

IB-DRR - Incremental Learning with Information-Back Discrete Representation Replay

Jian Jiang

Edoardo Cetin

Oya Celiktutan

Centre for Robotics Research, Department of Engineering, King's College London, UK

{jian.jiang, edoardo.cetin, oya.celiktutan}@kcl.ac.uk

Abstract

Incremental learning aims to enable machine learning models to continuously acquire new knowledge given new classes, while maintaining the knowledge already learned for old classes. Saving a subset of training samples of previously seen classes in the memory and replaying them during new training phases is proven to be an efficient and effective way to fulfil this aim. It is evident that the larger number of exemplars the model inherits the better performance it can achieve. However, finding a trade-off between the model performance and the number of samples to save for each class is still an open problem for replay-based incremental learning and is increasingly desirable for real-life applications. In this paper, we approach this open problem by tapping into a two-step compression approach. The first step is a lossy compression, we propose to encode input images and save their discrete latent representations in the form of ‘codes’ that are learned using a hierarchical Vector Quantised Variational Autoencoder (VQ-VAE). In the second step, we further compress ‘codes’ losslessly by learning a hierarchical latent variable model with bits-back asymmetric numeral systems (BB-ANS). To compensate for the information lost in the first step compression, we introduce an Information Back (IB) mechanism that utilizes raw exemplars for a contrastive learning loss to regularise the training of a classifier. By maintaining all seen exemplars’ representations in the format of ‘codes’, Discrete Representation Replay (DRR) outperforms the state-of-art method on CIFAR-100 by a margin of 4% average accuracy with a much less memory cost required for saving samples. Incorporated with IB and saving a small set of old raw exemplars as well, the average accuracy of DRR can be further improved by 2%.

1. Introduction

Deep neural networks leveraging large-scale annotated datasets have been shown to be powerful in many real-world

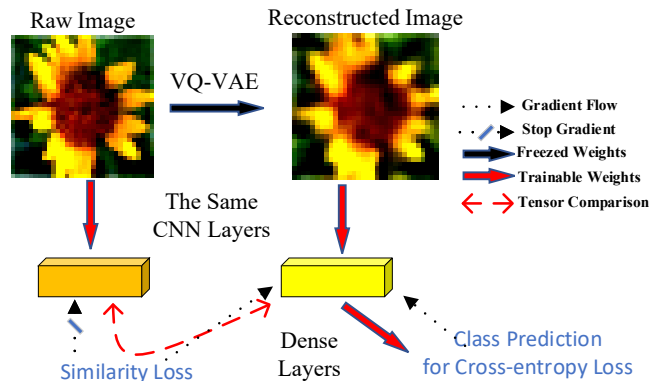


Figure 1. Information Back (IB) mechanism inside our classifier. IB maximizes the similarity between the latent representations of the raw image and its reconstructed view. With the help of IB, raw images are indirectly used for regularizing the optimisation of classification loss, i.e., the cross-entropy loss in our case.

tasks such as image classification. One downside is that these data-driven methods work under the assumption that all training samples (exemplars) are simultaneously available during the training phase [10, 19, 17]. However, due to the growing need for systems that can adapt to dynamic environments and can continually learn new tasks, current deep neural networks are not adequate as they suffer from catastrophic forgetting - when a model is continuously updated using novel incoming data, the updates can override knowledge acquired from previous classes. Incremental learning, also known as continual learning, never-ending learning or life-long learning, aims to design systems that can keep learning new knowledge while maintaining the performance for the previously learned tasks.

An effective and commonly used strategy in the current incremental learning methods is *replay* or *rehearsal*, which is mixing the new data with the old data when learning new information to maintain old knowledge. This can be achieved by saving a small subset of old data [24, 39, 26, 14, 36, 13, 31]. However, the success of this strategy highly depends on how many and which exemplars to

save [24, 36, 14]. Alternatively, a line of research has focused on learning a generator to produce pseudo old samples [23, 41, 32]. However, the generator needs to be updated using fake old data that may lead to a vicious cycle if some pseudo data of bad quality is produced. Recent work [23, 41, 32] has shown that image generator is unstable and unreliable in incremental learning for complex dataset such as CIFAR-100 [18] and only works on small and relatively simple datasets such as MNIST [20]. Although carefully selected exemplars or generated representative exemplars can help improve the performance up to some degree [36, 14], the larger number of exemplars a model inherits the better performance it can achieve. However, saving all available data is memory-expensive. The trade-off between model performance and the number of exemplars per class is hard to be optimised due to different applications requirement and varies among different distinctive tasks.

In this paper, we approach this open problem by introducing *Discrete Representation Replay* (DRR). The DRR module allows saving the maximum number of exemplars possible at a small memory cost with two-step compression. As shown in Figure 2, DRR benefits from a Vector-Quantised Variational Autoencoder (VQ-VAE) [29, 35] as well as a Bits-Back Asymmetric Numeral Systems (BB-ANS) that paired with hierarchical latent variables. The VQ-VAE reduces the memory required for 50,000 uncompressed training exemplars on the CIFAR-100 [18] dataset from 146.48 MB to 7.26 MB and BB-ANS further reduces it to 4.24 MB, resulting in an overall memory gain of 97%. To learn new classes, DRR combines latent representations (in the format of *codes*) of new training samples with those of old samples and reconstruct them back to the RGB image domain at run time. Our experiments on CIFAR-100 [18] showed that Discrete Representation Replay (DRR) outperformed the state-of-the-art replay-based method [24] in the multi-class incremental learning setting by a margin of 4% while reducing the memory size required for saving exemplars. To alleviate the loss of information caused by VQ-VAE, we design a negative-pairs-free contrastive learning component called Information Back (IB) inside our classifier that saves and uses raw images to regularize the training of the classifier, see Fig. 1. We assume the reconstruction processes of different images share similar transformation patterns as the same VQ-VAE is applied. Note that IB requires to save an additional subset of raw previously seen samples but they are not directly used for minimizing the classification loss. They serve as regularizers to force the classifier to get back some general information lost during image reconstruction. Incorporated with IB component, called IB-DRR, the testing accuracy of DRR increases by 2% on CIFAR-100. Additionally, our empirical results show that even without maintaining previously seen raw

