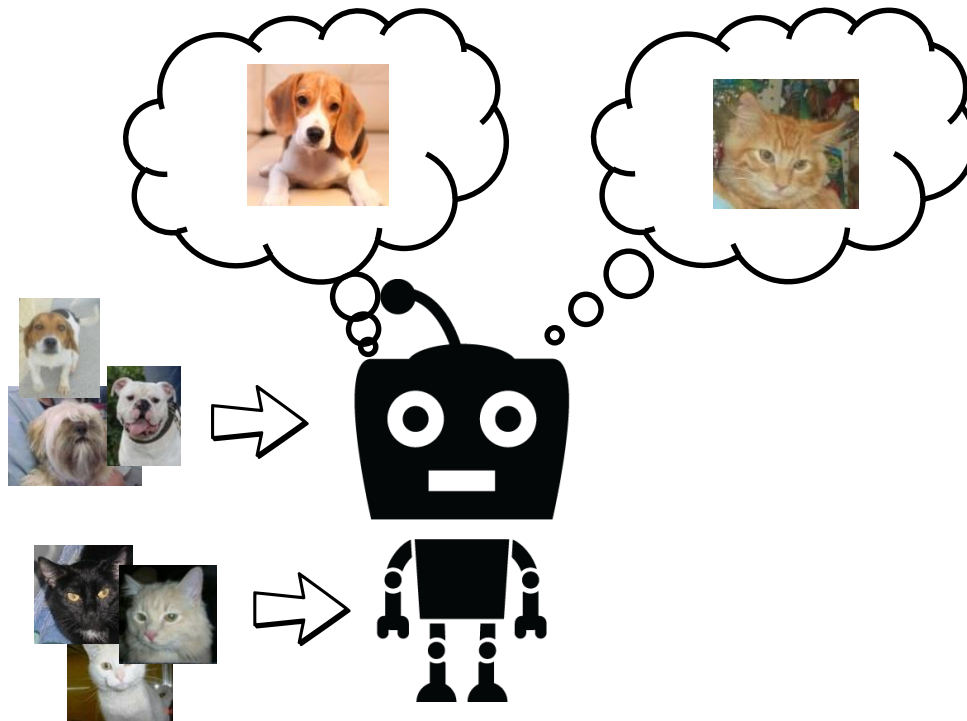


Learning and forgetting in image generation




BCN.ai

Luis Herranz


Computer Vision Center, UAB

December 2018

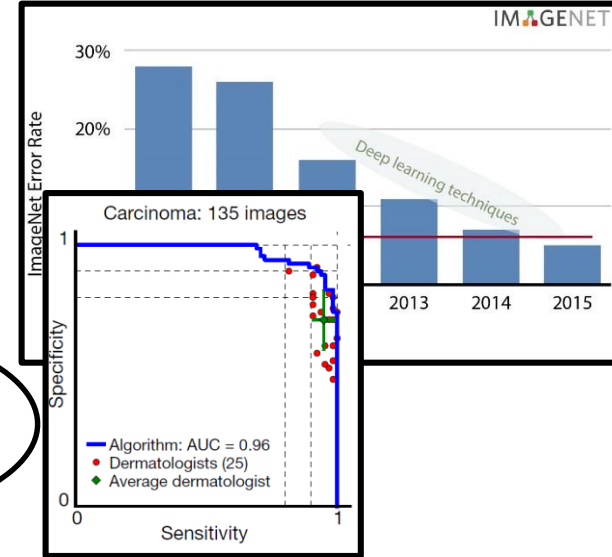



Successful AI in computer vision



Siamese cat



Malignant melanocytic lesion



Outperform humans!

...but

- Tons of data
- All data in advance
- Heavily supervised
- Discriminative tasks


Humans learn differently

Limited data and supervision:

Few-shot learning
Zero-shot learning
Unsupervised learning

Image generation
("imagination")

Continual learning



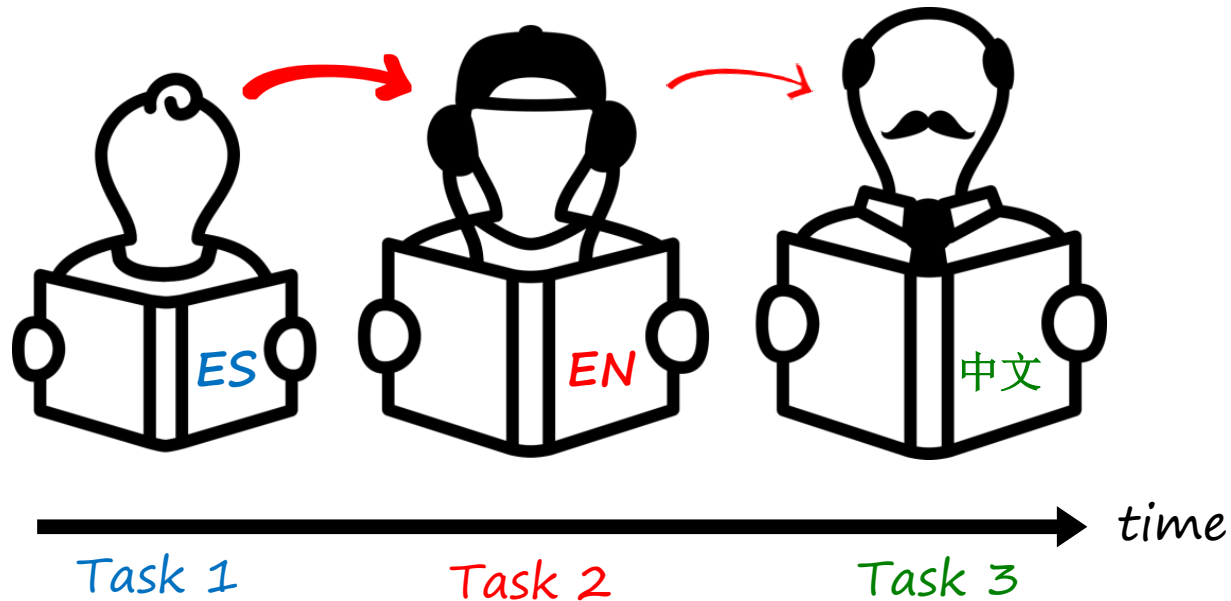
...without forgetting

Hey Jan, can you paint a girl with a pearl earring?

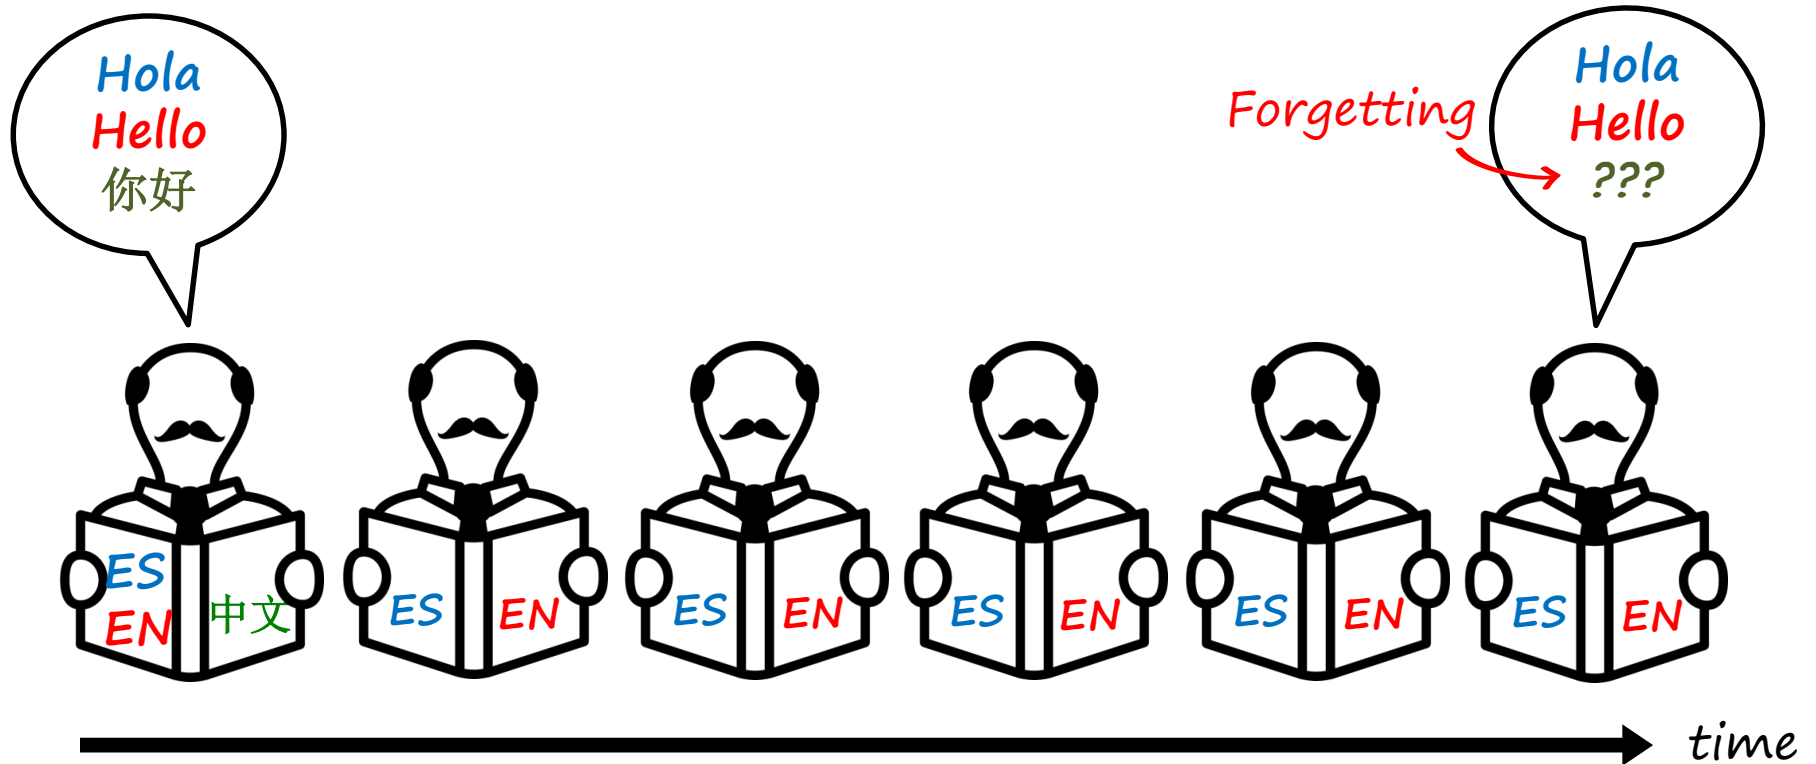


Continual learning in humans (a.k.a. lifelong/sequential learning)

- Reuse of past knowledge (i.e. knowledge transfer, transfer learning)
- Learn new skills for new tasks



...and forgetting



Why is continual learning important?

