DDGR: Continual Learning with Deep Diffusion-based Generative Replay

Rui Gao[1], Weiwei Liu[1]

[1]School of Computer Science, National Engineering Research Center for Multimedia Software,

Institute of Artificial Intelligence and Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University, Wuhan, China

Email: {grcswhu1002,liuweiwei863}@gmail.com

Motivation

We focus on the challenges of generative replay (GR) methods in continual learning:

- GR method is one of the most popular methods in continual learning, but most of them ignore the benefit offered by the classifier to the generator.
- Existing GR methods using VAE or GAN which is beaten by diffusion models [Ho et al., 2020].
- Most GR methods reuse generated samples which may produce low-quality samples for previous tasks [Shin et al., 2017].

- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In NeurIPS, 2020.
- Shin, H., Lee, J. K., Kim, J., and Kim, J. Continual learning with deep generative replay. In NeurIPS, pp. 2990–2999, 2017.

Our Method

Denoising diffusion probabilistic model

- Our work is based on DDPM [Ho et al., 2020] which is a typical diffusion model architecture and consists of a forward and a reverse processes. We adopt DDPM as the generator.

Learning $p_\theta(\mathbf{x}_n \mid \mathbf{x}_{n+1})$ to approximate $q(\mathbf{x}_n \mid \mathbf{x}_{n+1})$

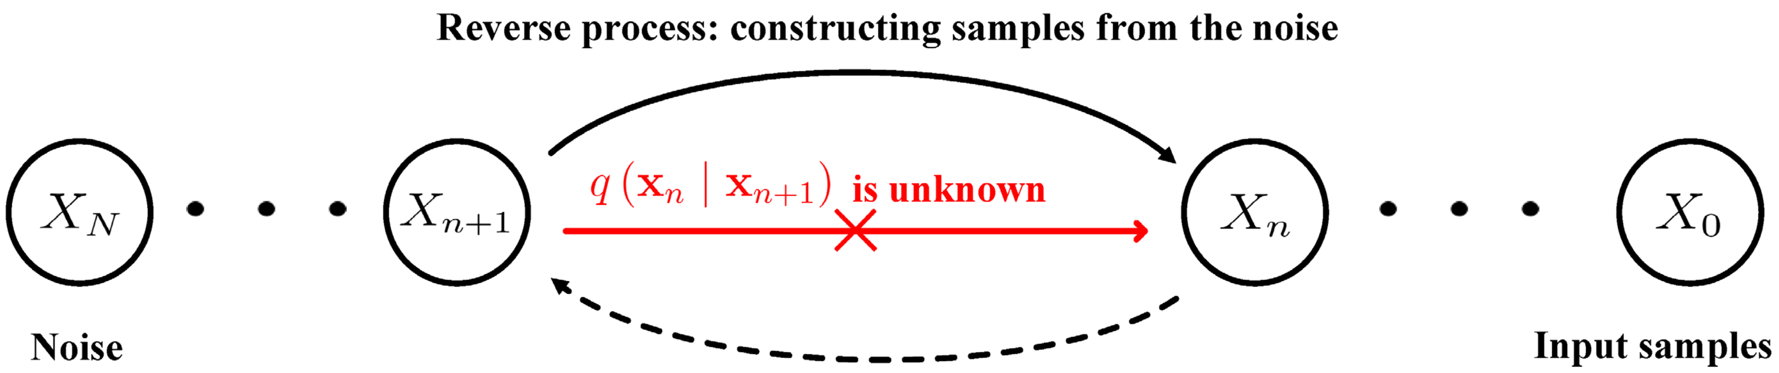

Forward process: adding Gaussian noise

$$\mathbf{x}_{n+1} = \sqrt{\bar{\alpha}_{n+1}}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_{n+1}}\boldsymbol{\epsilon}$$

Instruction-operator

- We use the classifier, which was trained in sequence on previous tasks, to instruct the sampling of DDPM. Specifically, we calculate the instruction-operator through the classifier .

$$\mathbf{G} = \nabla_{\mathbf{x}_n} \log p_\zeta(y \,|\mathbf{x}_n)\big|_{\mathbf{x}_n = \boldsymbol{\mu}_\theta} = \nabla_{\mathbf{x}_n} \ell(\boldsymbol{f}_\zeta(\mathbf{x}_n), y)\big|_{\mathbf{x}_n = \boldsymbol{\mu}_\theta}$$

- The instruction-operator can be regarded as a type of distillation of previous knowledge to classifier. Benefiting from the instruction of the classifier, DDGR improves the quality of the samples of previous tasks produced by the generator.

