Bookworm continual learning: beyond zero-shot learning and continual learning

Kai Wang, **Luis Herranz**, Anjan Dutta, Joost van de Weijer

Learning and machine perception (LAMP) group
Computer Vision Center
Universidad Autónoma de Barcelona

ContinualAI reading group, October 2020

Computer Vision Center (UAB campus)

Only Center in Europe fully devoted to Computer Vision

23 Years

+130 Staff

+20 Nationalities

M€2,3 Income / year

8 PhD thesis /year

+100 Intl publications / year

BOOKWORM AND GENERALIZED CONTINUAL LEARNING

Continual learning

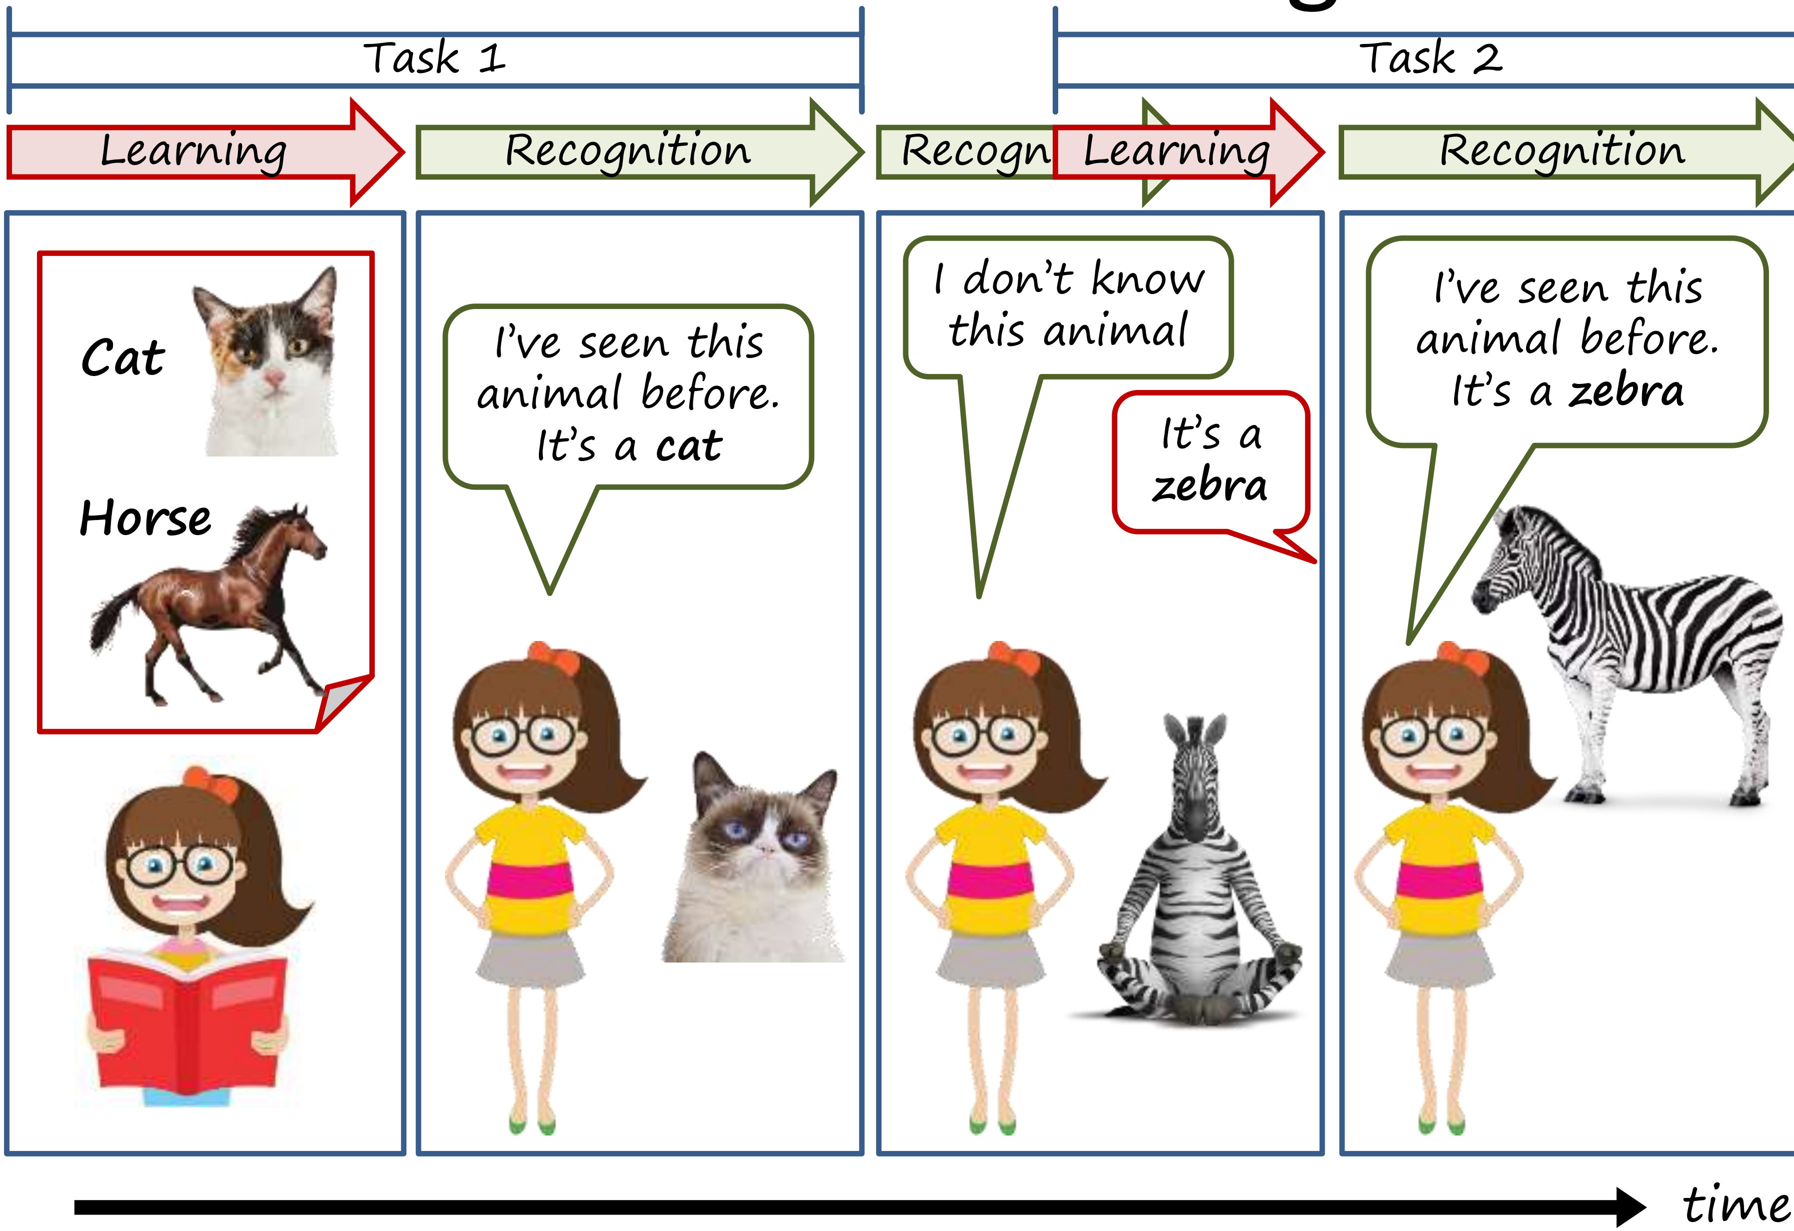

How to address continual learning?

- Regularization-based approaches
  - Prior focused (i.e. regularize parameters)
  - Data-focused (i.e. regularize features)
- Replay approaches
  - Rehearsal (e.g. exemplars)
  - Pseudorehearsal **(feature generation)**
- Parameter isolation approaches
  - Fixed architecture
  - Dynamic architecture

De Lange et al., T. Continual learning: A comparative study on how to defy forgetting in classification tasks, arxiv 2019.

Generative replay

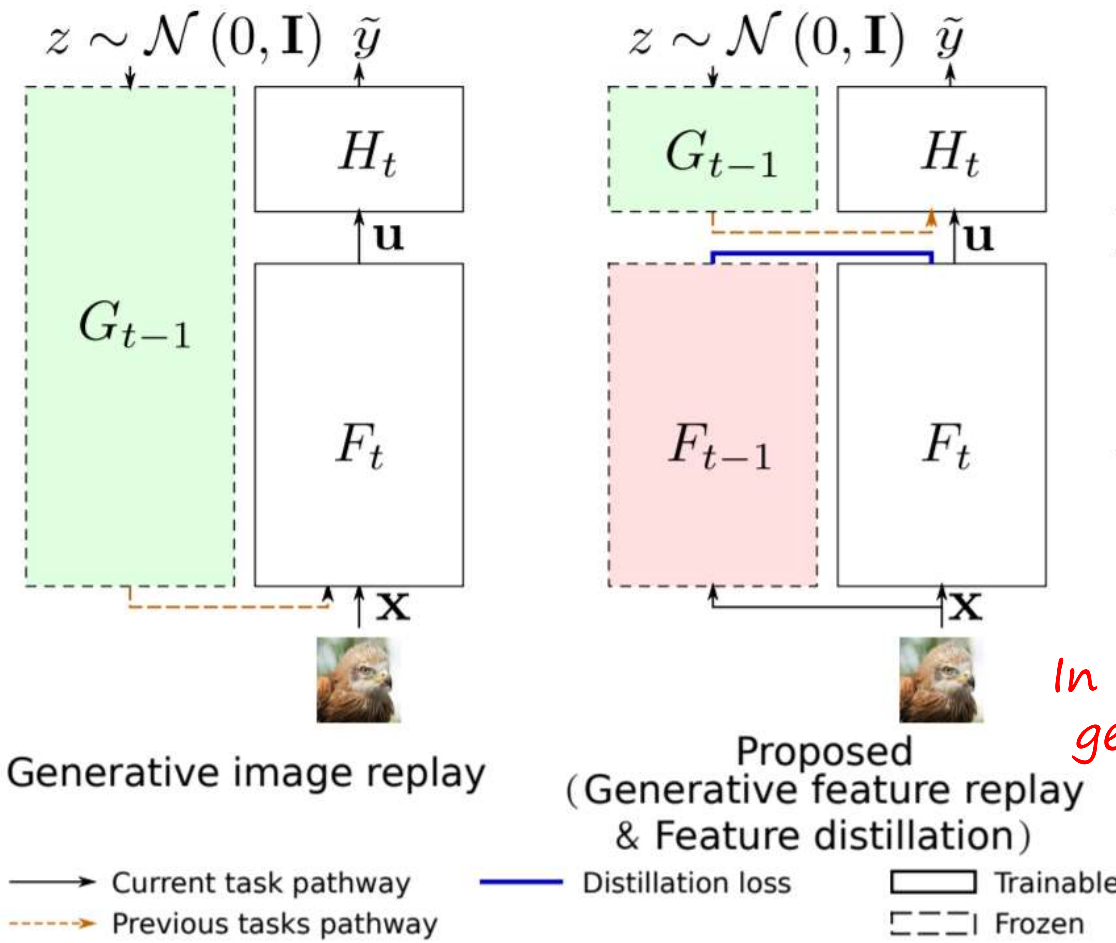

4 tasks on CIFAR-100
GAN-based replay

	T1	T2	T3	T4
Image (MeRGAN)	82.4	37.7	17.8	9.7
Block 1	80.7	41.6	26.5	20.1
Block 2		41.0	26.5	20.0
Block 3		51.1	37.0	26.6
Block 4 (Ours)		57.6	48.2	41.5

