Fairness Continual Learning Approach to Semantic Scene Understanding in Open-World Environments

Thanh-Dat Truong[1], Hoang-Quan Nguyen[1], Bhiksha Raj[2,3], Khoa Luu[1]

[1]CVIU Lab, University of Arkansas

[2]Carnegie Mellon University

[3]Mohammed Bin Zayed University of AI

https://uark-cviu.github.io/

Continual Semantic Segmentation (CSS)

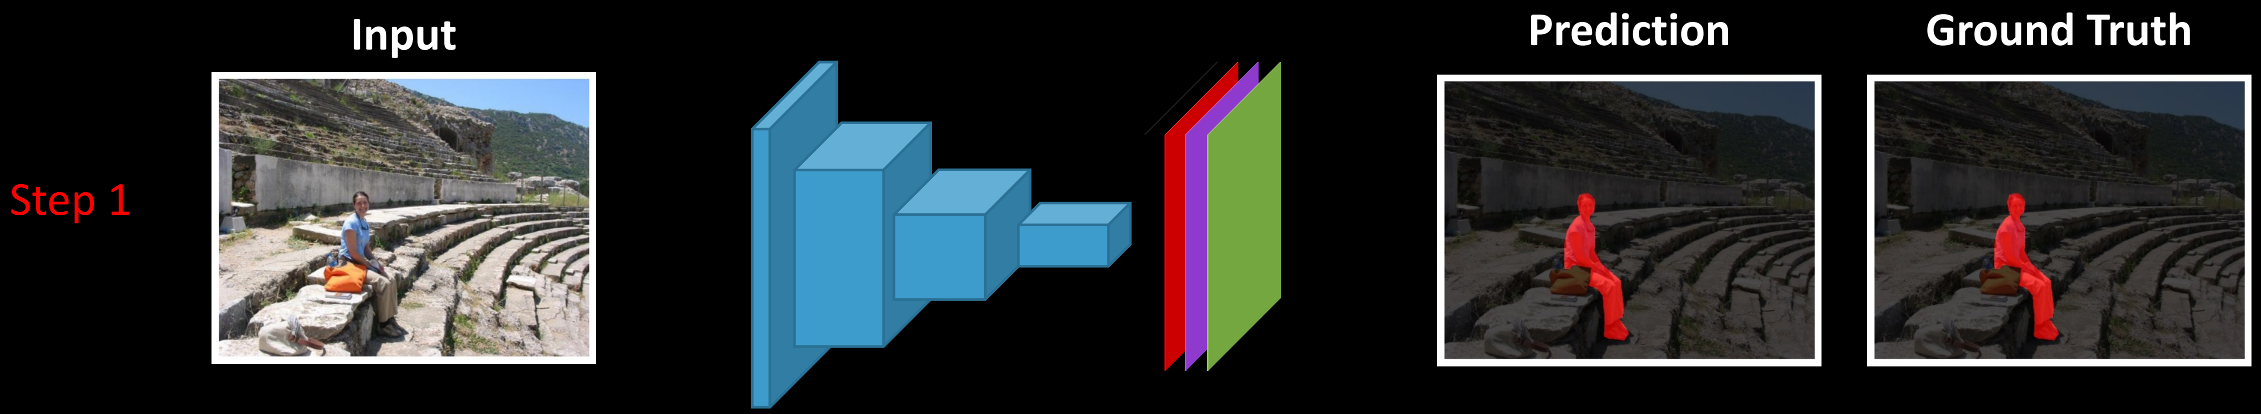

Continual Semantic Segmentation (CSS)

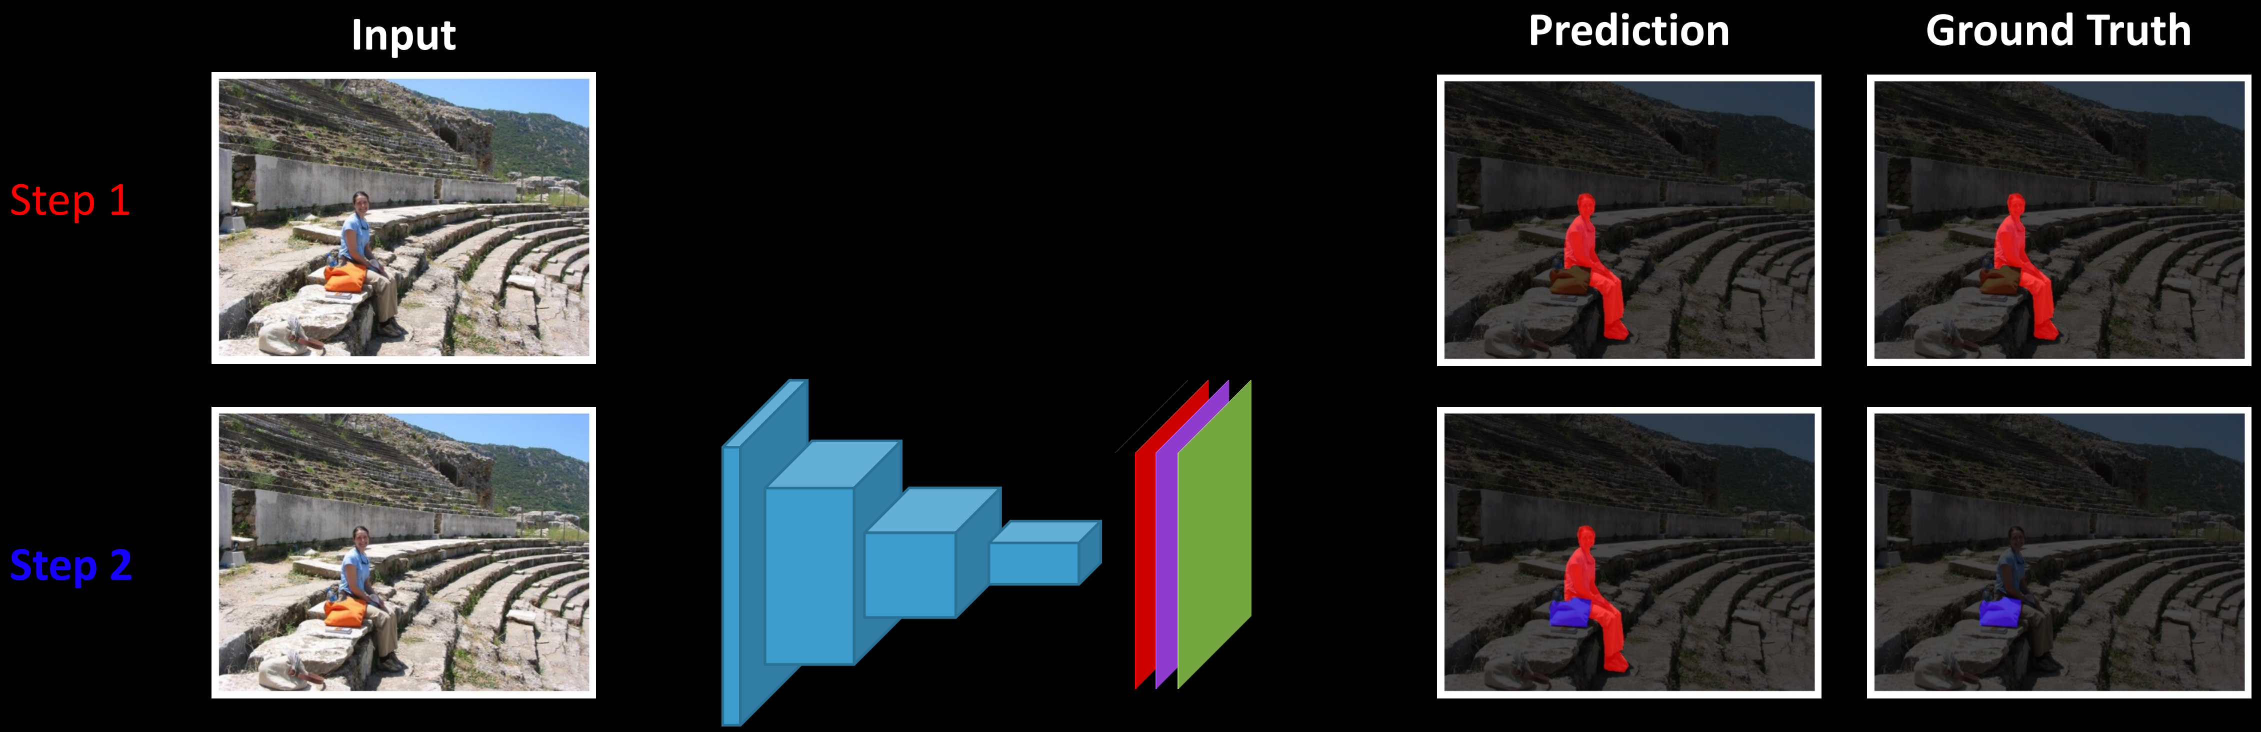

Continual Semantic Segmentation (CSS)

