Continual Learning

Nicolas Zucchet - Deep Reinforcement Learning Seminar - 19.05

Motivation

Why Continual Learning ?

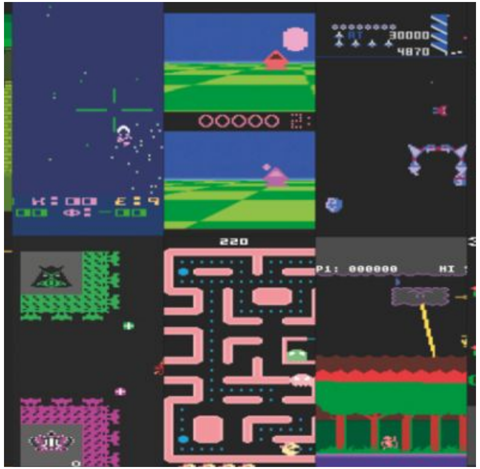

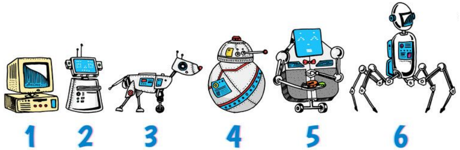

How Humans learn ?





tasks to learn

underlying tasks

How Humans learn ?

?

no catastrophic forgetting

?

reuse of important experiences

Definition

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks.

Source: https://nips.cc/Conferences/2018/Schedule?showEvent=10910

Definition

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks.

Source: https://nips.cc/Conferences/2018/Schedule?showEvent=10910

Definition

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks.

Source: https://nips.cc/Conferences/2018/Schedule?showEvent=10910

Definition

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks.

Source: https://nips.cc/Conferences/2018/Schedule?showEvent=10910

Definition

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks.

Similar to: sequential, incremental, lifelong learning

Related to: meta, multi-task, transfer, few-shot learning

Source: https://nips.cc/Conferences/2018/Schedule?showEvent=10910

Different approaches

- ➔ Increasing size networks
  - ◆ ***R**einforcement **C**ontinual **L**earning (RCL)*
- ➔ Fixed size networks
  - ◆ ***U**ncertainty Guided **C**L with **B**ayesian Neural Networks (UCB)*
- ➔ Mixing data from different tasks
  - ◆ *Experience Replay for CL*

