

Revisiting a kNN-based Image Classification System with High-capacity Storage

Kengo Nakata, Youyang Ng, Daisuke Miyashita, Asuka Maki,
Yu-Chieh Lin, and Jun Deguchi

Kioxia Corporation, Kawasaki, Japan
{kengo1.nakata, youyang.ng, daisuke1.miyashita, asuka.maki,
yuchieh.lin, jun.deguchi}@kioxia.com

Abstract. In existing image classification systems that use deep neural networks, the knowledge needed for image classification is implicitly stored in model parameters. If users want to update this knowledge, then they need to fine-tune the model parameters. Moreover, users cannot verify the validity of inference results or evaluate the contribution of knowledge to the results. In this paper, we investigate a system that stores knowledge for image classification, such as image feature maps, labels, and original images, not in model parameters but in external high-capacity storage. Our system refers to the storage like a database when classifying input images. To increase knowledge, our system updates the database instead of fine-tuning model parameters, which avoids catastrophic forgetting in incremental learning scenarios. We revisit a kNN (k-Nearest Neighbor) classifier and employ it in our system. By analyzing the neighborhood samples referred by the kNN algorithm, we can interpret how knowledge learned in the past is used for inference results. Our system achieves 79.8% top-1 accuracy on the ImageNet dataset without fine-tuning model parameters after pretraining, and 90.8% accuracy on the Split CIFAR-100 dataset in the task incremental learning setting.

Keywords: kNN classifier, continual learning, explainable AI

1 Introduction

Image classification systems using deep neural networks (DNN) have achieved superhuman recognition performance in computer vision tasks [32,30,42,22,15,13]. On the other hand, the knowledge for image classification is implicitly stored in model parameters and is not accessible to users. For example, users cannot retrieve knowledge related to only *cats* or *dogs* from the model parameters. Training datasets may contain a small amount of inappropriate data (e.g., incorrectly labeled or undesirably biased images) [36,1,43]. If the model parameters are optimized using such training datasets, then the image classification systems will implicitly contain false knowledge. In existing systems, however, users cannot selectively eliminate or correct only the false knowledge in model parameters.

To update knowledge (add, delete, or modify), users need to fine-tune the model parameters. Especially, in continual or incremental learning scenarios [31,23]

where new data is continually added to a training dataset, the cost for fine-tuning is incurred repeatedly as new data is added. If model parameters are fine-tuned with only newly added data in an attempt to reduce the cost, the model will acquire new knowledge for the added data while forgetting the knowledge learned in the past, which is often called catastrophic forgetting [18,31]. Although various methods have been proposed to mitigate the impact of catastrophic forgetting [39,35,7,27,40,4], as long as modifying model parameters, forgetting some knowledge is inevitable. Refs. [38,25] propose a zero-shot classifier that does not require any fine-tuning by sufficiently pretraining model parameters on large-scale datasets which contain a variety of images. Such classifier does not face the catastrophic forgetting even without fine-tuning, but it cannot additionally acquire user’s desired knowledge with the user’s own datasets.

Moreover, in existing systems, users cannot verify the validity of inference results or evaluate the contribution of knowledge to the results. For instance, users do not know how models utilize knowledge learned in the past to classify input images. Although we can identify unnecessary parameters through accuracy evaluation or some other parameter analyses [56,53,48], we cannot interpret why the parameters, that is, knowledge, are unnecessary (cf. explainable AI).

We investigate a system that stores image feature maps, labels, and original images of entire training datasets in external high-capacity storage as knowledge for image classification. Our image classification system refers to the external storage like a database when classifying input images. To increase knowledge, our system adds image feature maps and labels for new data to the database instead of fine-tuning model parameters using those new data. We employ a kNN (k-Nearest Neighbor) classifier [12], which is one of the most classic classification algorithms. Recently, many papers have employed kNN for the purpose of evaluating their proposed representation learning algorithms [5,55,33].

In this paper, we shed new light on the potentials of kNN classification system with high-capacity storage. A concurrent work [34] demonstrates that kNN retrieval improves long-tail recognition. Not only this, we empirically show that million scale kNN retrieval is affordable for practical image classification tasks, and that our system with kNN avoids catastrophic forgetting in continual learning scenarios and achieves better accuracy than conventional methods. Furthermore, by reviewing neighborhood samples referred by the kNN algorithm, we can verify the validity of inference results, such as whether those referred samples contain incorrectly labeled data. If those referred samples contain incorrectly labeled data, our system can correct the false knowledge and improve the accuracy by eliminating only the incorrectly labeled data from the database.

The main contributions of this paper are as follows:

- We investigate a large-scale kNN system that stores knowledge for image classification in high-capacity storage and refers to the storage when classifying images, and empirically demonstrate its effectiveness and applicability on various image classification datasets and in continual learning scenarios.
- We also show that a large-scale kNN system has a capability of verifying the validity of inference results and selectively correcting only false knowledge.

2 Related Work

2.1 Data-driven Image Classification

DNNs, hundreds of millions to billions parameters [15,19,13,46] of which are optimized by supervised learning with labeled images, have achieved state-of-the-art results in visual understanding and image classification tasks. In recent years, unsupervised visual clustering, which does not rely on a labeled dataset for learning the feature representation and classifier, has been developed [50,9,21,10,47,20,17]. Unsupervised pretraining strategies on large-scale datasets that include unlabeled or noisily labeled images have demonstrated the potential of large-scale open-domain knowledge for improving the accuracy of closed-domain image classification tasks [15,25,38]. In addition, although the applicable dataset scale is still limited, Ref. [28] presented an architecture that utilizes an entire dataset (e.g., all training images) to classify an input image during the inference process. Data has never been more important in the quest to further improve the usability of image classification. Inspired by the data-driven approaches in classification strategies, our work focuses on the application of knowledge retrieval to image classification. We present a method that utilizes the potential of both the trained parametric representation model and the available datasets during the inference process.