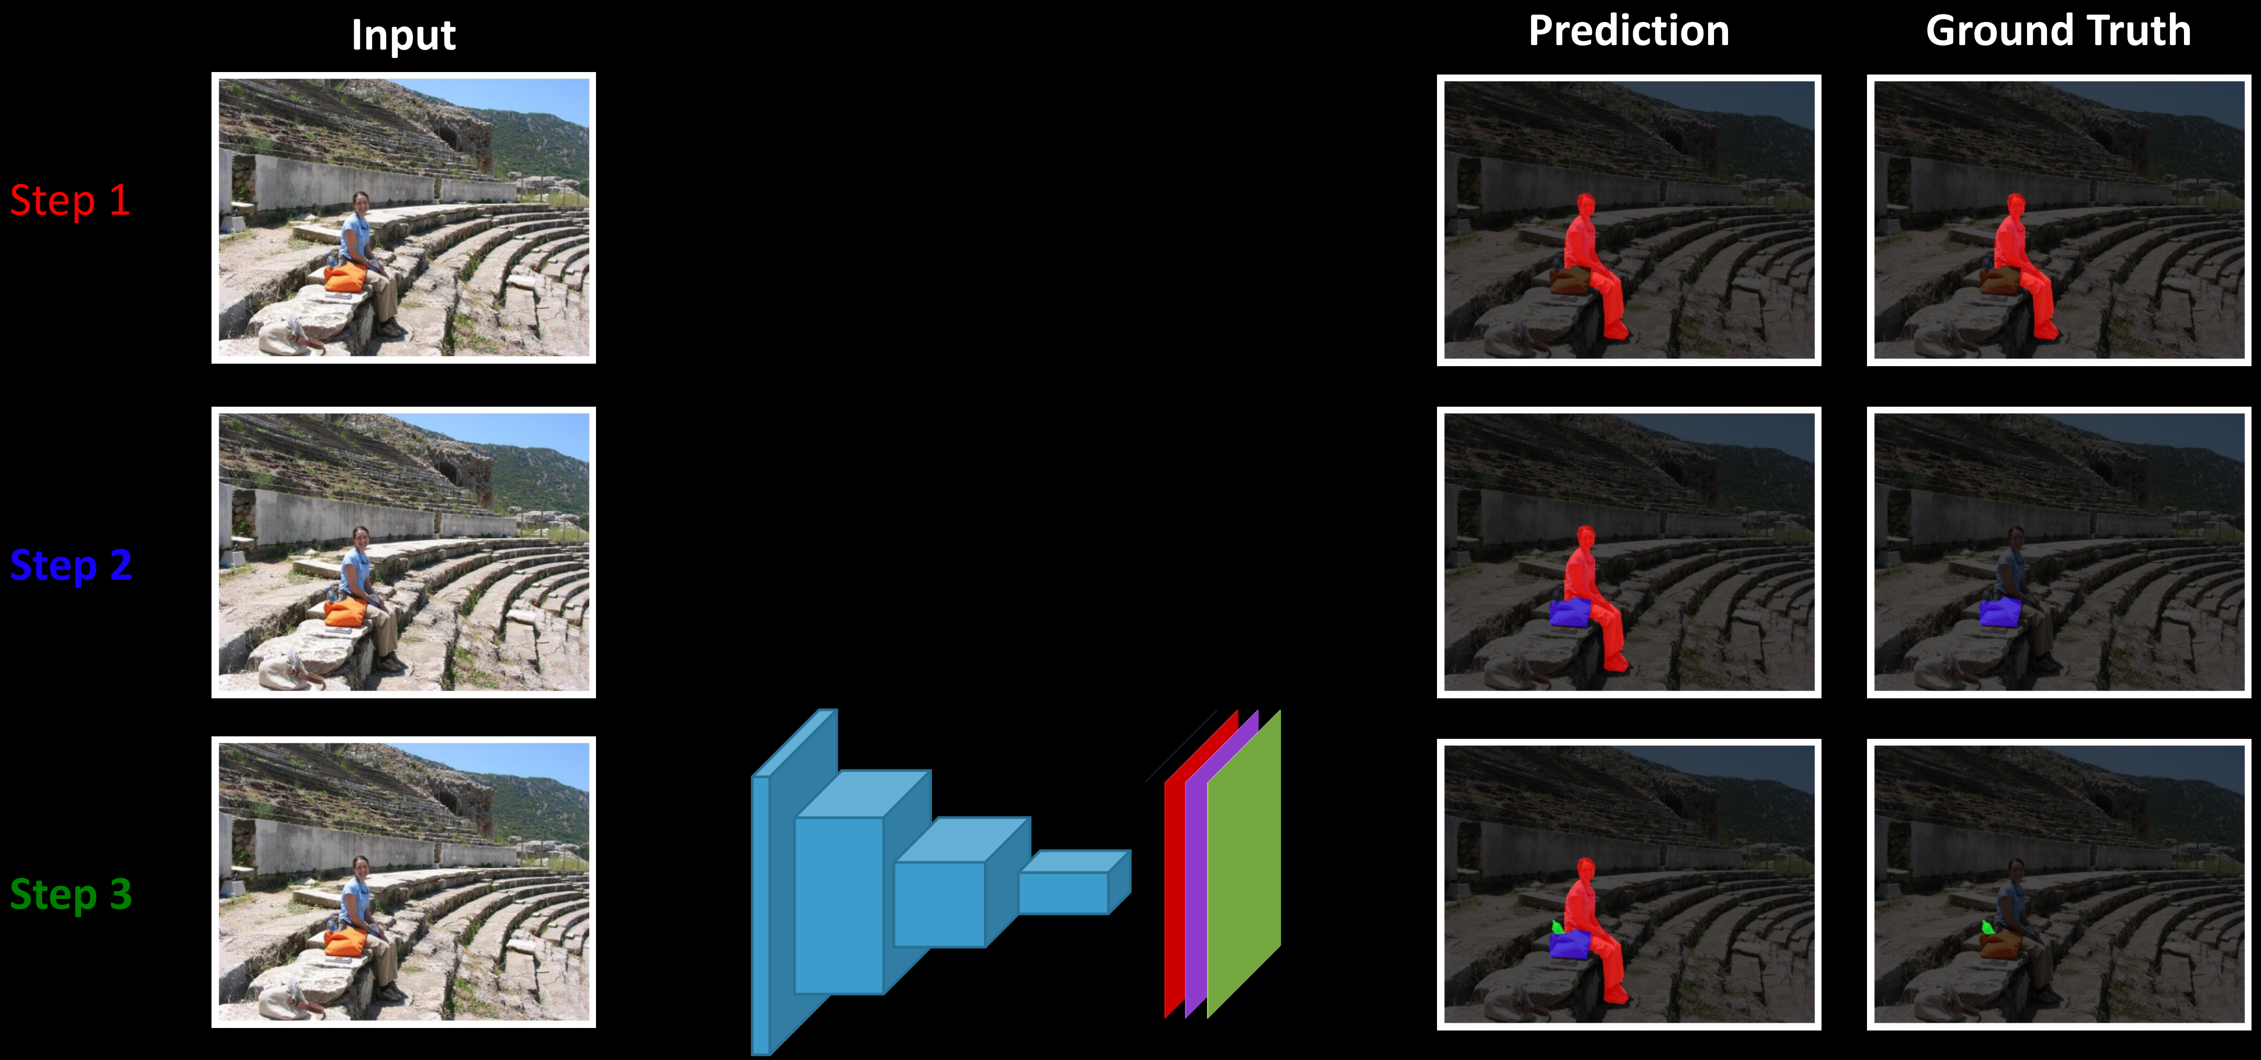

Continual Semantic Segmentation (CSS)