Desired properties

Presence of transfer

Online learning

Bounded size

Resistance to catastrophic forgetting

Reinforced Continual Learning

J. Xu et al., 2018

Source: https://arxiv.org/abs/1805.12369

Motivation

Increasing network size

RL to find the increase

Network construction

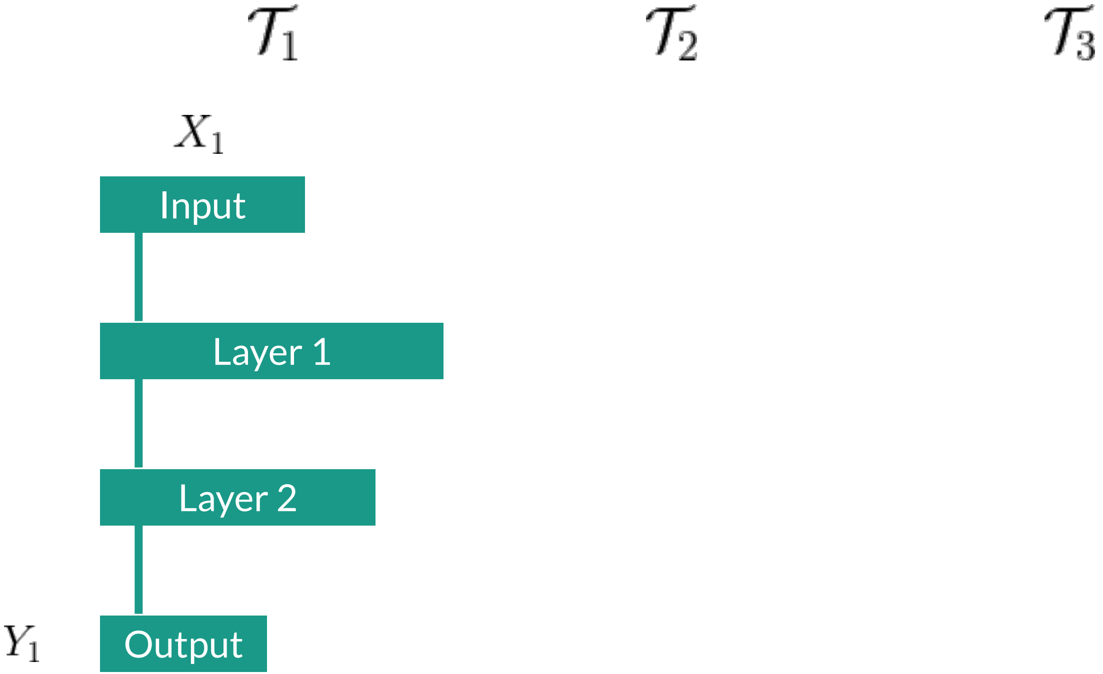

Network construction

$\mathcal{T}_1$ $\mathcal{T}_2$ $\mathcal{T}_3$

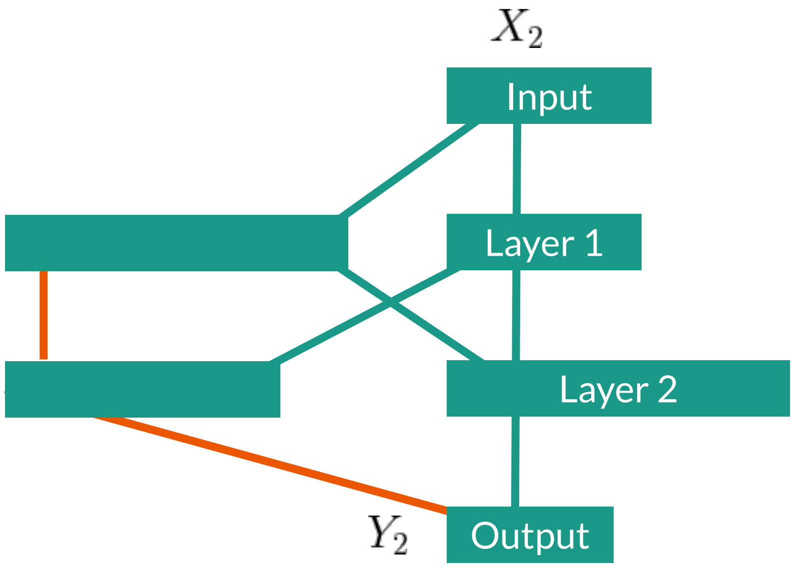

Network construction

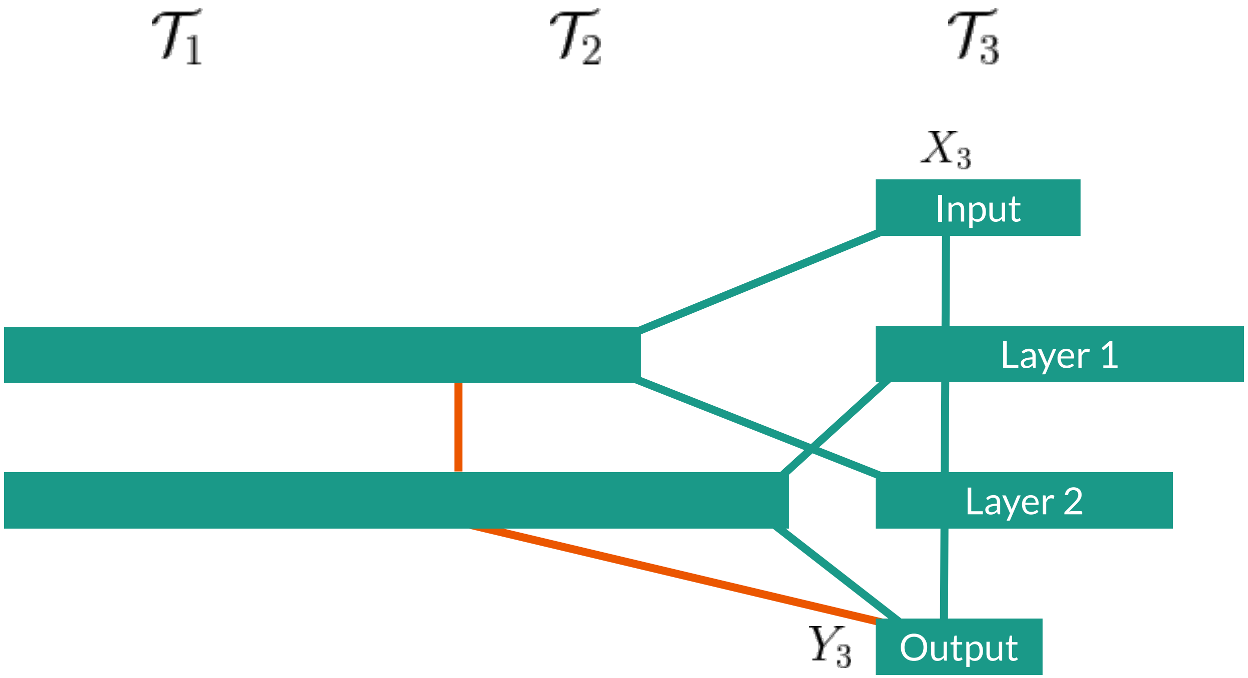

Final network

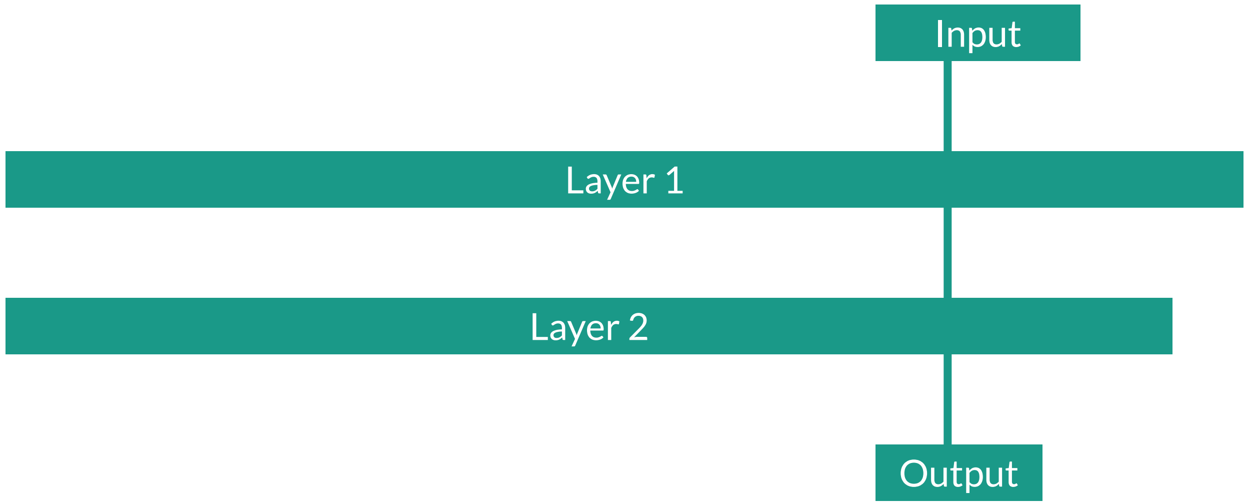

Expansion

Objective: best performing model without paying too much

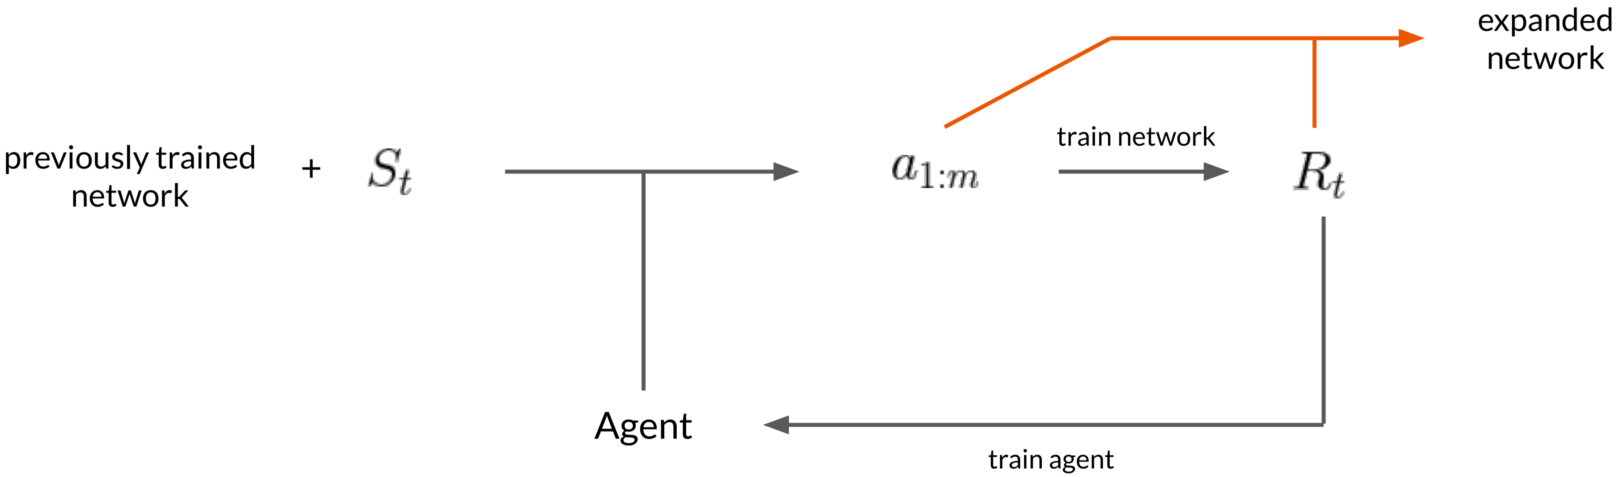

S_t embedding of the task

$a_{1:m}$ number of neurons to add for each layer

Agent: Actor-Critic architecture

$$R_t = \mathrm{Acc}_t(S_t, a_{1:m}) + \mathrm{Comp}_t(S_t, a_{1:m})$$

Results

- No catastrophic forgetting
- Economic

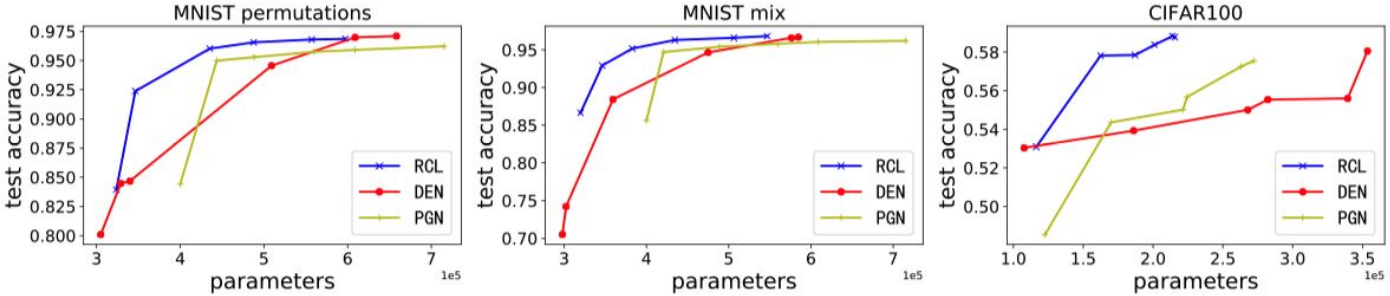