Efficient processing of data streams

Learning with privacy (e.g. discard data after processing)

cat



bird



dog




cat



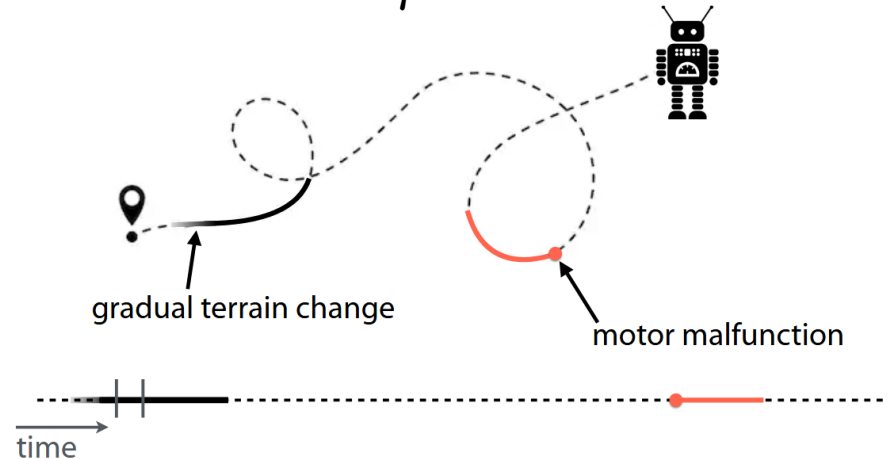
dog




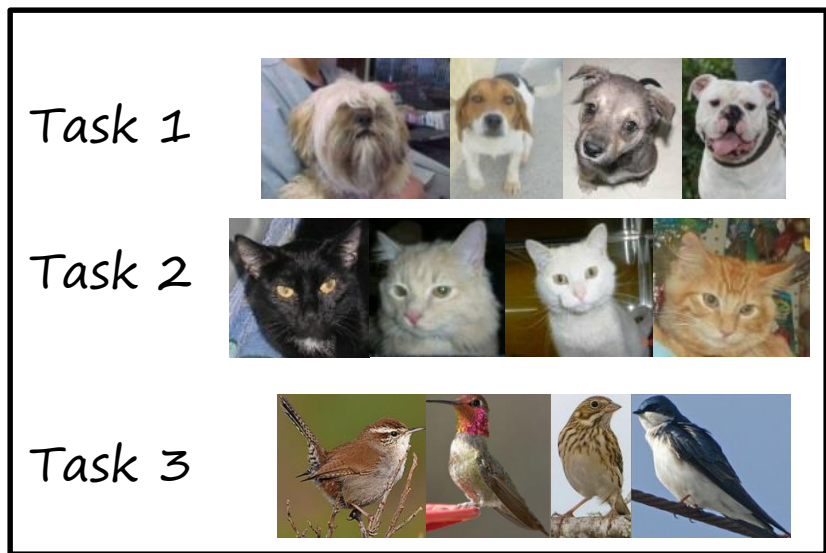
Human-robot interaction



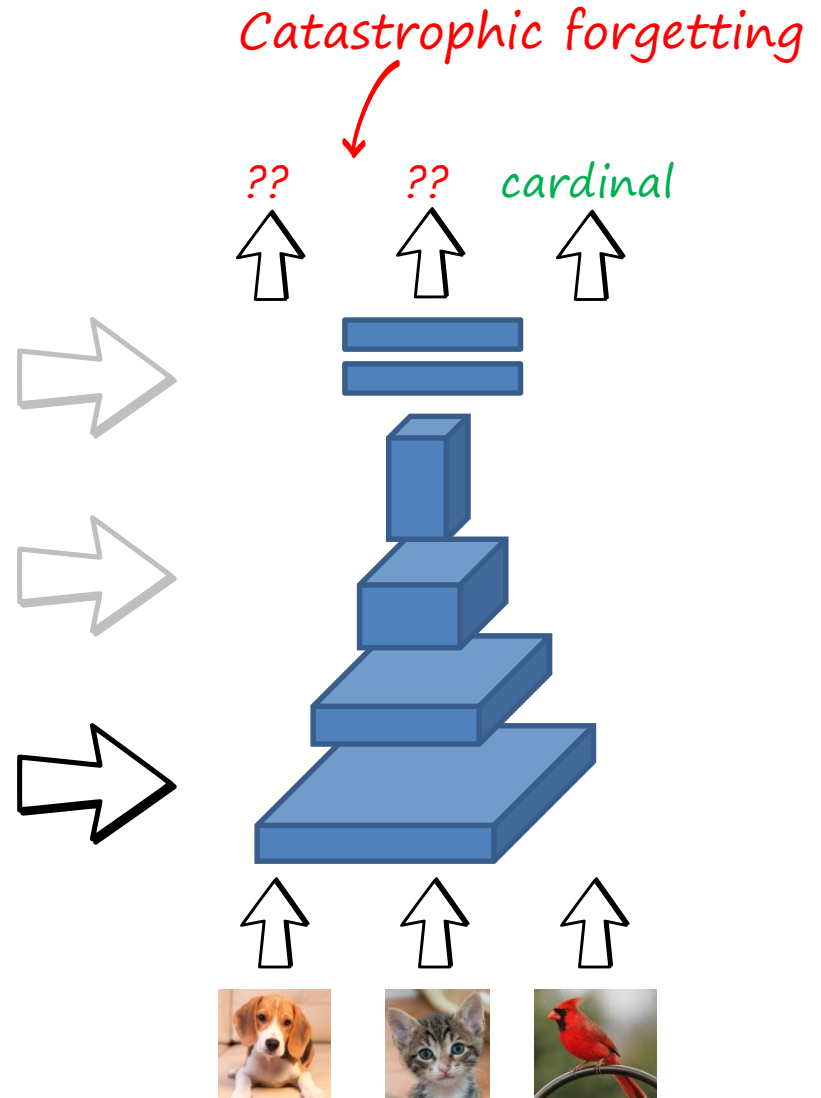
Deal with changing environments and problems



Learning in neural networks



Continual learning in neural networks



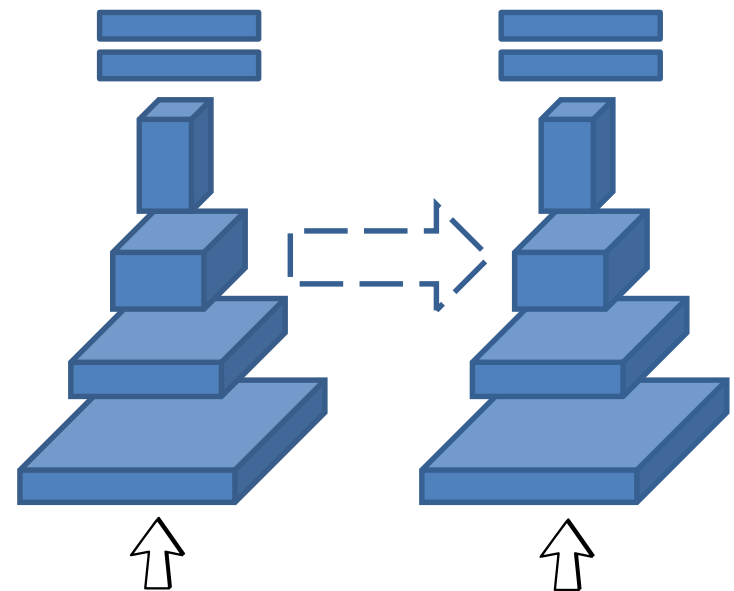

Transfer learning and continual learning

Forgets source task, i.e. catastrophic forgetting (who cares?)

Forgets task 1 (big deal!!)

Transfer+adaptation

Continual learning



Source task


Target task

Task 1

Task 2

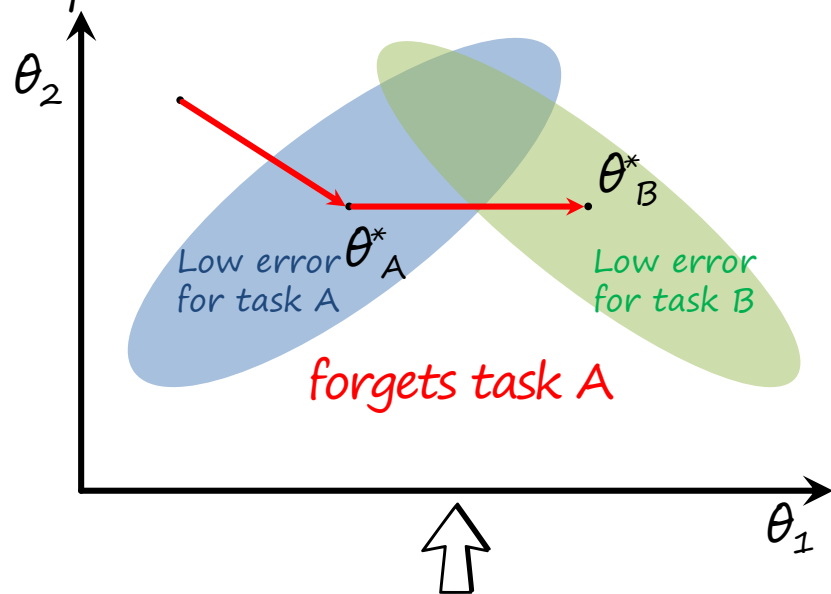
Continual learning =
transfer learning - (catastrophic) forgetting