data IB-DRR can still boost performance in some settings, especially, the settings where the number of new classes to learn is large. In summary, our main contributions are as follows:

- We formalize incremental learning as a representation compression problem and introduce *Discrete Representation Replay* (DRR) strategy for this task. More explicitly, we propose a two-step compression which allows us to save training exemplars at a competitively small memory cost. This is enabled by a VQ-VAE that is used to save and reconstruct images during new training phases and hierarchical latent variable models with BB-ANS that further reduces the memory cost.
- We introduce an Information Back (IB) mechanism that utilizes a subset of raw images for old classes and all raw images from new classes via a contrastive learning loss to alleviate the information loss and regularise the training.
- Finally, we show that on the CIFAR-100 dataset [18] our DRR surpasses the state-of-the-art method [24] by a margin of 4% without saving any raw examples. And by saving a small number of raw samples, IB-DRR yields the best performance with a margin of 6% as compared to the state-of-the-art method [24]. Our preliminary results on the ImageNet [7] also indicate the viability of this approach across different datasets.

2. Related Work

There is a significant body of work on incremental learning, which can be divided into three broad categories, namely, regularization-based (e.g., [17, 45]), parameter-isolation-based (e.g., [33, 30]) and replay-based (e.g., [39, 14, 27, 4, 24, 23]). Our proposed approach falls into the category of replay-based methods. Replay, also known as rehearsal, is the most competitive and popular strategy in incremental learning and achieves state-of-the-art performance on many benchmark datasets including CIFAR-100 [18], MNIST [20] and ImageNet [40]. We summarise these replay-based methods in three categories, namely, image-level replay, feature-level replay, and compression for replay.

Image-level Replay. Storing and replaying exemplars in the memory to solve catastrophic forgetting was first introduced in 1990s [38]. Recently, Isele *et al.* [14] reintroduced replay strategy in incremental learning and proposed a ranking function to select and store exemplars for each class in a long-term memory called episodic memory. To control the memory required for replay, Rebuffi *et al.* [36] proposed a method name iCaRL that used a fixed size memory for all classes. To maintain the most important exemplars, they used a buffer-reconstruction procedure called Herding to select a number of exemplars from a new class to save and dis-

card the least preferred exemplars from previous classes to prevent memory from growing. To address the imbalance between a small number of exemplars for old classes and substantially more samples of new classes, Hou *et al.* [13] proposed a method named (LUCIR) that treated both old and new classes uniformly by introducing cosine normalization on latent representations of samples of new classes and old classes. They also selected samples to save by leveraging an online mining method. Most recently, Liu *et al.* [24] developed a method called Mnem, where exemplars to save were considered as learnable parameters. The parameterized exemplars of old classes were optimized in an end-to-end manner during the training, and by utilizing strategies used in [13] their model achieved state-of-the-art performance as compared to existing replay-based methods on both CIFAR-100 and ImageNet. On the other hand, using exemplars directly for training classifiers may be prone to over-fitting [27, 4]. Therefore, Gradient Episodic Memory (GEM) [27] was proposed to indirectly utilize saved exemplars. Their model computed gradients given new data as well as a set of task-specific gradients using episodic memory, and it projected the gradients in the feasible region defined by task gradients. Like GEM, our Information-Back Discrete Representation Replay (IB-DRR) also used saved raw exemplars indirectly for classification. To avoid saving original inputs, a number of works [41, 32, 46] utilised a generative model such as Generative Adversarial Network (GAN) [8] or Variational Autoencoder (VAE) [15] to produce fake data that is similar to data from old classes. In this way, they aimed to transfer incremental learning problem into image generating problem. However, the generative model itself needs to maintain the ability to generate fake data for old tasks, and this approach still has to deal with catastrophic forgetting that occurs in the generative model rather than the classifier. This has been demonstrated by a recent work [22] pointing out that a GAN itself has a catastrophic forgetting problem and the training of GAN can be regarded as an incremental learning problem. In addition, fake data may not reflect the real data distribution, which can potentially lead to a vicious cycle. Therefore, training a generative model in incremental learning settings is unstable and unreliable and recent generative-replay-based methods [41, 32, 46] work on comparatively simple datasets such as MNIST [20] and SVHN [28] datasets.