Feature replay performs better than image replay

In this work we also use feature generation, but based on VAE

[Liu et al., Generative Feature Replay For Class-Incremental Learning, CLVISION workshop at CVPR 2020](#)

Zero-shot learning

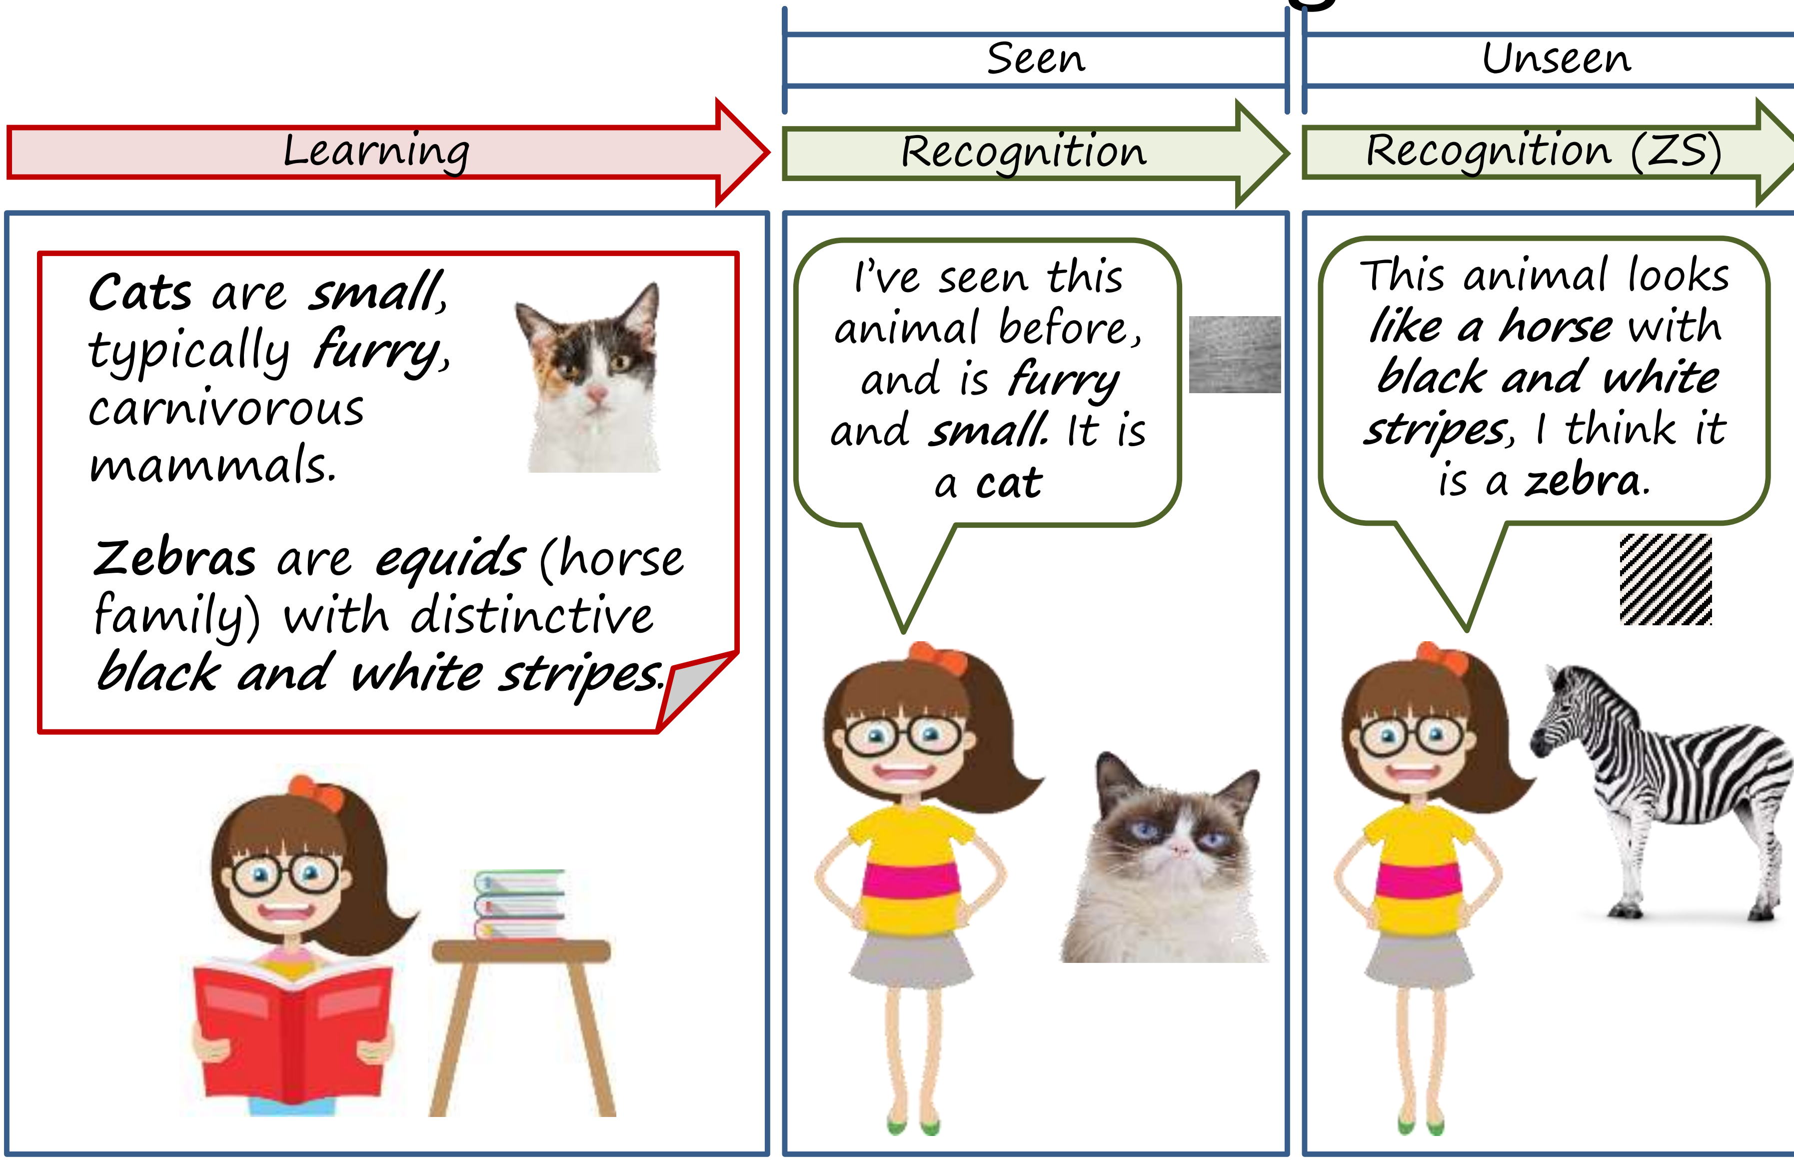

Bookworm continual learning

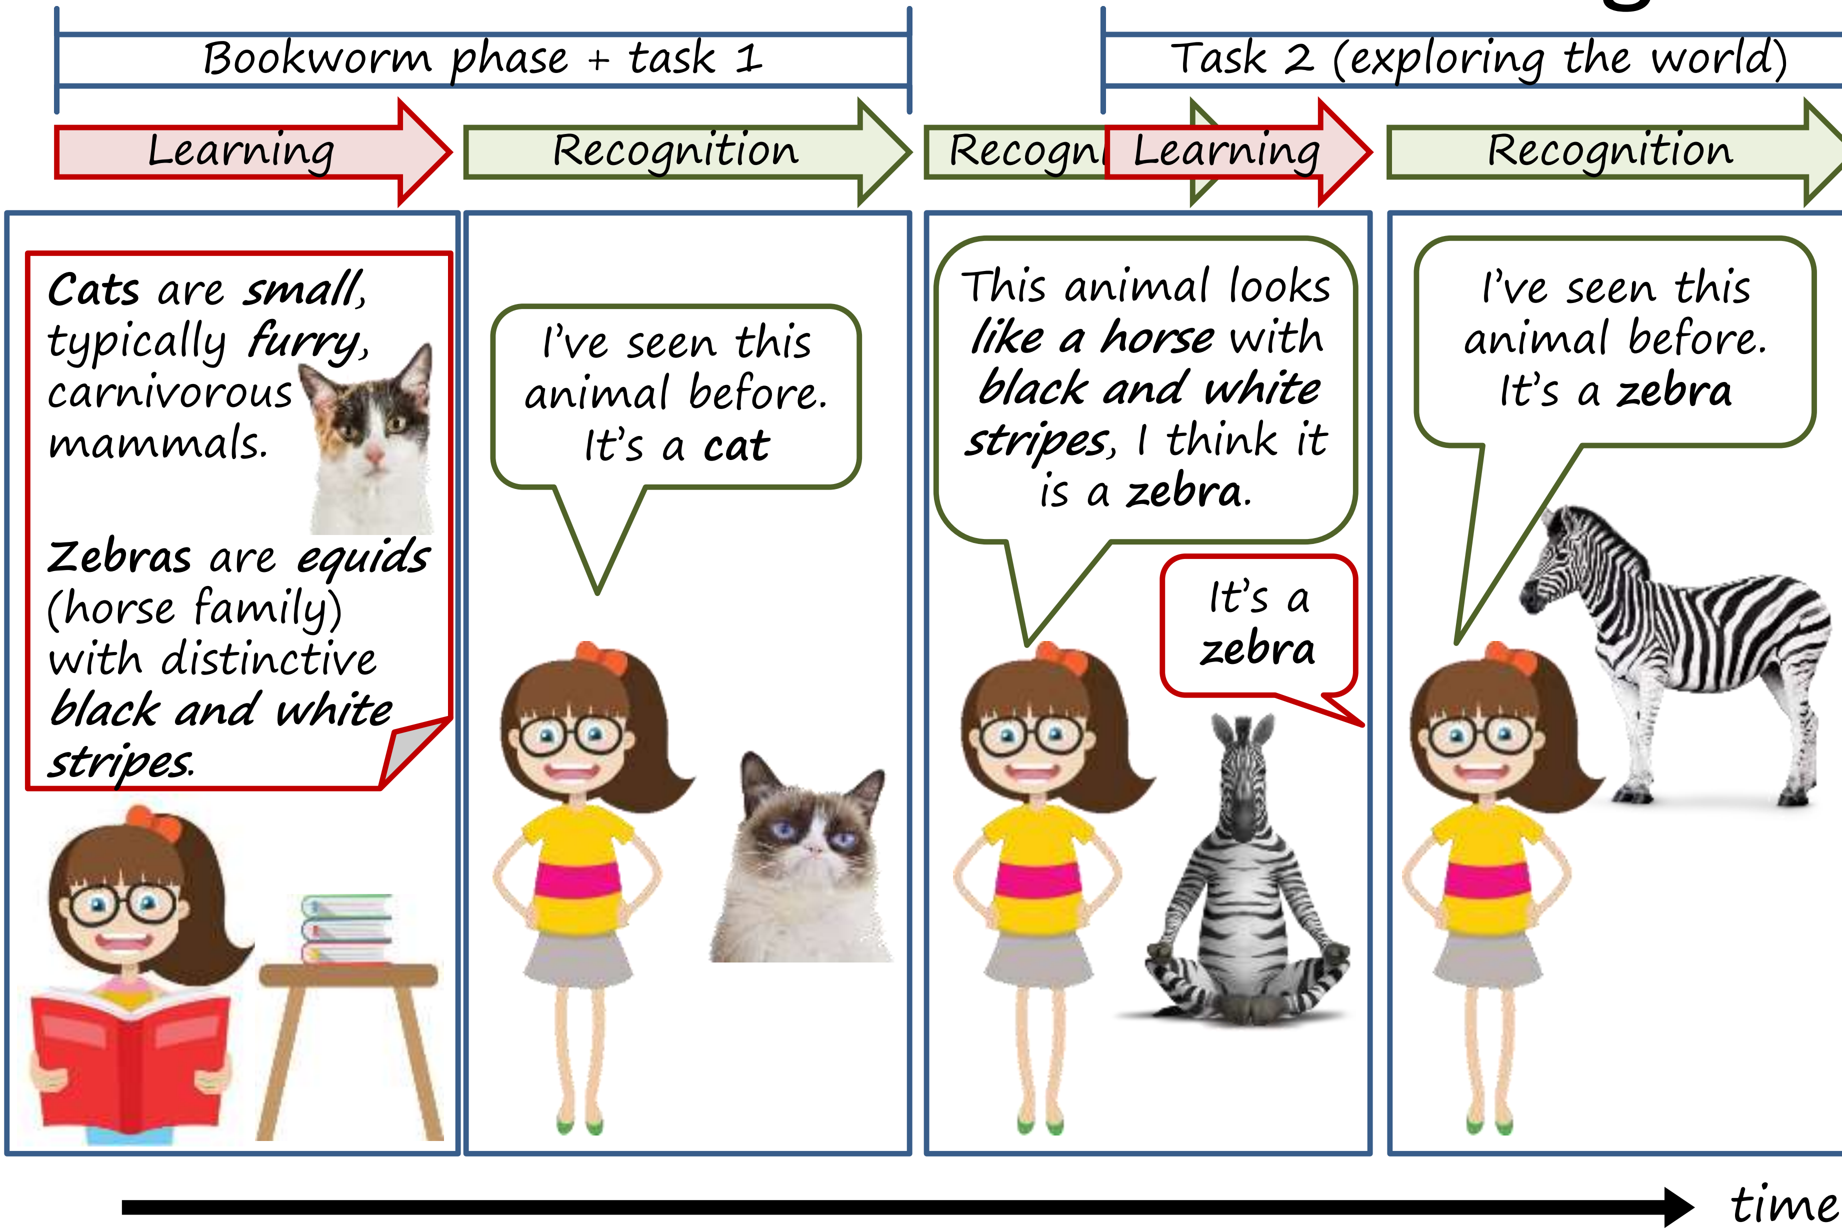

Zero-shot classification

Training

Dog, cat, polar bear, horse
(seen categories)

CNN

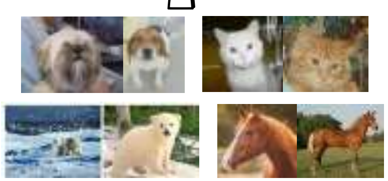

Test

Zebra
(unseen category)