Catastrophic interference and forgetting




Training/fine tuning

$$\mathcal{L}(\theta) = \mathcal{L}_B(\theta)$$

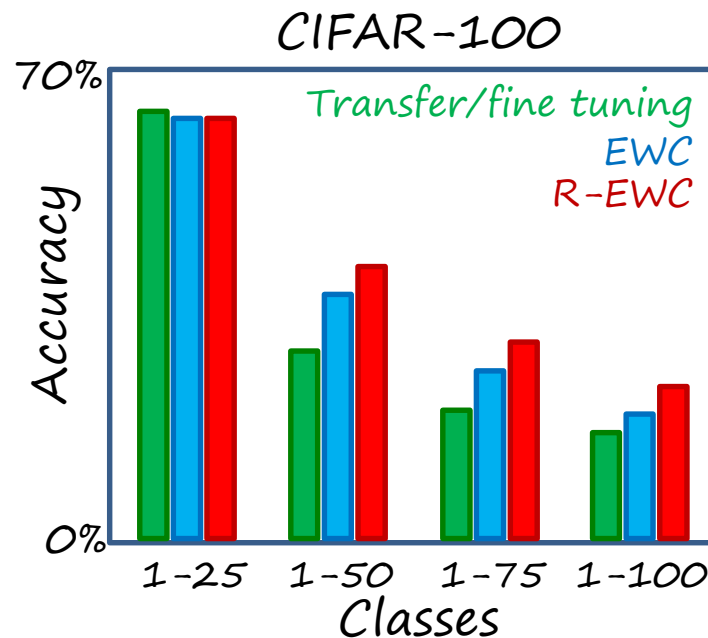
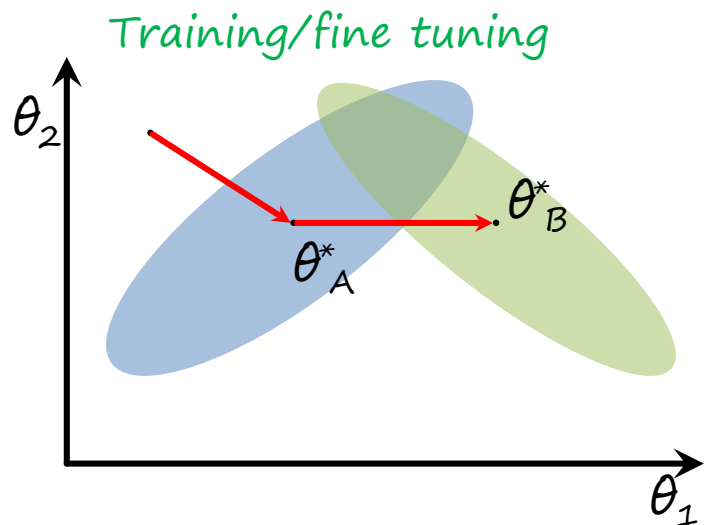


Elastic weight consolidation (EWC)


$$\mathcal{L}(\theta) = \mathcal{L}_B(\theta) + \frac{\lambda}{2} (\theta - \theta_A^*)^\top F_A (\theta - \theta_A^*)$$




Catastrophic interference and forgetting



Elastic weight consolidation (EWC)




Rotated Elastic weight consolidation (R-EWC)




Generative models: networks that imagine

Training data
(e.g. $64 \times 64 \times 3 \approx 12\text{K}$ dims)



Learning
→

Sampling



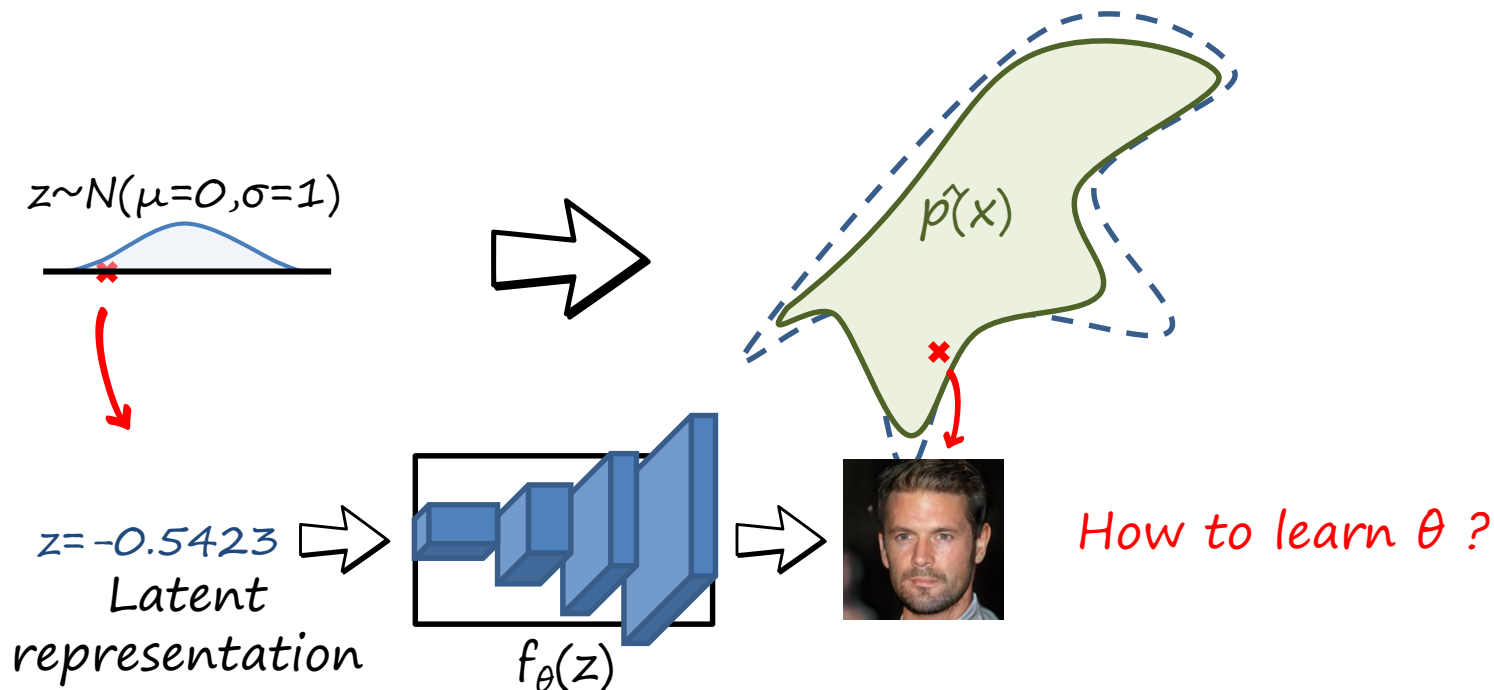
- Different approaches
- Density estimation
 - Variational autoencoders
 - Autoregressive models