Input | Prediction | Ground Truth

Step 1

Step 2

Step 3

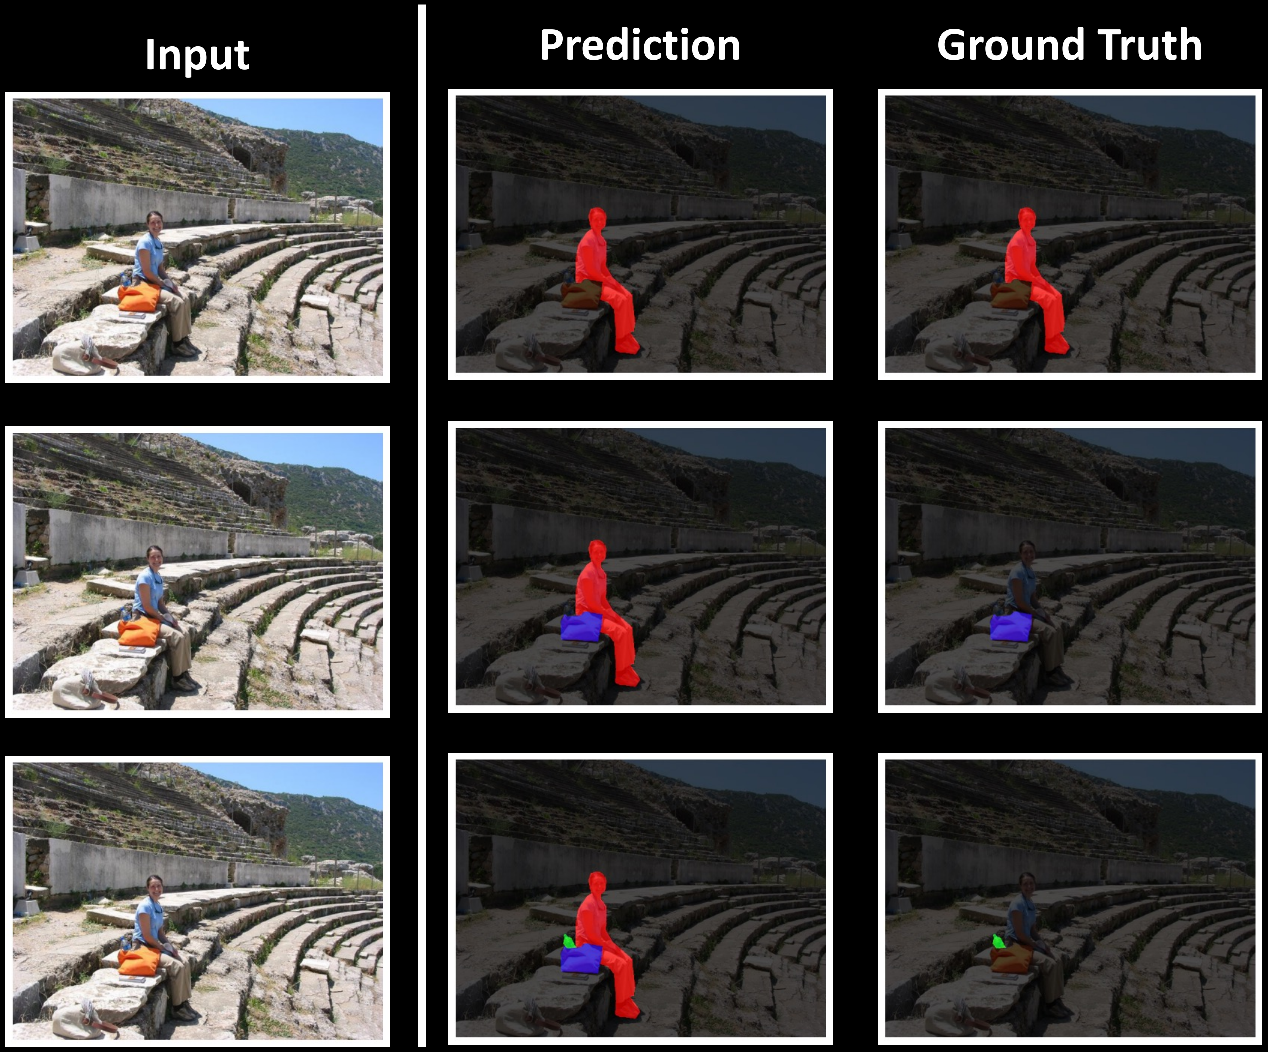

Major Challenges in CSS

Continual Semantic Segmentation (CSS)

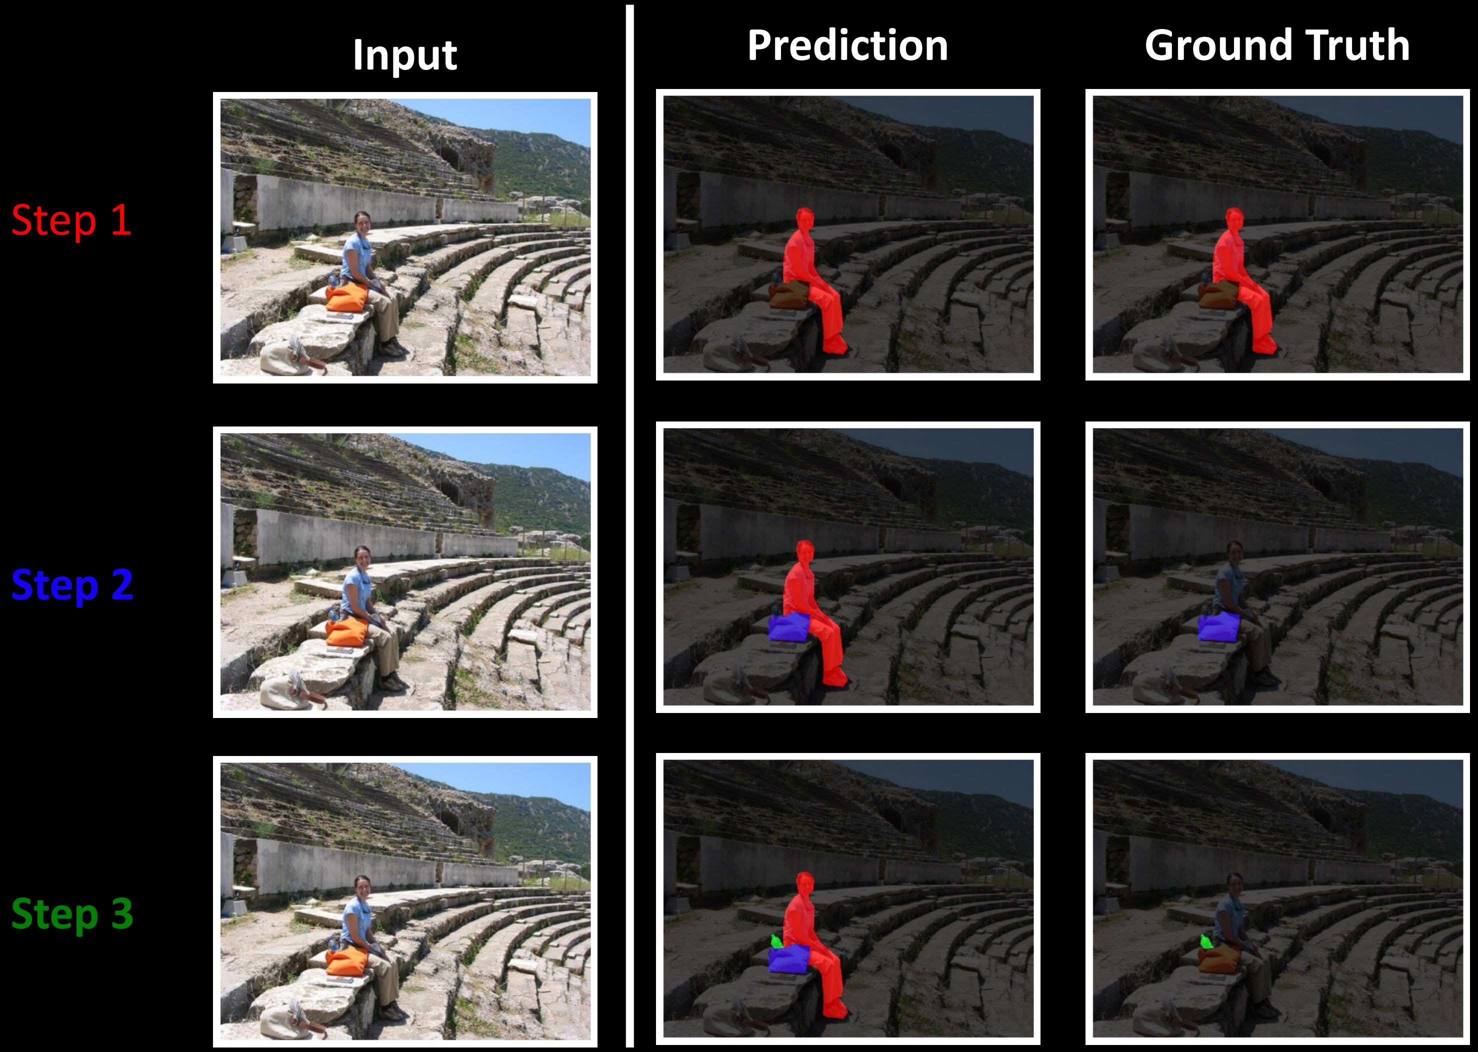

Major Challenges in CSS

1 **Catastrophic Forgetting**

The model partially or completely forgets the knowledge learned in the previous steps

Continual Semantic Segmentation (CSS)