Deep diffusion-based generative replay

- (1) Generate samples from DDPM with instruction-operator.
- (2) Train the classifier with generated samples and current tasks' data.
- (3) Update DDPM.

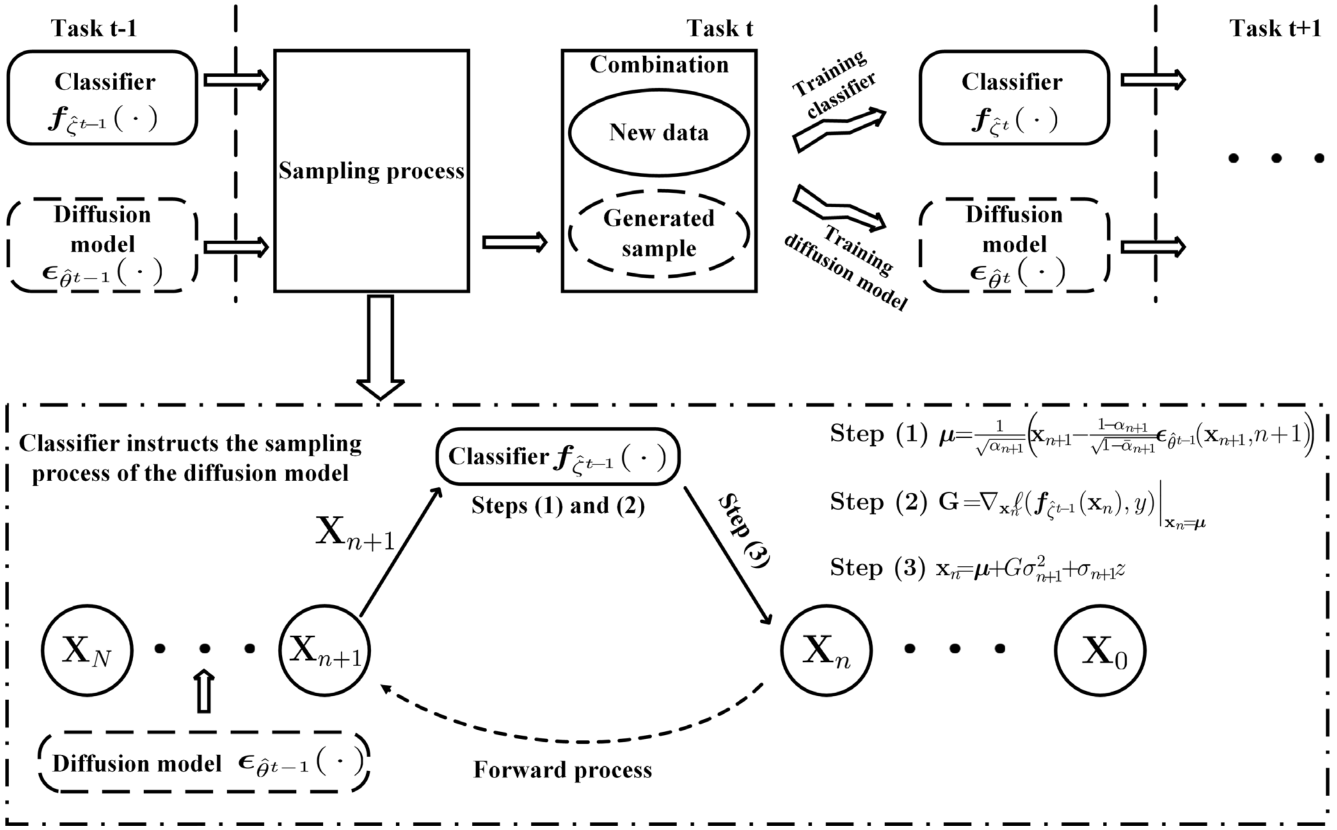

Experiments

Datasets

In this paper, we consider two scenarios commonly encountered in CL, namely class incremental (CI) [van de Ven et al., 2019] and class incremental with repetition (CIR) [Cossu et al., 2022].