Feature-level Replay. Another line of work has focused on saving the latent representations instead of saving original image data or generating fake data. Indeed, in some incremental learning scenarios [23, 1, 41, 21], training data from a certain task is no longer available after a model is trained for the task. Because saving data firstly leads to privacy and security issues and secondly results in a significant increase in the memory as the model gradually learns new classes. Pellegrini *et al.* [31] pretrained a classifier on

ImageNet [40]. Then its low-level layers were kept and frozen to serve as a feature extractor for other datasets like Core50 [25]. They extracted and saved a small subset of latent representations (namely, the output of the feature extractor) of inputs and replayed them during new training phases. They trained and updated a classifier using both the features of new data and replayed features. Recently, Xialei *et al.* [23] combined a generative model with a feature-level replay strategy and used a GAN to produce fake latent representations and achieved competitive performance on the CIFAR-100. Their model called generative feature replay (GFR) consisted of a modified classifier that took latent representation as inputs, a feature extractor and a generator. The modified classifier was jointly trained using cross-entropy loss with the feature extractor which encodes original images to corresponding latent representations. After trained on new data, they were frozen and the generator was updated using latent representations produced by them given corresponding data. Besides, they applied the knowledge distillation technique to the feature extractor to avoid forgetting. Feature-level replay methods generally require less memory cost compared to image-level counterparts when used to save the same amount of samples. Our approach can also be regarded as a feature-level replay method. Because we save features rather than images and the saved features are replayed and reconstructed to the image domain by the decoder in our model for training a classifier.

Compression for Replay. There are three most related continual learning works [2, 3, 37] similar to our DRR. Ayub *et al.* [2] used a pretrained autoencoder structure that saved compressed latent representations of previous samples and reconstructed them back to image-domain for replaying. Caccia *et al.* [3] used stacked VQ-VAEs and updated the codebooks to control drifting representations during the incremental learning whereas our DRR used a hierarchical VQ-VAE and our codebooks were frozen after the VQ-VAE was pretrained. Caccia *et al.* reported that saving more samples might degrade the image quality due to the drift that occurred in their VQ-VAE. Thus, in their approach storing more samples did not necessarily improve classification performance. Our results showed that our VQ-VAE with frozen codebooks did not suffer from data shift problem much. Matthew *et al.* [37] used a discrete VAE model with a ‘codes’ buffer. However, unlike our DRR whose latent space was able to represent infinite latent representations of samples, their latent space had a restricted size that was proportional to the VAE capacity (VAE capacity was a major hyperparameter in their work). As shown in their work, the smaller capacity their VAE had, the higher compression it could achieve but the fewer maximum samples they could save. Besides, to make their discrete VAE fit new data, the VAE was updated using reconstructed images via

codes in the buffer as well as real images from new tasks. It potentially transferred the problem of catastrophic forgetting from classifiers into generators like other generative models we discussed before. In contrast, our trained DRR has a stable high compression rate with generally good reconstructed image quality and can be generalized to other datasets (see Appendix for reconstructed images via DRR).

3. Background

3.1. Vector Quantised Variational AutoEncoder

Vector Quantised Variational AutoEncoder (VQ-VAE) [29] was originally designed for reducing the computation workload in PixelCNN [44] for the task of image generation. PixelCNN-based methods [44, 6] can generate high-quality fake images, however, due to the nature of auto-regressive models, they need to be trained over each pixel leading to a considerably high computation cost. The VQ-VAE encodes an RGB image into a ‘codes’ matrix with a smaller size. Instead of original images, the codes are then used as an input to auto-regression models, resulting in a significant computation reduction. The VQ-VAE relies on vector quantisation that encodes inputs into discrete representations rather than continuous ones and as a compressor it has been shown to be effective in image generation and image compression [29, 35]. Similarly, in this paper, we benefit from the VQ-VAE not only as a good feature extractor but also as a powerful compressor. Considering that a PixelCNN is also a generative model that has to deal with incremental learning as we discussed in Section 2, instead of generating fake images using an auto-regressive generator, our model reconstructs the images via saved codes using the VQ-VAE only. The VQ-VAE defines a discrete embedding space called ‘codebook’ $\mathbf{c} \in \mathbb{R}^{K \times d}$, where K is the size of codebook and d is the dimension of each embedding vector $c_k \in \mathbb{R}^d, k \in 1, 2, 3, \dots, K$. The encoder $E(x)$ learns a non-linear mapping that encodes the input $\mathbf{x} \in \mathbb{R}^M$ to the latent representation $\mathbf{z}_e \in \mathbb{R}^d$, where $d \ll M$. Then vector quantisation $VQ(\cdot)$ is applied over \mathbf{z}_e and, given the codebook (the prototype set), \mathbf{z}_e is replaced with a set of embedding vectors in the codebook using nearest neighbor search as formalised by

$$\mathbf{z}_q = VQ(\mathbf{z}_e) = \mathbf{c}_k, \text{ where } \mathbf{k} = \arg \min_k \|\mathbf{z}_e - c_k\|_2 \quad (1)$$

The quantised output \mathbf{z}_q is then fed into the decoder to obtain the reconstructed data. The objective of VQ-VAE can be defined as

$$\mathcal{L}(\mathbf{x}, D(\mathbf{c}_k)) = \|\mathbf{x} - D(\mathbf{c}_k)\|_2^2 \quad (2)$$

$$+ \|sg[E(\mathbf{x})] - \mathbf{c}_k\|_2^2 \quad (3)$$

$$+ \beta \|sg[\mathbf{c}_k] - E(\mathbf{x})\|_2^2 \quad (4)$$

where the term in Eq. 2 is the *reconstruction error* of a typical autoencoder (e.g., mean squared error). And $D(\cdot)$ refers to the decoder of the VQ-VAE. The term in Eq. 3 is called *codebook loss*, which brings the selected embedding vectors \mathbf{c}_k close to \mathbf{z}_e . Besides, the gradients produced by Eq. 3 are applied to the codebook only. The last term in Eq. 4, called *commitment loss*, enables encoder to produce similar values to the chosen embedding vectors and it is only applied to the encoder weights. A hyperparameter β is used to balance the learning rates of different terms and $sg[\cdot]$, namely stop-gradient, is an operation that prevents gradients to propagate to its argument. For example, in Eq. 3 $sg[E(x)]$ blocks the gradients calculated by this loss term from flowing into the weights of the encoder $E(\cdot)$. The visualized vector quantisation operation can be found in Fig. 2.

