

Politecnico



# Recognizing New Visual Categories in Real Open World Environments

**Doctoral Examination Committee: Prof. Bottino Andrea Prof. Ciliberto Carlo Prof. locchi Luca Prof. Montrucchio Bartolomeo Prof. Zanuttigh Pietro** 





Supervisor: **Prof. Caputo Barbara** 

Candidate: **Fontanel Dario** 



Closed world assumption

### Training



Closed world assumption

### Training



### Deployment



Closed world assumption

### Training



### Deployment



Open world



Train







Open world



Train



Test









Open world



Train



Test









Open world



Train



Test







Knowledge (t)



*Unknown* category found



Open world



Train



Test







Knowledge (t)





Train

*Unknown* category found



Open world



Train



Test





found



Knowledge (t)







Train



Unknown category







Open world





**Unkno** found

Knowledge (t)

Unknown category





Open world





Unknown category found

Knowledge (t)





Open world





Unknown category found

Knowledge (t)











## Outline Recognizing New Visual Categories in Real Open World Environments

## Detection of the unknown "PAnS: Prototopycal Anomaly

Segurane Fabia Cermelli, Massimiliano Mancini, Barbara Caputo

## Incremental learning

"Weakly supervised Incremental Learning for Semantic Section Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

## Unified framework "Boosting Deep Open World Recognition by Clustering"

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo IROS 20

## Impact of domain shift

"On the Challenges of Open World Recognition under Shifting Visual Boria Fantagel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo ICRA 21



Open world





Unknown category found

Knowledge (t)











Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2021, Nashville, USA

## Anomaly Segmentation

It refers to the task of identifying the pixels in an image that belong to unknown categories.





## Anomaly Segmentation

It refers to the task of identifying the pixels in an image that belong to unknown categories.







Æ

/!\





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches





Image

Fully-Convolutional Network

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches





Segmentation output

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches



Network



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches



Image

Fully-Convolutional Network



probability

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches



Image

Fully-Convolutional Network



probability

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches





Image

Fully-Convolutional Network



Segmentation prediciton

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Traditional approaches





Image

Fully-Convolutional Network



Synthesizer



Synthesized image



## Segmentation prediciton



## "PAnS: Prototopycal Anomaly Segmentation" Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Traditional approaches





Image

## **Fully-Convolutional Network**



Synthesizer



Synthesized image

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo **Traditional approaches** 

## Due to semantic segmentation scenes complexity

## Generative approaches tend to produce artifacts while synthesizing pixels of known classes

Reconstruction (top) of SPADE [1] on an image from StreetHazards dataset (middle) [2]. 28











Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework

Confidence values independent for each class

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework

Confidence values independent for each class



No additional computation required

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework

Confidence values independent for each class



No additional computation required Scores bounded so that a threshold can be applied



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework

Confidence values independent for each class







Image

Fully-Convolutional Classifier Per pixel probability Network

No additional computation required Scores bounded so that a threshold can be applied





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

## Framework

Confidence values independent for each class





No additional computation required Scores bounded so that a threshold can be applied



Cosine classifier



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework



Scores bounded so that a threshold can be applied



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Framework




Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo Experimental setting

# AUPR

# AUROC

# area under the Precision-Recall curve

## Architecture

ResNet-50 architecture [67] as backbone and PSPNet [26] as head module.



area under the True Positive Rate and False Positive Rate curve

# **FPR95**

# measures the False Positive Rate (FPR) at 95% of recall

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Results

## **StreetHazards**

D.Hendrycs, et al. A benchmark for Anomaly Segmentation (2019)





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Results



D.Hendrycs, et al. A benchmark for Anomaly Segmentation (2019)

Method	AUPR	AUROĈ	FPR95
AE	2.2	66.1	91.7
Dropout	7.5	69.9	79.4
MSP	6.6	87.7	33.7
MSP + CRF	6.5	88.1	29.9
SynthCP	9.3	88.5	28.4
PAnS	8.8	91.1	23.2





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Results

## **StreetHazards**

D.Hendrycs, et al. A benchmark for Anomaly Segmentation (2019)

Method	AUPIR	AUROĈ	FPR95
AE	2.2	66.1	91.7
Dropout	7.5	69.9	79.4
MSP	6.6	87.7	33.7
MSP + CRF	6.5	88.1	29.9
SynthCP	9.3	88.5	28.4
PAnS	8.8	91.1	23.2