- **Generative adversarial networks (GANs)**

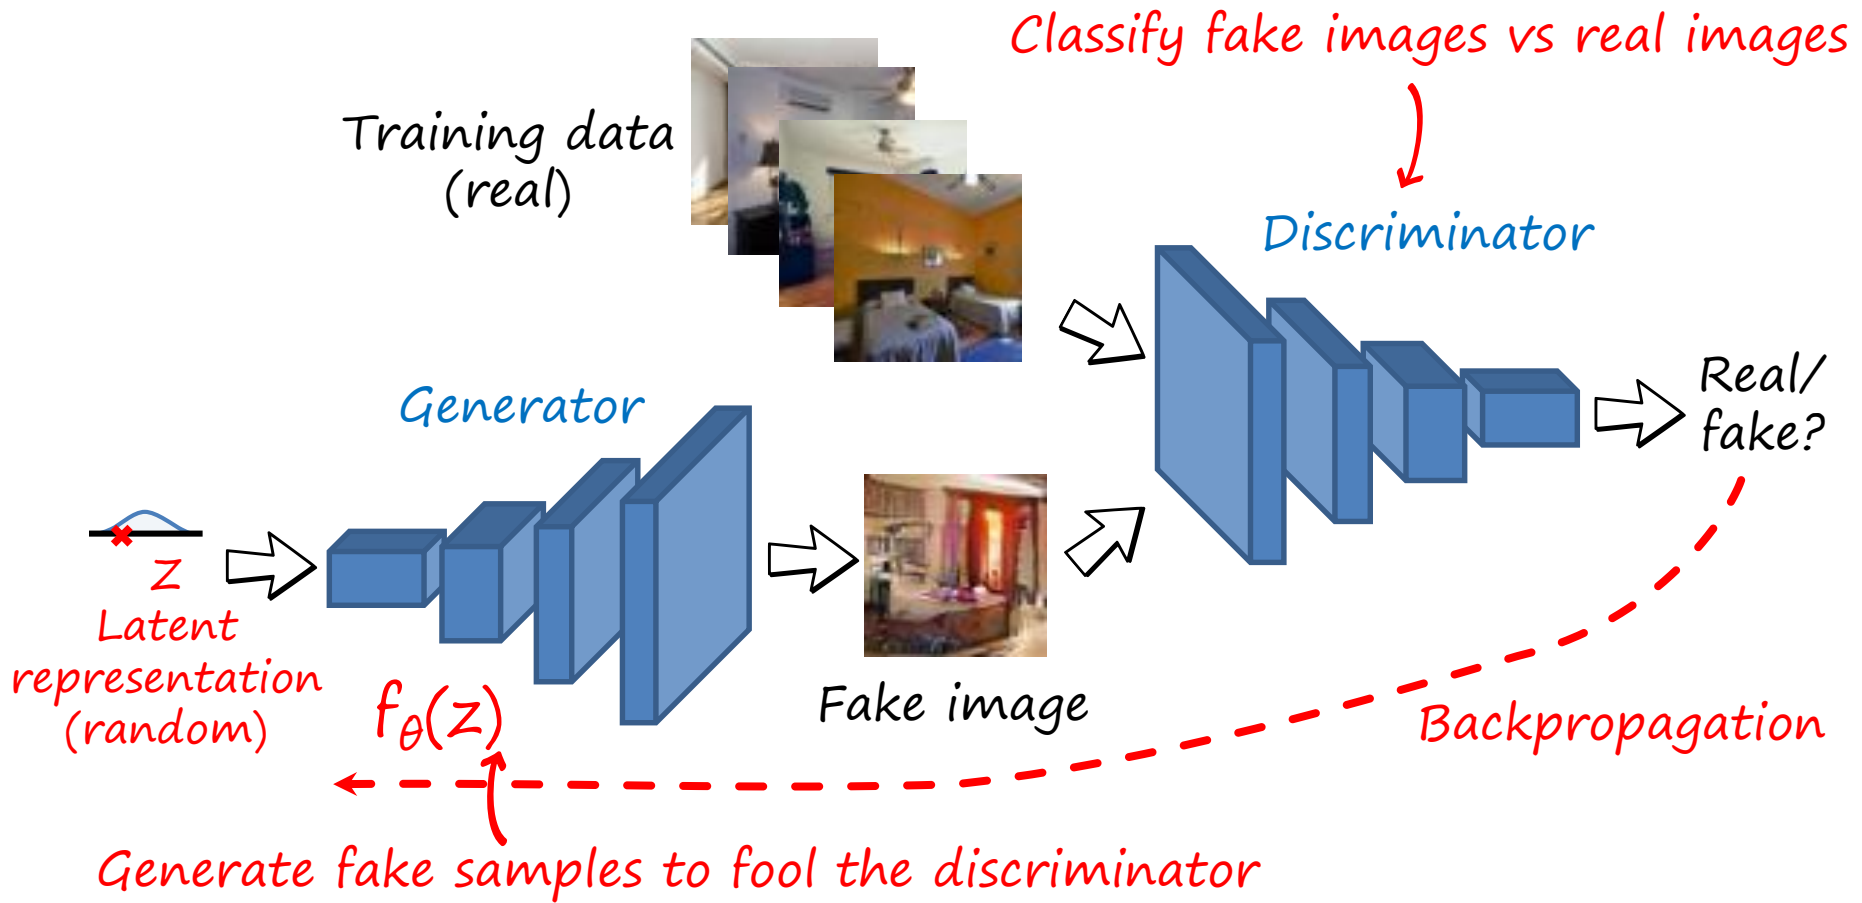
Indirect generative sampling

Learning and sampling from a complex distribution directly is very difficult

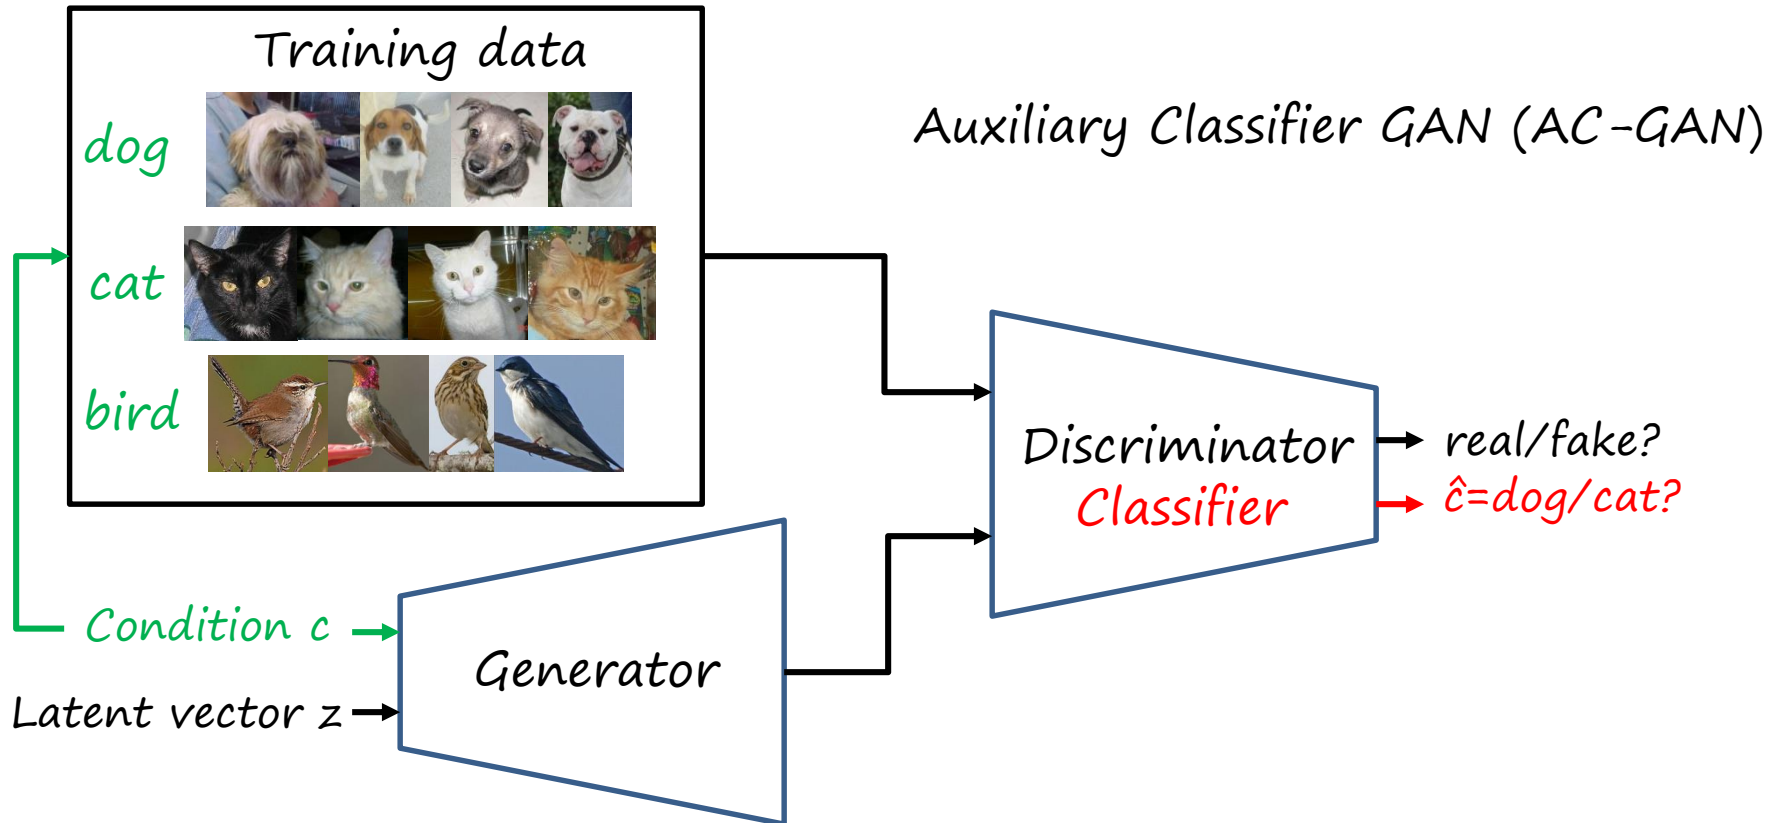
Idea: *sample from a simple distribution and learn a transformation*



Generative Adversarial Networks (GANs)



Conditional Generative Adversarial Networks



Generative Adversarial Networks

Wasserstein GAN (WGAN-GP)



