Discriminating Surface. In *IJCAI*, 2016.
[4] Jiao et al. Deep Low-Resolution Person Re-Identification. In *AAAI*, 2018.
[5] Wang et al. Cascaded SRGAN for Scale-Adaptive Low Resolution Person Re-Identification. In *IJCAI*, 2018.