## Image

## Semantic labels

PAnS



## Preliminaries

Open world





Unknown category found

Knowledge (t)



Knowledge (t+1)









Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2022, New Orleans, USA

Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

Task overview

Suppose we are given a model pretrained on a set of classes, using expensive pixel-level annotations

## Train on {person, motorbike, car}



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

Task overview

Suppose we are given a model pretrained on a set of classes, using expensive pixel-level annotations

Our goal is to incrementally segment new classes over time using only image-level labels

## Train on {person, motorbike, car}



44





Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

Experimental setting

## Comparisons

Fully supervised IL methods

Standard fully supervised incremental learning approaches serving as upper bounds

Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

## Experimental setting

## Comparisons

Fully supervised IL methods

Standard fully supervised incremental learning approaches serving as upper bounds

Weakly-supervised methods Used to generate pseudo-labels offline and train the same IL model

Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

## Experimental setting

## Comparisons

Fully supervised IL methods

Standard fully supervised incremental learning approaches serving as upper bounds

Weakly-supervised methods Used to generate pseudo-labels offline and train the same IL model

## Architecture

ResNet-101 architecture [67] as backbone and Deeplav v3 [26] as head module.

Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

## Experimental setting

## Comparisons

#### Fully supervised IL methods

Standard fully supervised incremental learning approaches serving as upper bounds

#### Weakly-supervised methods

Used to generate pseudo-labels offline and train same IL model the

### Architecture

ResNet-101 architecture [67] as backbone and Deeplav v3 head module

[26] as

## 20 semantic classes + background 10k images



The PASCAL Visual Object Classes (VOC) Challenge. Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J. and Zisserman, A, IJCV10

## Pascal VOC 2012

15-5 15 classes first, 5 classes added

10-10 10 classes first, 10 classes added

Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

## Experimental setting

## Comparisons

#### Fully supervised IL methods

Standard fully supervised incremental learning approaches serving as upper bounds

#### Weakly-supervised methods

Used to generate pseudo-labels offline and train same IL model the

#### Architecture

ResNet-101 architecture [67] as backbone and Deeplav v3 head module.

[26] as



Pascal VOC 2012

15-5

20 semantic classes + background 10k images

15 classes first,

5 classes added

## 10-10

10 classes first,

10 classes added



Fabio Cermelli, Dario Fontanel, Antonio Tavera, Marco Ciccone, Barbara Caputo

Results

15 - 5





image

CAM

Old classes: car New classes: sheep, motorbike

10-10

GT

54

SS

SEAM

EPS

WILSON



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2020, Las Vegas, USA IEEE Robotics and Automation Letters 2020



## Preliminaries

Open world





Unknown category found

Knowledge (t)



Knowledge (t+1)





















Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo Global clustering loss

> Scores based on distances between samples features and respective class centroids

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo Global clustering loss

> Scores based on distances between samples features and respective class centroids



low score major update



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo Global clustering loss

> Scores based on distances between samples features and respective class centroids

The network learns to keep representations of images as close as possible to the centroids of the respective class



low score major update



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo Global clustering loss

> Scores based on *distances* between samples features and respective class centroids

The network learns to keep representations of images as close as possible to the centroids of the respective class

**Cross entropy** loss



# high score minor update

major update

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Local clustering loss





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Local clustering loss

Scores based on *distances* between features of samples of different classes





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Local clustering loss

Scores based on *distances* between features of samples of different classes

The network learns to keep samples of the same class closer, while pushing away samples of different classes





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Local clustering loss

Scores based on *distances* between features of samples of different classes

Local clustering loss

The network learns to keep samples of the same class closer, while pushing away samples of different classes





## Preliminaries

Open world





Unknown category found

Knowledge (t)



Knowledge (t+1)









## "Boosting Deep Open World Recognition by Clustering (B-DOC)" Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Class-specific rejection thresholds** 



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Class-specific rejection thresholds** 



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Class-specific rejection thresholds** 



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set


Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Class-specific rejection thresholds** 



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Class-specific rejection thresholds** 



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set

> Maximum distance loss

If the **distance** between a sample of class y and its class centroid is greater than  $\Delta_{v}$ , the loss increases  $\Delta_{v}$ .