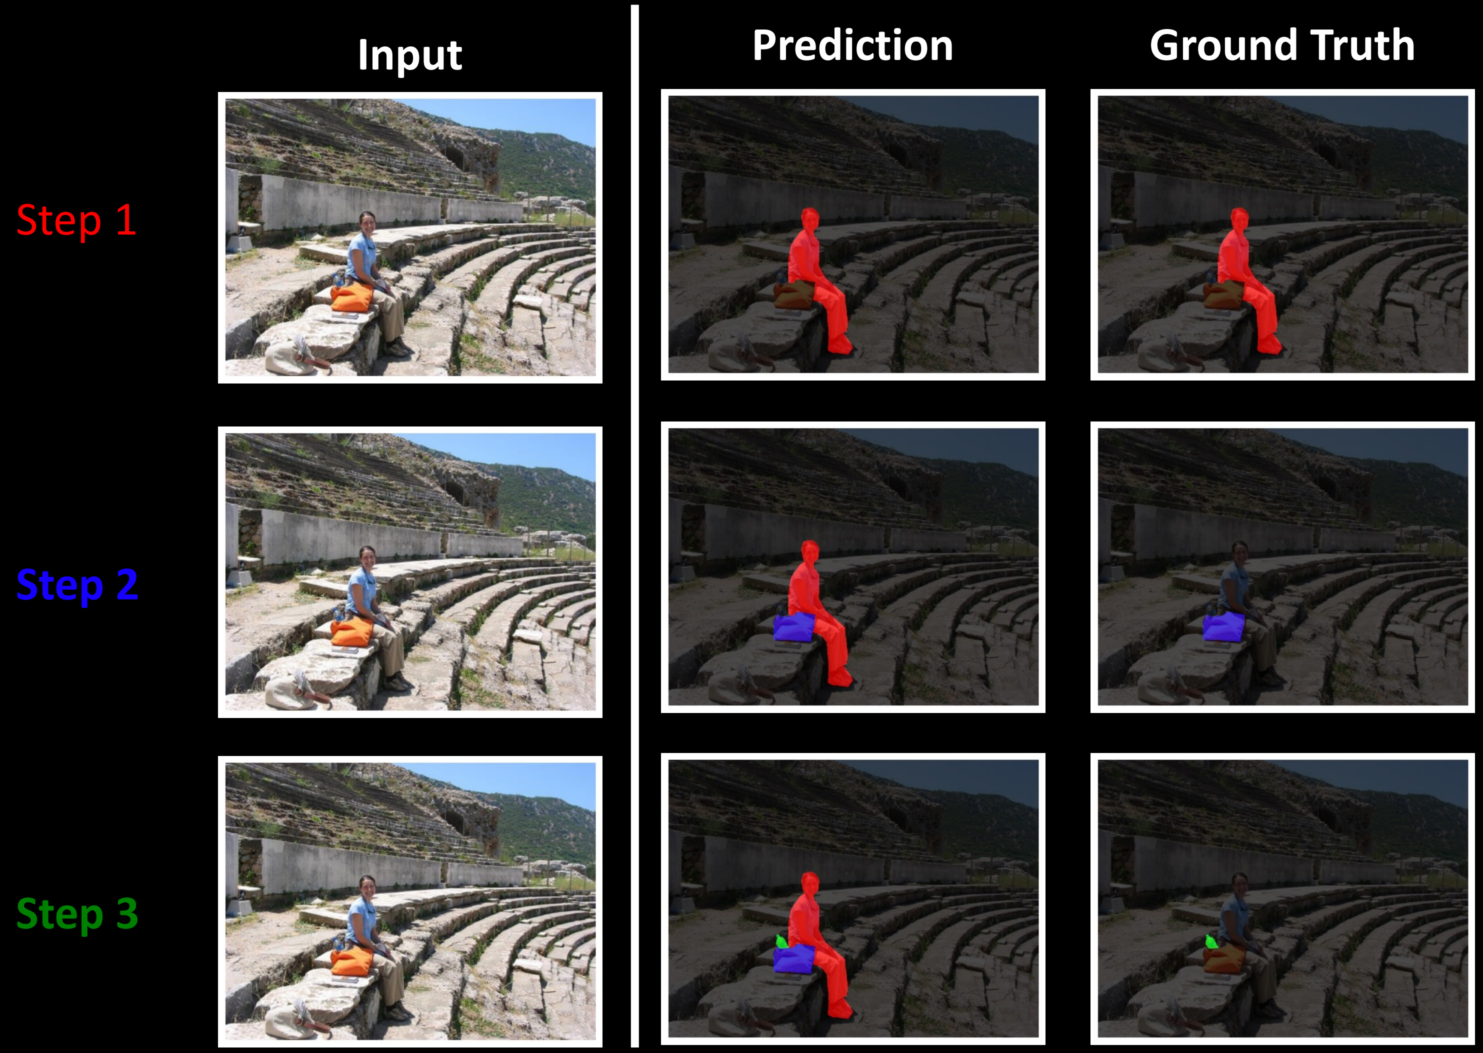

Major Challenges in CSS

1. **Catastrophic Forgetting**
   The model partially or completely forgets the knowledge learned in the previous steps

2. **Background Shifting**
   The labels in the previous and future steps has been collapsed into a background class

Continual Semantic Segmentation (CSS)

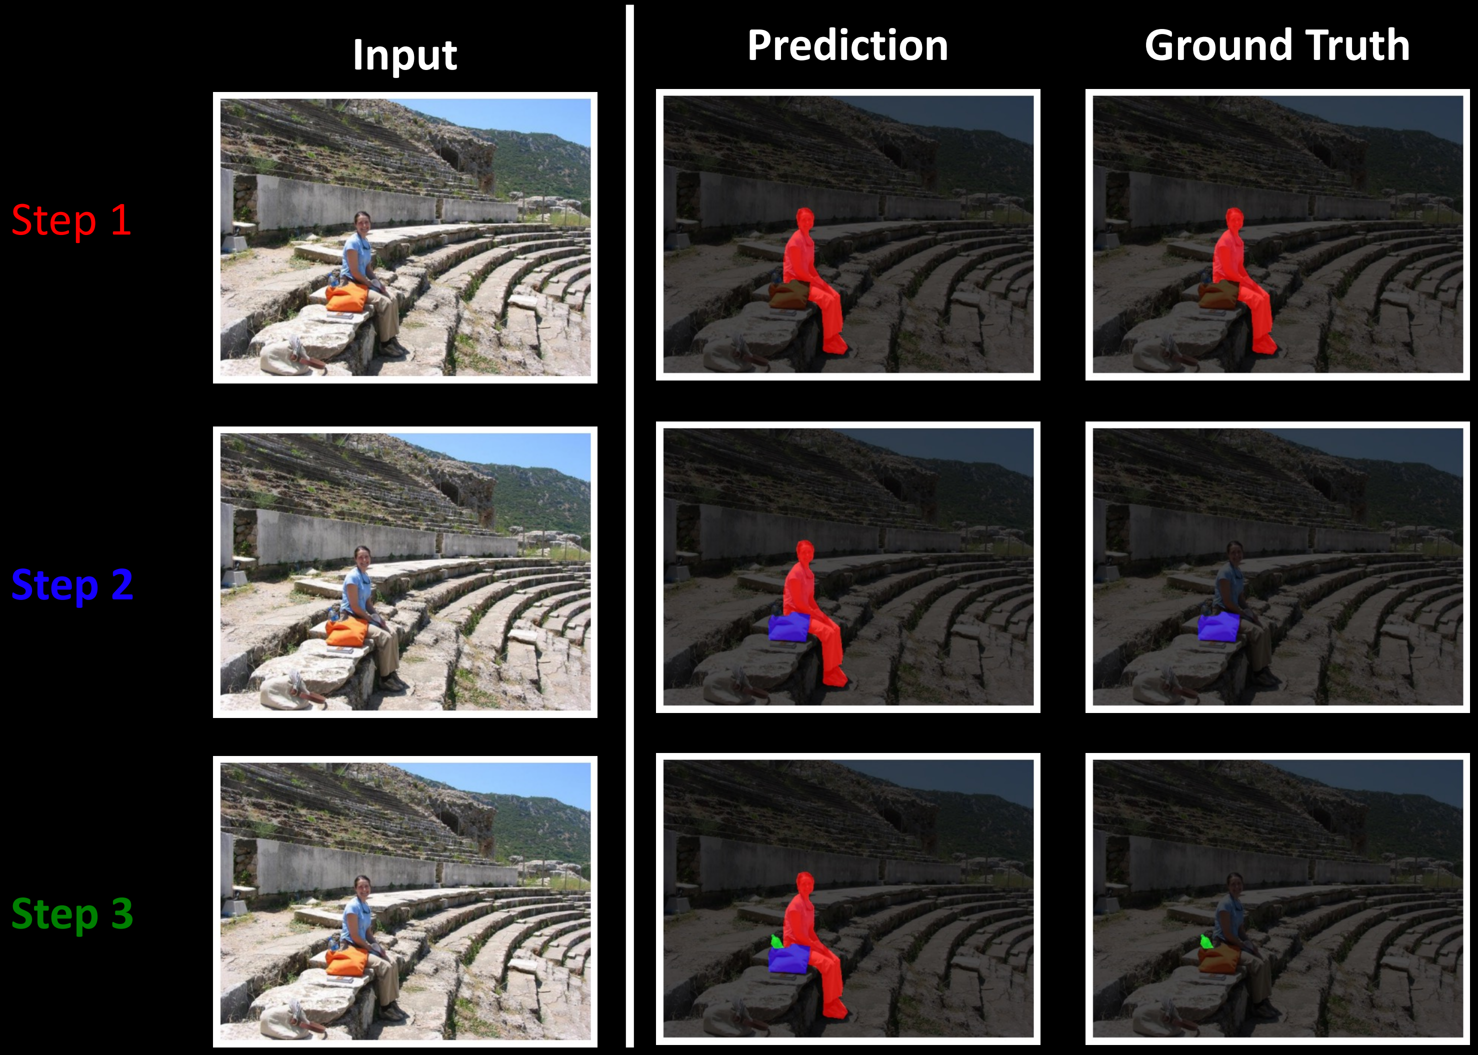

Major Challenges in CSS

1. **Catastrophic Forgetting**
   The model partially or completely forgets the knowledge learned in the previous steps

2. **Background Shifting**
   The labels in the previous and future steps has been collapsed into a background class

3. **Fairness**
   The distributions of the classes have been imbalanced toward a specific group of classes