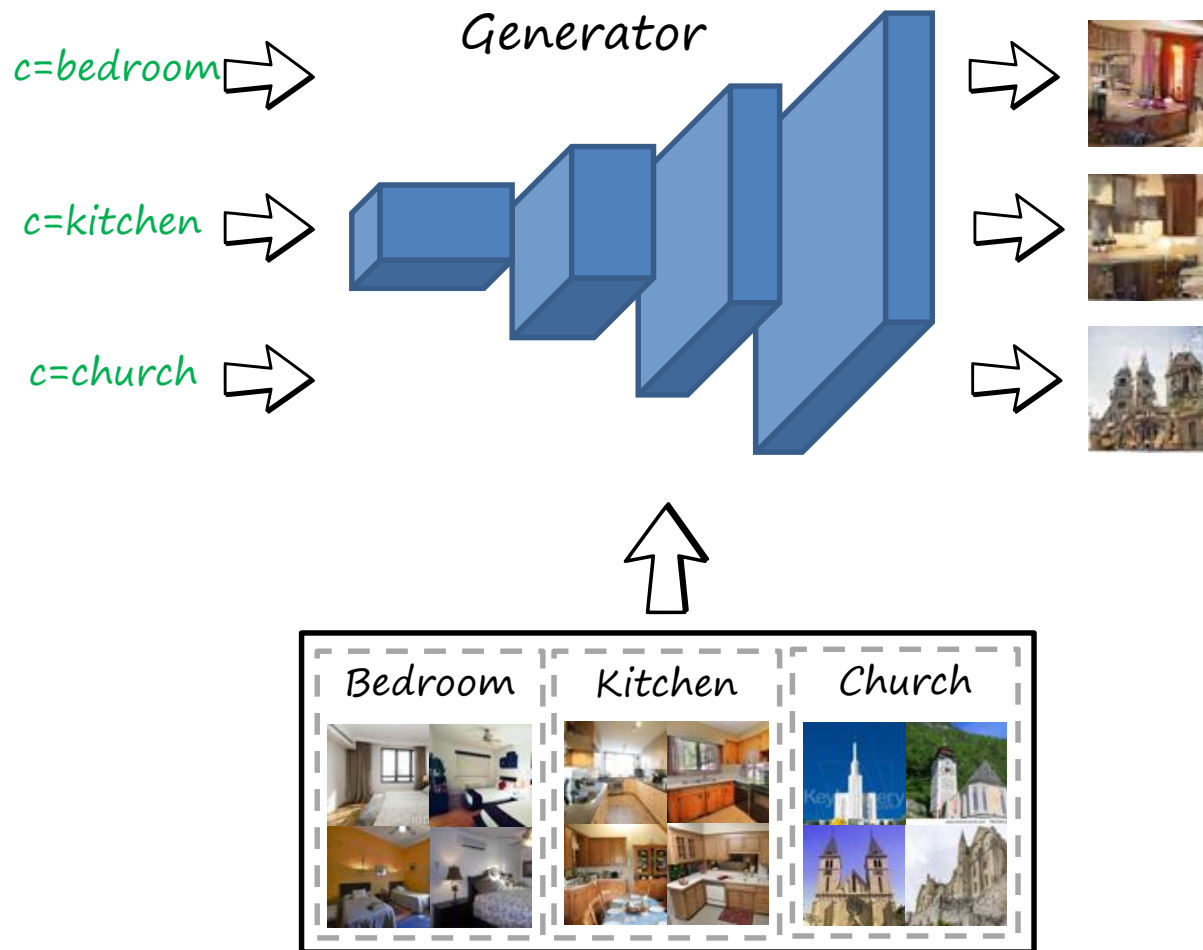
Progressive growing of GANs



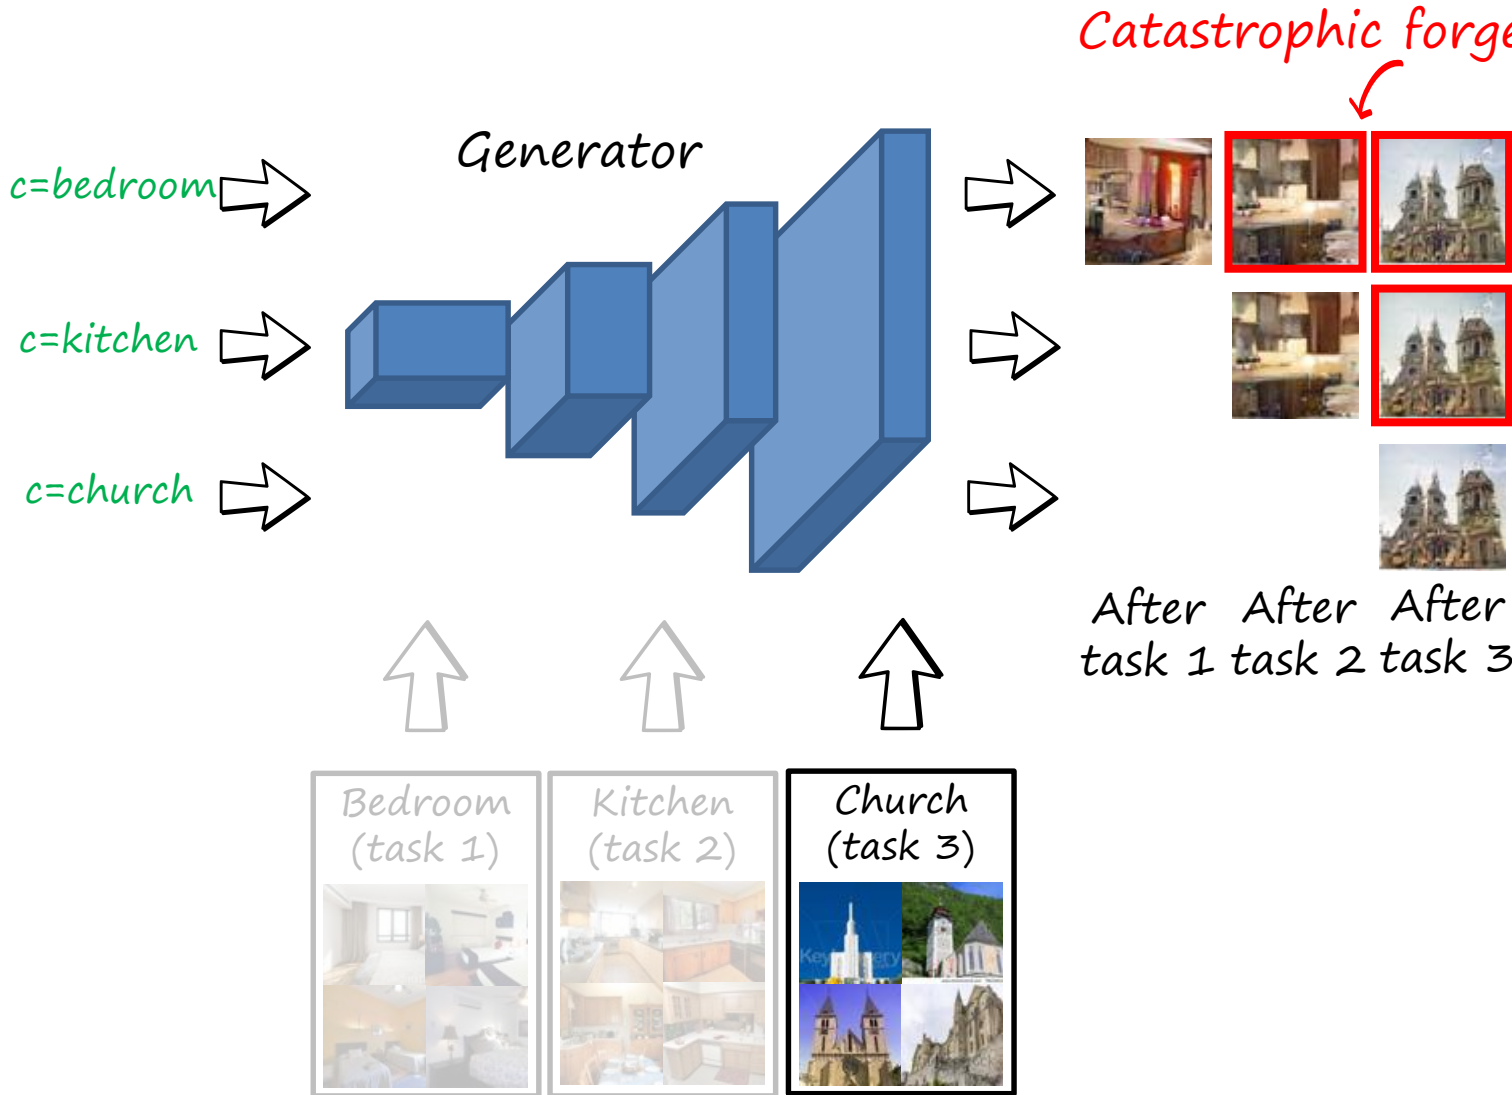
BigGAN (conditional GAN)



Conditional image generation



Continual learning in image generation



Sequential learning for image generation

Go to <http://www.lherranz.org/2018/10/29/mergans> to play the videos

MNIST 10 categories (10 tasks)

$c=0$ $c=1$ $c=2$ $c=3$ $c=4$ $c=5$ $c=6$ $c=7$ $c=8$ $c=9$



LSUN 4 categories (4 tasks)