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

#### **Class-specific rejection thresholds**



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set

> Maximum distance loss

If the distance between a sample of class y and its class centroid is greater than  $\Delta_v$ , the loss increases  $\Delta_{v}$ .





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

#### **Class-specific rejection thresholds**



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set

#### Maximum distance loss

If the distance between a sample of class y and its class centroid is greater than  $\Delta_v$ , the loss increases  $\Delta_v$ . If the **distance** between a sample *not* in class y and y class centroid is **less** than  $\Delta_v$ , the loss





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

#### **Class-specific rejection thresholds**



To exploit the clustered feature space, we introduce class-specific rejection thresholds  $\Delta$ based on distances, *learned* on an held out set

#### Maximum distance loss

If the distance between a sample of class y and its class centroid is greater than  $\Delta_v$ , the loss increases  $\Delta_v$ . If the **distance** between a sample *not* in class y and y class centroid is **less** than  $\Delta_v$ , the loss





### "Boosting Deep Open World Recognition by Clustering (B-DOC)" Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo Experimental settings



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Experimental settings

Without Unknown (no rejection)



Test set does **not** contain **unknown** classes

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

Experimental settings

Without Unknown (no rejection)



Test set does **not** contain **unknown** classes Without Unknown (with rejection)

Test set does **not** contain **unknown** classes but the model has the possibility to **reject** 

#### "Boosting Deep Open World Recognition by Clustering (B-DOC)" Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Experimental settings** 

Without Unknown (no rejection)



Test set does **not** contain unknown classes

Open World Recognition (Avg)

Average between rejection) and

Without Unknown (with With Unknown (with rejection)

Without Unknown (with rejection)

Test set does **not** contain unknown classes but the model has the possibility to **reject** 

#### "Boosting Deep Open World Recognition by Clustering (B-DOC)" Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

**Experimental settings** 

Without Unknown (no rejection)



Test set does **not** contain unknown classes

Open World Recognition (Avg)

Average between Without Unknown (with rejection) and

With Unknown (**with rejection**)

Without Unknown (with rejection)

Test set does **not** contain unknown classes but the model has the possibility to **reject** 

Open World Recognition (Harm)

Harmonic mean between Without Unknown (with rejection) and With Unknown (with rejection)

Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

#### Results

#### **RGB-D Object**

Categories: 26 known + 25 unknown *Incremental steps: 11 first + 5 each* 

K. Lai, L. Bo, X. Ren, D. Fox, A large-scale hierarchical multi-view rgb-d object dataset, ICRA-11



Without Unknown (no rejection)

#### Without Unknown (with rejection)



Open World Recognition (Harm)



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo', Elisa Ricci, Barbara Caputo

#### Results

#### **CORE-50**

#### Categories: 5 known + 5 unknown Incremental steps: 2 first + 1 each



V. Lomonaco and D. Maltoni, Core50: a new dataset and benchmark for continuous object recognition, CoRL-17



Without Unknown (no rejection)

#### Without Unknown (with rejection)



Open World Recognition (Harm)



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo IEEE International Conference on Robotics and Automation (ICRA) 2021, Xi'an, China IEEE Robotics and Automation Letters 2021

Closed domain assumption

#### Training



#### Deployment



Closed domain assumption

#### Training



#### Deployment







Closed domain assumption

#### Training



#### Deployment





Closed domain assumption

#### Training



#### Deployment

2 . . .

# Are OWR models robust to domain shift?



Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Benchmark

- objects
- Different acquisition conditions
- incremental step) Synkuu

"On the Challenges of Open World Recognition under Shifting Visual Domains"

# • Three datasets containing the same 51 daily-life

# Standard protocol: 26 (known) + 25 (unknown) 11 (initial step) + 5 (each









Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

#### Benchmark

# synROD



Markus, RAL 20

### Synthetic dataset generated utilizing publicly freely available 3D models.

## **Reproduction of realistic lighting.**

Unsupervised Domain Adaptation through Inter-modal Rotation for RGB-D Object Recognition, Loghmani Mohammad Reza, Robbiano Luca, Planamente Mirco, Park Kiru, Caputo Barbara and Vincze





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Benchmark

### **Objects captured in a** controlled scenario.

No clutter or changes in illumination or background, only varying camera angles.