Figure 3: Average test accuracy *v.s.* model complexity for RCL, DEN and PGN.

Pros and cons

Pros

- No catastrophic forgetting
- Final network economical (number of neurons)

Cons

- Huge training time
- Drawbacks of RL without its advantages (state is not used)

Uncertainty guided CL with bayesian neural networks

S. Ebrahimi et al., 2019

Source: https://openreview.net/forum?id=HklUCCVKDB

Principle

Probabilistic weights

Posterior $\mathbb{P}(w|\mathcal{D}) = \frac{\mathbb{P}(w)\mathbb{P}(\mathcal{D}|w)}{\mathbb{P}(\mathcal{D})}$ **intractable!**

$\rightarrow$ $w_i \sim \mathcal{N}(\underbrace{\mu_i}_{\text{value}}, \underbrace{\sigma_i}_{\text{incertitude}}), \quad \theta = (\mu, \sigma)$ $q(w|\theta)$

$$\theta^* = \text{argmin}_\theta D_{KL}(q(w|\theta)||\mathbb{P}(w|\mathcal{D}))$$

Algorithm (Bayes by Backprop)

M.C. approximation using N samples $w^{(i)} \sim q(w|\theta)$

$$\mathcal{L}_{BBB}(\theta, \mathcal{D}) \approx \sum_{i=1}^{N} \underbrace{\log(q(w^{(i)}|\theta))}_{-\left\|\frac{w^{(i)}-\mu}{\sigma}\right\|^2} - \underbrace{\log(\mathbb{P}(w^{(i)}))}_{\text{prior}} - \underbrace{\log(\mathbb{P}(\mathcal{D}|w^{(i)}))}_{\text{network}}$$

For each task:

$$(w^{(i)}) \sim q(w|\theta_t)$$

$$\theta_t \longrightarrow \mathcal{L}_{BBB}(\theta_t, \mathcal{D}) \xrightarrow{\text{gradient descent step}} \theta_{t+1}$$

until loss plateaus

Algorithm

Specific learning rate for each parameter

Update the learning rate between each task

$$\alpha_\mu^{(t+1)} = \alpha_\mu^{(t)} \sigma$$

$\alpha_\mu^{(t)}$
1
0.5
0.5

μ	σ
2	0.1
4	3
1.5	1

$\longrightarrow$

$\alpha_\mu^{(t+1)}$
0.1
1.5
0.5

Learning on a new task:

$$\begin{matrix} \theta^{(t)} \\ \alpha_\mu^{(t)} \end{matrix} \longrightarrow \alpha_\mu^{(t+1)} = \alpha_\mu^{(t)} \sigma \longrightarrow \theta^{(t+1)} = BBB(\theta^{(t)}, \alpha_\mu^{(t+1)})$$