CNN

How to model the **semantic relations** between seen and unseen categories?

A **semantic model**:
- Attributes (e.g. white, black, stripes, four legs, ...)
- Language (via word embeddings)
- ...

Zero-shot learning via visual-semantic alignment

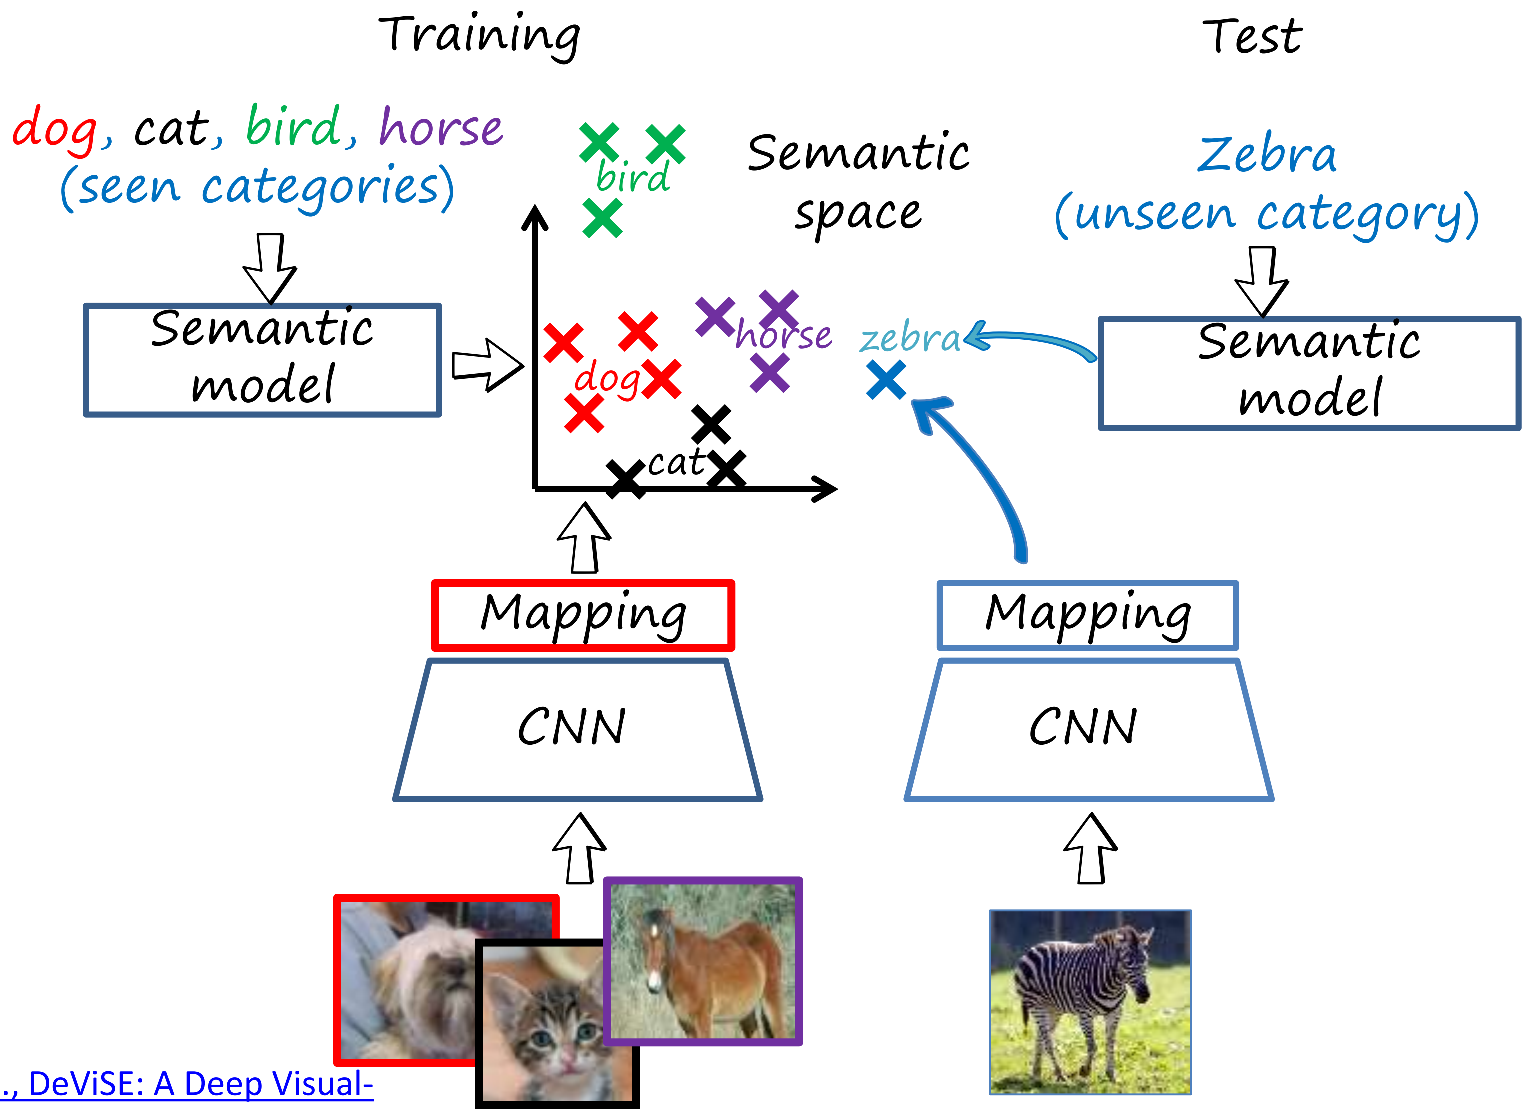

Frome et al., DeViSE: A Deep Visual-Semantic Embedding Model, NIPS 2013

Zero-shot learning with feature generation

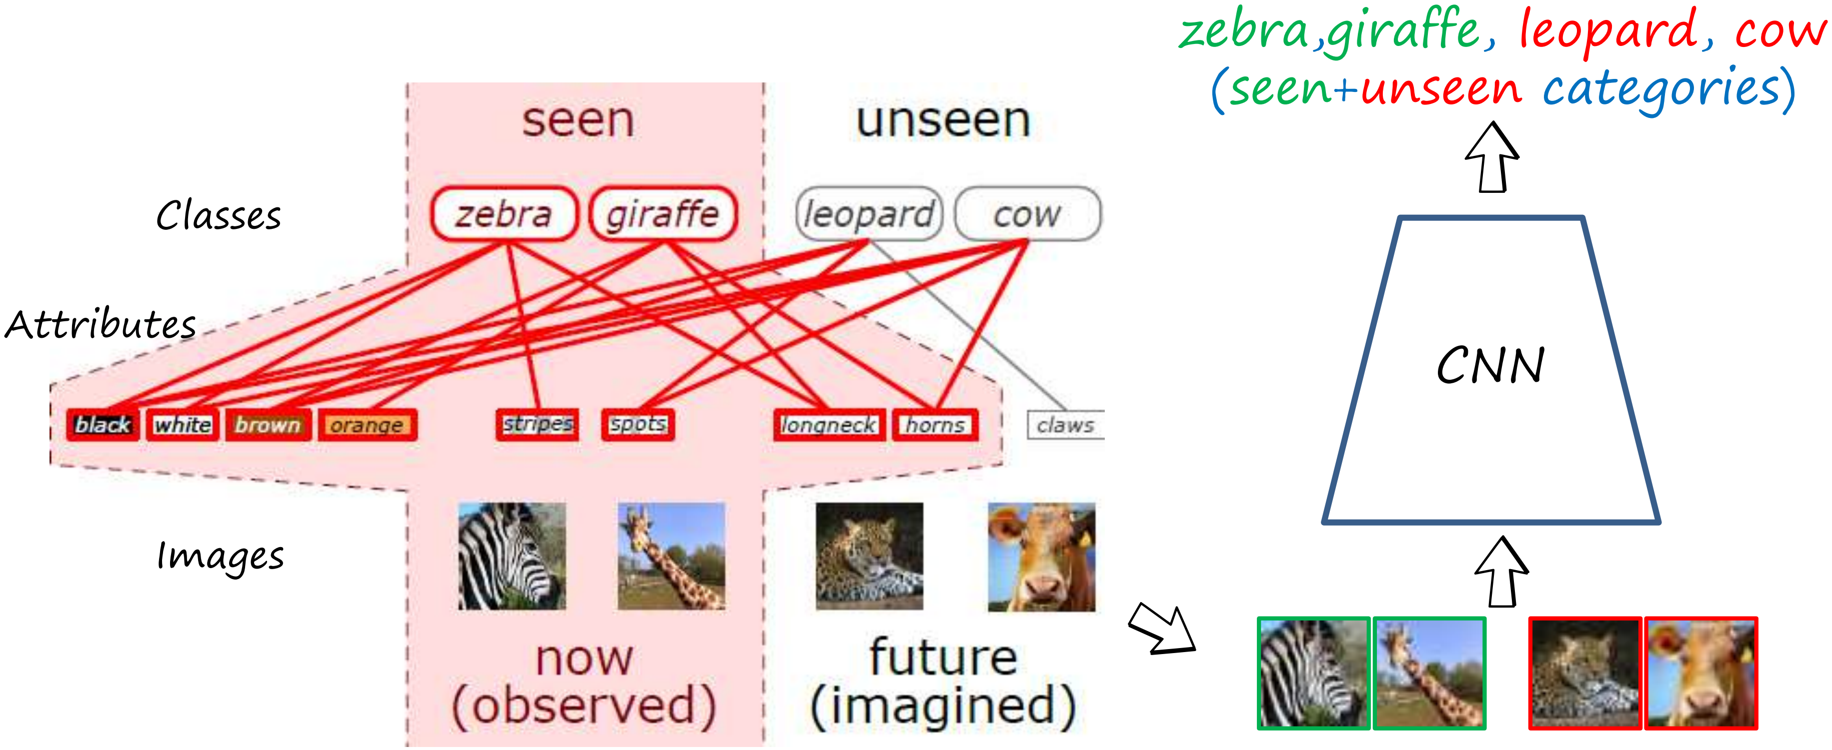

Zero-shot learning as continual learning (learning)