3.2. Bits-Back Asymmetric Numeral Systems

Asymmetric Numeral Systems (ANS) is used to compress sequences of discretely distributed symbols (a ‘symbol’ is a one dimension data point) into a sequence of bits and can recover those bits back to symbols. Bits-back with ANS (BB-ANS) [42, 43, 16] tries to approximate true data distribution $p_{data}(\mathbf{x})$ by $p_\theta(\mathbf{x})$ and encode datapoint into a number of bits equal to its negative log probability assigned by a latent variable model. BB-ANS models use latent variable models with a marginal distribution $p_\theta(\mathbf{x})$ defined by $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \int p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}$, and they utilize an inference model $q_\theta(\mathbf{z}|\mathbf{x})$ of variational autoencoders (VAE) [15] to approximate the posterior $p_\theta(\mathbf{z}|\mathbf{x})$. The marginal likelihood $p_\theta(\mathbf{x})$ is rewritten as:

$$\log p_\theta(\mathbf{x}) = \mathbb{E}_{q_\theta(\mathbf{z}|\mathbf{x})} \log \frac{p_\theta(\mathbf{x}, \mathbf{z})}{q_\theta(\mathbf{z}|\mathbf{x})} + \mathbb{E}_{q_\theta(\mathbf{z}|\mathbf{x})} \log \frac{q_\theta(\mathbf{x}, \mathbf{z})}{p_\theta(\mathbf{z}|\mathbf{x})} \quad (5)$$

where the first term in Eq. 5 is the Evidence Lower Bound (ELBO) that is jointly optimized using the reparameterization trick with the inference model and generative model of VAE. However, for latent variable models, to encode \mathbf{x} , it is necessary to encode \mathbf{z} as well, inducing an extra message length for the prior $p(\mathbf{z})$ equals to $-\log p(\mathbf{z})$. To get those extra bits back, BB-ANS firstly initializes ANS with a bit-stream of N random bits. Then it performs three steps (denoted as S1, S2, S3) as follows: S1 Decode \mathbf{z} from bit-stream using the inference model $q_\theta(\mathbf{z}|\mathbf{x})$; S2 Encode \mathbf{x} using generative model $p_\theta(\mathbf{x}|\mathbf{z})$; S3 Encode \mathbf{z} to bit-stream using the prior $p(\mathbf{z})$. The S2 and S3 add $-\log p_\theta(\mathbf{x}|\mathbf{z})$ bits and $-\log p(\mathbf{z})$ respectively, but BB-ANS gets $-\log q_\theta(\mathbf{z}|\mathbf{x})$ bits back from the S1. The net message length then becomes $\log q_\theta(\mathbf{z}|\mathbf{x}) - \log p_\theta(\mathbf{x}|\mathbf{z}) - \log p(\mathbf{z})$, which is on average equal to the negative ELBO. The cost of initial N random bits is negligible for encoding long sequences. A visualized BB-ANS is shown in Fig. 2.

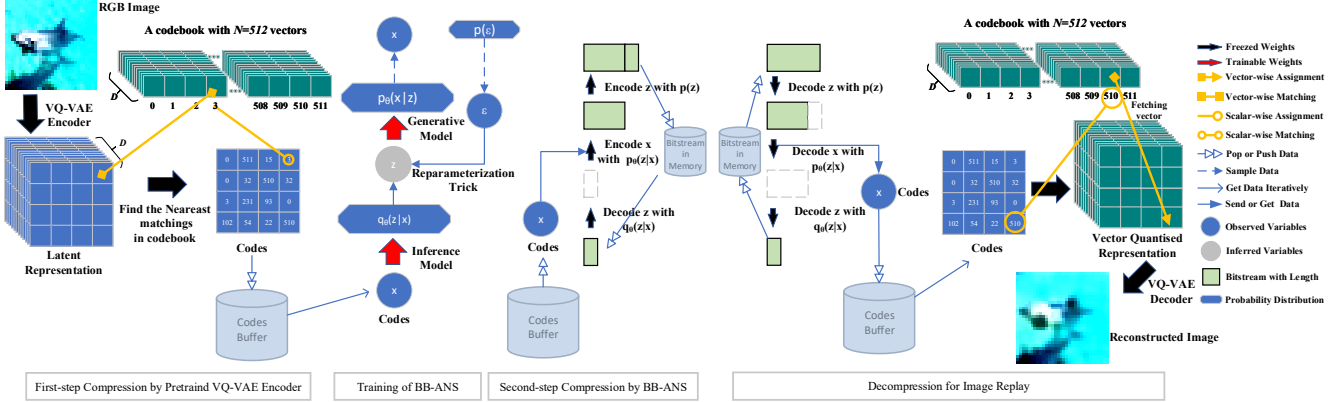


Figure 2. DRR uses a two-step compression. In the first step, a pretrained VQ-VAE encodes RGB images into codes that are saved in a buffer. During each new phase, BB-ANS is fine-tuned with codes from new images and those previous saved. Updated BB-ANS is used to perform the second step compression, i.e., compress codes and save them into bit-streams in the memory. Discrete latent representations are replayed by first decompressing bit-streams to codes and then decompressing the codes back to the image domain. This data flow only shows the core idea of DRR. In practice, a hierarchical VQ-VAE and a BB-ANS with hierarchical latent variables are used.