Parameter evolution

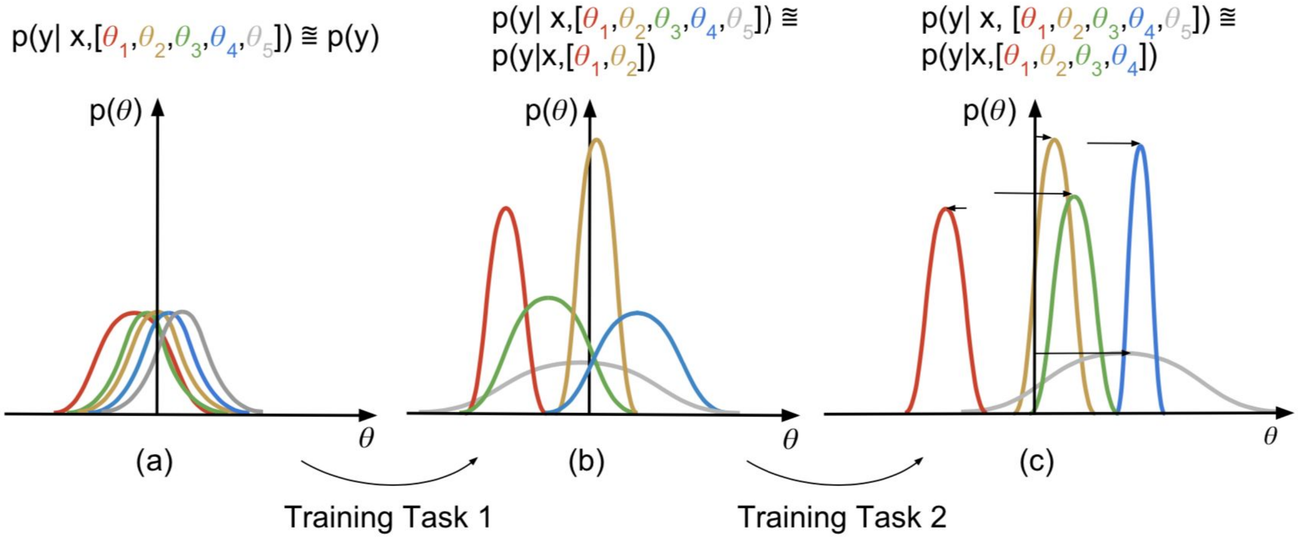

Results

Each task = 2 digits

Each task = one permutation

Each task = 2/20 classes

Catastrophic forgetting measure

Mean accuracy on all the tasks

(a) 5-Split MNIST, 5 tasks.

Method	BWT	ACC
VCL-Vadam†	-	99.17
VCL-GNG†	-	96.50
VCL	-0.56	98.20
IMM	-11.20	88.54
EWC	-4.20	95.78
HAT	0.00	99.59
ORD-FT	-9.18	90.60
ORD-FE	0.00	98.54
BBB-FT	-6.45	93.42
BBB-FE	0.00	98.76
UCB-P (Ours)	-0.72	99.32
UCB (Ours)	0.00	**99.63**
ORD-JT*	0.00	99.78
BBB-JT*	0.00	99.87

(b) Permuted MNIST, 10 permutations.

Method	#Params	BWT	ACC
SI ‡	0.1M	-	86.0
EWC ‡	0.1M	-	88.2
HAT ‡	0.1M	-	91.6
VCL-Vadam†	0.1M	-	86.34
VCL-GNG†	0.1M	-	90.50
VCL	0.1M	-7.90	88.80
UCB (Ours)	0.1M	-0.38	91.44
LWF	1.9M	-31.17	65.65
IMM	1.9M	-7.14	90.51
HAT	1.9M	0.03	97.34
BBB-FT	1.9M	-0.58	90.01
BBB-FE	1.9M	0.02	93.54
UCB-P (Ours)	1.9M	-0.95	97.24
UCB (Ours)	1.9M	0.03	**97.42**
BBB-JT*	1.9M	0.00	98.12

(c) Alternating CIFAR10/100

Method	BWT	ACC
PathNet	0.00	28.94
LWF	-37.9	42.93
LFL	-24.22	47.67
IMM	-12.23	69.37
PNN	0.00	70.73
EWC	-1.53	72.46
HAT	-0.04	78.32
BBB-FE	-0.04	51.04
BBB-FT	-7.43	68.89
UCB-P (Ours)	-1.89	77.32
UCB (Ours)	-0.72	**79.44**
BBB-JT*	1.52	83.93

Joint training

Pros and cons

Pros

- Stability
- Know when the network cannot learn more
- General framework

Cons

- Huge training time
- Low plasticity

Experience Replay for CL

D. Rolnick et al., 2018

Source: https://papers.nips.cc/paper/8327-experience-replay-for-continual-learning

Motivation

Hippocampal replay:

- reduce catastrophic forgetting
- improve generalisation

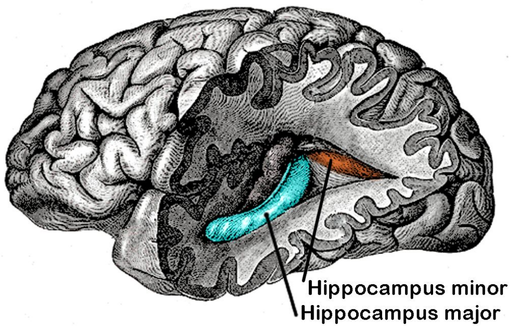

Experience replay

Experience replay

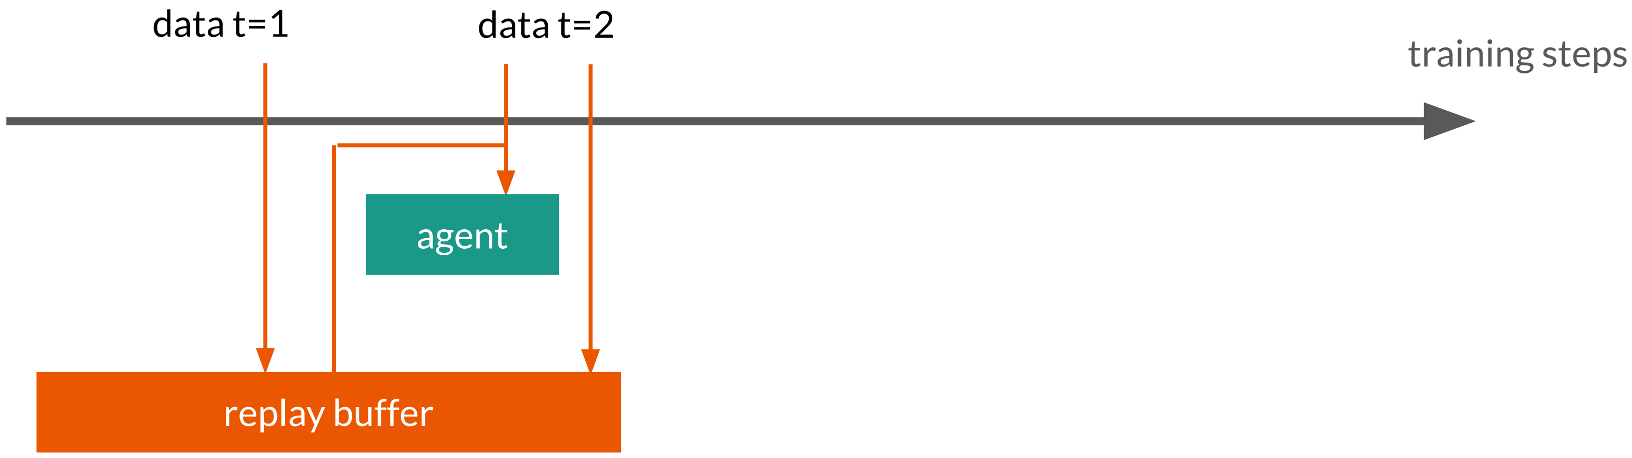

Experience replay

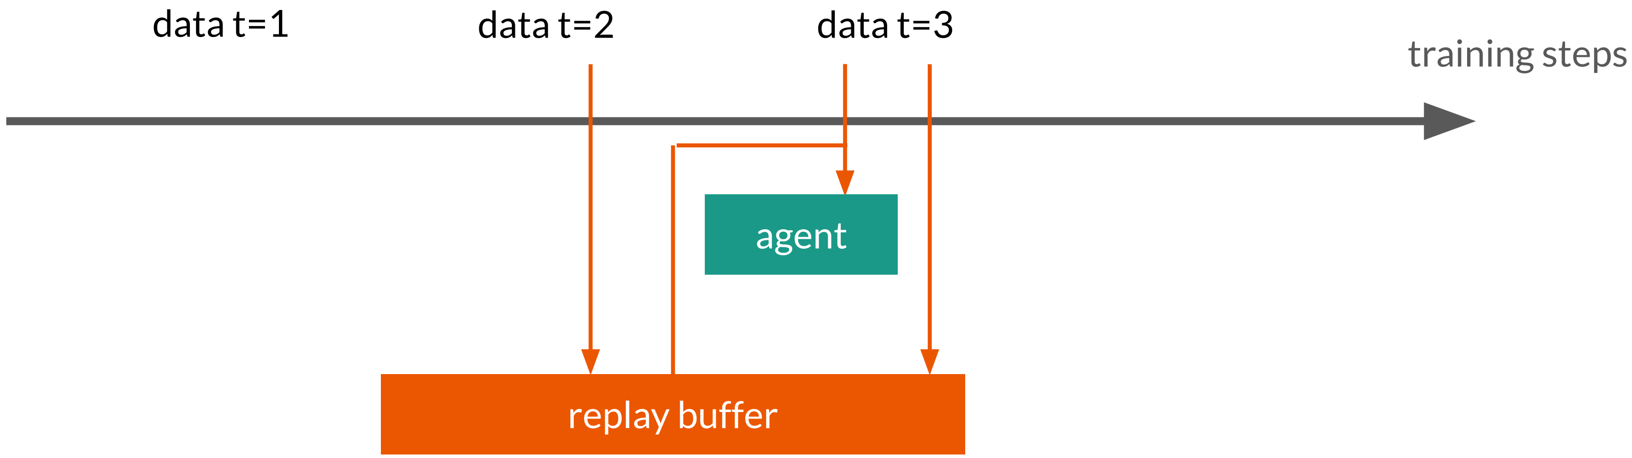

Training modes