$c=$ bedroom

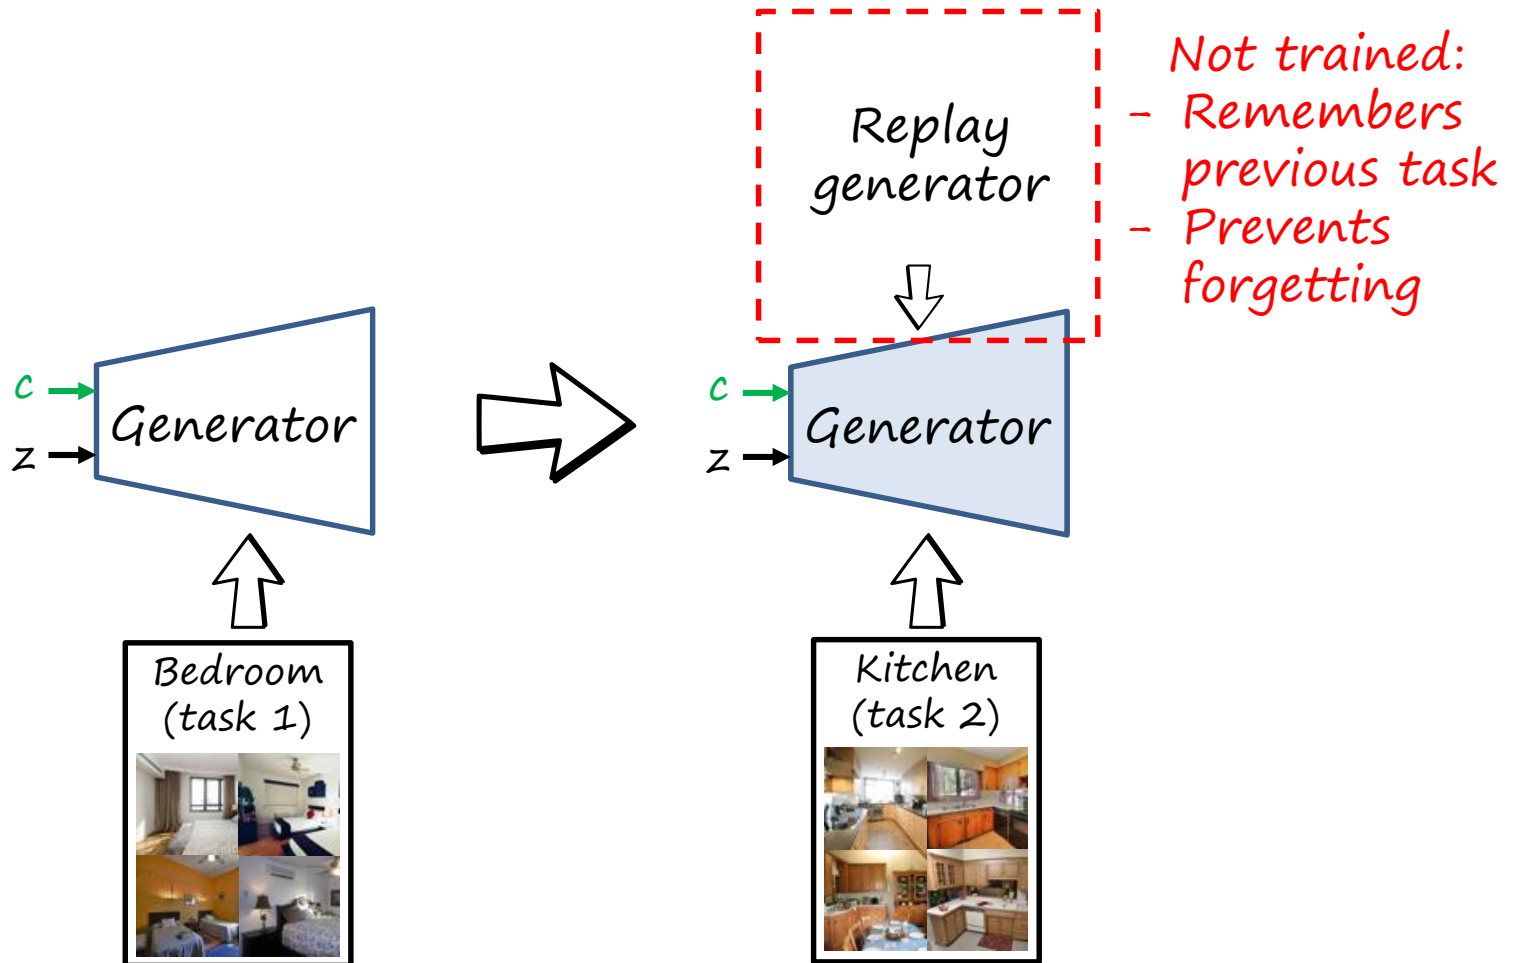
$c=$ kitchen

$c=$ church

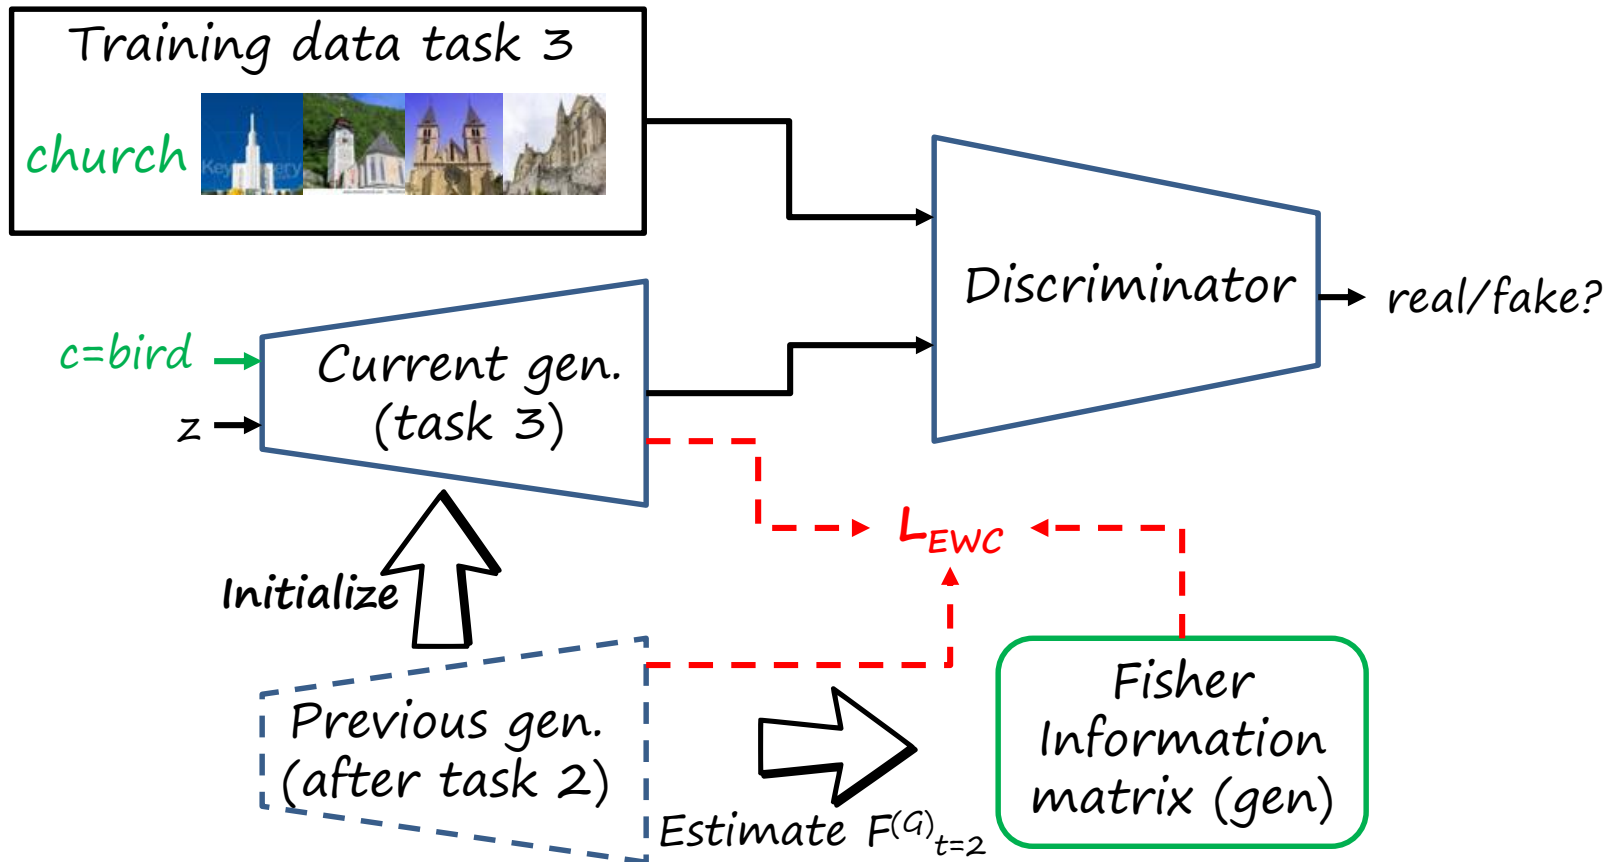
$c=$ tower



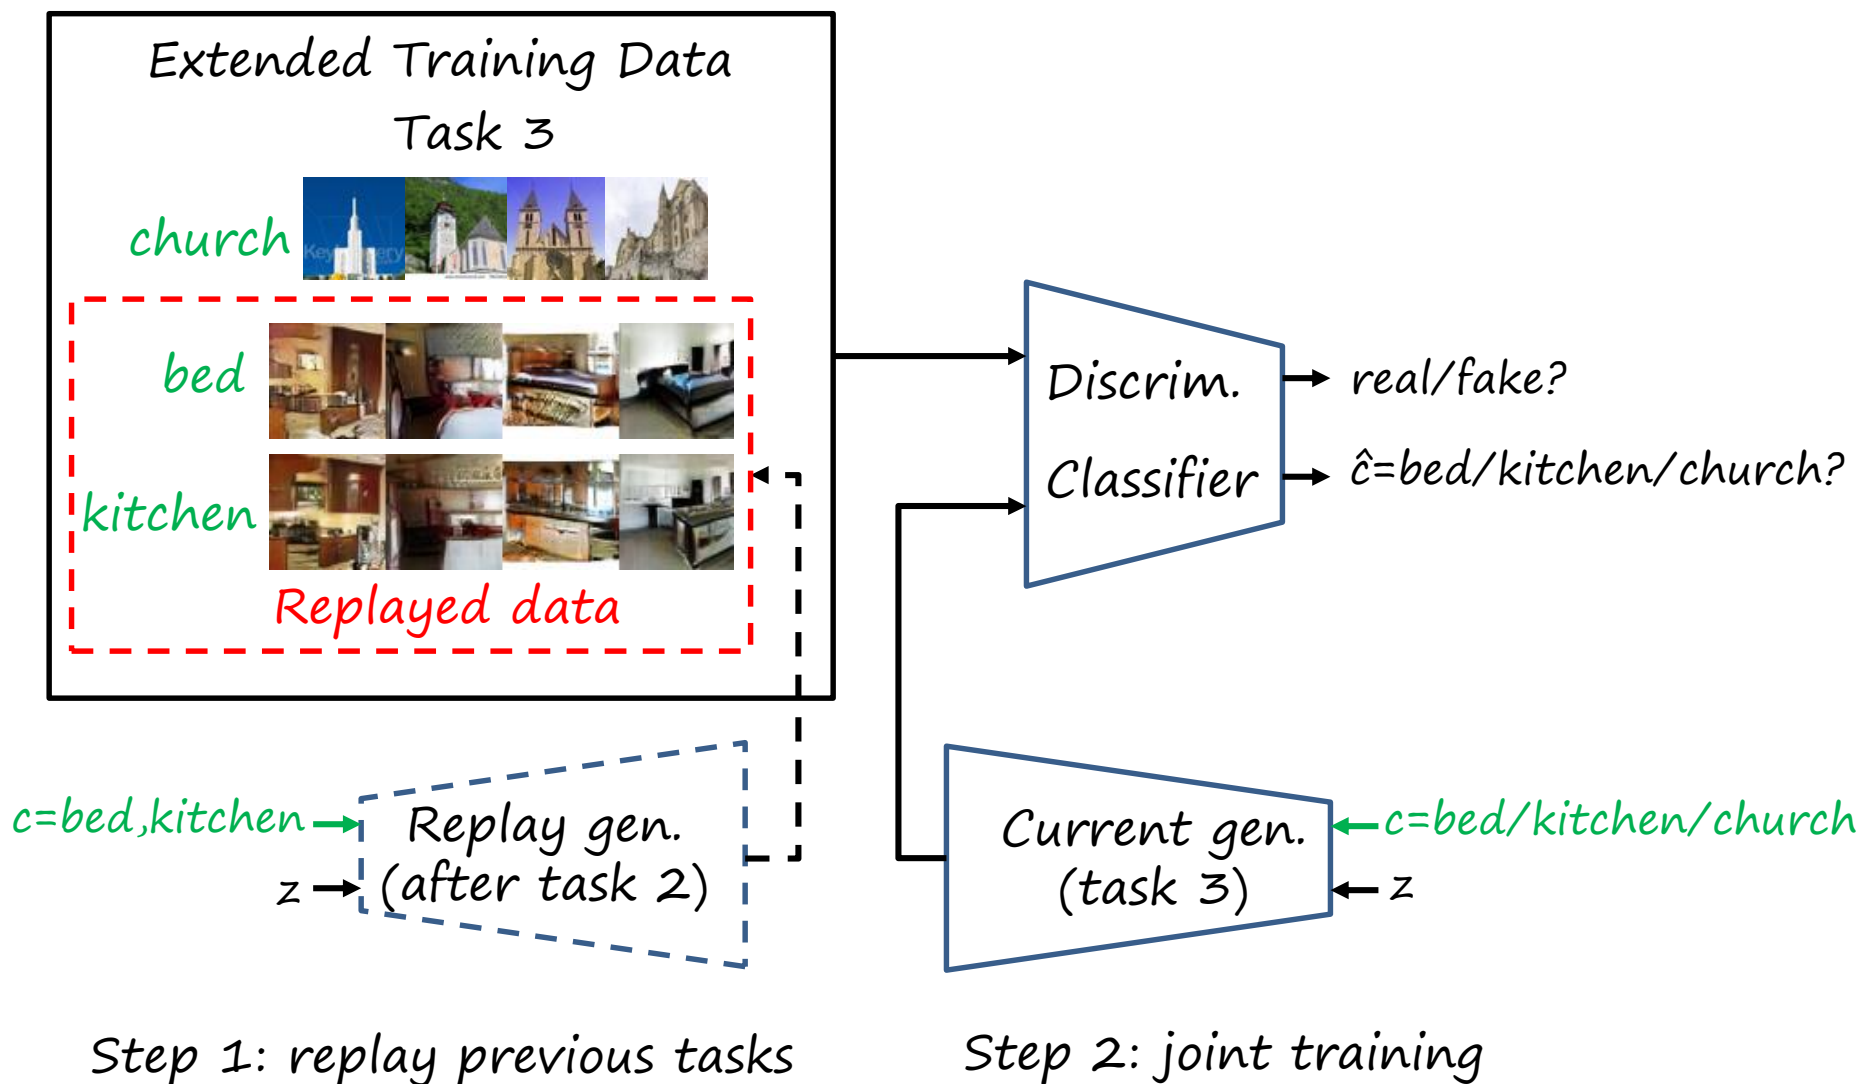
Memory Replay GANs (MeRGANs)