Fairness in Continual Learning

The Class Distribution based on the Number of Pixels

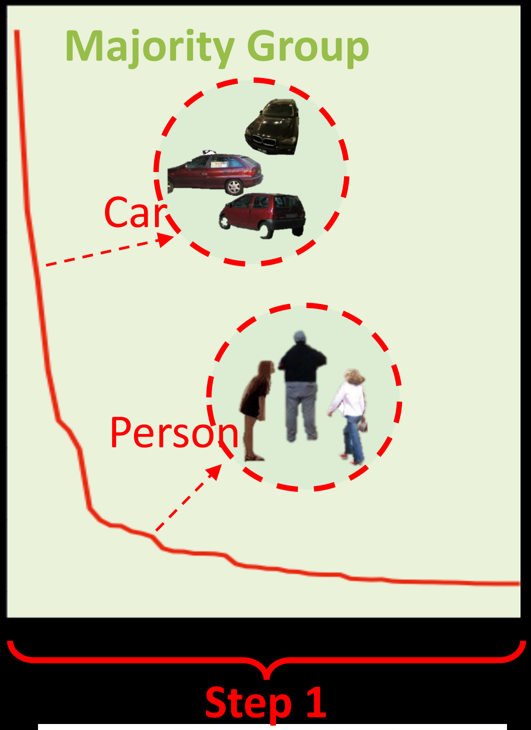

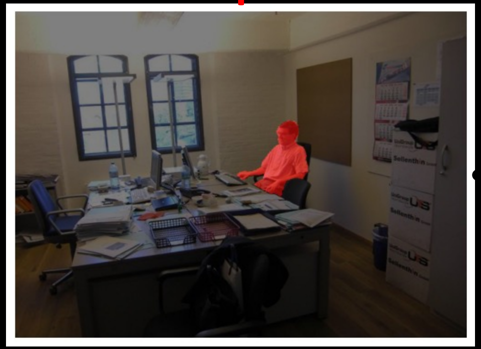

Fairness in Continual Learning

The Class Distribution based on the Number of Pixels

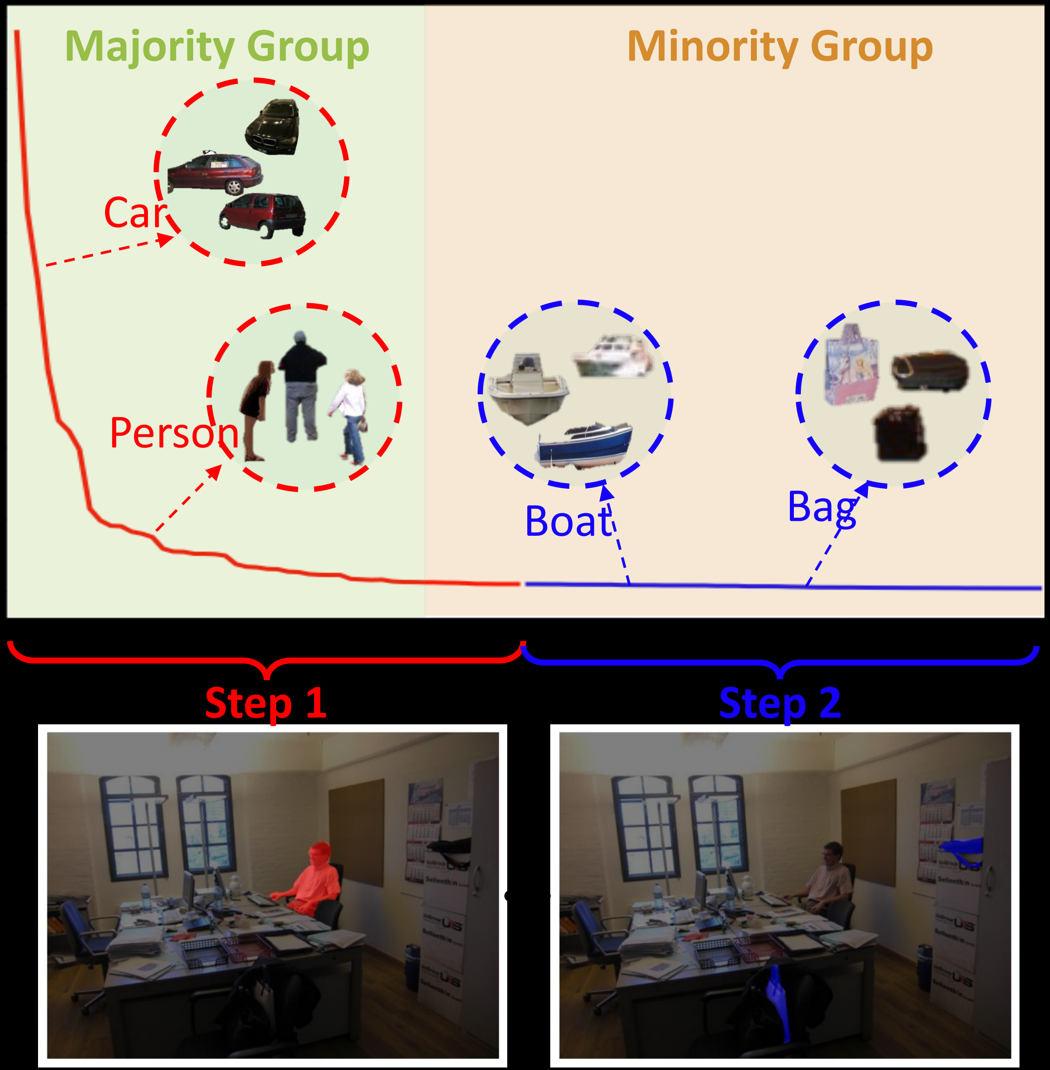

Fairness in Continual Learning

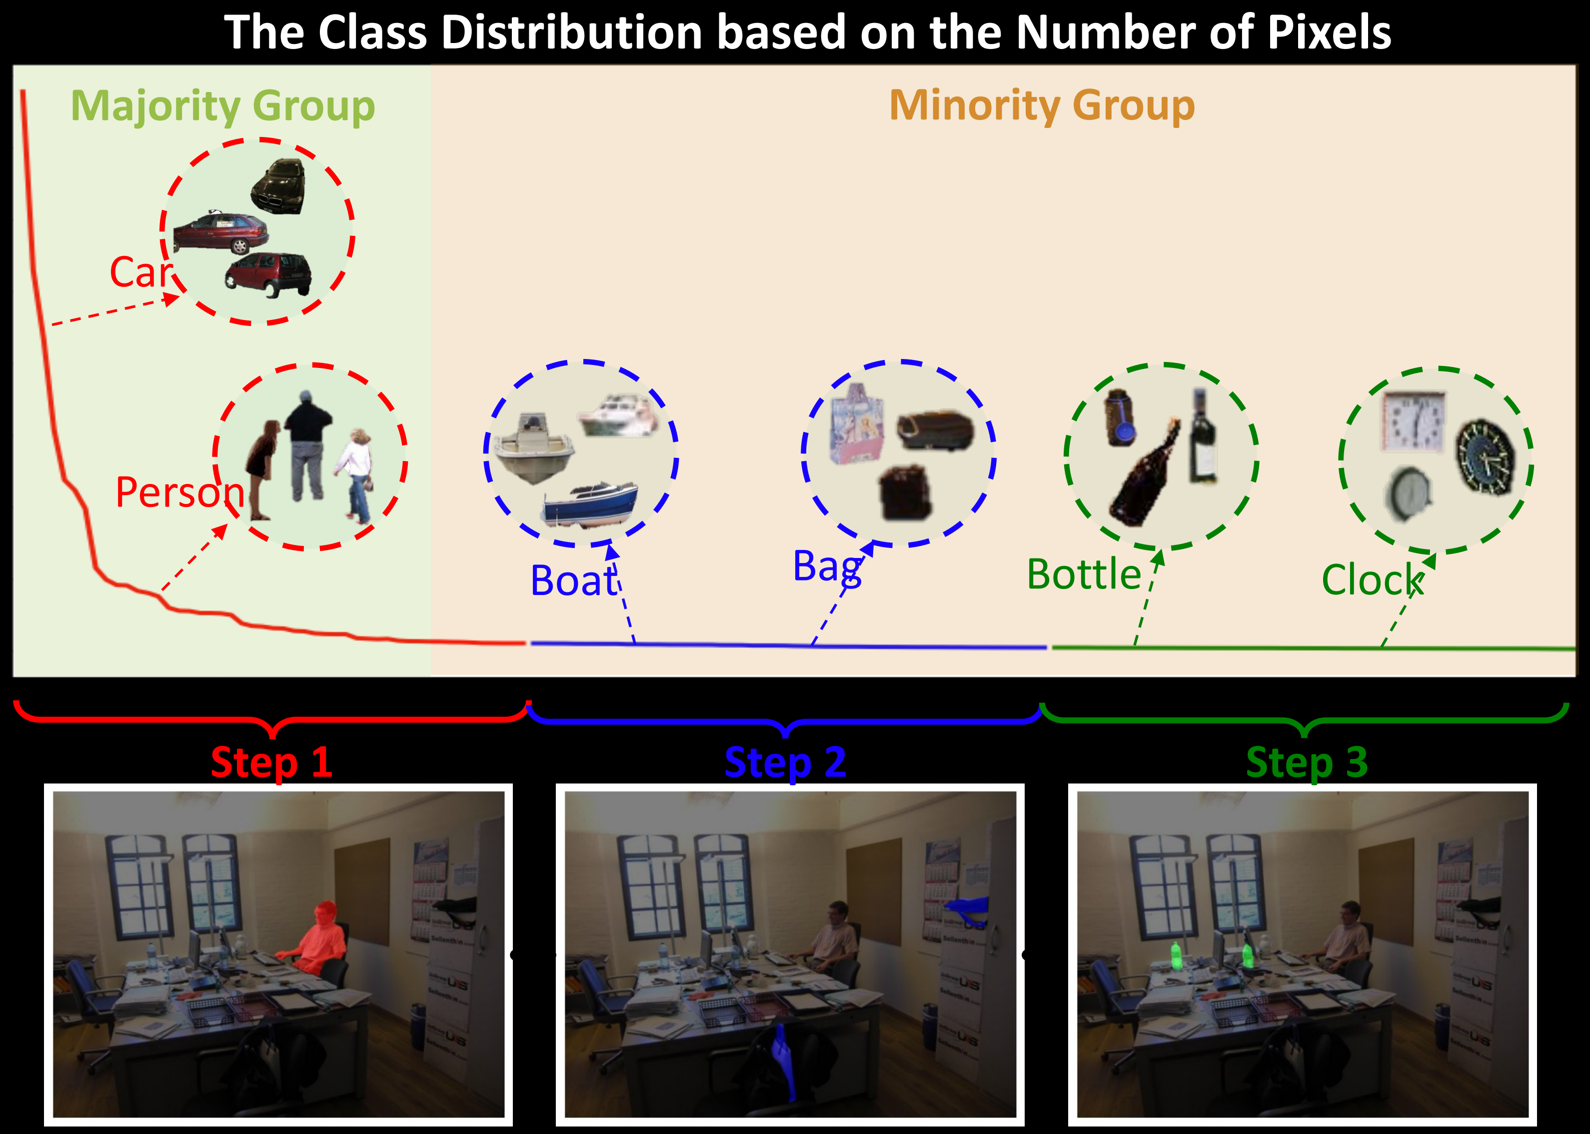

Fairness in Continual Learning

The Class Distribution based on the Number of Pixels

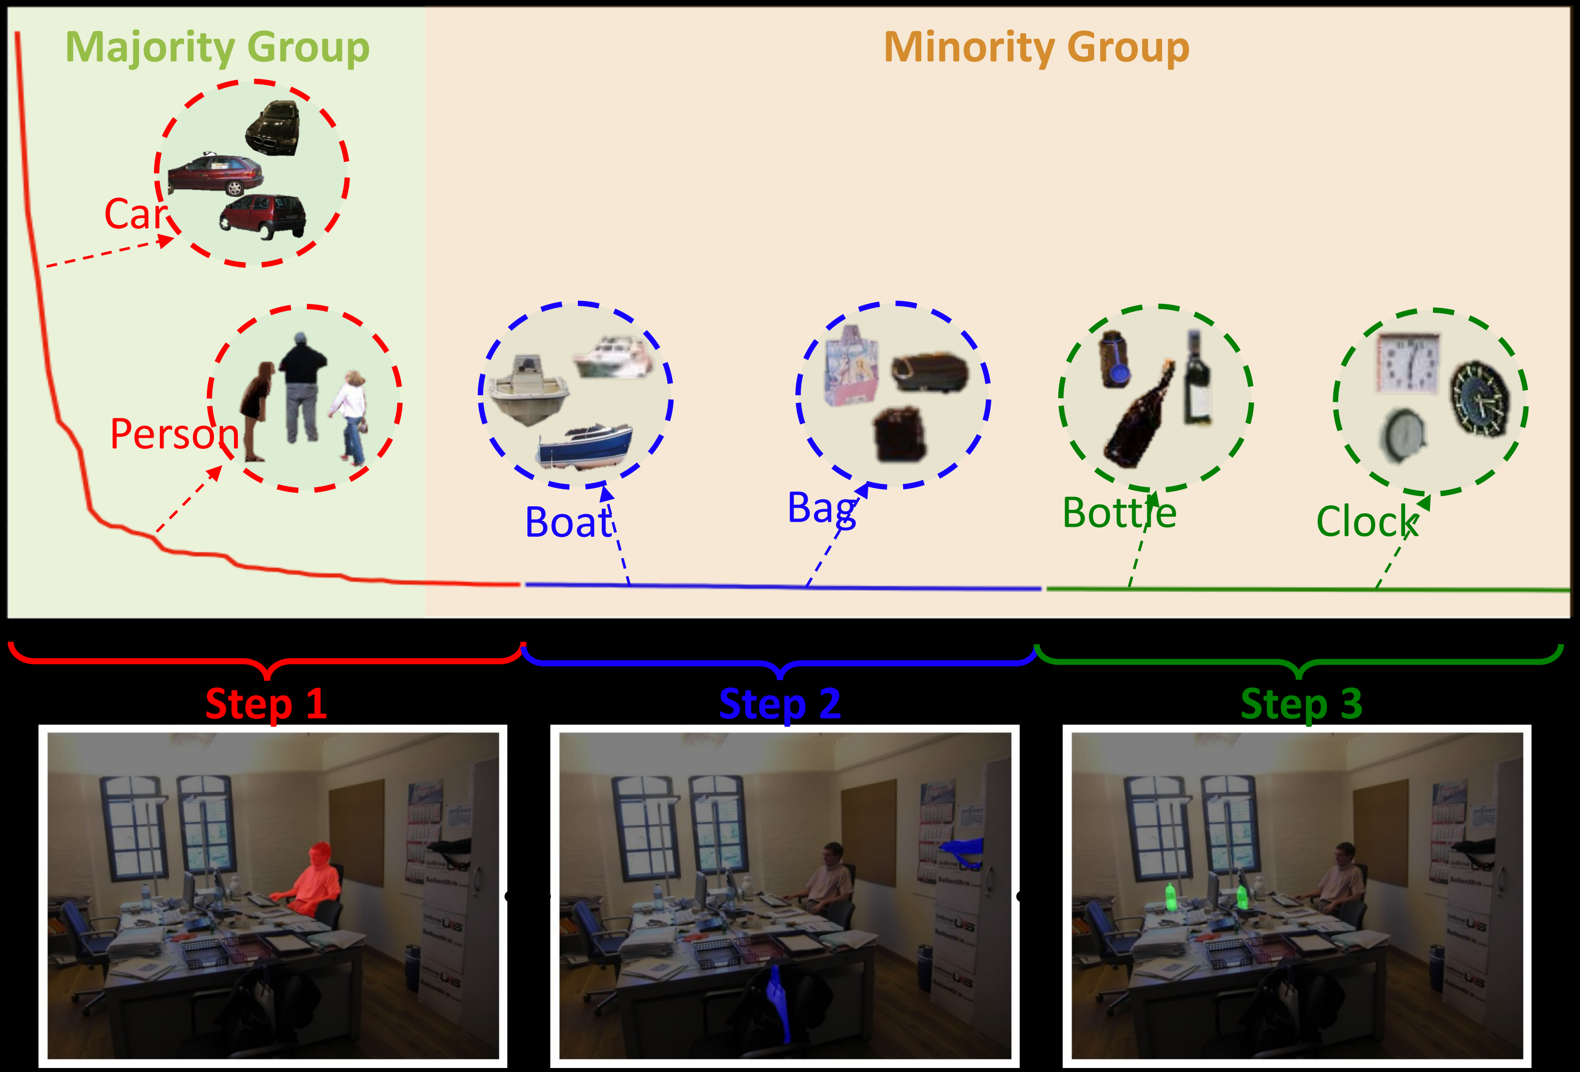

The distribution of classes in the majority group in the task dominates the ones in the minority groups in the later tasks

The model behaves unfairly among classes

Contributions

- Introduce a novel fairness metric for continual semantic segmentation
- Propose new Fairness Continual Learning approach to Semantic Segmentation
  - Promote fairness by a new fairness loss based on the class distribution
  - Impose consistency of segmentation maps by a Conditional Structural Consistency Loss
  - Model the catastrophic forgetting and background shift problems via the new Prototypical Contrastive Clustering loss
  - Proved as a new, generalized continual learning paradigm of knowledge distillation
- Achieve State-of-the-Art Performance on Continual Semantic Segmentation benchmarks and Promote Fairness of the model predictions

Fairness Objective

$$\theta^* = \operatorname*{argmin}_{\theta} \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}}\, \mathcal{L}(\boldsymbol{y}, \widehat{\boldsymbol{y}})$$

Subject to:

$$\underbrace{\max_{c_a, c_b} \left| \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}} \sum_{i,j} \mathcal{L}\left(y_{i,j} = c_a\right) - \mathbb{E}_{x,y} \sum_{i,j} \mathcal{L}\left(y_{i,j} = c_b\right) \right| \leq \epsilon}_{\text{Impose the Fair Behavior of the Model by Maintaining the Small Difference of Error Rates Between Classes}}$$