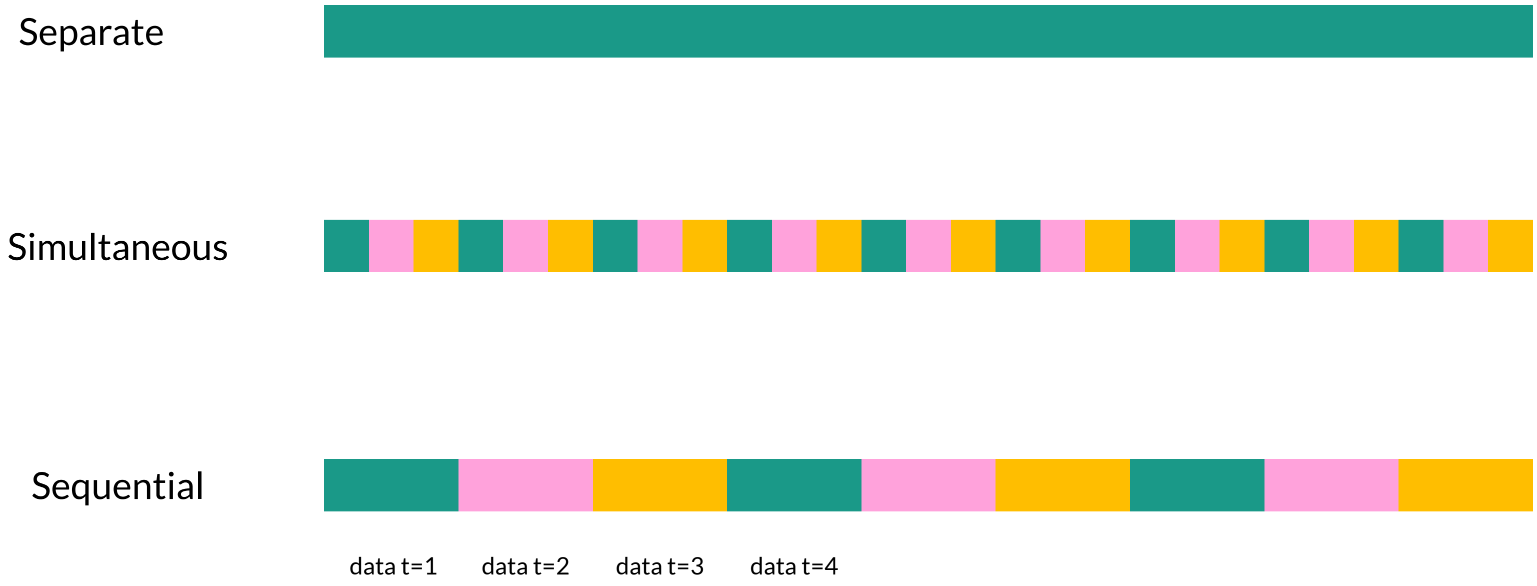

Without replay

minimal interference BUT catastrophic forgetting

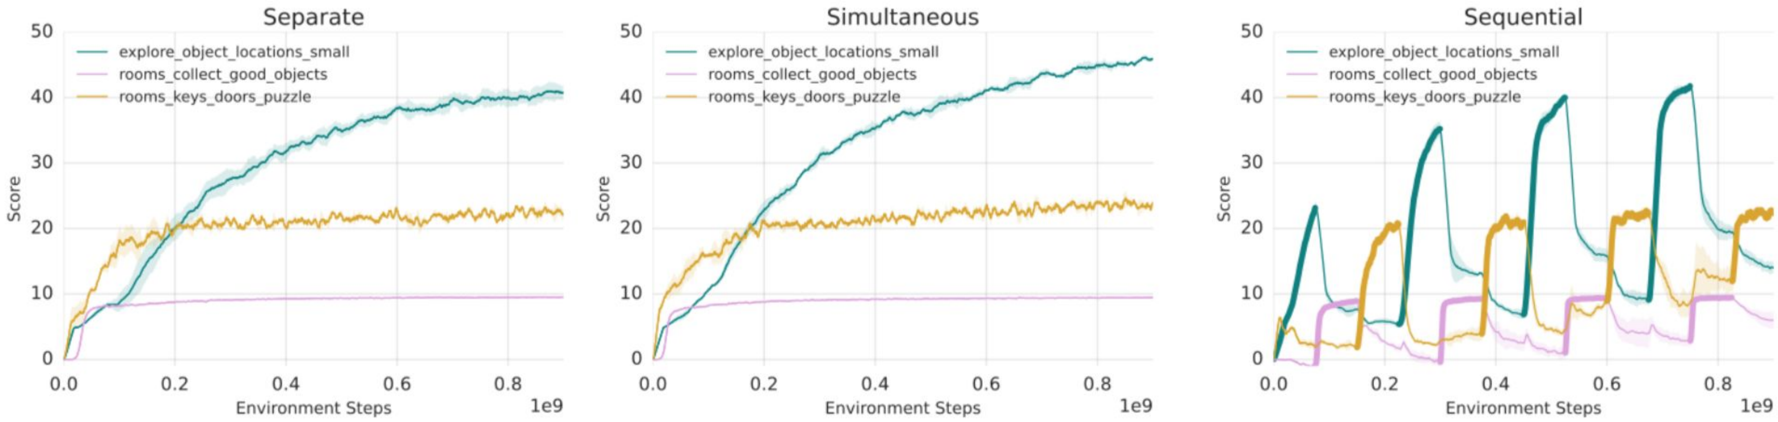

with V-Trace learning algorithm

Benefits of experience replay

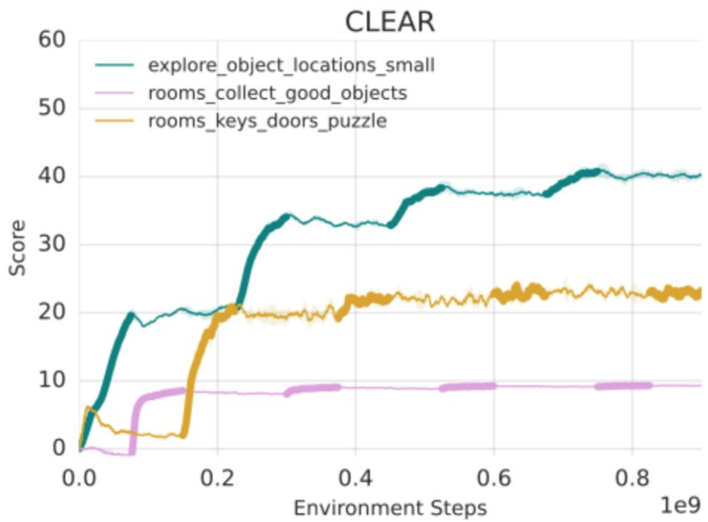