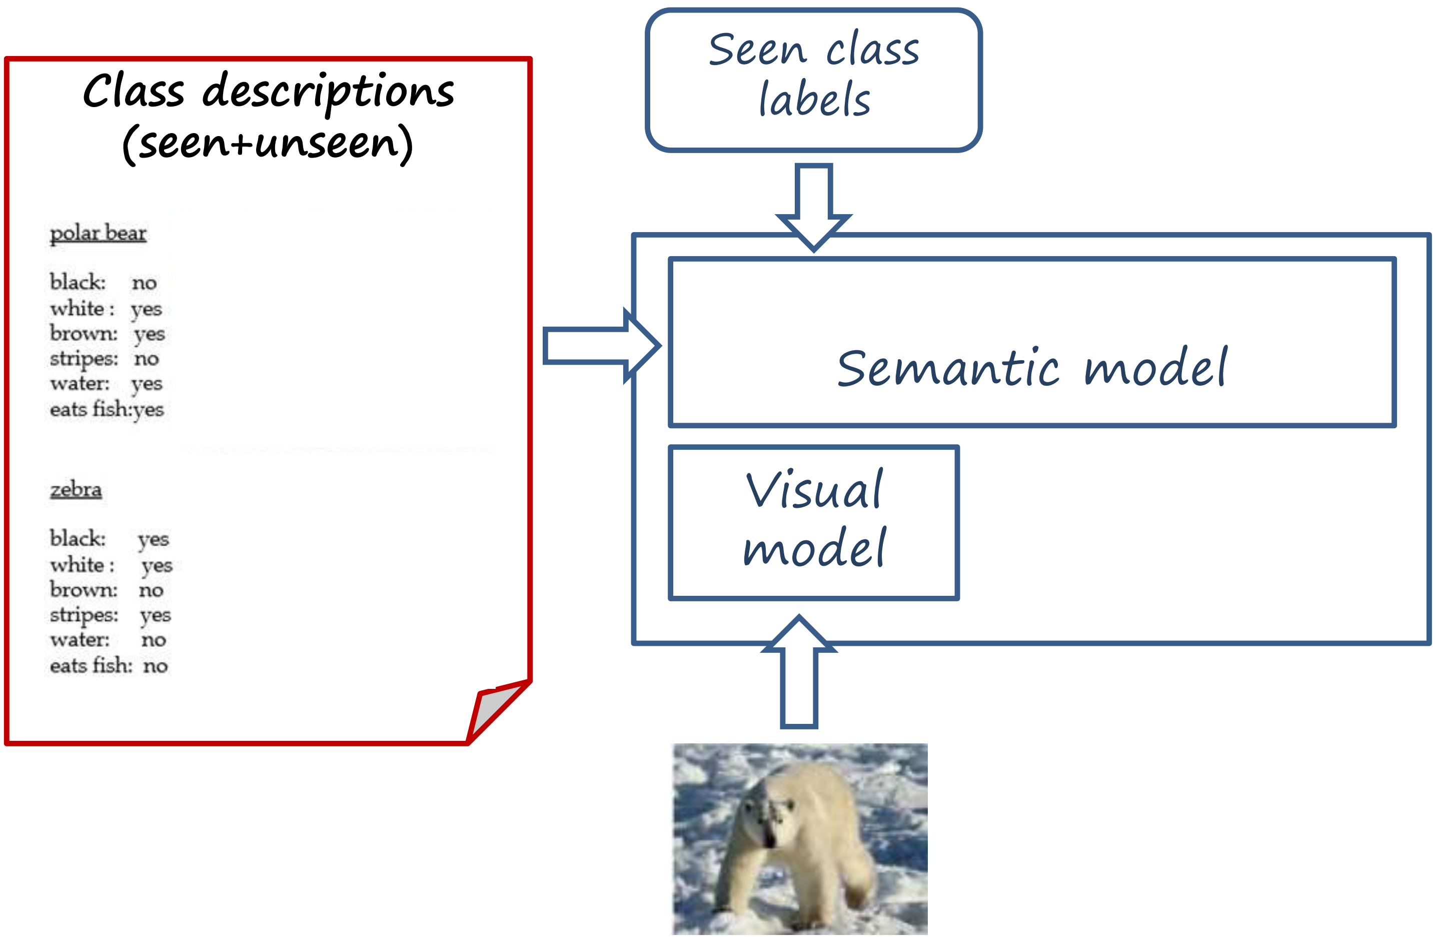

Zero-shot learning as continual learning (inference+evaluation)

ZSL (only unseen)

Generalized ZSL
(both seen+unseen)

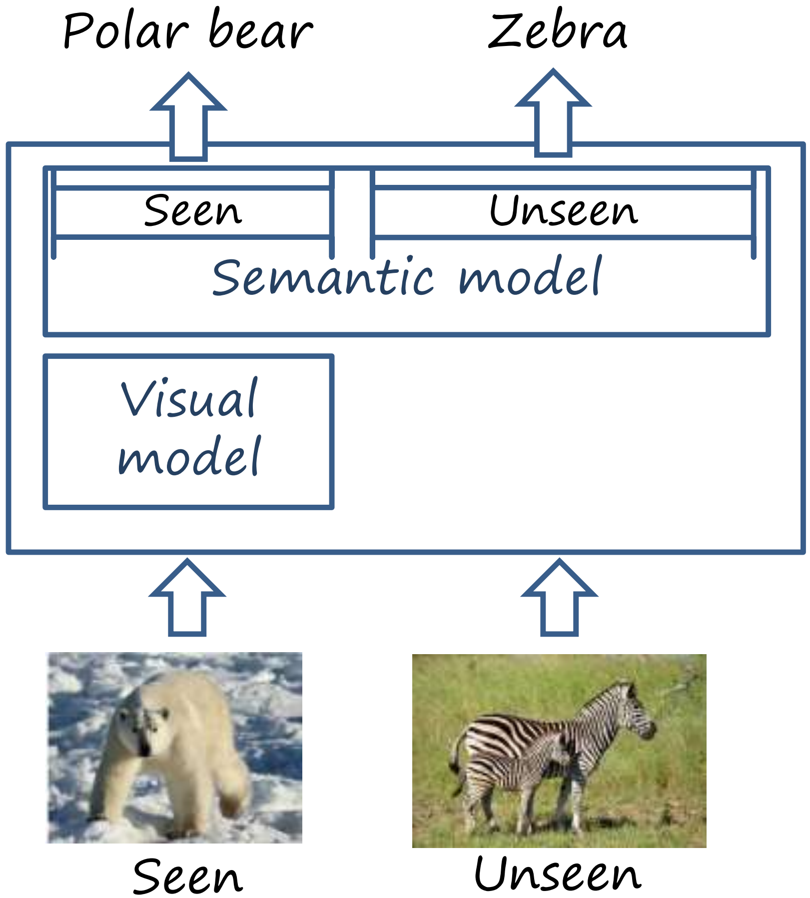

Why don't we
keep updating
the **visual model**
(continually)?

Bookworm continual learning (learning)