4. Discrete Representation Replay (DRR)

In this section, we formalize the incremental learning as a representation compression problem, and propose Discrete Representation Replay (DRR) by a two-step compression as shown in Fig. 2.

Lossy Compression by VQ-VAE. In the first step of compression, we utilize a VQ-VAE to process the inputs and save their discrete latent representations in a format called ‘codes’ that can be later used to reconstruct the original inputs. The VQ-VAE provides a lossy compression where ‘reconstructed’ images in our approach are a blurrier version of the original ones - naturally caused by an undercomplete autoencoder that gradually reduces the dimensionality of its intermediate layers and results in a smaller size of latent representation compared to the original input. This lossy compression contributes to the majority of the total compression rate of DRR. In our proposed model, we use a hierarchical VQ-VAE introduced in [35] (see Appendix for the network structure). Because our preliminary results have shown that the original VQ-VAE [29] cannot reconstruct the images adequately for our task. In our hierarchical VQ-VAE, vector quantisation is applied 2 times after different intermediate layers. It aims to learn hierarchical discrete latent representations such that a top-level representation models global information while a bottom-level representation captures local details. Two levels are trained jointly inside the hierarchical VQ-VAE and two independent codebooks are used respectively. Additionally, we adopt a fully convolutional design for our VQ-VAE, i.e., all of the layers are either convolutional or element-wise functions. In this way, the VQ-VAE is able to process images with arbitrary sizes. Our experiments show that our VQ-VAE trained on CIFAR-100 can reconstruct images with a higher resolution

well for classification like images from ImageNet. After we pretrain the VQ-VAE, we freeze and use it as a feature extractor and save the resulting latent representations in the form of codes, which significantly decreases the memory cost.

Lossless Compression by BB-ANS. In the second step, we utilize a BB-ANS with hierarchical latent variables called Bit-Swap introduced in [16] that is more expressive than original BB-ANS [42], i.e., can more closely model real data distribution. Instead of using a fixed prior for z , Bit-Swap assumes that z is generated by a latent variable z_1 , and z_1 is generated by another latent variable z_2 and so on. In this way, the sampling process of both the generative model and the inference model obeys a Markov chain dependency between the stochastic variables [16]. In a lossless manner, we use Bit-Swap to compress ‘codes’ produced by the VQ-VAE. We train from scratch a Bit-Swap model at first and keep fine-tuning it for future new data distribution of codes.

In DRR, a new classifier is trained with reconstructed images both from new classes and old classes. Our preliminary results show that training from scratch the classifier every time it is required to learn new classes always performs better than fine-tuning the previously learned classifier. Though it is widely accepted that the learned model weights contain information of old data and the learned weights are an unwritten definition for the concept ‘past knowledge’ in the increment learning domain, intuitively, if given access to old distribution, training the model from scratch is more flexible, leading a better fit to the new data distribution (a mixture of old and new data). In fact, He *et al.* [11] show that in some cases training from scratch is more robust and usually converges a solution *on par with or even better than* the fine-tuning counterpart.

5. Information-Back DRR (IB-DRR)

The pseudo data distribution produced by either generative models or reconstruction models has a drift towards original real data distribution. In our DRR, we lose the information of real data because of the nature of under-complete autoencoders as well as vector quantisation operations. This causes a considerably large decrease in classification score as compared to using original data, e.g., around an overall 12% decrease for our DRR on the CIFAR-100 dataset. We can regard the reconstructed images as augmented data that is transformed by some blurring transformation, though this kind of augmentation may not necessarily be a good augmentation for classification. To remedy this problem, we propose to utilize contrastive learning by treating an image and its reconstructed version as a positive pair. Inspired by a recent negative-pairs-free contrastive learning method called SimSiam [5], we propose an Information Back (IB) mechanism for DRR, see Fig. 1. Specifically, our IB-DRR maximizes the cosine similarity of the latent representations (produced by the classifier) of real images and their reconstructed views produced by VQ-VAE. We denote the convolution layers of a classifier before its dense layers (fully connected layers) as f , and the output of f , i.e., latent representations of raw input x and its reconstructed view \hat{x} as $r_1 \triangleq f(x)$ and $r_2 \triangleq f(\hat{x})$, respectively. The objective of IB is to minimize the negative cosine similarity between r_1 and r_2 :

$$\mathcal{L}_{sim} = -sg \left[\frac{r_1}{\|r_1\|_2} \right] \cdot \frac{r_2}{\|r_2\|_2}$$

where $sg[\cdot]$ prevents gradients to propagate to its argument. Then IB-DRR jointly optimises the \mathcal{L}_{sim} and a cross-entropy loss \mathcal{L}_{ce} , i.e., $\mathcal{L} = \lambda \mathcal{L}_{sim} + \mathcal{L}_{ce}$ where λ is a scalar used to balance the two loss terms. Note that IB-DRR requires maintaining two memory set, one is codes of all old training samples, the other is a small subset of raw training exemplars of old classes. Like GEM [27, 4], the saved raw old exemplars are not directly used for minimizing the classification loss. They serve as regularizers to force the classifier to get back some lost information that is important for classification. More explicitly, it forces the classifier to treat reconstructed images and the real ones in a similar way. The general training process of IB-DRR is similar to DRR except that IB-DRR jointly optimises the two losses as follows: (1) the contrastive learning loss between raw exemplars and their reconstructed views; and (2) the cross-entropy loss computed using reconstructed images from new exemplars and old codes.

