CVPR VIRTUAL JUNE 19-25

Continual Semantic Segmentation via Repulsion-Attraction of Sparse and Disentangled Latent Representations

Umberto Michieli and Pietro Zanuttigh

umberto.michieli@dei.unipd.it

University of Padova

Continual Learning in Semantic Segmentation

Our Focus: Class-Incremental Continual Learning in Semantic Segmentation

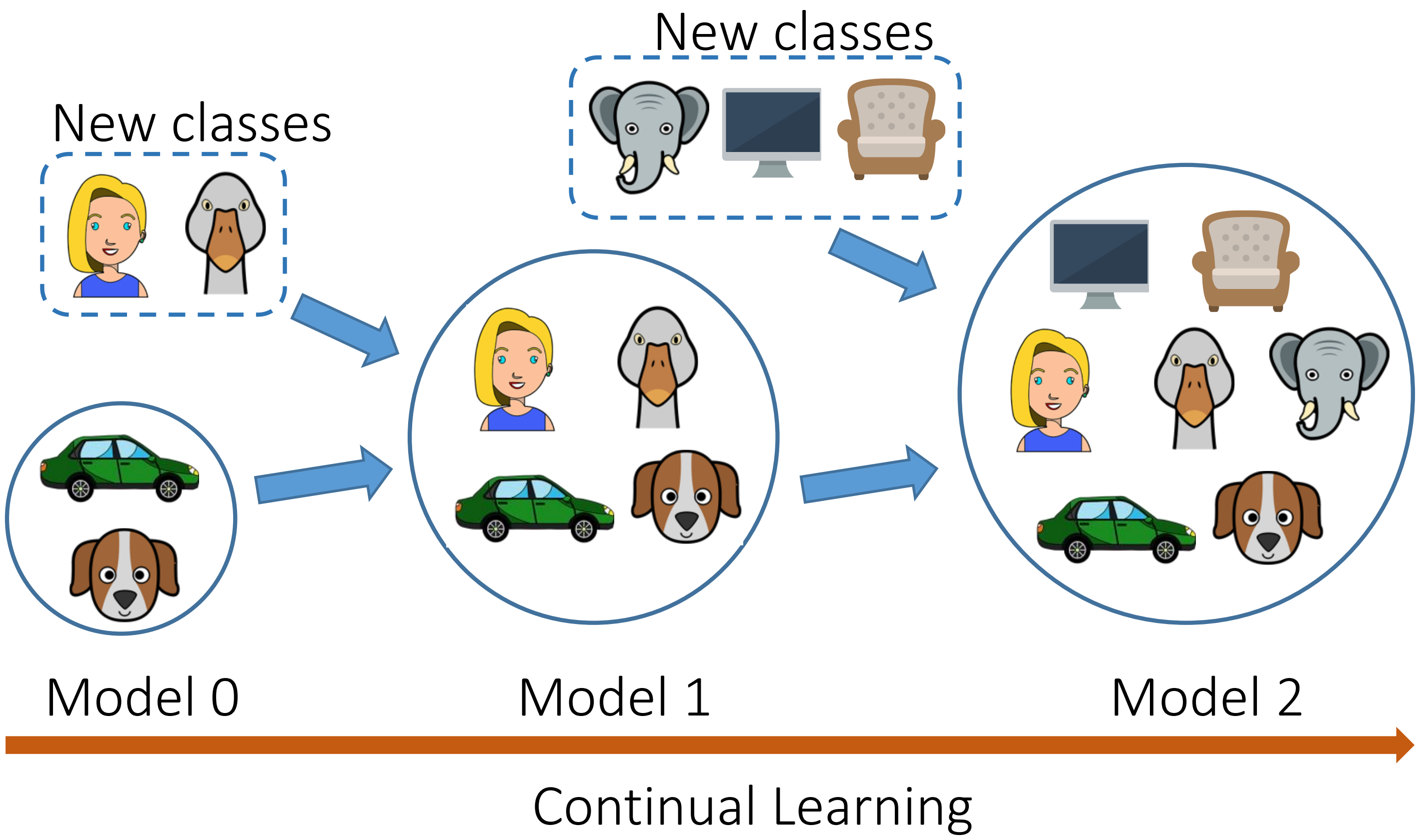

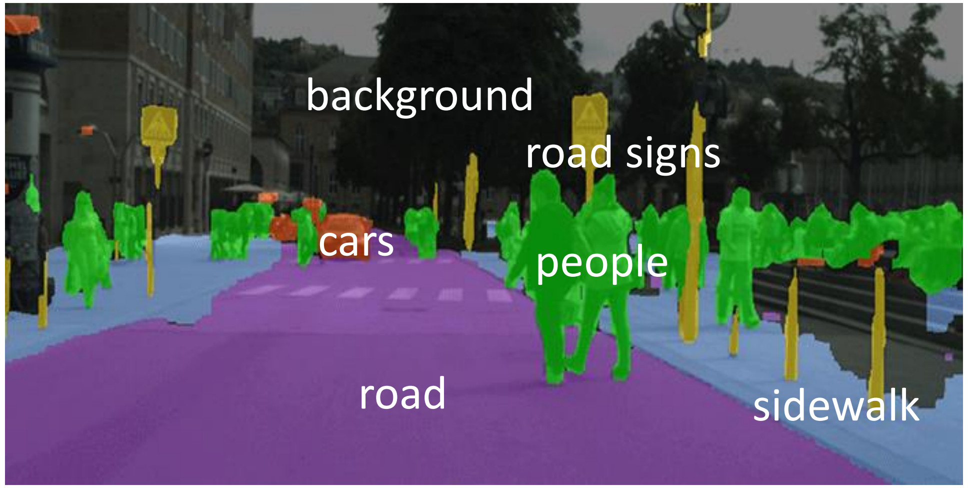

Continual Segmentation - Different Setups