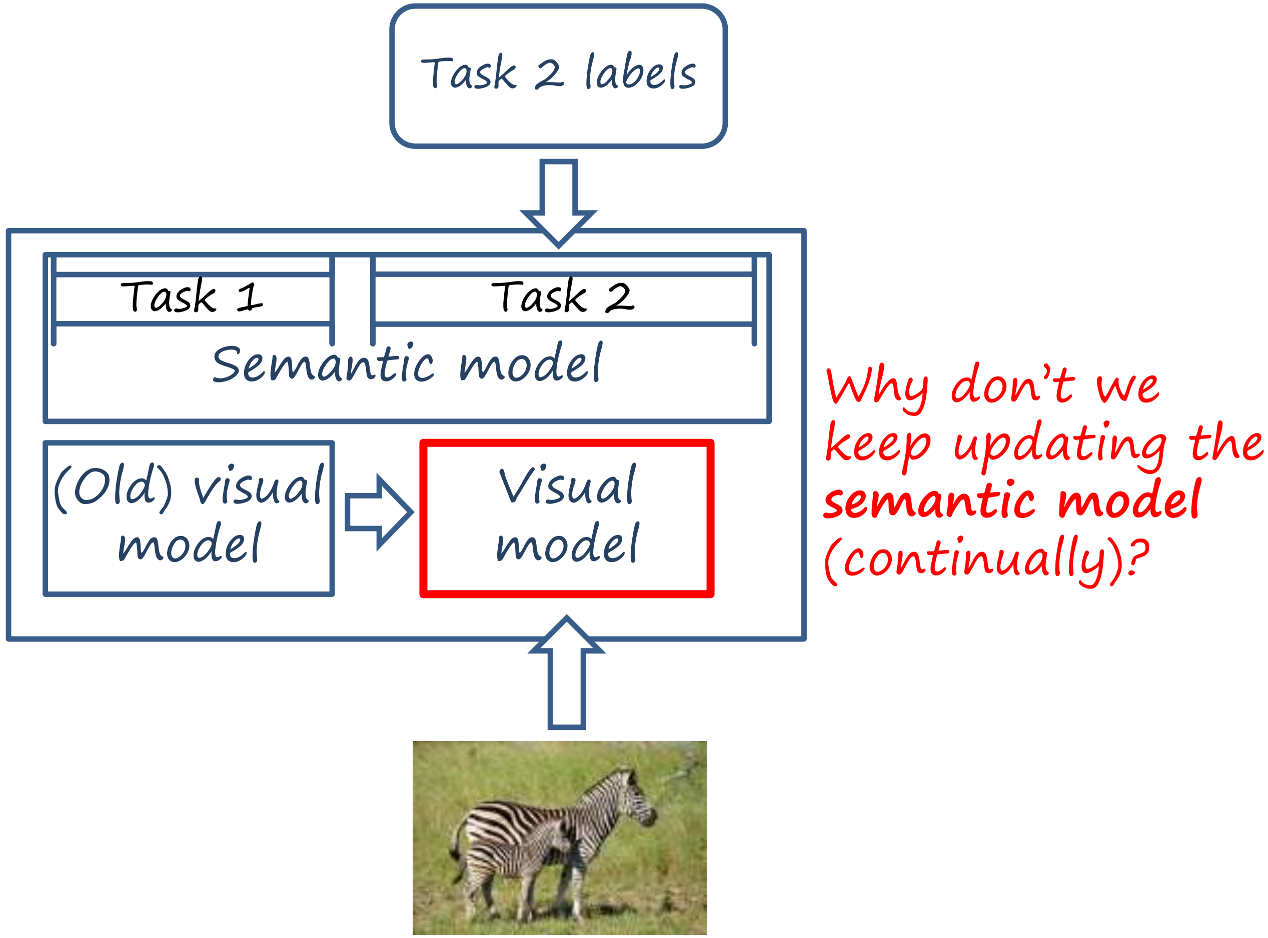

CL vs (G)ZSL vs BCL

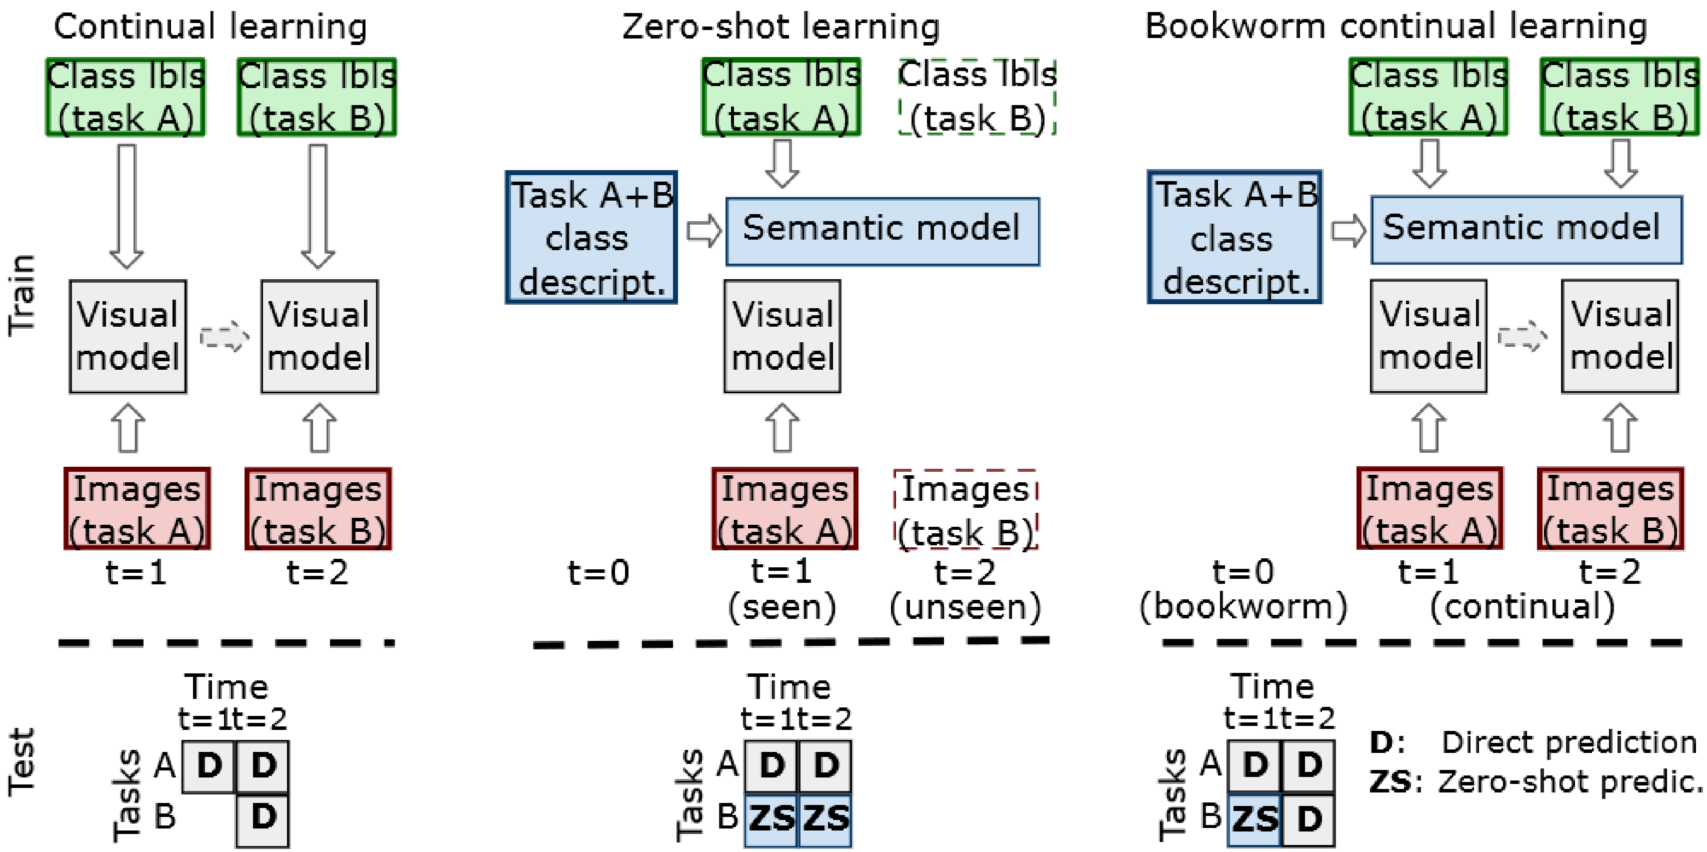

Generalized continual learning

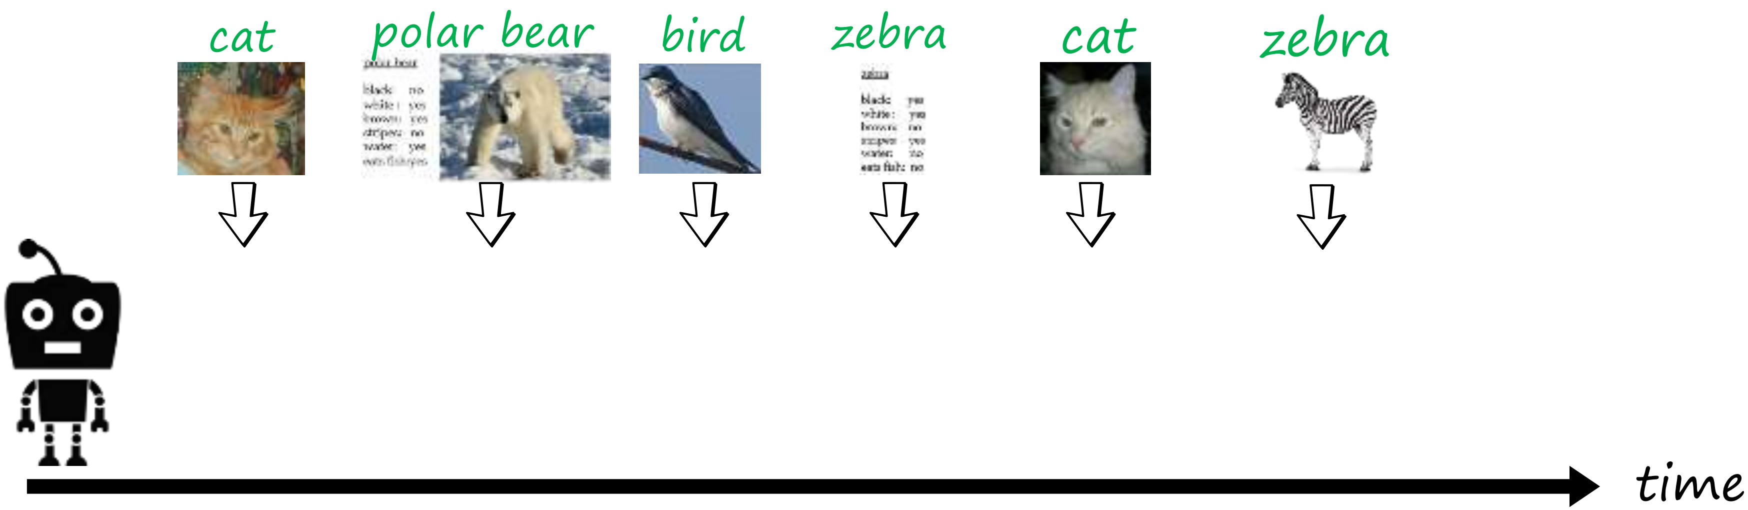

Summary

Setting	Sub-models Vis.	Sub-models Sem.	Continual Vis.	Continual Sem.	Predictions Seen	Predictions Unseen
JT	✓	✗	✗	-	✓	✗
CL	✓	✗	✓	-	✓	✗
ZSL	✓	✓	✗	✗	✗	✓
GZSL	✓	✓	✗	✗	✓	✓
BCL	✓	✓	✓	✗	✓	✓
GCL	✓	✓	✓	✓	✓	✓

Settings exploiting semantic relations

JT: joint training
CL: continual learning
ZSL: zero-shot learning
GZSL:generalized ZSL
BCL: bookworm CL
GCL: generalized CL

AN APPROACH TO BCL: BIDIRECTIONAL IMAGINATION

Overview of the approach

- Requirements: alleviate catastrophic forgetting and enable zero-shot inference
- Our framework is based on feature generation
  - Same mechanism (i.e. replay/imagination) in both continual learning to prevent forgetting and ZSL to infer unseen classes
  - Relatively easy to extend to BCL
- Approach: Bidirectional Imagination (*BImag*)
  - We generate both past and future classes
- We use VAE as generator

Continual learning

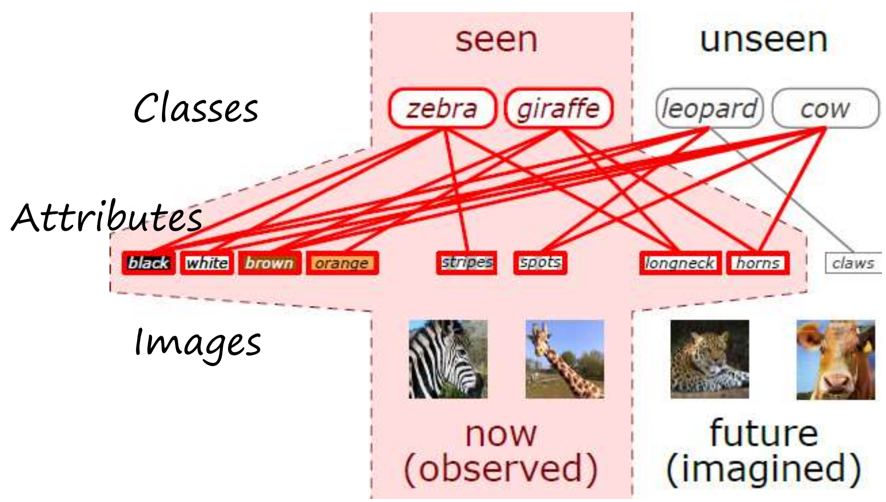

(Generalized) zero-shot learning

Naive approach: *attr-BImag*

Attribute-conditional generation

- Generative feature replay
- Hierarchical generative replay
  - Attribute generator+feature generator
- Semantic model: class descr. matrix A