6. Experiments

Our experiments are designed for class-incremental classification and we consider the single-headed classifier

scheme where there is only a single unified output layer. In other words, a unified classifier learns novel classes and classifies all the classes that it has seen so far. We compare our approach with baseline methods and the state-of-the-art methods that use the replay strategy.

6.1. Experimental Setup

Following [13, 24, 23, 36], we evaluate our proposed strategies on the CIFAR-100 dataset [18], and we train three variants for ablation study: 1) DRR: discrete representation replay with saving latent representations (in the format of codes) of all seen exemplars; 2) IB-DRR: information back discrete representation replay with an extra requirement of saving a small set of raw original exemplars of previously seen classes; and 3) IB-DRR*: a variant of IB-DRR without saving raw data, so it can only utilize raw (original) exemplars of new coming classes for IB. More explicitly, IB-DRR* is IB-DRR without saving raw samples for old classes; and DRR is obtained by ablating Information Back mechanism from IB-DRR*.

Dataset. The CIFAR-100 dataset contains 100 classes and each class has 500 training samples and 100 test samples with the image size of $32 \times 32 \times 3$.

Architecture. For our hierarchical VQ-VAE, we use two codebooks (two level of codes) each with a size of 512 (number of codes) and set the embedding dimension to 64. We use a Bit-Swap with 8 hierarchical latent variables with a Markov chain structure for each level of codes. In our preliminary experiments, we find the Resnet-32 used in previous replay-based methods [24, 34, 36] causes underfitting of our DRR in some cases. We use a ResNet-18 [12] used in [23] as a classifier (see Appendix for further discussion).

Hyperparameters and Configuration. Following the *class incremental* settings introduced in [13, 24, 23], our VQ-VAE and Bit-Swap are pretrained given the half of the classes from CIFAR-100 at the initial training phase 0 and the rest of the classes are added gradually in the future phases. That is, the training has 1 initial phase for the classification of the first 50 classes and N incremental phases for the remaining 50 classes. We freeze the VQ-VAE but keep fine-tuning the Bits-Swap for future phases. Note that our classifier is trained with reconstructed images but tested with raw images. Our preliminary results show that training the classifier with the mixture of raw images of new classes and reconstructed images of old classes degrades the classification results. We evaluate the performance of our incremental learning method by setting N to 25, 10 and 5, meaning that 2, 5, and 10 classes are gradually added at each incremental phase, respectively. The VQ-VAE is trained by an Adam optimizer with a constant learning rate 0.0003 for 1400 epochs. We train or fine-tune the Bit-Swap by an Adam optimizer with a constant learning rate 0.002 until its loss no longer decreases. For DRR, IB-DRR and

Method	25 phases	10 phases	5 phases
<i>UB</i>	78.60	78.92	78.86
GFR [23]	54.01	60.14	60.18
Mnem [24]	60.96	60.78	60.76
DRR	64.37	64.04	64.00
IB-DRR*	62.90	64.85	65.36
IB-DRR	65.65	66.40	66.71

Table 1. Comparison in terms of the average overall accuracy (%) for varying N incremental phases on CIFAR-100.

IB-DRR*, we train the classifier with the image replay strategy as explained in Section 4. We use a stochastic gradient descent (SGD) optimizer with 0.9 momentum and 0.0005 weight decay parameters for the classifier. The initial learning rate is set to 0.1 and divided by 5 after 60, 120, and 160 epochs. A warmup scheduler [9] is applied for the first 5 epochs. The classifier is trained from scratch during a new training phase for 200 epochs. If IB is used, we set the hyperparameter $\lambda = 0.005$ for the contrastive learning loss \mathcal{L}_{sim} .

Memory Budget. For DRR and IB-DRR* we only save codes for seen classes. For IB-DRR, we follow [24, 34, 36] and save additionally 20 raw exemplars per class for seen classes.

Baselines. We compare our methods with the following 5 baseline methods: *LWF* [21], *GFR* [23], *iCaRL* [36], *LU-CIR* [13], *Mnem* [24] and a **Upper Bound (UB)**. The first two baselines do not require to save exemplars while others (including ours) need to save old exemplars. *iCaRL*, *LU-CIR* and *Mnem* save 20 samples per class while **UB** saves all (50,000) samples. We follow their original implementation settings (see Appendix for further information).

6.2. Evaluation Metrics

For each new training phase i , we evaluate the performance of the model on the test set from the new classes and old classes (henceforth, phase-wise accuracy). We use two evaluation metrics. Following the previous work [13, 23, 36], the first metric is the *average overall accuracy* that computes the mean of phase-wise accuracy over all training phases, namely from the initial phase to phase i . However, in the initial phase N_0 , there is no catastrophic forgetting; therefore, we exclude the test accuracy of the initial phase from the *average overall accuracy*. Let A_i be the performance of the model on the held-out test set of phase i , the average overall accuracy is then defined as: $\bar{A} = \frac{1}{N} \sum_{i=1}^N A_i$ where N is the total number of phases (in our experiments $N = 5, 10, \text{ and } 25$). The second metric is the *last phase accuracy* that reports the phase-wise accuracy of the last training phase only. Please note that, in the last phase, all classes in the dataset are encountered, thus the accuracy is given for all 100 classes.