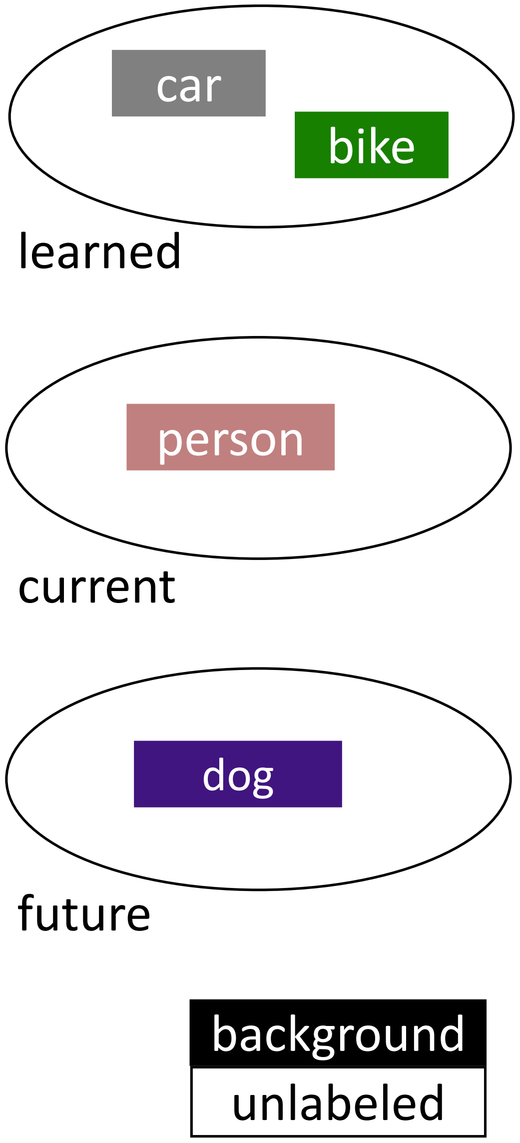

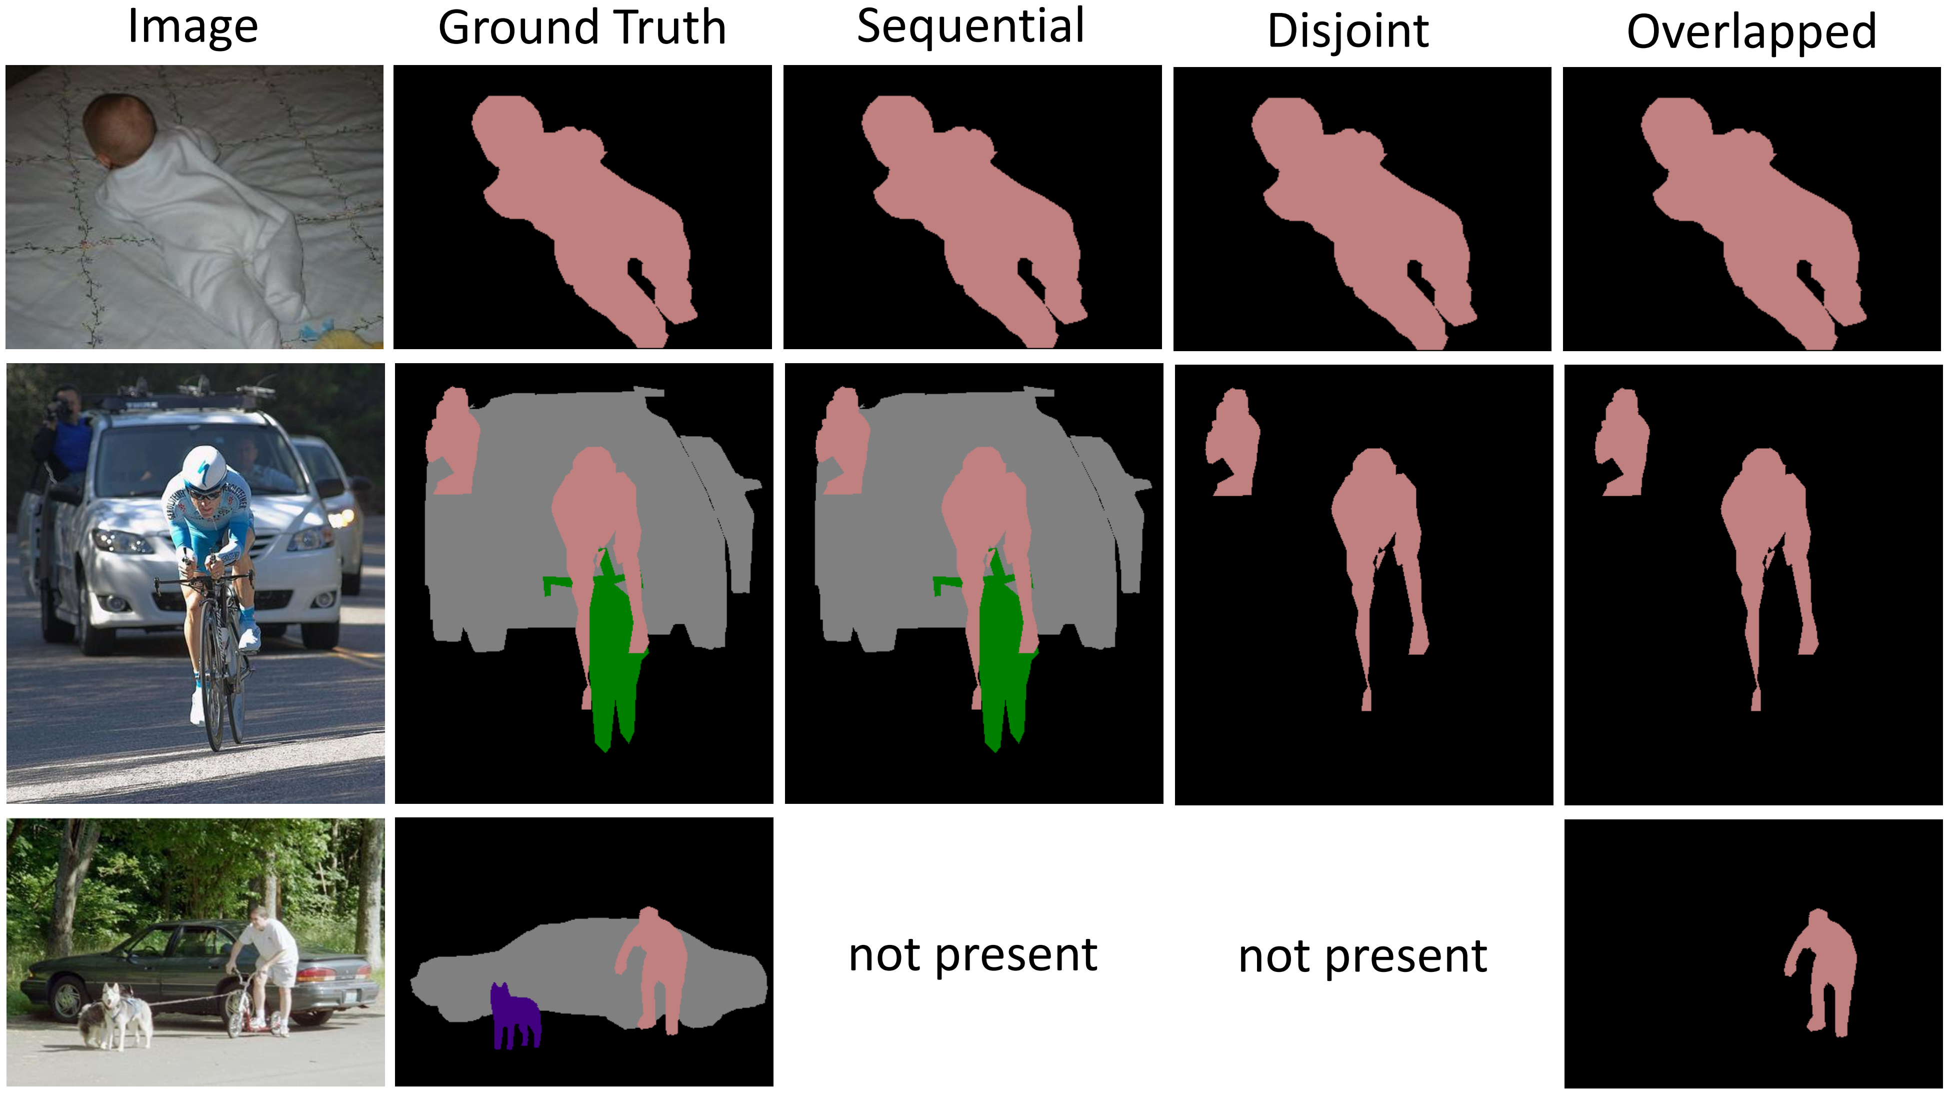

SDR Architecture

SDR: **S**parse and **D**isentangled **R**epresentations

We combine task-related cross entropy loss with **4 constraints**:

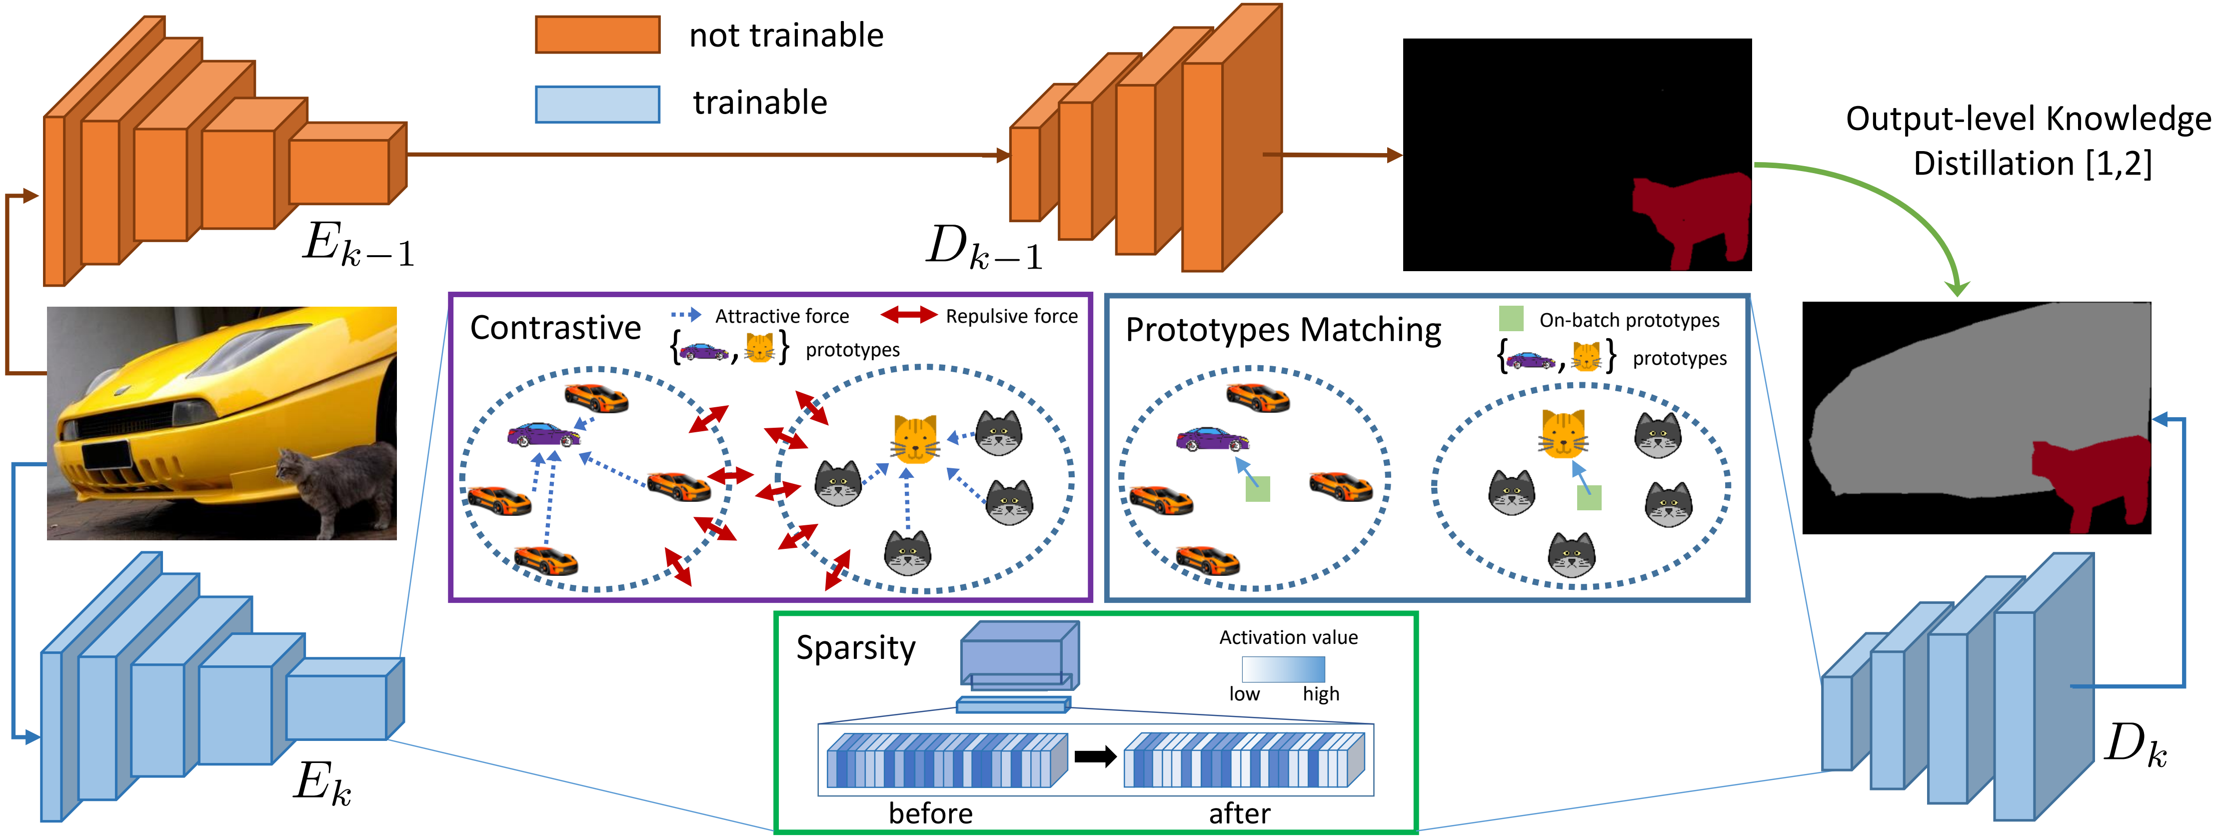

[1] Michieli, U., et al., "Incremental learning techniques for semantic segmentation" ICCVW, 2019.

[2] Cermelli, F. et al., "Modeling the background for incremental learning in semantic segmentation" CVPR, 2020.

SDR Architecture

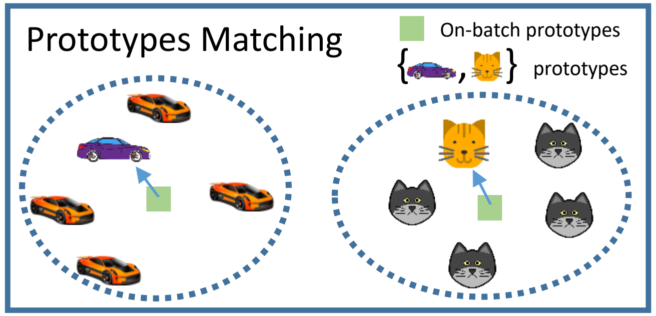

Set of previous classes

$$\mathcal{L}_{pm} = \frac{1}{|\mathcal{C}_{k-1}|} \left\| \hat{\boldsymbol{p}}_c - \boldsymbol{p}_c \right\|_F \qquad c \in \mathcal{C}_{k-1}$$

→ On-batch prototypes constrained to be close to representations learned from previous steps

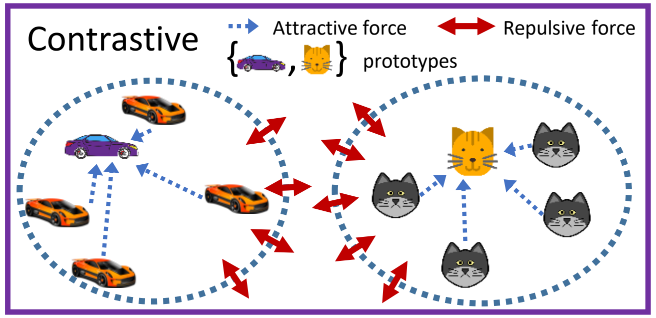

- Attractive: $$\mathcal{L}_{cl}^{a} = \frac{1}{|c_j \in \boldsymbol{y}_n^*|} \sum_{c_j \in \boldsymbol{y}_n^*} \sum_{\boldsymbol{f}_i \in \boldsymbol{F}_n} \left\| \left(\boldsymbol{f}_i - \boldsymbol{p}_{c_j}\right) \mathbb{I}[y_i^* = c_j] \right\|_F$$

Features of the same class tightly clustered around prototype

- Repulsive: $$\mathcal{L}_{cl}^{r} = \frac{1}{|c_j \in \boldsymbol{y}_n^*|} \sum_{c_j \in \boldsymbol{y}_n^*} \sum_{c_k \in \boldsymbol{y}_n^*, c_k \neq c_j} \frac{1}{\|\hat{\boldsymbol{p}}_{c_j} - \hat{\boldsymbol{p}}_{c_k}\|_F}$$

Features of different classes separated from each other

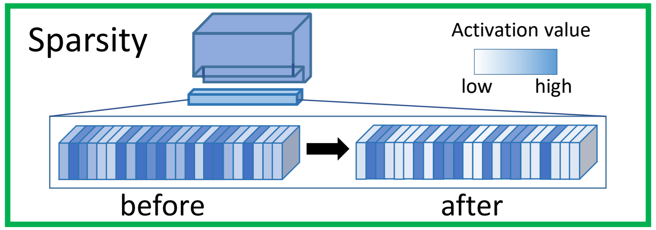

features normalized with respect to the class-conditional maximum value

$$\mathcal{L}_{sp} = \frac{1}{|\boldsymbol{f}_i \in \boldsymbol{F}_n|} \sum_{\boldsymbol{f}_i \in \boldsymbol{F}_n} \frac{\sum_j \exp(\bar{f}_{i,j})}{\sum_j \bar{f}_{i,j}}$$