$$\boldsymbol{a} = G(\boldsymbol{y}) = A\boldsymbol{y}$$

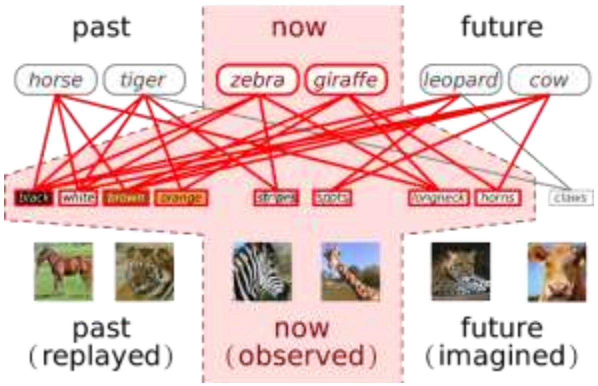

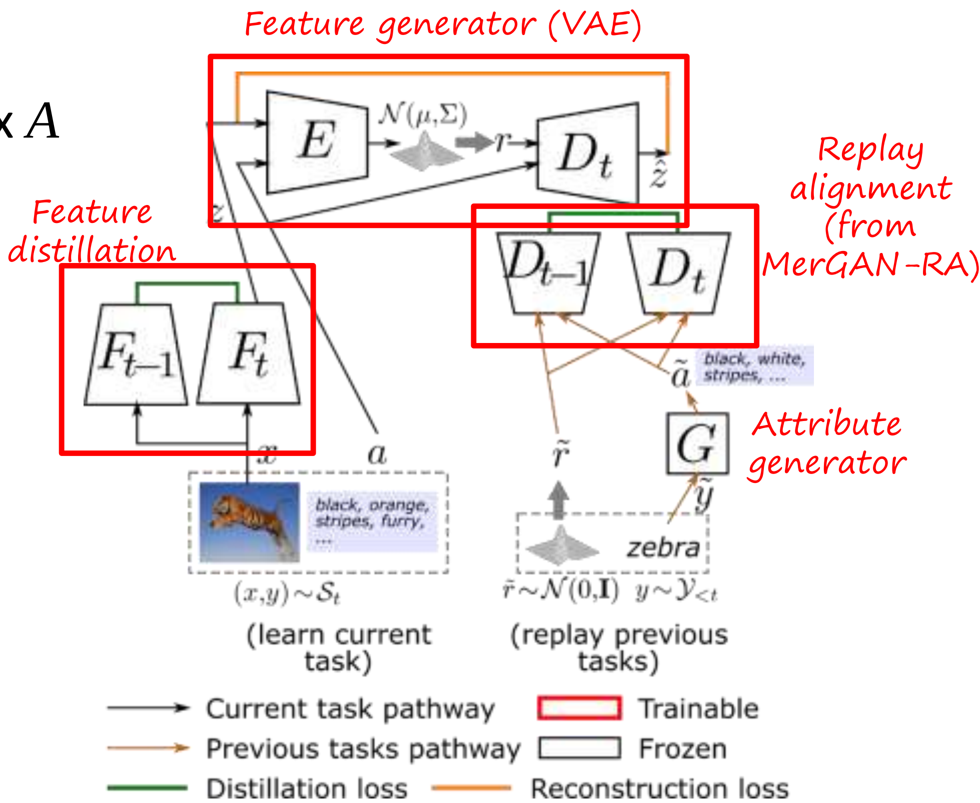

[X. Liu et al., Generative Feature Replay For Class-Incremental Learning, CLVISION@CVPR 2020.]
[C. Wu *et al.*, "Memory Replay GANs: learning to generate images from new categories without forgetting", NeurIPS 2018]

attr-Bimag. Training and inference

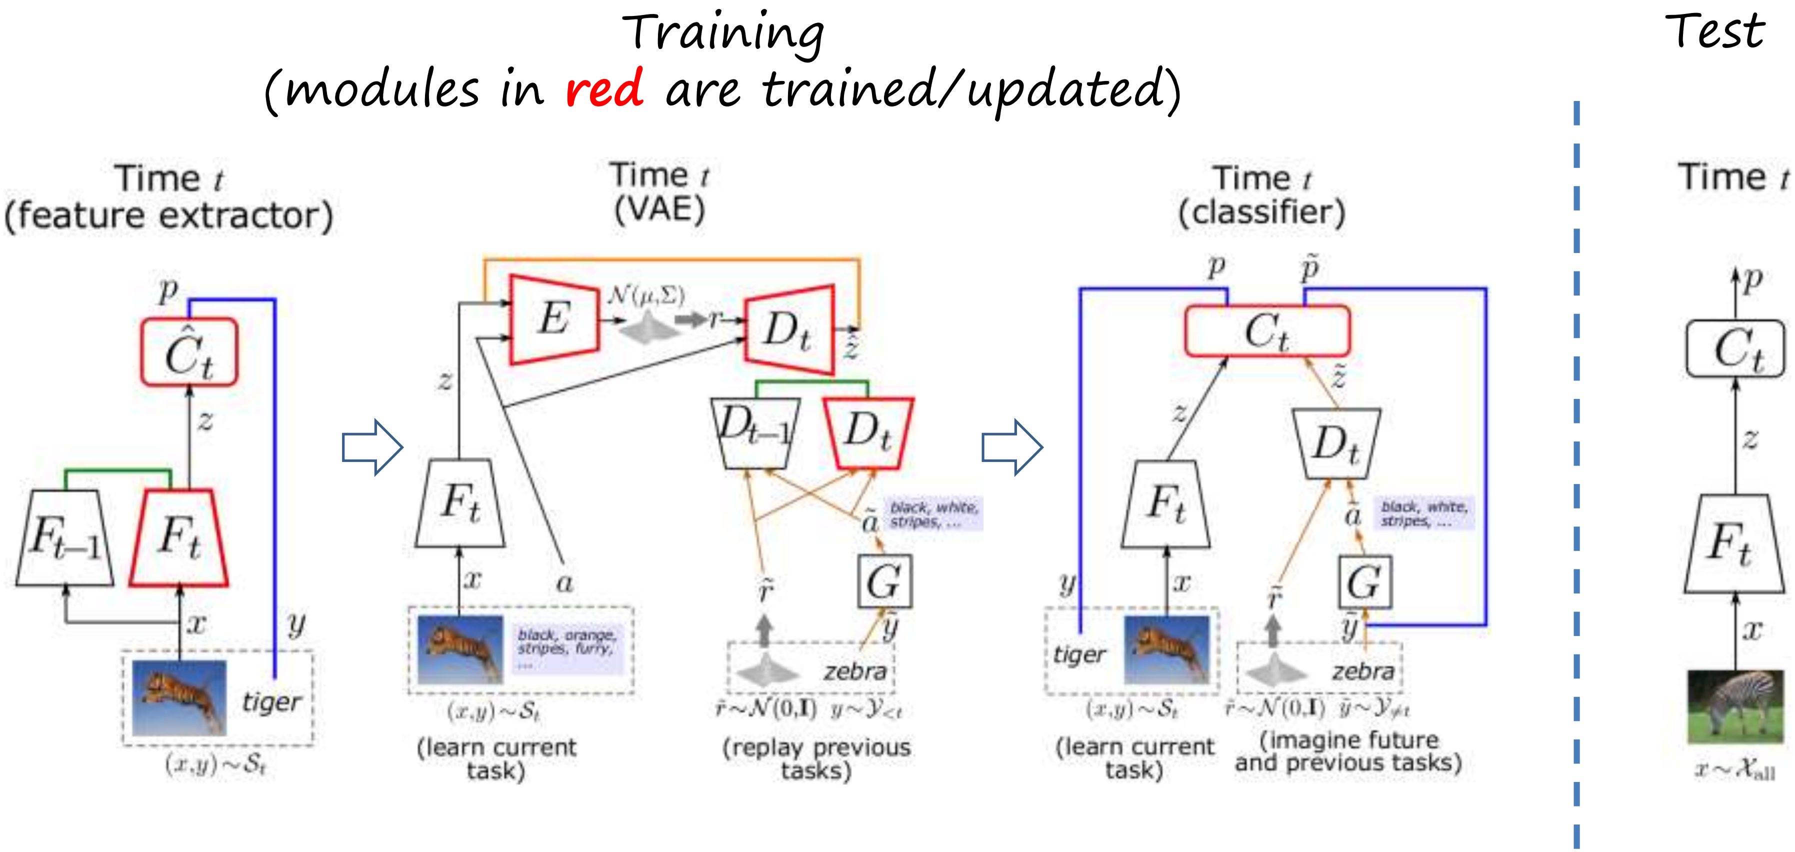

class-BImag

- Non-informative semantic model $A = \mathbf{I}$
  - Then $\boldsymbol{a} = \mathbf{y}$, i.e. continual learning model
  - *class-BImag*: our continual learning baseline
- *attr-Bimag* performs **worse** than *class-Bimag* preventing forgetting
  - The semantic model seems to harm generative replay of past tasks

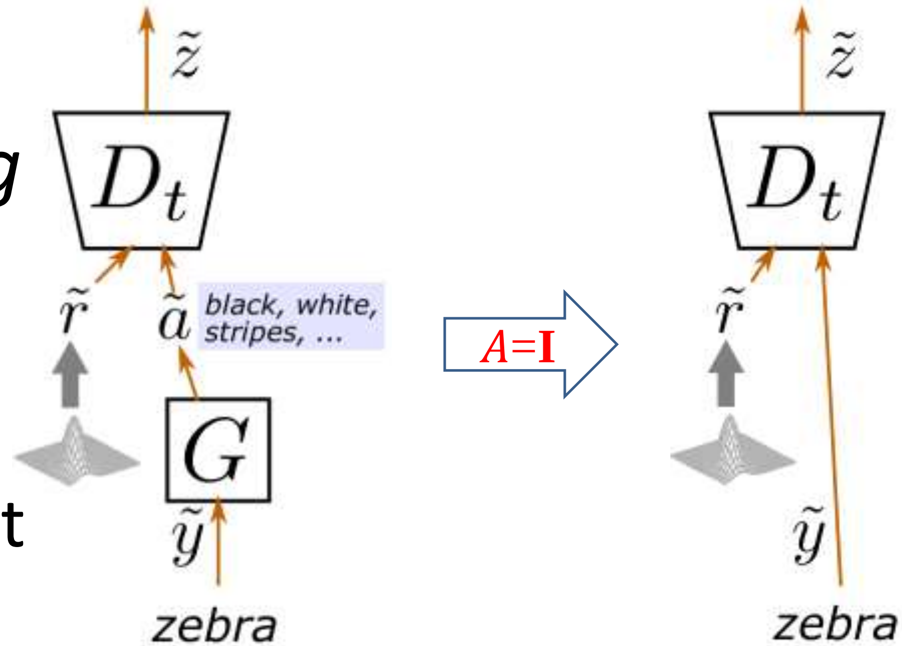

A practical example

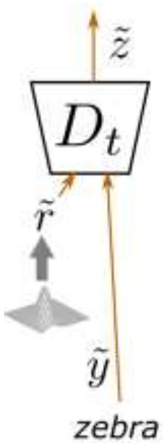

Class cow (note the high diversity)

class-Bimag. VAE observes the visual information directly.

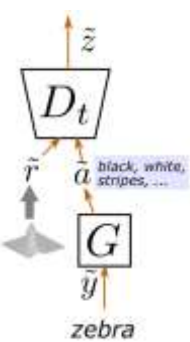

attr-Bimag. VAE doesn't observe the visual information directly but attributes.

cow →

black: 40%
white: 60%
brown: 40%
patches: 40%

This average description is not representative in the case of high diversity.