- CI: CIFAR-100 and ImageNet [Deng et al., 2009]
- CIR: CORe50 [Lomonaco et al., 2017].

	Tasks	Classes/task	Train data/task	Task selection
CIFAR-100	5	$\{50, 5\}$ or $\{50, 10\}$	$\{25000, 2500\}$ or $\{25000, 5000\}$	random class
ImageNet	5	$\{500, 50\}$ or $\{500, 100\}$	$\{650000, 65000\}$ or $\{650000, 130000\}$	random class
CORe50	79	$\{10, 5\}$	$\{3000, 1500\}$	random class

- van de Ven, G. M. and Tolias, A. S. Three scenarios for continual learning. CoRR, abs/1904.07734, 2019.
- Cossu, A., Graffieti, G., Pellegrini, L., Maltoni, D., Bacciu, D., Carta, A., and Lomonaco, V. Is class-incremental enough for continual learning? Frontiers Artif. Intell., 5: 829842, 2022.
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In CVPR, pp. 248–255, 2009.
- Lomonaco, V. and Maltoni, D. Core50: a new dataset and benchmark for continuous object recognition. In CoRL, 2017.

Results in CI

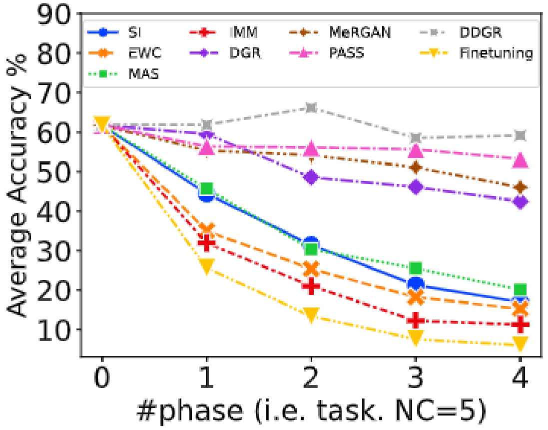

(a) CIFAR-100 AlexNet NC=5

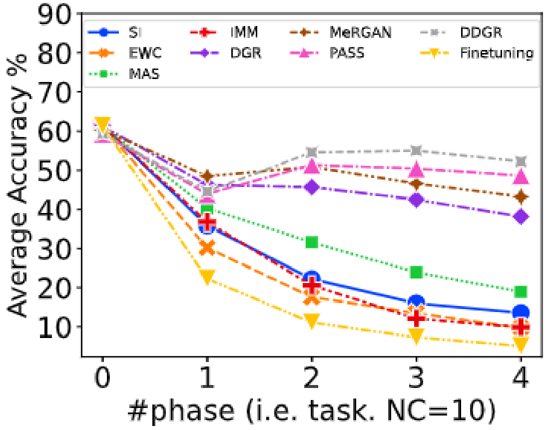

(b) CIFAR-100 AlexNet NC=10

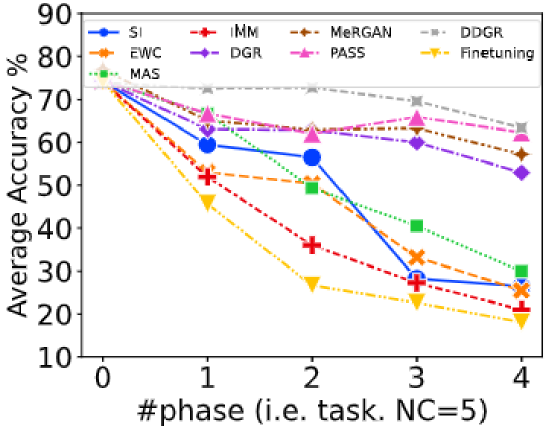

(c) CIFAR-100 ResNet NC=5

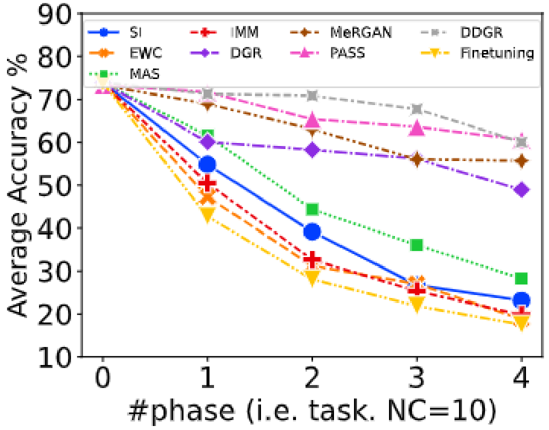

(d) CIFAR-100 ResNet NC=10

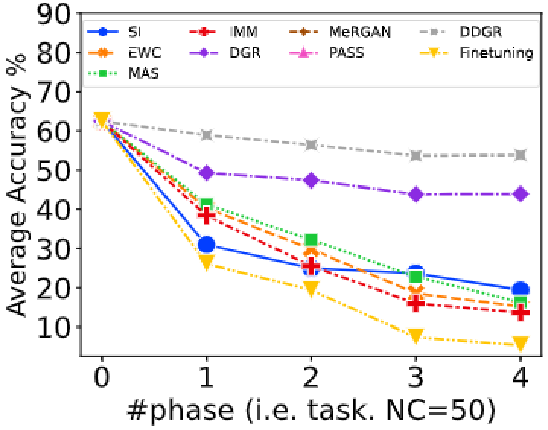

(e) ImageNet AlexNet NC=50

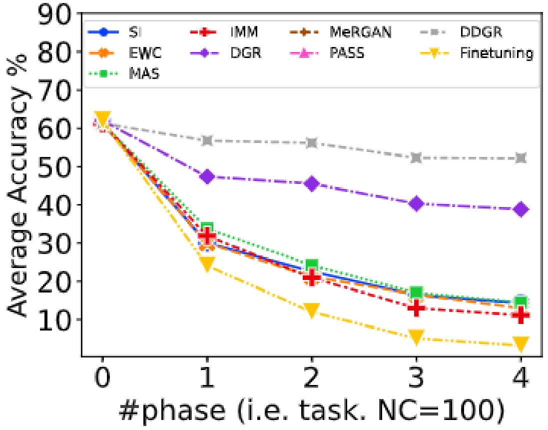

(f) ImageNet AlexNet NC=100

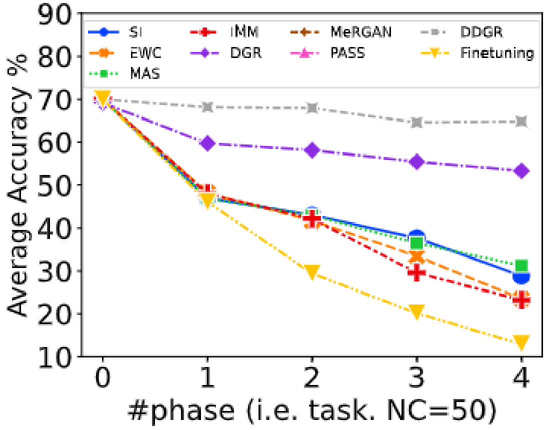

(g) ImageNet ResNet NC=50

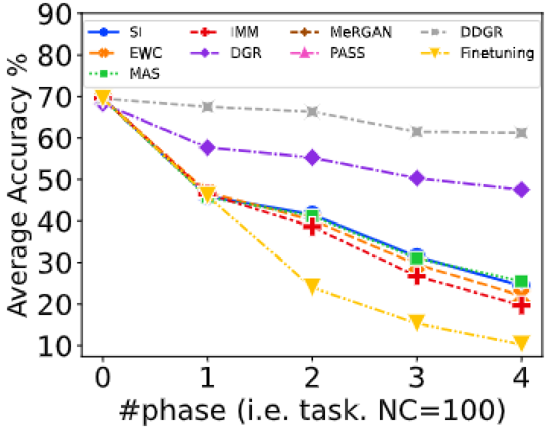

(h) ImageNet ResNet NC=100

Results in CIR

In class incremental scenario with repetition, we refer to each task as a batch.