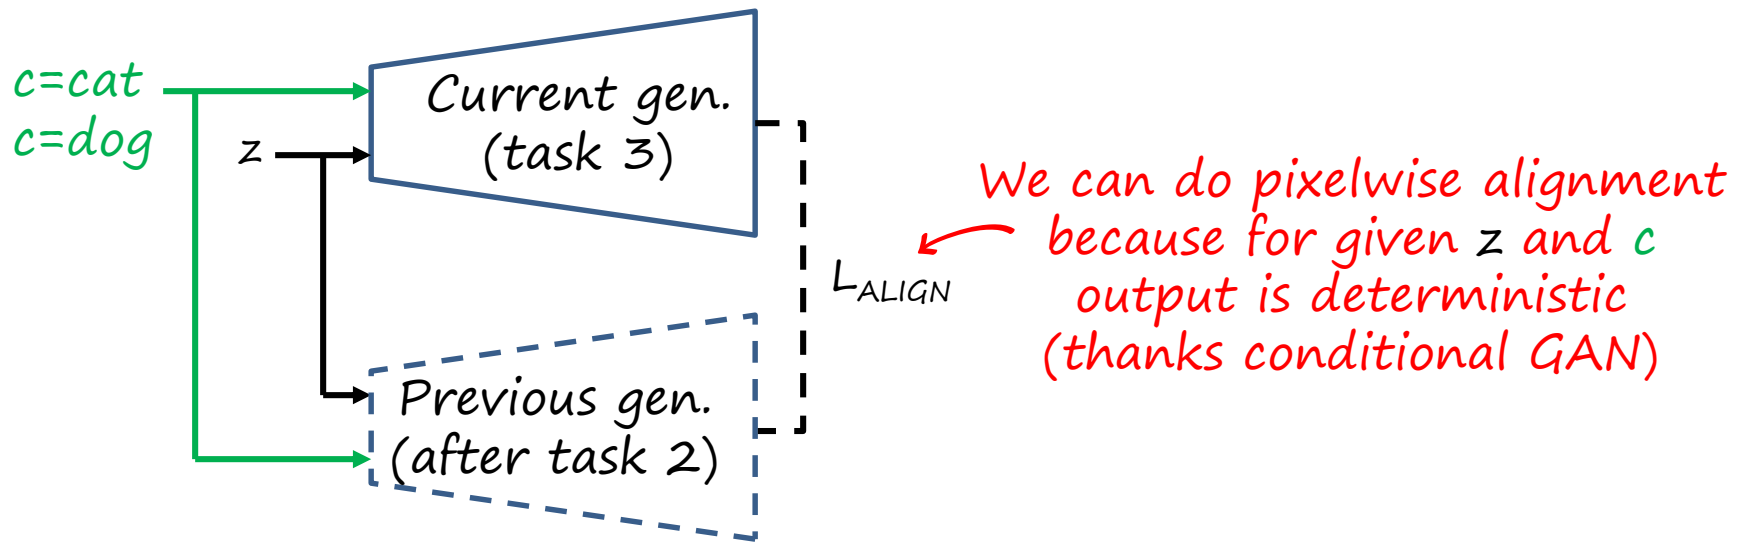
GAN with EWC



MeRGAN-JTR: joint training w/ replay

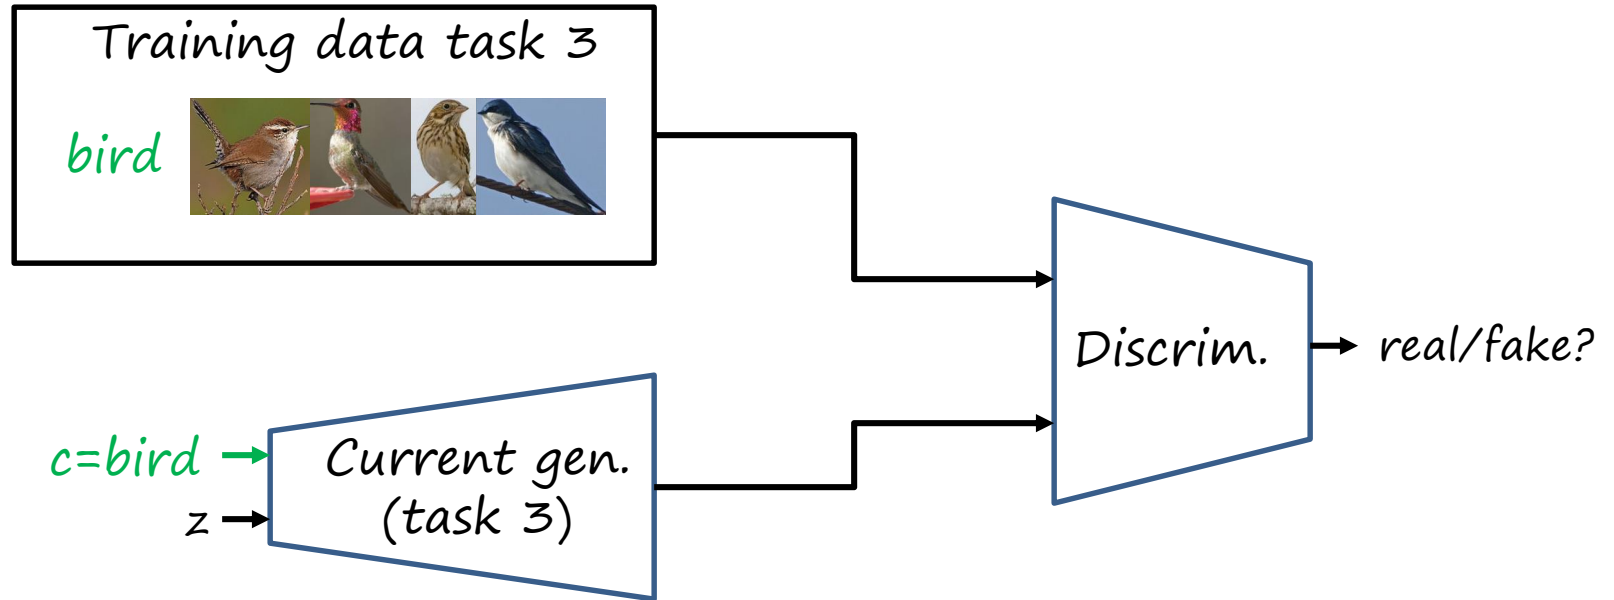


MeRGAN-RA: replay alignment



Step 1: replay previous tasks and align

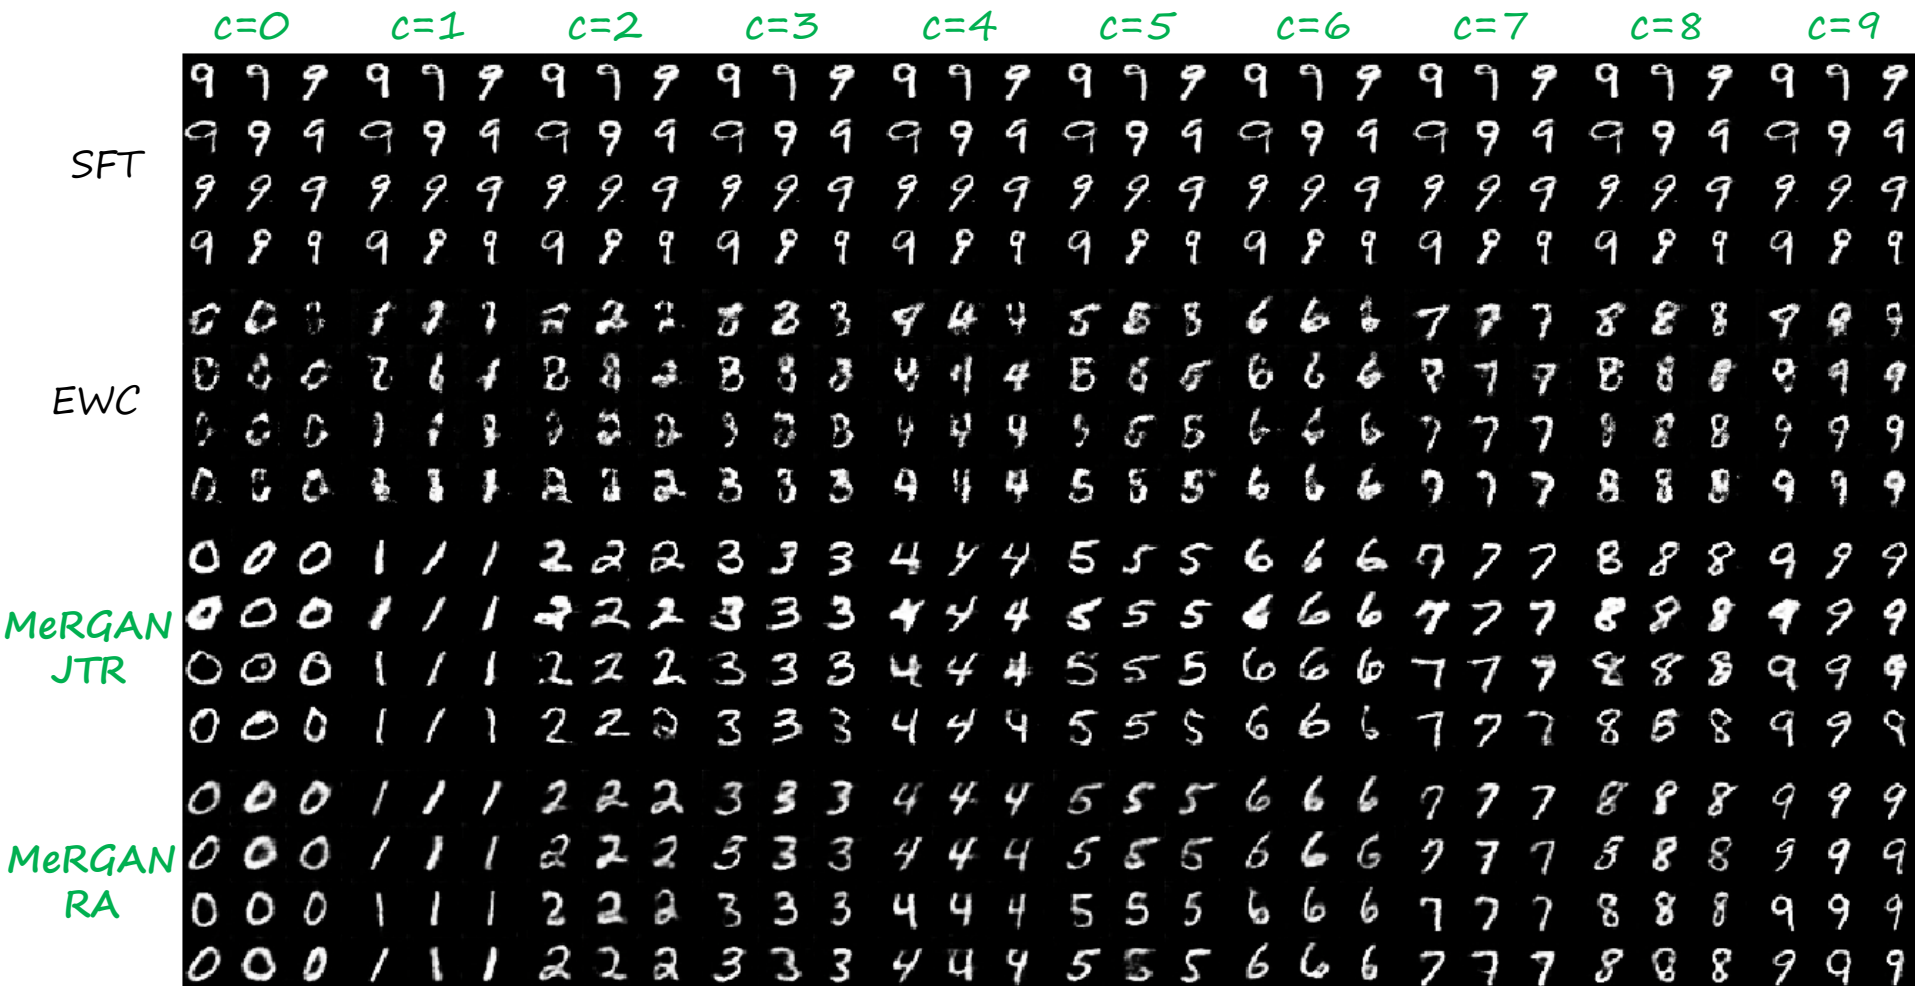
MeRGAN-RA: replay alignment



Step 2: learning new task

Digit generation (MNIST)

MNIST 10 categories (10 tasks)



Go to <http://www.lherranz.org/2018/10/29/mergans> to play the video

Scene generation (LSUN)


4 tasks (4 categories) 18-layer ResNet generator

Sequential fine tuning
Task 1 Task 2 Task 3 Task 4



Different bedrooms!


Task 1 Task 2 Task 3 Task 4



Remembers the category

Same bedroom!