Possibles solutions

- No direct connection between visual and class information
  - Idea: include class as condition in VAE (class-attr-BImag)
- Asymmetric generation
  - Problem: the VAE has more information about the past (visual and attributes) than about the future (attributes)
  - Idea: decouple forward and backward replay directions (asym-BImag)
- Deterministic semantic models cannot capture all the visual diversity in classes
  - Idea: use stochastic models instead of deterministic

class-attr-BImag

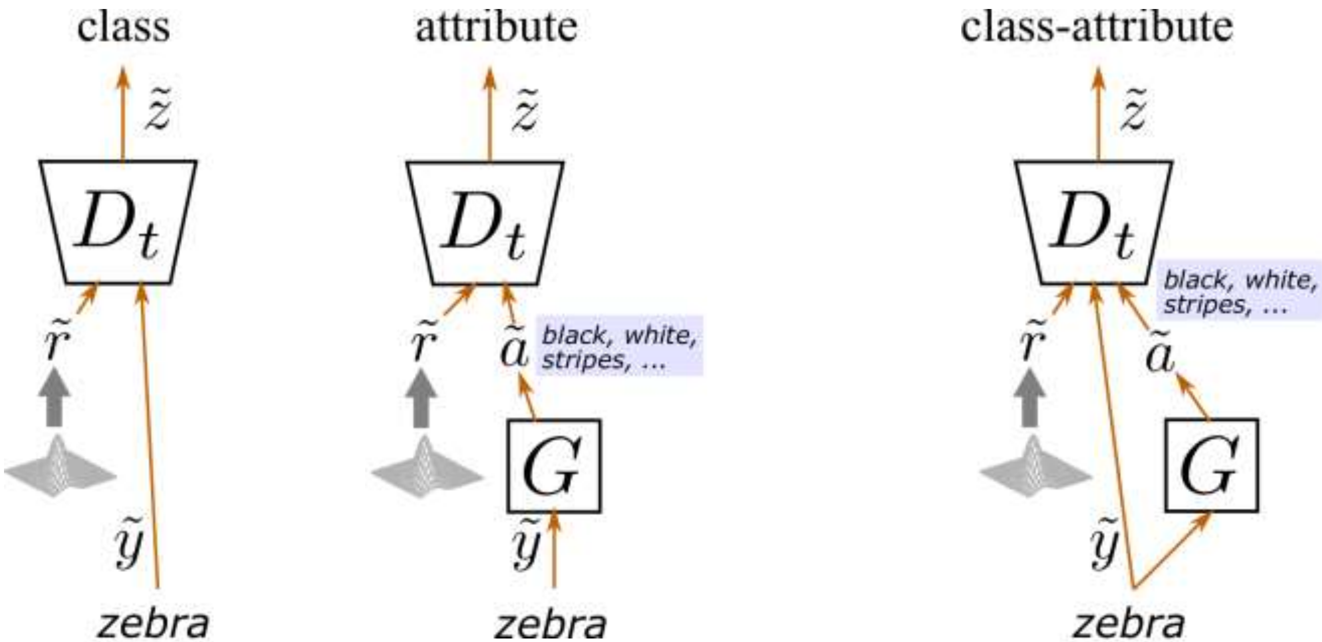

In practice we concatenate class and attribute vectors

asymmetric-BImag

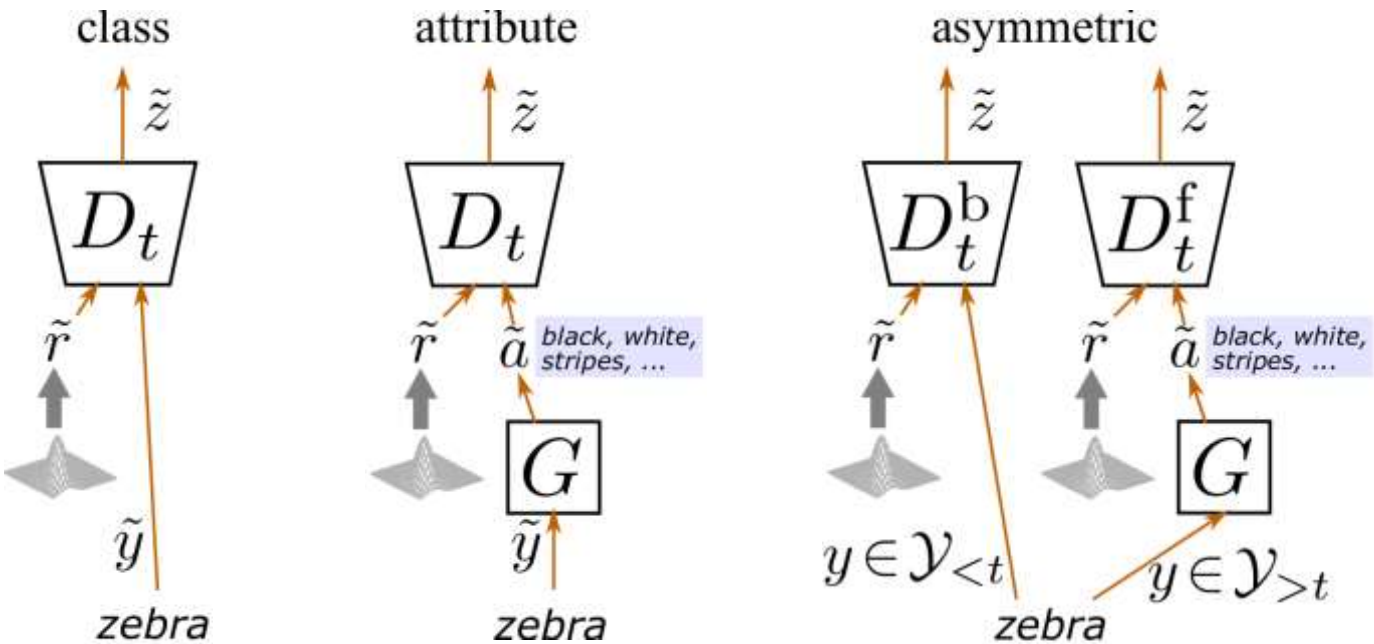

One VAE for future In practice we concatenate class and attribute vectors

Stochastic semantic model

Problem: not clear how to get instance descriptions of unseen classes

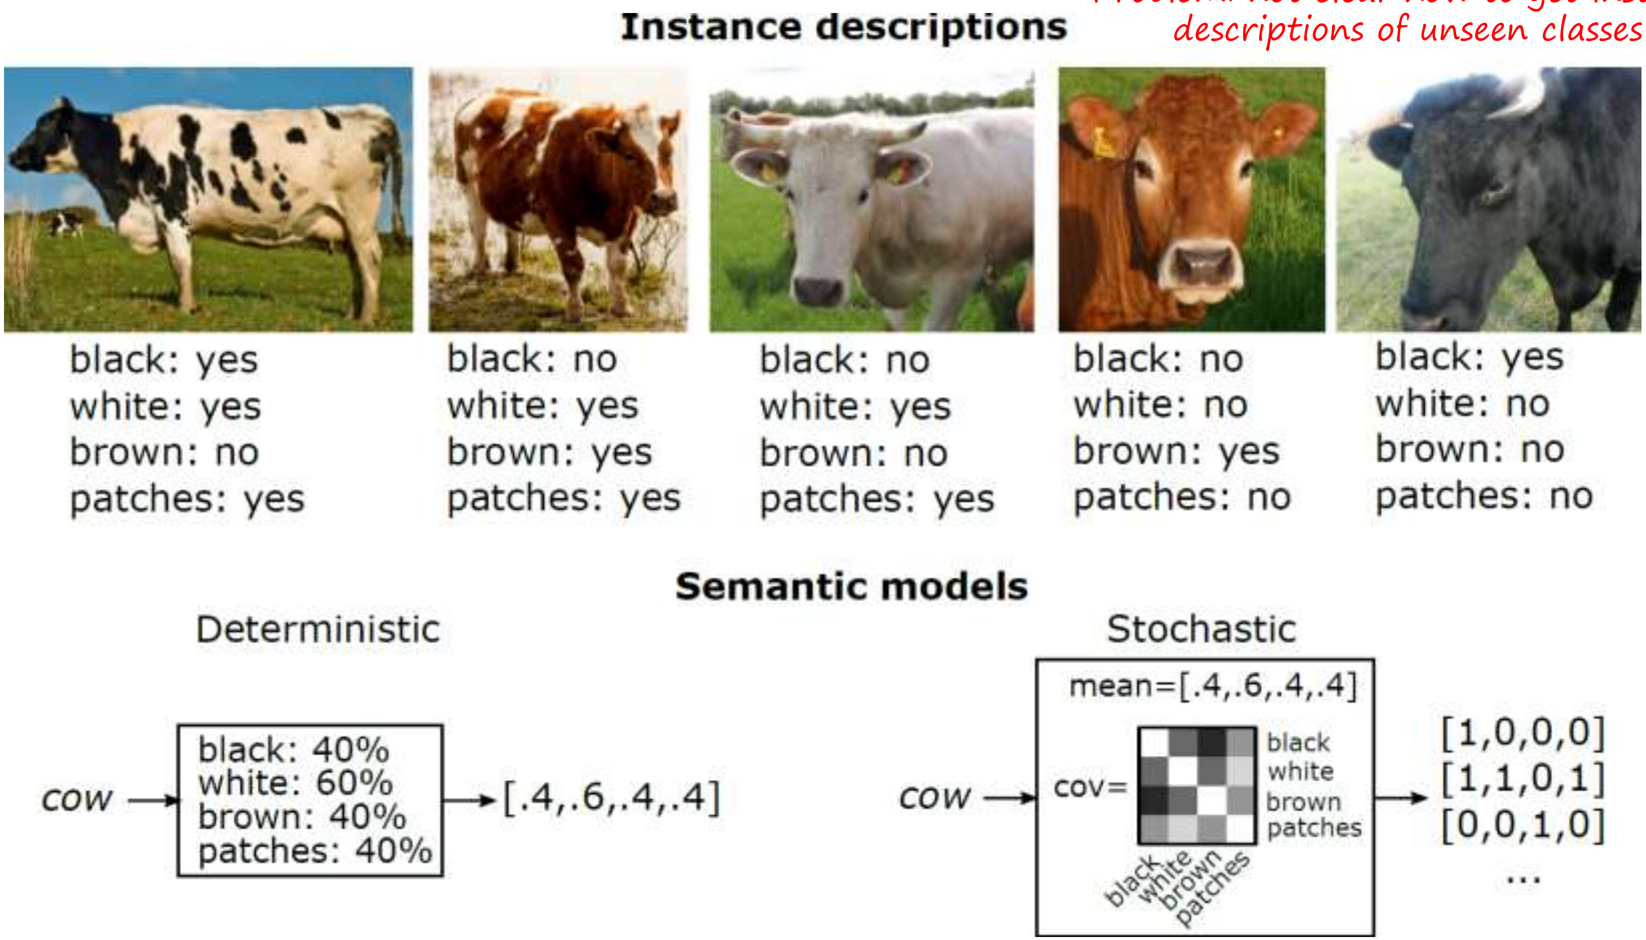

BImag with a stochastic semantic model

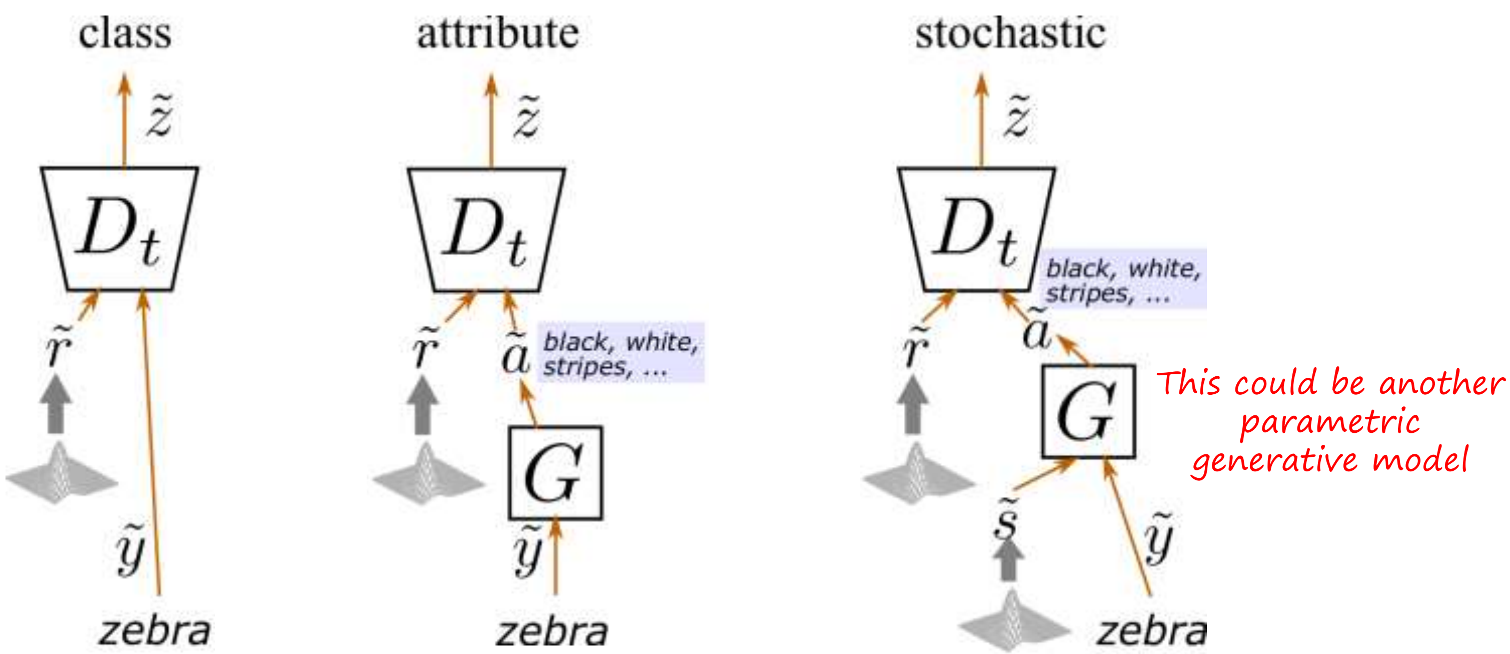

Experiments

- Datasets
  - Animal with attributes 2. Two (40/10) and three (30/10/10) task splits.
  - CUB. Two (150/50) and three (100/50/50) task splits.
  - SUN. Two task splits (645/72).
- Evaluation metrics
  - AUSUC (used in GZSL, less sensitive to calibration than harmonic mean)

    $$\hat{y} = \arg\max f_c(x) - \gamma \mathbb{I}[c \in \mathcal{U}]$$

    - Adapted to two tasks as AUTAC (area under the task-accuriacies curve).
    - Adapted to three tasks as VUTAS (volume under the task-accuriacies surface).
  - Overall accuracy
- Settings
  - GZSL (BCL for t=1)
  - BCL

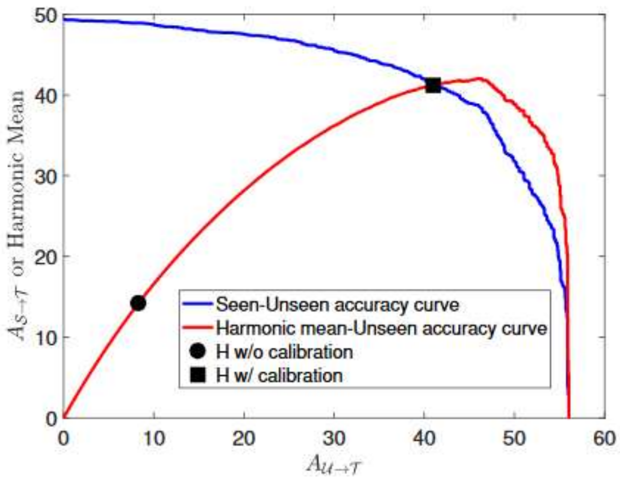

Experiments on GZSL

- Here attr-BImag is a typical GZSL baseline
- Surprisingly, joint class-attribute conditioning (class-attr-BImag) improves performance in GZSL

Method	Gen.	FE	CUB				AWA			
			A (%)	B (%)	H (%)	AUSUC	A (%)	B (%)	H (%)	AUSUC
attr-BImag	VAE	fix	60.84	39.70	48.05	0.347	72.28	**62.02**	**66.76**	0.540
	VAE	ft	**77.74**	41.30	53.94	0.484	73.83	59.97	66.18	0.555
cls-attr-BImag	VAE	fix	59.28	40.97	48.45	0.349	74.20	54.18	62.63	0.453
	VAE	ft	73.57	45.09	55.91	**0.515**	**76.93**	51.40	61.63	**0.578**
Mishra *et al.*	VAE	fix	-	-	34.5	-	-	-	51.2	-
f-CLSWGAN	GAN	fix	57.7	43.7	49.7	-	61.4	57.9	59.6	-

Experiments on two tasks

	CUB 150/50				AWA 40/10			
	CL	GZSL/BCL			CL	GZSL/BCL		
	class	attr	cls-att	asym	class	attr	cls-att	asym
$t=1$ (GZSL)	0.018	0.484	$\boxed{\mathbf{0.515}}$	0.484	0.039	0.555	$\boxed{\mathbf{0.578}}$	0.555
$t=2$	0.691	0.670	$\boxed{0.685}$	0.691	0.917	0.914	$\boxed{\mathbf{0.923}}$	0.917
Mean	0.355	0.577	$\boxed{\mathbf{0.600}}$	0.588	0.478	0.735	$\boxed{\mathbf{0.750}}$	0.736

Twp tasks experiments (AUSUC) on CUB 150/50 and AwA 40/10.

Experiments on three tasks

	CUB 100/50/50						AWA 30/10/10			
	CL	GZSL/BCL					CL	GZSL/BCL		
		Class-level			Instance-level			Class-level		
	class	attr	class-att	asym	epis	class-epis	class	attr	class-att	asym