Impose the Fair Behavior of the Model by
Maintaining the Small Difference of Error Rates Between Classes

Fairness Objective

$$\underbrace{\max_{c_a, c_b} \left| \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}} \sum_{i,j} \mathcal{L}\left(y_{i,j} = c_a\right) - \mathbb{E}_{x,y} \sum_{i,j} \mathcal{L}\left(y_{i,j} = c_b\right) \right| \leq 2C\, \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}}\, \mathcal{L}(\boldsymbol{y}, \hat{\boldsymbol{y}})}_{\text{The Fairness Objective Is Also Imposed by the Loss}}$$

Fairness Objective

$$\theta^* = \underset{\theta}{\operatorname{argmin}}\, \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}}\, \mathcal{L}(\boldsymbol{y}, \widehat{\boldsymbol{y}})$$

$$= \underset{\theta}{\operatorname{argmin}} \int \mathcal{L}\,(\boldsymbol{y}, \widehat{\boldsymbol{y}})p(\boldsymbol{y}), p(\widehat{\boldsymbol{y}})\boldsymbol{dy}d\widehat{\boldsymbol{y}}$$

$$= \underset{\theta}{\operatorname{argmin}} \int \mathcal{L}(y_{i,j}, \hat{y}_{i,j})p(y_{i,j})p(\boldsymbol{y}_{\setminus(i,j)}|y_{i,j})p(\widehat{\boldsymbol{y}})\boldsymbol{dy}d\widehat{\boldsymbol{y}}$$

Suffer Imbalance Distributions

The Gradients Produced in the Majority Group Largely Dominant the Ones in the Minority Group

The Class Distribution based on the Number of Pixels

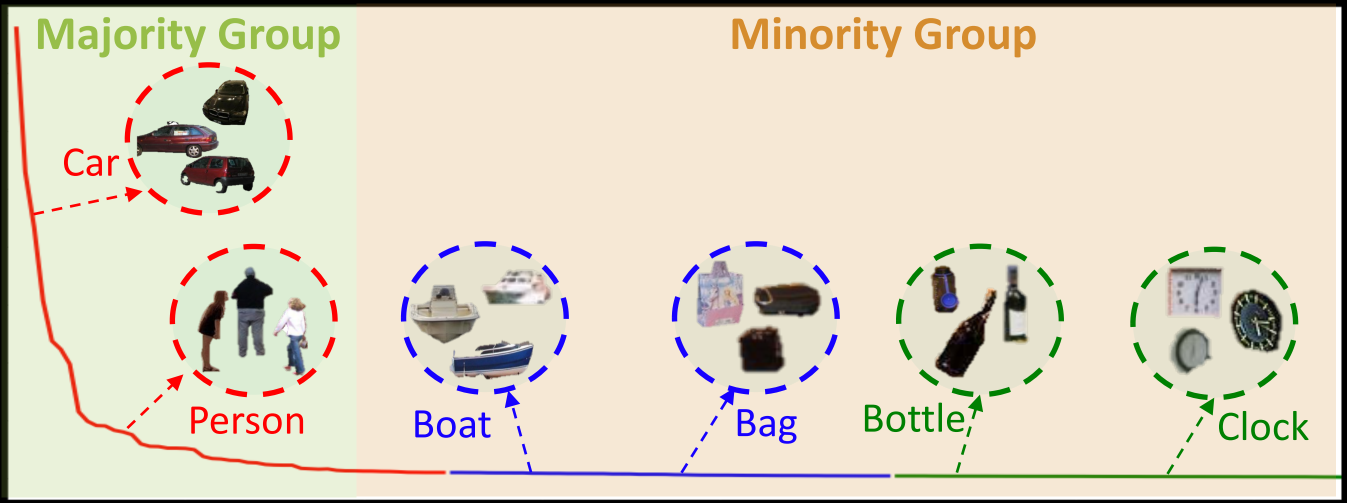

Fairness Continual Learning Approach

$$\theta^* = \operatorname*{argmin}_{\theta} \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}}\, \mathcal{L}(\boldsymbol{y}, \widehat{\boldsymbol{y}}) \frac{q(y_{i,j})q(\boldsymbol{y}_{\setminus(i,j)}|y_{i,j})}{p(y_{i,j})p(\boldsymbol{y}_{\setminus(i,j)}|y_{i,j})}$$