Method	25 phases	10 phases	5 phases
<i>UB</i>	77.91	77.91	77.91
GFR [23]	40.39	51.29	53.34
Mnem [24]	50.78	51.53	54.32
DRR	60.87	60.87	60.87
IB-DRR*	60.50	62.29	62.83
IB-DRR	62.86	63.65	63.97

Table 2. Comparison in terms of last phase accuracy (%) (after models trained for 100 classes) on CIFAR-100 for varying N incremental phases.

6.3. Results and Analyses

In Figure 3, we compare our methods with the state-of-the-art replay-based methods that use image-level replay [24] and feature-level replay [23] as well as other baseline methods [36, 21, 13] in terms of *phase-wise accuracy*. It is obvious that methods using replay have better performance than LwF [21] which is one of the most competitive methods without replay, demonstrating the superiority of replay-based methods. IB-DRR outperforms others except for the upper bound (UB) in all experimental settings. Looking at the most challenging 25-phase setting, our models start with a lower phase-wise accuracy at the initial training phase (with 50 classes) due to the fact that reconstructed images are of low quality as compared to original images. However, despite the fact it is trained with reconstructed, low-quality images, our models can maintain the knowledge over old classes when learning new classes with improvement in the average accuracy by a margin of 6% and 4% by IB-DRR and DRR respectively (see Table 1). In addition, IB-DRR and DRR surpass other baselines by a margin of at least 9% and 3% according to the *last phase accuracy* (see Table 2). In 5-phase and 10-phase settings, IB-DRR* performs better than DRR while in 25-phase setting we see the opposite situation. Recall that, both of them only save codes for previously seen classes. Intuitively, the Information Back (IB) mechanism benefits more diverse raw samples from more new classes to regularize the classifier to get back some information that is useful for classification. From this point of view, in 5-phase or 10-phase IB-DRR* has a large enough number of new classes to utilize. As IB-DRR also saves old raw samples, it can perform well in all the settings we tested. Moreover, our methods behave similarly to the upper bound with a relatively ‘constant’ decrease in accuracy.

6.3.1 Comparison of Memory Cost

Table 3 provides the memory required for saving extra exemplars as well as assistant model weights on the CIFAR-100. Saving discrete latent representations in codes costs




Figure 3. Phase-wise accuracy (%) on Cifar100 based on 5-phase (10 class per phase), 10-phase (5 class per phase), and 25-phase (2 class per phase) settings, respectively.

Method	Exemplar cost (MB)	Assistant model cost (MB)
UB	148 (50,000 exemplars)	NA
GFR [23]	NA	4.5 (GAN)
Mnem [24] or LUCIR [13]	6.2 (2,000 exemplars)	NA
DRR or IB-DRR*	4.24 (50,000 exemplars)	5.81 (VQ-VAE) + 13.4 (Bit-Swap)
IB-DRR	6.2 + 4.24	5.81 (VQ-VAE) + 13.4 (Bit-Swap)

Table 3. Comparison of memory required for saving exemplars and assistant models, e.g., generators and autoencoders on CIFAR-100.

less memory while saving 25 times more exemplars as compared to other replay-based methods [24, 13]. Please note that we are not able to compare the cost of assistant models. Because many incremental learning methods (especially those that use knowledge distillation techniques) save a copy of old models. For example, Mnem [24] saves the model weights Θ_0 after the 1st phase and uses knowledge distillation loss to encourage Θ_{i-1} and Θ_i to maintain the same prediction ability on old classes. GFR [23], LUCIR [13] and Mnem [24] all need to save a copy for the latest model and train a new model with the help of the old model copy, therefore, each of them should be part of the total memory cost of assistant models.

6.4. Can DRR generalise to other datasets?

To validate that our VQ-VAE trained on CIFAR-100 has a good generalization ability when applied to a different dataset even with a higher resolution, we presented preliminary quantitative results on Subset-ImageNet that contained random 100 classes from ImageNet-1K. For our method, we used the pretrained VQ-VAE on CIFAR-100 as described previously in Section 6.1 to process images from ImageNet. We compared our DRR with Mnem [24] which is the SOTA method. Both approaches used images with the size of $224 \times 224 \times 3$ as inputs and Resnet-18 as the classifier. Mnem saved 20 raw samples per class (330 MB in total) while DRR saved codes of all samples. Results of the 5-phase setting showed DRR and Mnem achieved an average overall accuracy of 78.05% and 72.58% respectively. In ad-

dition, DRR resulted in a memory cost of 474 MB for saving 128,856 raw exemplars which originally cost 18,494 MB.

7. Conclusion

In this paper, we formalised incremental learning as a representation compression problem and proposed a novel approach to this problem. Our proposed approach, Discrete Representation Replay (DRR), performs a two-step compression using a Vector-Quantised Variational Autoencoder (VQ-VAE) and a Bits-back Asymmetric Numeral Systems with hierarchical latent variables (Bit-Swap). Our experimental results showed that DRR outperforms the state-of-the-art approaches on the CIFAR-100 dataset in terms of accuracy and memory size required for saving exemplars. In addition, we introduced an Information Back (IB) mechanism that utilized raw exemplars to regularize the training of the classifier and IB further boosted the performance of the DRR by saving a small set of raw exemplars of previously seen classes. Our preliminary results showed implications that our approach could be generalised to other datasets such as ImageNet. As future work, we will extend IB-DRR by introducing extra network components for contrastive learning to alleviate information loss caused by the VQ-VAE.