$t = 1$ (GZSL)	0.0002	0.1207	0.1231	0.1207	**0.1429**	0.0680	0.0049	**0.2364**	0.1580	**0.2364**
$t = 2$	0.0078	0.2205	0.2246	**0.2454**	0.2239	0.2225	0.0560	0.3231	0.3105	**0.3375**
$t = 3$	**0.3949**	0.3106	0.3762	**0.3949**	0.2459	0.3840	**0.7306**	0.7296	**0.7306**	**0.7306**
Mean	0.1343	0.2173	0.2413	**0.2537**	0.1613	0.2248	0.2638	0.4297	0.3997	**0.4348**

Three tasks experiments (VUTAS) on CUB 100/50/50 and AwA 30/10/10.

Conclusions (Role of semantic relations in CL)

- **Semantic models** should be considered as integral part (together with visual models) for human-like continual learning
- Two new settings (visual+semantic): **generalized continal learning** (GCL) and **bookworm continual learning** (BCL)
  - BCL: generalizes both CL and ZSL. Semantic model fixed
  - GCL: generalizes BCL. Continual semantic model
- Many open questions (settings?, datasets? evaluation? training?)

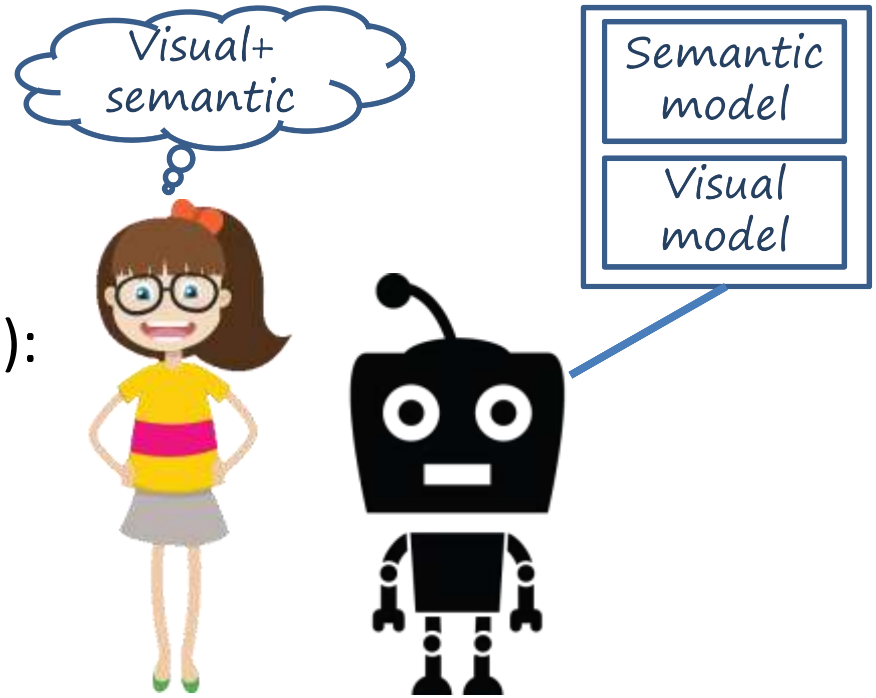

Setting	Sub-models		Continual		Predictions	
	Vis.	Sem.	Vis.	Sem.	Seen	Unseen
JT	✓	✗	✗	-	✓	✗
CL	✓	✗	✓	-	✓	✗
ZSL	✓	✓	✗	✗	✗	✓
GZSL	✓	✓	✗	✗	✓	✓
BCL	✓	✓	✓	✗	✓	✓
GCL	✓	✓	✓	✓	✓	✓

Conclusions (Bidirectional imagination)

- **BImag** uses feature generation to prevent catastrophic forgetting and infer future classes
- Integrating semantic information in replay generators is not trivial and can **interfere** with they performance (attr-BImag).
- Possible solutions
  - Class-attribute conditioning
  - Asymmetric generation
  - Stochastic semantic models?
- A lot of space for improvement

Works in CL and ZSL at LAMP-CVC

Continual learning

- **R-EWC**. X. Liu et al., Rotate your Networks: Better Weight Consolidation and Less Catastrophic Forgetting, ICPR 2018
- **MerGANs**. C. Wu et al., Memory Replay GANs: learning to generate images from new categories without forgetting, NeurIPS 2018.
- L. Yu et al., Semantic Drift Compensation for Class-Incremental Learning, CVPR 2020
- X. Liu et al., Generative Feature Replay For Class-Incremental Learning, CLVISION@CVPR 2020.
- M. Masana et al., Ternary Feature Masks: continual learning without any forgetting, arxiv 2020
- X. Liu et al., Continual Universal Object Detection, arxiv 2020

Zero-shot learning

- Song et al., Generalized Zero-shot Learning with Multi-source Semantic Embeddings for Scene Recognition, ACM Multimedia 2020
- Yang et al., Simple and effective localized attribute representations for zero-shot learning, arxiv 2020

Generalized/bookworm continual learning

- **Wang et al., Bookworm continual learning: beyond zero-shot learning and continual learning, arxiv 2020 (short version at TASK-CV@ECCV 2020)**

Kai Wang

Anjan Dutta

Joost van de Weijer

Bookworm continual learning
https://arxiv.org/abs/2006.15176

www.lherranz.org
lherranz@cvc.uab.es

More details about our work in CL at
http://www.lherranz.org/category/continual-learning