# ROD





Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Benchmark

### More challenging environment with objects depicted against various backgrounds, scales, views, lighting, and occlusions.

Recognizing Objects In-the-wild: Where Do We Stand?, Loghmani Mohammad Reza, Caputo Barbara and Vincze Markus, ICRA 18

# ARID









Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Are OWR models Robust to Domain Shift?

Synthetic-to-Real









Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Are OWR models Robust to Domain Shift?



### **Constrained-to-Unconstrained**





Closed domain assumption

#### Training



D D

#### Deployment

# Can single source domain generalization solve the issue?

0 10



Single source domain generalization

#### Source domain



Single source domain generalization

#### Source domain



#### Target domain



















Single source domain generalization

#### Source domain



#### Target domain s























Single source domain generalization

#### Source domain



#### Target domains





















-













Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

#### Benchmark

#### Data augmentation with transformation sets (RSDA)

An evolutionary-based algorithm select the set of transformations that results in the worst model performance and randomly applies a subset of them.



Visual content taken from Volpi, R., & Murino, V.. Addressing model vulnerability to distributional shifts over image transformation sets. (ICCV 19)







#### Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

#### Benchmark

Data augmentation with transformation sets (RSDA)

An evolutionary-based algorithm select the set transformations that results in the worst model performance and randomly applies a subset of

them.



#### Self-supervised learning with relative rotations (RR)

By employing an auxiliary self-supervised task, the model concentrates on discriminative invariances and regularities, enhancing generalization to new domains



Visual content taken from Gidaris, S., Singh, P., & Komodakis, N.. Unsupervised representation learning by predicting image rotations. (ICLR18)





#### Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

#### Benchmark

#### Data augmentation with transformation sets (RSDA)

An evolutionary-based algorithm select the set transformations that results in the worst model performance and randomly applies a subset of



#### Self-supervised learning with relative rotations (RR)

By employing an auxiliary self-supervised magget moncentrates on discriminative and regulasities, enhancing generalization to demains.









Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Can DG methods address the problem?



### Synthetic-to-Real

#### Source: synROD Source: synROD 30 **B-DOC** DeepNNO 25 25.2 25.3 DeepNNO-RSDA 23.2 DeepNNO-RR Accuracy 20 DeepNNO-SC 19.1 17.3 15 14.6 10 9.9 10.0 9.3 8.0 5 4.9 4.3 0 ROD ARID ARID Open World Recognition (Harm) Without Unknown (with rejection)







Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Barbara Caputo

Can DG methods address the problem?

### **Constrained-to-Unconstrained**





### Conclusion

Recognizing New Visual Categories in Real Open World Environments

Detection of the unknown

We introduced PAnS that learns class-specific prototypes overcoming the limitations of the softmax operation in anomaly segmentation.

#### Incremental learning

We introduced WILSON that learns to generate pseudo labels to train a segmentation network from image level labels only, reducing annotation costs.

#### Unified framework

We introduced B-DOC that adopts a global-to-clustering training loss objective, and a learnable rejection threshold per each class.

### Impact of domain shift

We introduced the first benchmark to assess OWR methods under shifting visual domains and laid the foundations for further research.

### Future works

Recognizing New Visual Categories in Real Open World Environments

**Open World Recognition methods** OWR approaches need to collect and label a new set of images every time a new class is discovered, usually involving a human in the loop.

### Future works

Recognizing New Visual Categories in Real Open World Environments

**Open World Recognition methods** OWR approaches need to collect and label a new set of images every time a new class is discovered, usually involving a human in the loop.

labelled

• Adopt an active learning pipeline to identify the most meaningful data to be
## Future works

Recognizing New Visual Categories in Real Open World Environments

**Open World Recognition methods** OWR approaches need to collect and label a new set of images every time a new class is discovered, usually involving a human in the loop.

- labelled

• Adopt an active learning pipeline to identify the most meaningful data to be

• Investigate how to generate pseudo-labels to avoid the labelling process

## Future works

Recognizing New Visual Categories in Real Open World Environments

**Open World Recognition methods** OWR approaches need to collect and label a new set of images every time a new class is discovered, usually involving a human in the loop.

- labelled

• Adopt an active learning pipeline to identify the most meaningful data to be

• Investigate how to generate pseudo-labels to avoid the labelling process

Adopt a few shot approach that requires limited amounts of labelled data



# Thank you for the attention!