Ideal Distributions
Where the Learned Model Behave Fairly

$$\theta^* \cong \operatorname*{argmin}_{\theta} \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}} \left\{ \mathcal{L}(\boldsymbol{y}, \widehat{\boldsymbol{y}}) + \frac{1}{N}\sum_{i,j}\left[\log\frac{q(y_{i,j})}{p(y_{i,j})} + \log\frac{q(\boldsymbol{y}_{\setminus(i,j)}|y_{i,j})}{p(\boldsymbol{y}_{\setminus(i,j)}|y_{i,j})}\right]\right\}$$

Continual Learning Objective

Fairness Objective From Class Distribution

Conditional Structural Consistency Objective

Fairness Continual Learning Framework

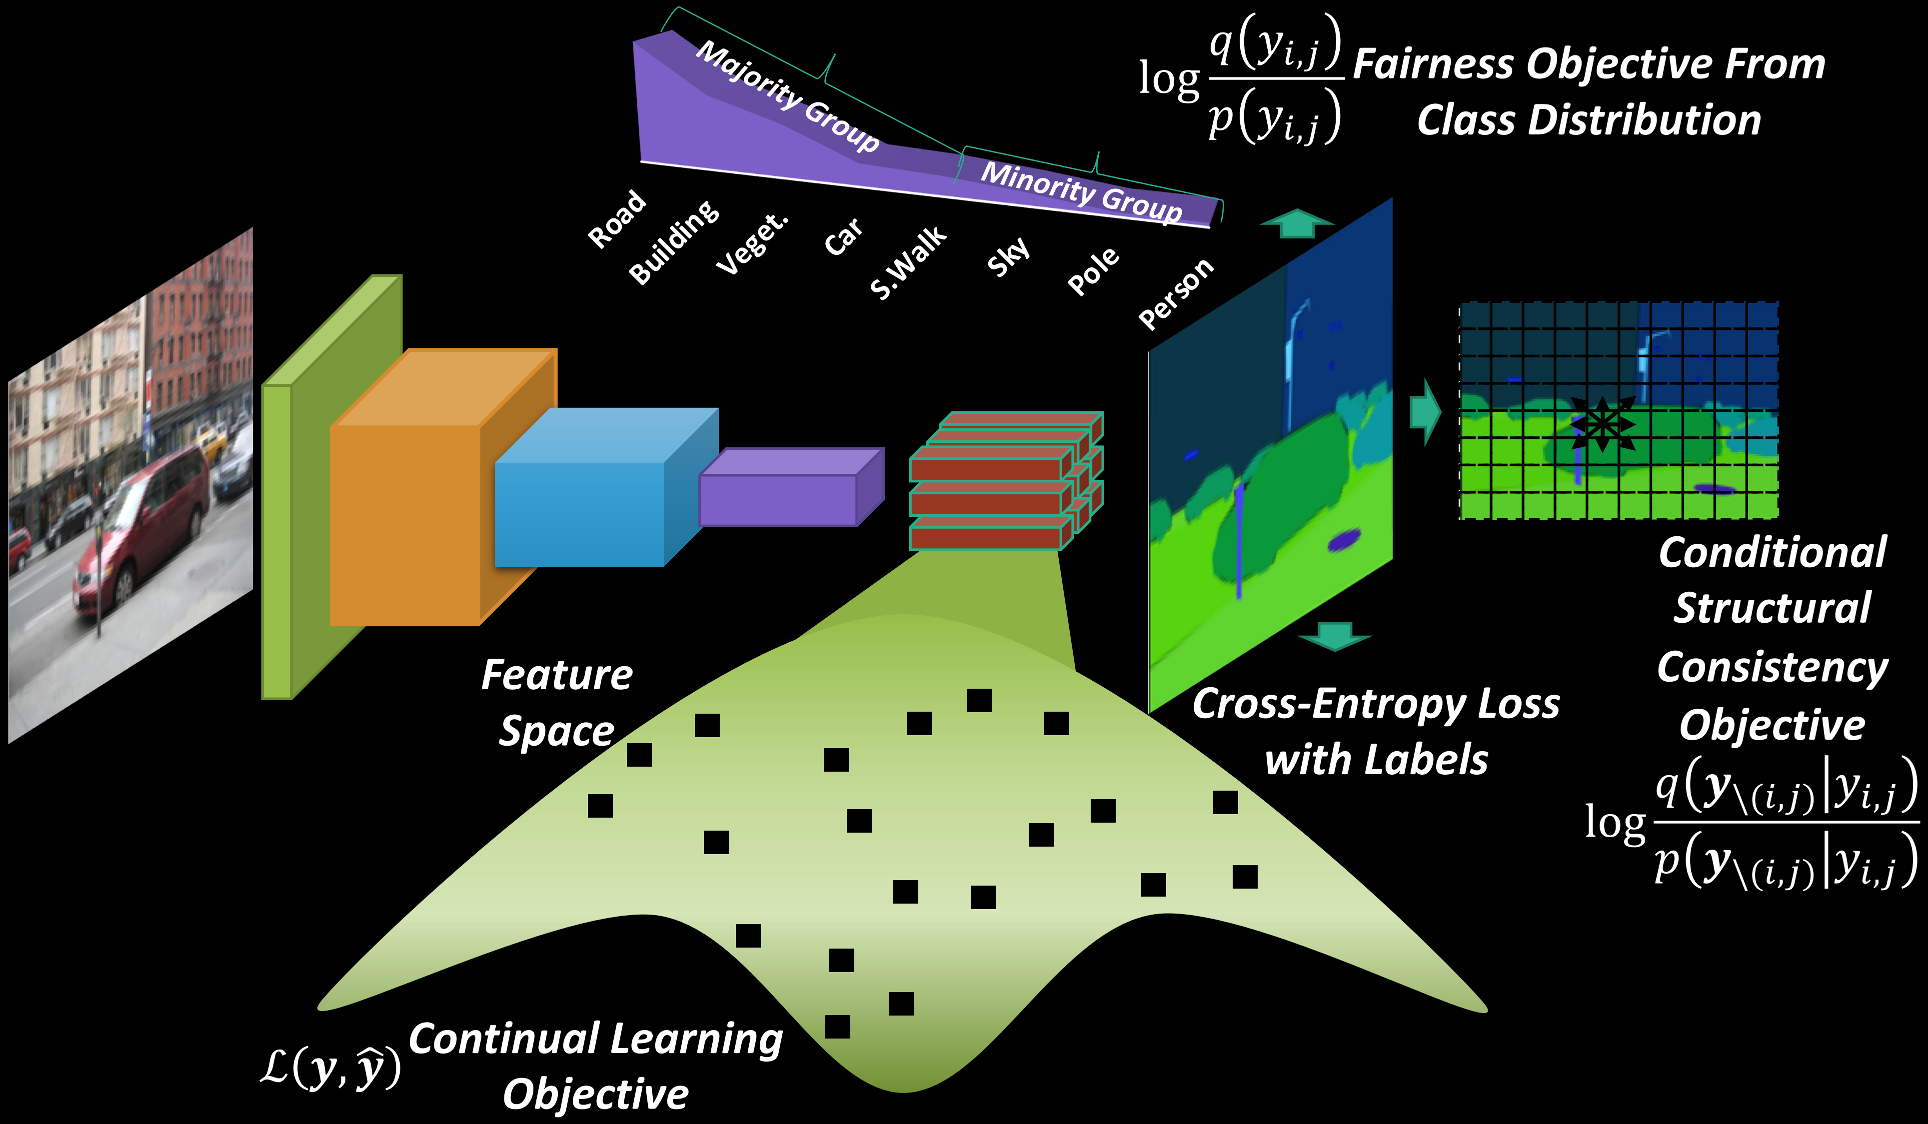

Fairness Continual Learning Framework