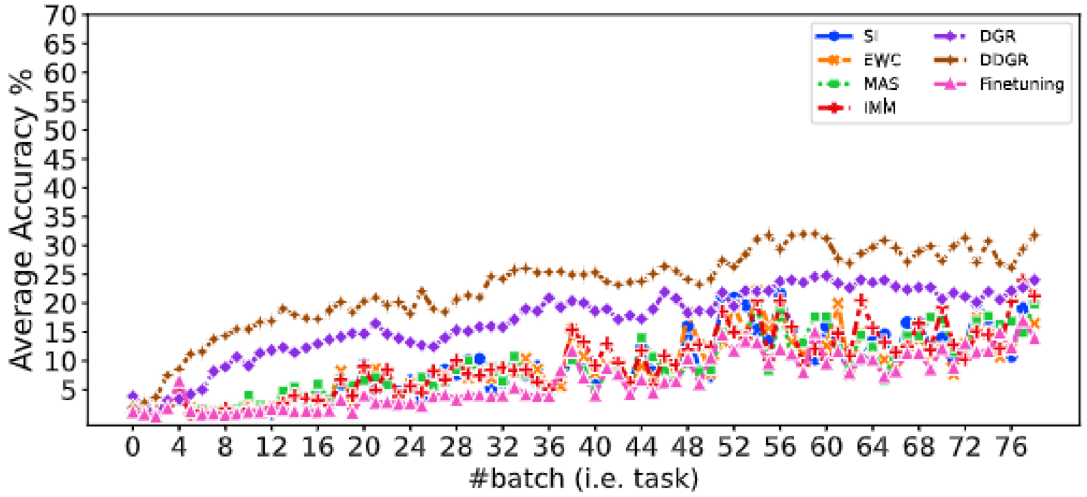

(a) AlexNet

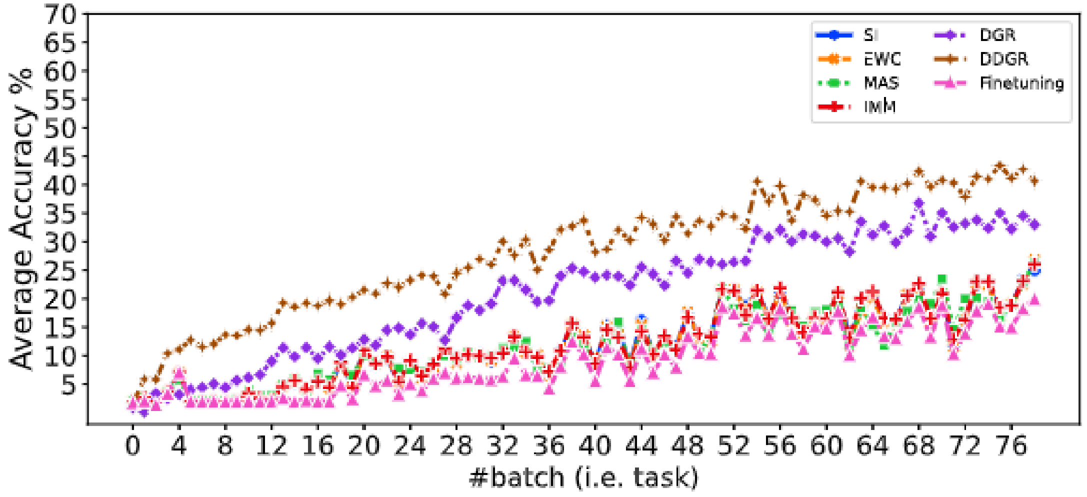

(b) ResNet

Conclusion

Main Contributions

Our main contributions can be summarized as follows:

- We focus on sampling process of a diffusion model and explore how this process might be instructed by a pretrained classifier. Specifically, we calculate instruction-operator through classifier at each time step of diffusion model to guide the generation of samples.
- The novel DDGR is proposed based on a diffusion model, where the classifier uses the instruction-operator to instruct the sampling process of the diffusion model. Benefiting from the instruction-operator, DDGR significantly improves the quality of generated samples for previous tasks.
- Extensive experimental results under class incremental (CI) and class incremental with repetition (CIR) settings demonstrate the advantages of DDGR.

Thanks for watching!