Set of active channels is narrowed, letting room for the representation of upcoming classes

CVPR VIRTUAL JUNE 19-25

Results – Pascal VOC2012

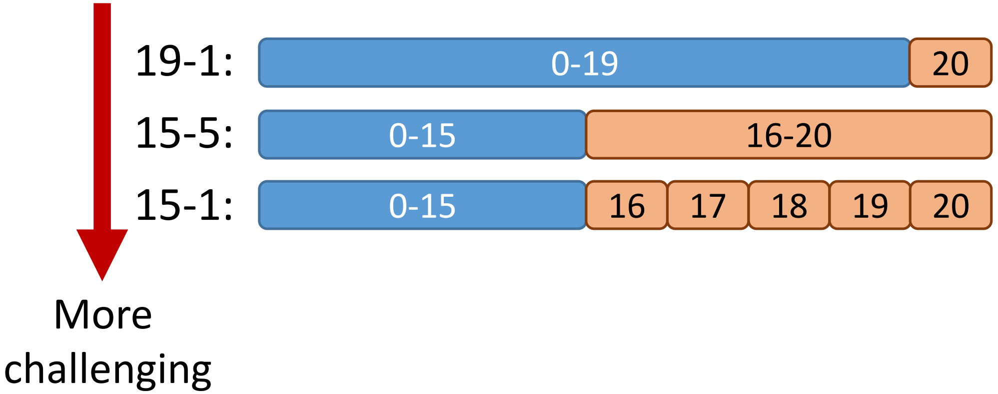

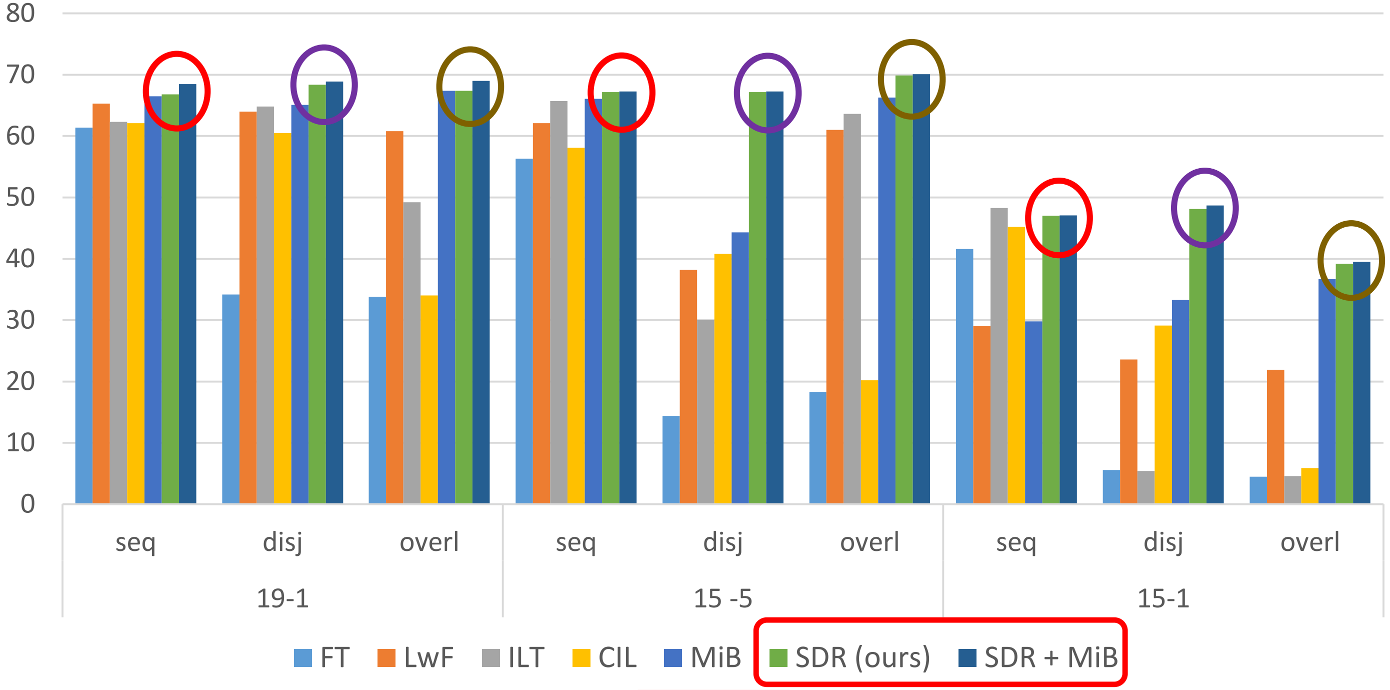

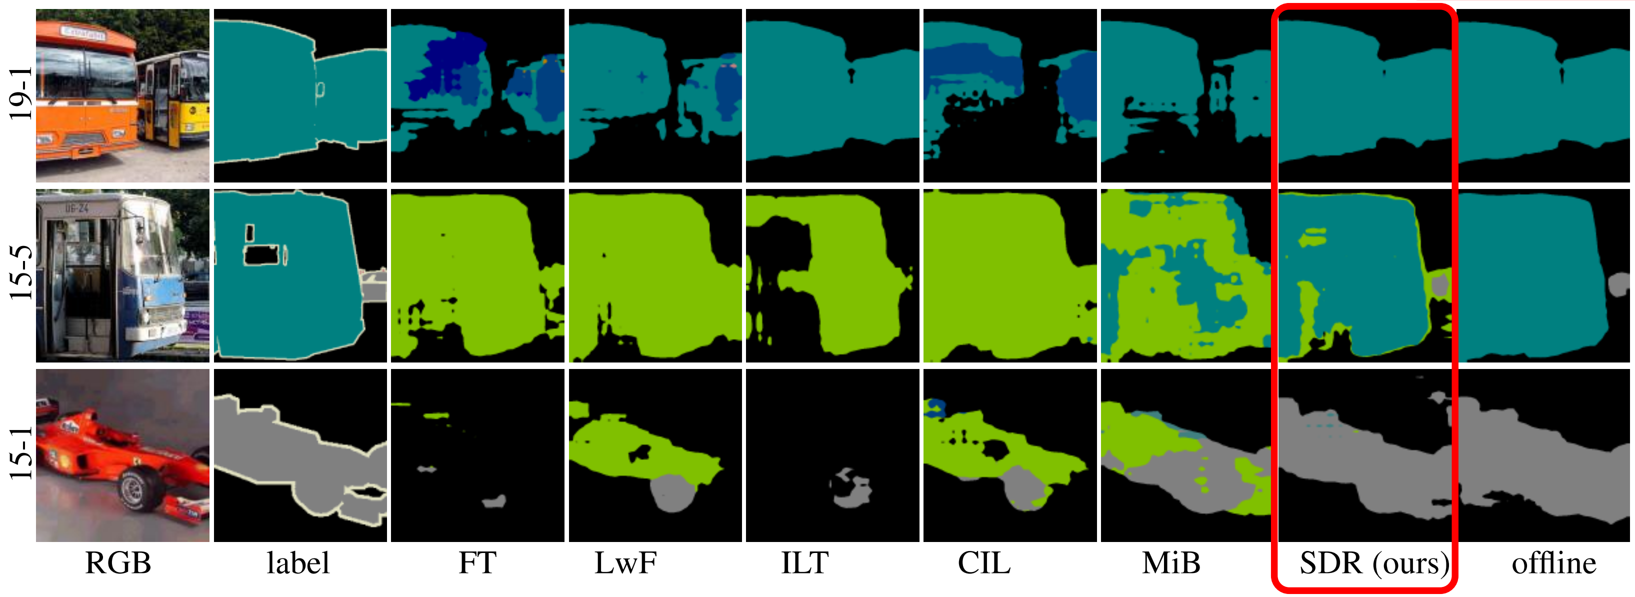

Results – ADE20K