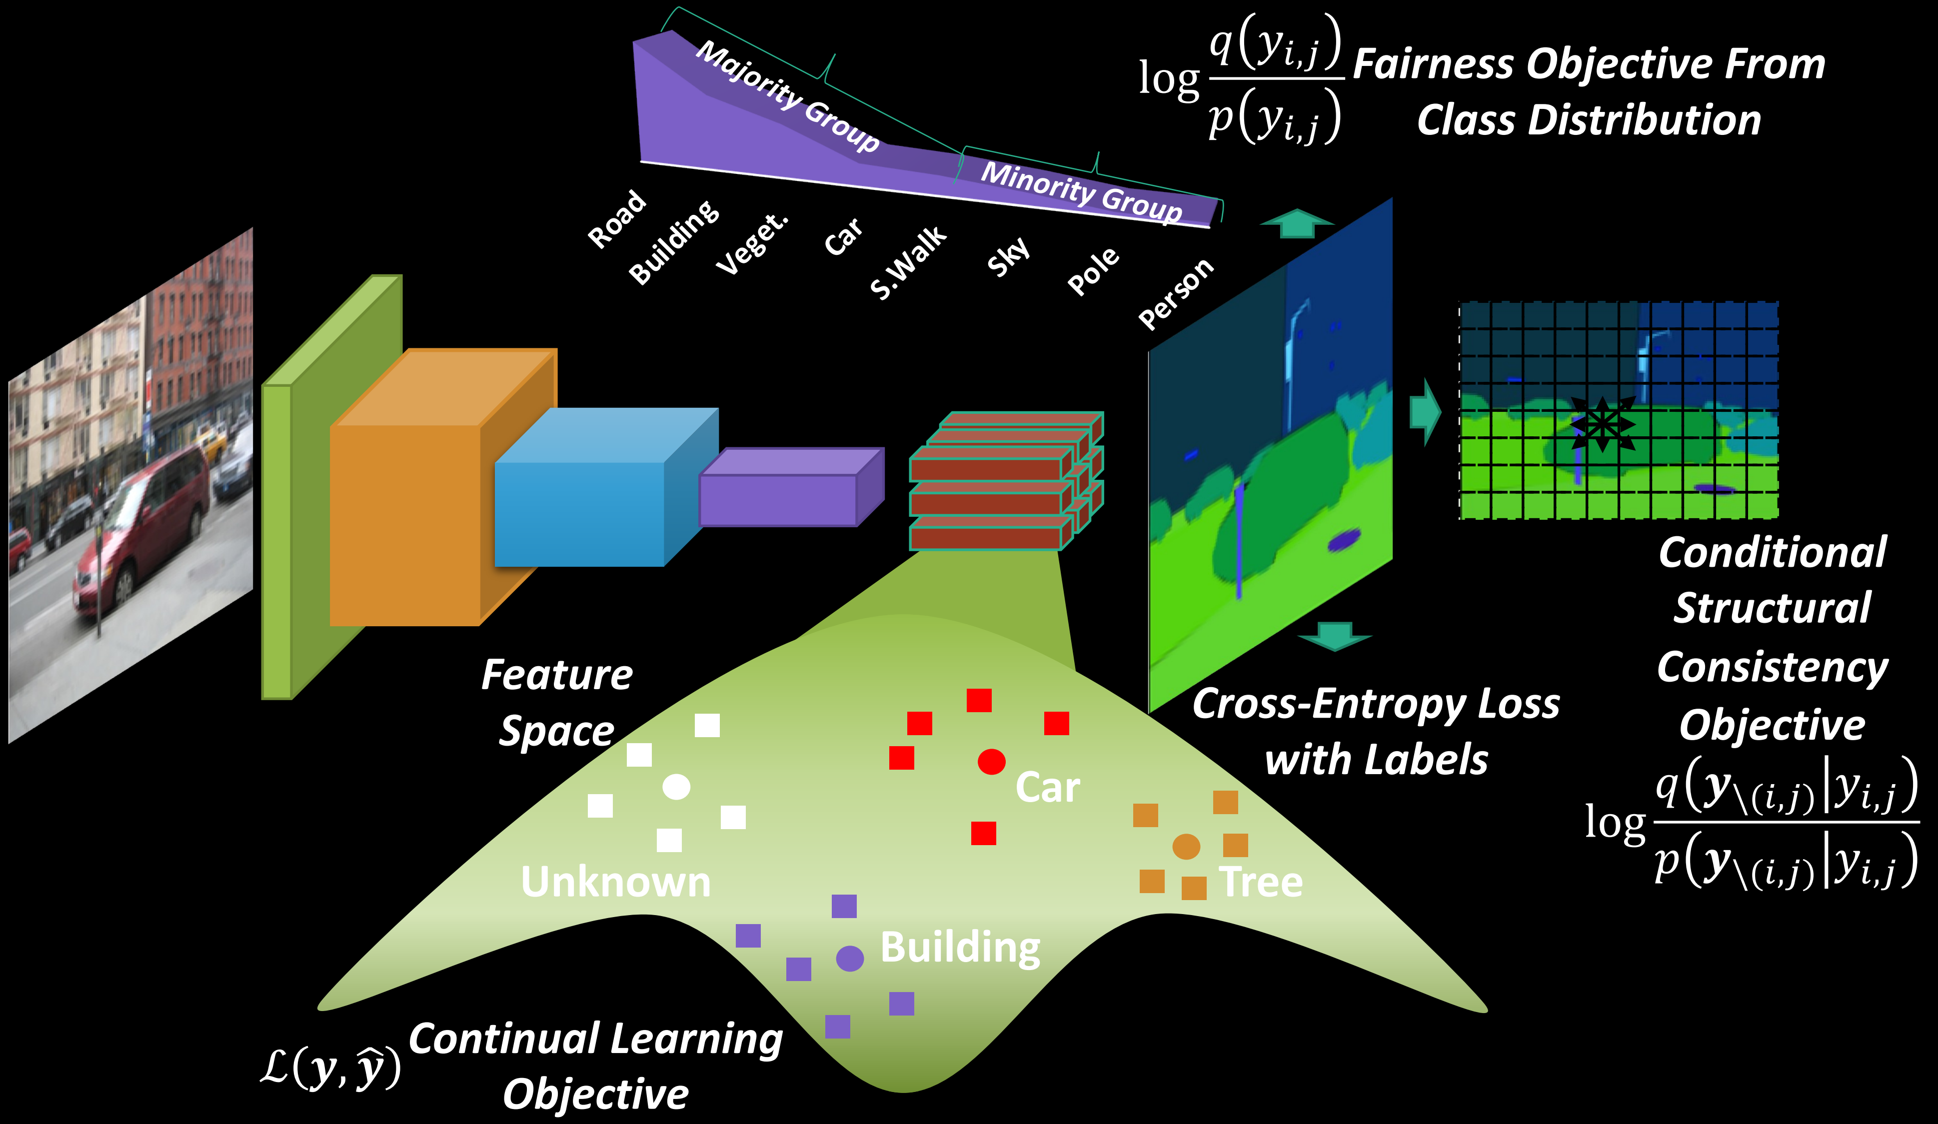

Fairness Continual Learning Framework

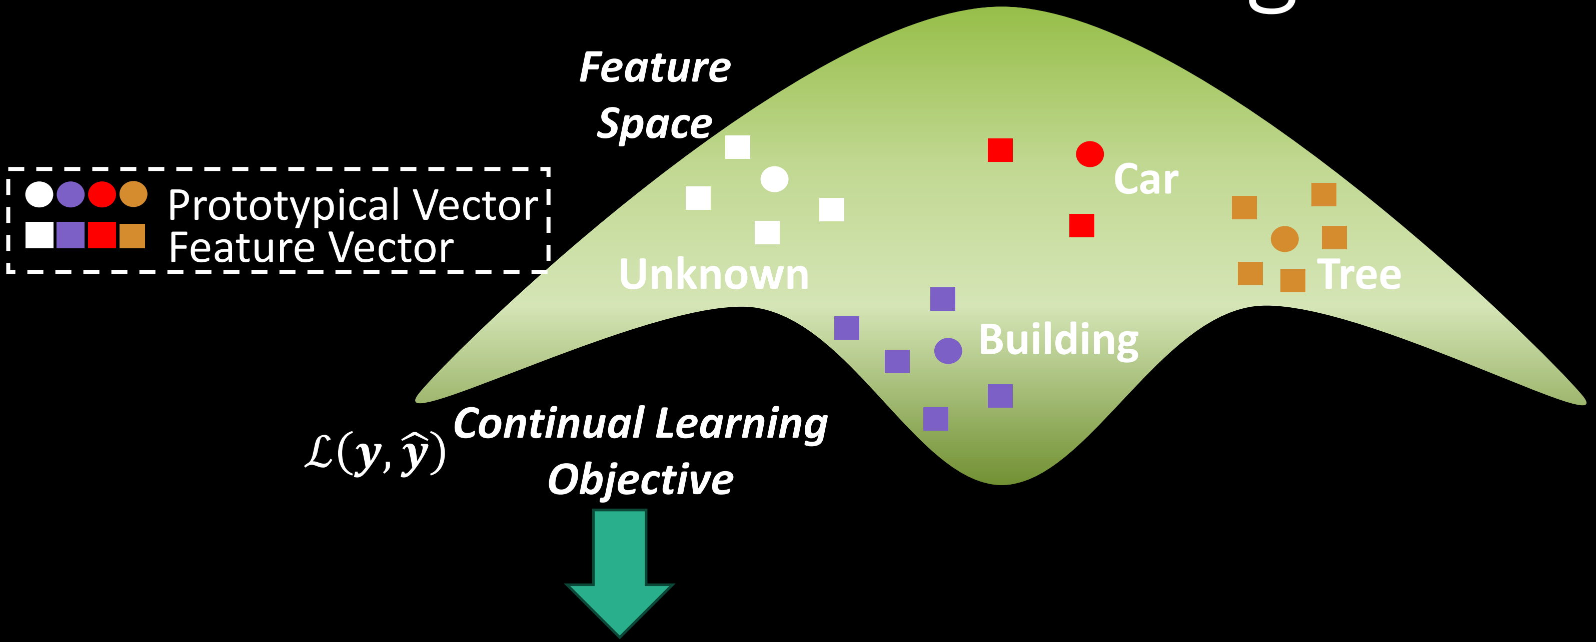

Prototypical Contrastive Clustering Loss

$$\mathcal{L}_{cluster}(\boldsymbol{x}^t, \mathcal{F}, \theta_t) = \sum_{i,j} \sum_{c} \mathcal{D}(\boldsymbol{f}^t_{i,j}, \boldsymbol{p}_c)$$

where

$$\mathcal{D}(\boldsymbol{f}^t_{i,j}, \boldsymbol{p}_c) = \begin{cases} \ell(\boldsymbol{f}^t_{i,j}, \boldsymbol{p}_c) & \text{If } \boldsymbol{f}^t_{i,j} = c \\ \max\left(0, \boldsymbol{\Delta} - \ell(\boldsymbol{f}^t_{i,j}, \boldsymbol{p}_c)\right) & \text{Otherwise} \end{cases}$$

Thank You For Watching