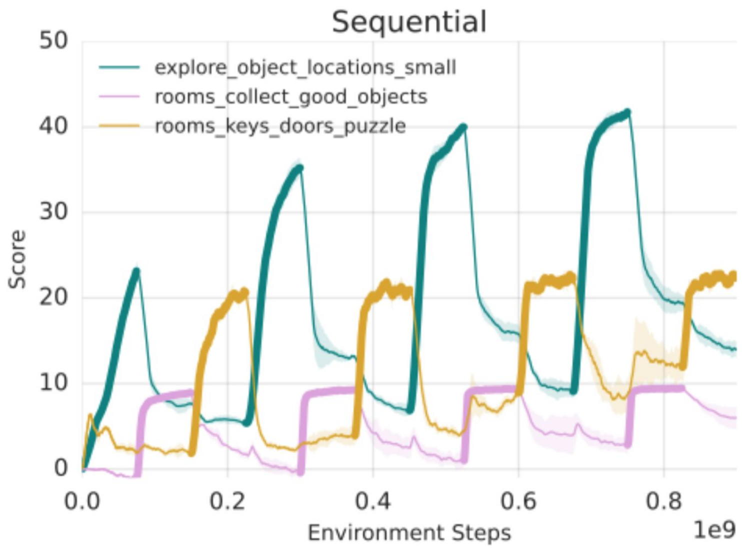

Mix new/old experience

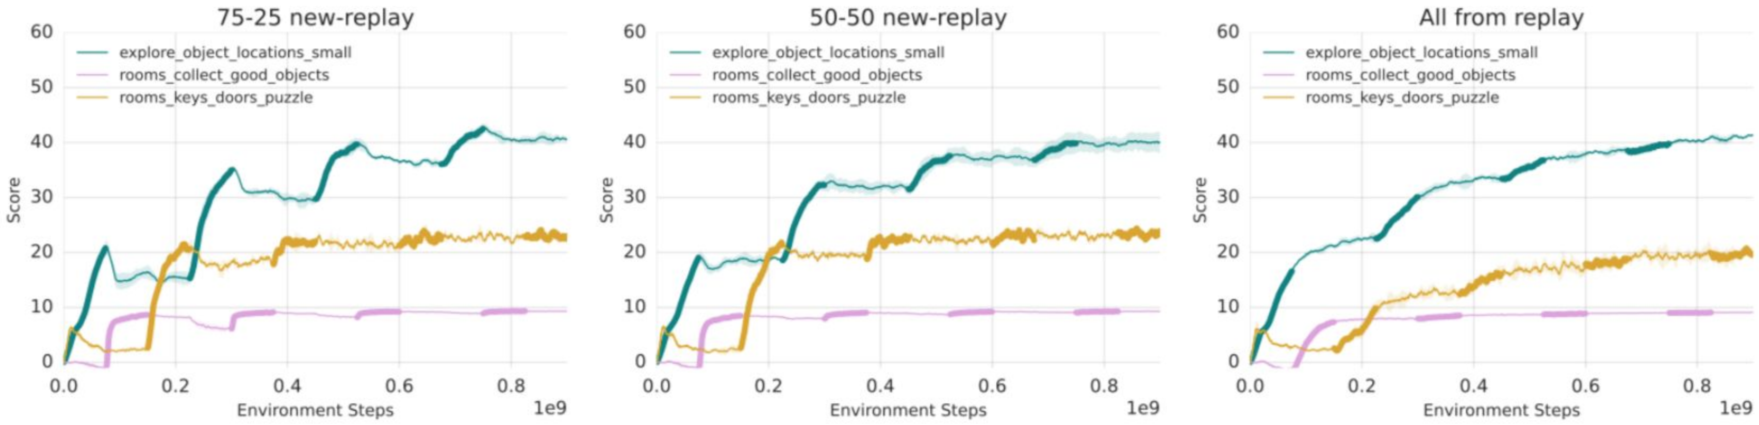

Pros and cons

Pros

- Insights on why experience replay works
- Nice example of Neuroscience justifying AI algorithms

Cons

Thank you!