Task 1 Task 2 Task 3 Task 4



Remembers the instance

Visualizing catastrophic interference


4 tasks (4 categories) 18-layer ResNet generator

c=bedroom

c=kitchen

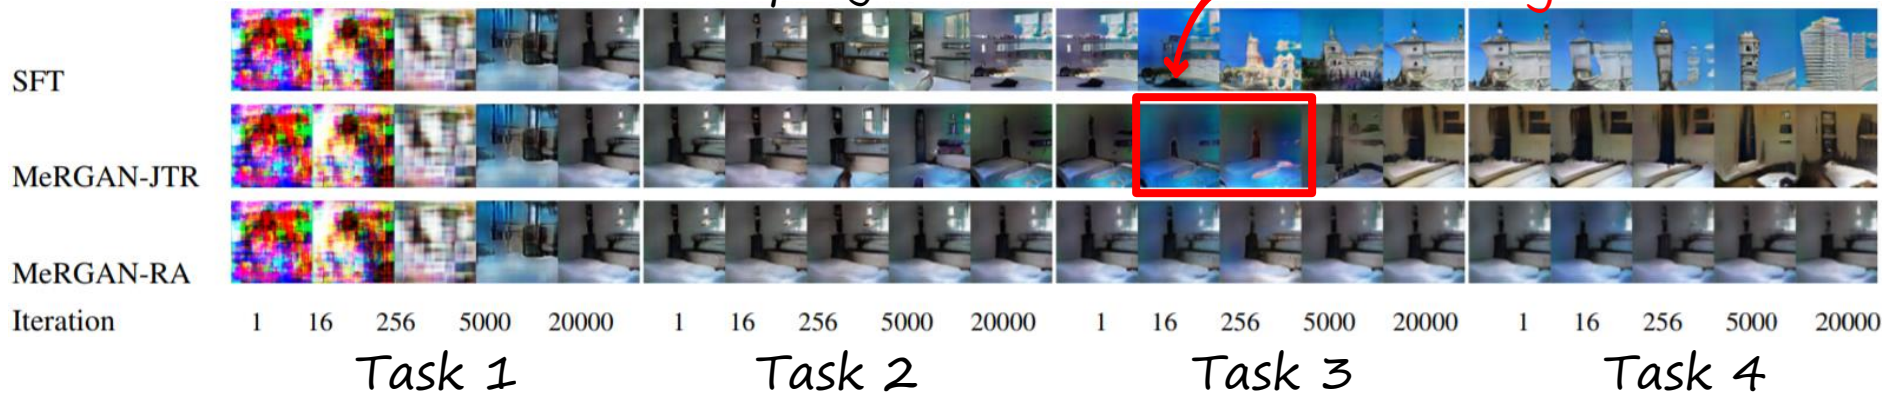
c=church

c=tower



Learning church interferes with remembering bedroom

Sampling bedrooms



Learning task 4 (tower) epoch 5000

bedroom

kitchen

church

tower

Go to <http://www.lherranz.org/2018/10/29/mergans>
to play the video



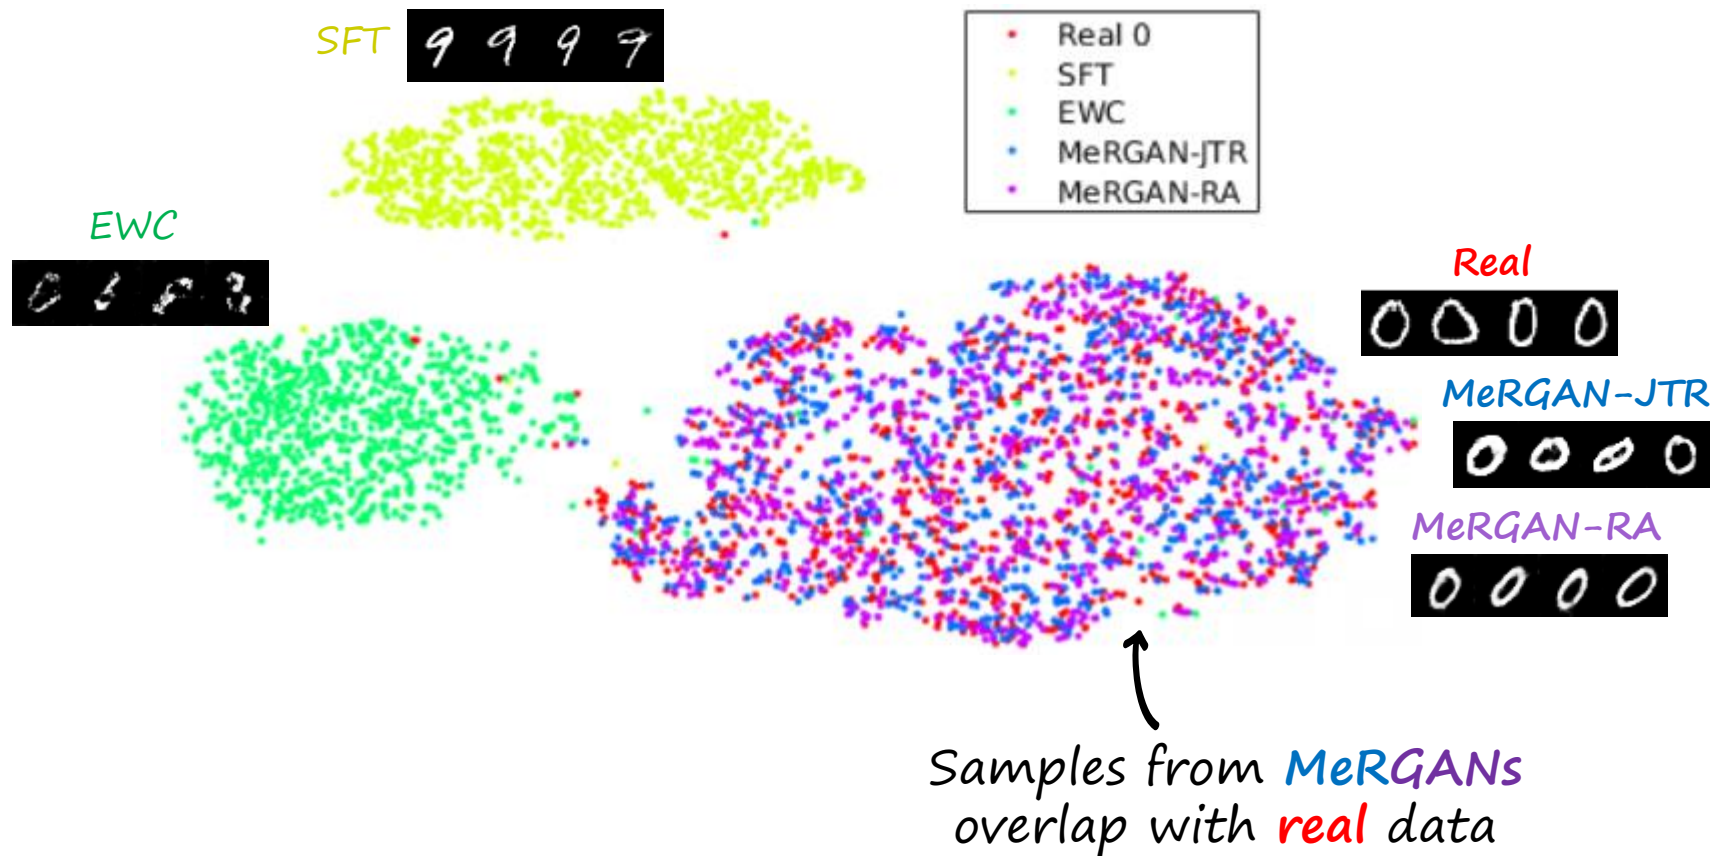
FT

JTR (ours)

RA (ours)

t-SNE visualizations (MNIST)


Generating digit 0 (i.e. first task) after learning 10 tasks



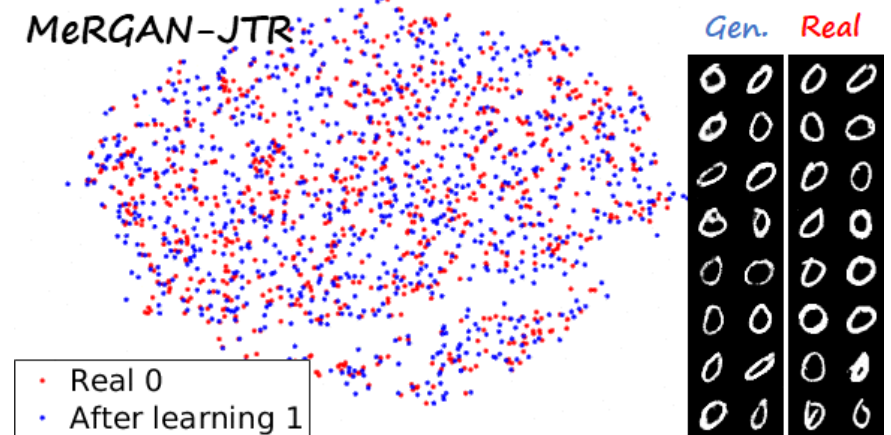
Learning and forgetting in t-SNE (MNIST)

Generating digit 0 (i.e. task 1)


SFT



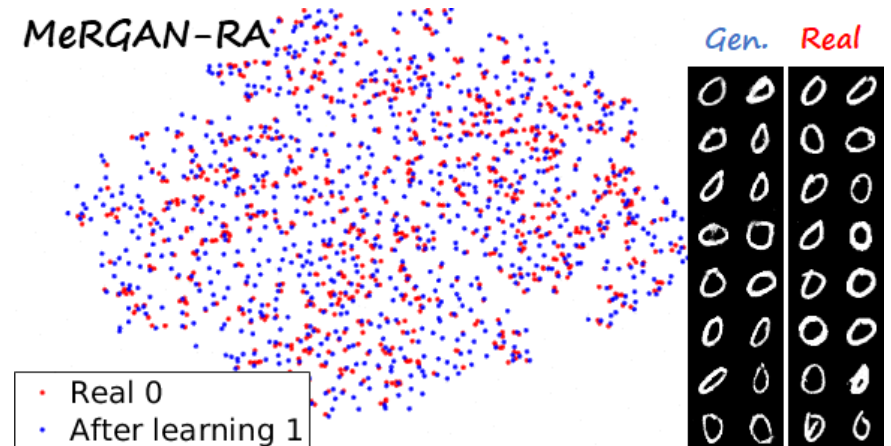
MeRGAN-JTR



EWC



MeRGAN-RA



Collaborators



Chenshen Wu



Xialei Liu



Yaxing Wang



Marc Masana



*Joost van de
Weijer*



*Bogdan
Raducanu*




*Antonio
López*



*Andrew
Bagdanov*

Thank
you!



www.cvc.uab.es/lamp
lherranz@cvc.uab.es

More details at <http://www.lherranz.org/category/continual-learning>
MeRGANs, NeurIPS 2018, <https://arxiv.org/abs/1809.02058>
R-EWC, ICPR 2018 <https://arxiv.org/abs/1802.02950>