→ SDR outperforms competitors, especially when multiple steps are involved

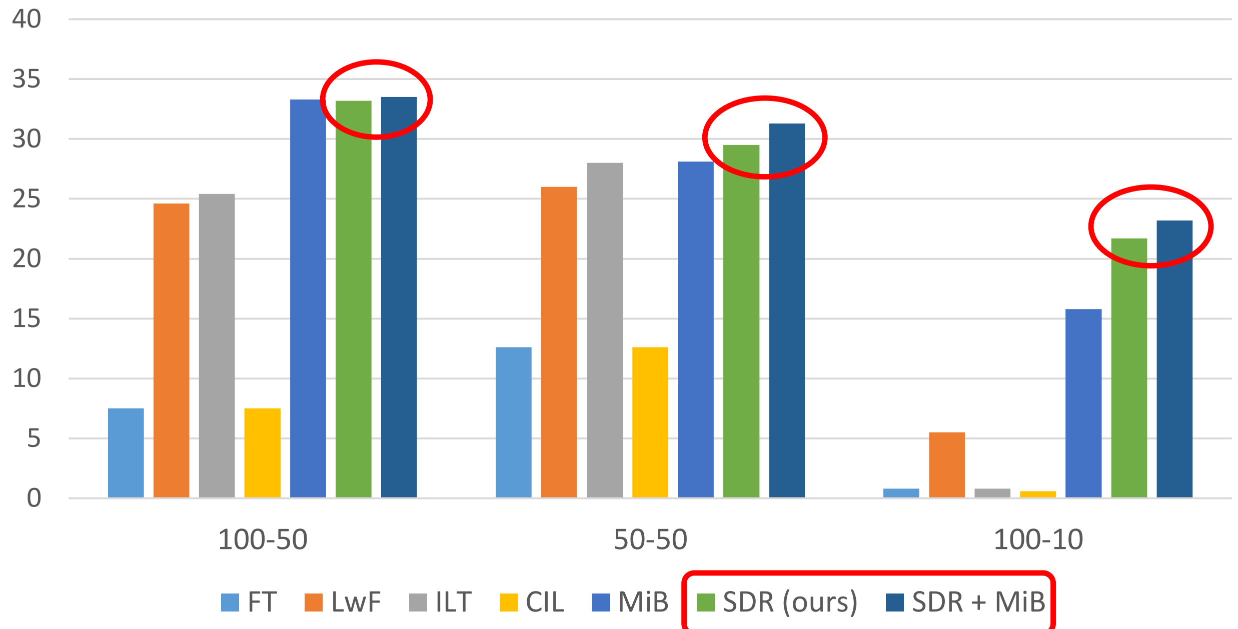

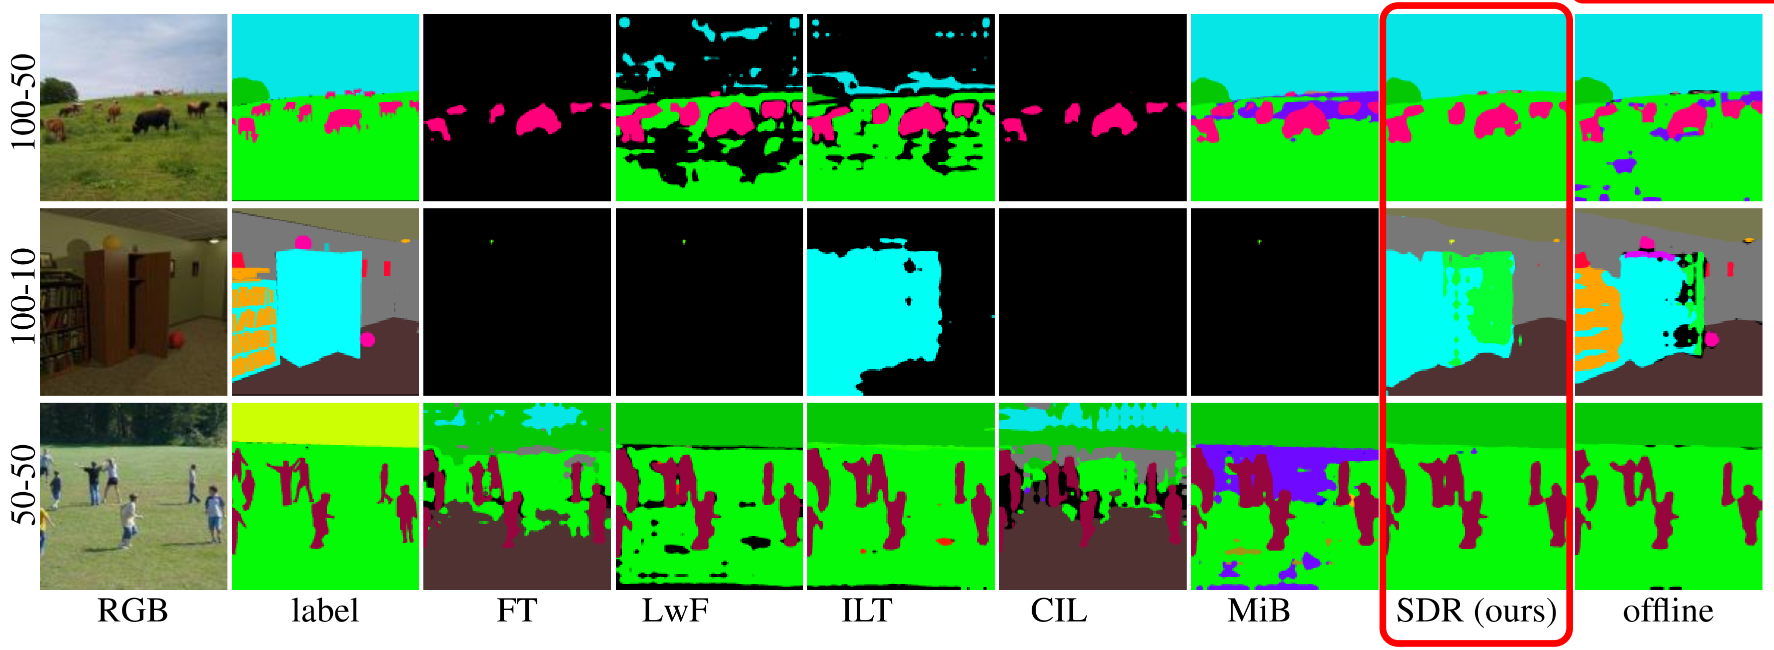

Conclusion

- ✓ We propose 3 novel latent space shaping techniques to avoid forgetting and promote learning of new concepts:
  - **prototype matching**
  - **contrastive learning**
  - **sparsity**
- ✓ We jointly tackle sequential, disjoint and overlapped scenarios
- ✓ We achieve state-of-the-art results on a variety of tasks and datasets

Paper website: https://lttm.dei.unipd.it/paper_data/SDR/

Code available: https://github.com/LTTM/SDR