Acknowledgement

The work of Jian Jiang has been supported by the King’s China Scholarship Council (K-CSC) PhD Scholarship programme.

References

- [1] Rahaf Aljundi. Continual learning in neural networks. *arXiv preprint arXiv:1910.02718*, 2019. [3](#)
- [2] Ali Ayub and Alan R Wagner. Eec: Learning to encode and regenerate images for continual learning. *arXiv preprint arXiv:2101.04904*, 2021. [3](#)
- [3] Lucas Caccia, Eugene Belilovsky, Massimo Caccia, and Joelle Pineau. Online learned continual compression with adaptive quantization modules. In *International Conference on Machine Learning*, pages 1240–1250. PMLR, 2020. [3](#)
- [4] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018. [2](#), [3](#), [6](#)
- [5] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020. [6](#)
- [6] Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. PixelSNAIL: An Improved Autoregressive Generative Model. *35th International Conference on Machine Learning, ICML 2018*, 2:1364–1372, dec 2017. [4](#)
- [7] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A Downsampled Variant of ImageNet as an Alternative to the CIFAR datasets. *arXiv*, jul 2017. [2](#)
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [3](#)
- [9] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. [7](#)
- [10] Yanming Guo, Yu Liu, Ard Oerlemans, Songyang Lao, Song Wu, and Michael S Lew. Deep learning for visual understanding: A review. *Neurocomputing*, 187:27–48, 2016. [1](#)
- [11] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4918–4927, 2019. [5](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [6](#)
- [13] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. [1](#), [3](#), [6](#), [7](#), [8](#)
- [14] David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [1](#), [2](#)
- [15] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*, (MI):1–14, 2014. [3](#), [4](#)
- [16] Friso Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *International Conference on Machine Learning*, pages 3408–3417. PMLR, 2019. [4](#), [5](#)
- [17] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, 2017. [1](#), [2](#)
- [18] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. [2](#), [6](#)
- [19] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015. [1](#)
- [20] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. [2](#), [3](#)
- [21] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, jun 2018. [3](#), [7](#)
- [22] Kevin J Liang, Chunyuan Li, Guoyin Wang, and Lawrence Carin. Generative adversarial network training is a continual learning problem. *arXiv preprint arXiv:1811.11083*, 2018. [3](#)
- [23] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D Bagdanov, Shangling Jui, and Joost van de Weijer. Generative feature replay for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 226–227, 2020. [2](#), [3](#), [6](#), [7](#), [8](#)
- [24] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12245–12254, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [25] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In *Conference on Robot Learning*, pages 17–26. PMLR, 2017. [3](#)
- [26] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, volume 2017-Decem, pages 6468–6477, 2017. [1](#)
- [27] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, volume 2017-Decem, pages 6468–6477, 2017. [2](#), [3](#), [6](#)
- [28] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. [3](#)
- [29] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *arXiv preprint arXiv:1711.00937*, 2017. [2](#), [4](#), [5](#)
- [30] German I. Parisi, Jun Tani, Cornelius Weber, and Stefan Wermter. Lifelong learning of spatiotemporal representations with dual-memory recurrent self-organization. *Frontiers in Neurobotics*, 12(November), may 2018. [2](#)

- [31] Lorenzo Pellegrini, Gabriele Graffieti, Vincenzo Lomonaco, and Davide Maltoni. Latent Replay for Real-Time Continual Learning. dec 2019. 1, 3
- [32] Jason Ramapuram, Magda Gregorova, and Alexandros Kalousis. Lifelong generative modeling. *arXiv preprint arXiv:1705.09847*, 2017. 2, 3
- [33] Amal Rannen, Rahaf Aljundi, Matthew B. Blaschko, and Tinne Tuytelaars. Encoder based lifelong learning. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 2
- [34] Dushyant Rao, Francesco Visin, Andrei A Rusu, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Continual unsupervised representation learning. *arXiv preprint arXiv:1910.14481*, 2019. 6, 7
- [35] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *arXiv preprint arXiv:1906.00446*, 2019. 2, 4, 5
- [36] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 1, 2, 6, 7
- [37] Matthew Riemer, Tim Klinger, Djallel Bouneffouf, and Michele Franceschini. Scalable recollections for continual lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1352–1359, 2019. 3
- [38] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. 2
- [39] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P Lillicrap, and Greg Wayne. Experience replay for continual learning. *arXiv preprint arXiv:1811.11682*, 2018. 1, 2
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 2, 3
- [41] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, volume 2017-Decem, pages 2991–3000, 2017. 2, 3
- [42] James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. *arXiv preprint arXiv:1901.04866*, 2019. 4, 5
- [43] James Townsend, Thomas Bird, Julius Kunze, and David Barber. Hilloc: Lossless image compression with hierarchical latent variable models. *arXiv preprint arXiv:1912.09953*, 2019. 4
- [44] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016. 4
- [45] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3987–3995. JMLR. org, 2017. 2
- [46] Mengyao Zhai, Lei Chen, Frederick Tung, Jiawei He, Megha Nawhal, and Greg Mori. Lifelong gan: Continual learning for conditional image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2759–2768, 2019. 3