2.2 Knowledge Retrieval

Knowledge retrieval has seen substantial advancement in recent years, particularly in DNN-based natural language processing (NLP). DPR [26] applied a dense representation to passage retrieval in open-domain question answering tasks and achieved better retrieval performance than traditional sparse vector space models, such as TF-IDF and BM25. KEAR [51] brought external knowledge into the predicting process of Transformer [45] to reach human parity in a challenging commonsense task [44]. RETRO [3] introduced a frozen kNN retriever into the Transformer architecture in the form of chunked cross-attention to enhance the performance of auto-regressive language models. External world knowledge has been retrieved to assist in solving various NLP tasks. Our work looks to extend the adoption of knowledge retrieval beyond the modality of NLP. We introduce an image classification architecture based on knowledge retrieval, which is a data-driven paradigm that is centralized on available large-scale data resources and is supported by a trained representation model and a kNN classifier.

2.3 Continual Learning

When new data or tasks are added continually after a training process, fine-tuning is used to update the previously trained models with new knowledge. This continual learning, also known as incremental learning, online learning, or lifelong learning [31,23], is similar to natural human learning processes and is a key challenge for achieving artificial general intelligence. However, the simple

procedure of iterative fine-tuning using only new knowledge suffers from catastrophic forgetting [18,31]. To mitigate the impact of catastrophic forgetting, GEM [35] introduced episodic memory to store a subset of the previously learned data, enabling an external memory-driven continual learning strategy. ER [40], inspired by the human brain’s ability to replay past experiences, introduced the experience replay of memory. Episodic memory and experience replay have seen great developments in recent years [7,4,6]. iCaRL [39] evaluated the effect of catastrophic forgetting under a class incremental learning setting, where new classes are incrementally added to classifiers, and proposed a rehearsal approach with exemplar images of old classes being stored in memory. Inspired by the aforementioned works in adopting memory and the replay mechanism, we incorporate direct knowledge retrieval to solve the catastrophic forgetting problem in continual learning scenarios. Instead of a model-based approach, we introduce a data-based approach, leveraging available datasets as knowledge sources to adapt to the incrementation of tasks and classes.

2.4 Explainable AI

DNN models have been a technological breakthrough for various computer vision tasks. However, the explainability of DNN models remains a challenge and has led to slower-than-expected deployments in critical infrastructures. Attempts [54,41,8] have been made to interpret image classification models. Grad-CAM [41] proposed a gradient-based localization technique for rendering attention maps on the input images. Ref. [8] went beyond attention visualization by adopting gradients and the propagation of relevancy scores. These analysis methods can visualize the areas of the input images that DNN models focus on during the inference process, but they cannot analyze how the DNN models use knowledge acquired in training when classifying the input image. For example, if a training dataset contains false data (e.g., incorrectly labeled images), then these methods will not be able to selectively retrieve knowledge related to the false data from the trained models nor evaluate the impact of the false knowledge on the classification results. Our data-driven image classification architecture visualizes how knowledge is used to classify input images and enables selective modifications to specific knowledge without fine-tuning.

3 Approach

Fig. 1 provides an overview of our image classification system. Our system has three phases: pretraining, knowledge storing, and inference. The first step is to pretrain an image encoder model on a large-scale dataset containing a variety of images. This pretraining can use unlabeled or noisily labeled datasets. In the knowledge storing phase, the pretrained image encoder model extracts feature maps from support sets (e.g., a training dataset of a user’s desired downstream task). Extracted feature maps are paired with their corresponding labels and registered in the external storage like a database. In the inference phase, the

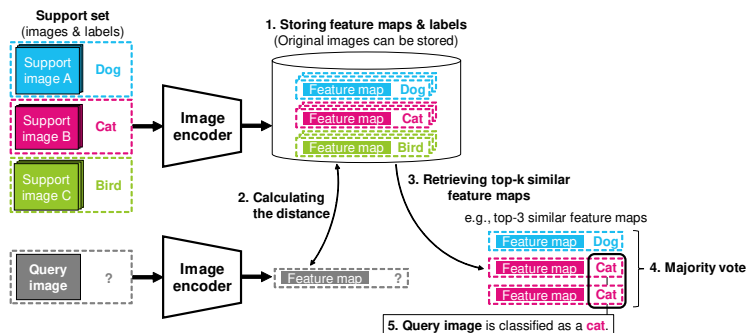


Fig. 1. Overview of our image classification system. Our system stores feature maps extracted from support images with the corresponding labels to the external storage. When classifying a query image, our system retrieves feature maps similar to the query one from the storage by calculating the distance based on cosine similarity. The query image is classified by majority vote on the labels of the top-k similar feature maps.

pretrained image encoder model extracts a feature map of a query image. Our system retrieves feature maps similar to the query one from the storage by calculating the distance based on cosine similarity. Then, the query image is classified by majority vote on the paired labels of the top-k similar feature maps retrieved. The following sections describe each phase in detail.

3.1 Pretraining

We pretrain an image encoder model to learn feature representations. We assume that this process is conducted on a large-scale computer system, such as a data center (not on the user side) [2]. Similarity-based retrieval in the latent space is performed in the following inference phase (see Section 3.3 for details). Therefore, transformations by the pretrained image encoder models should be such that semantically similar images are mapped to the neighborhood in the latent space. In this section, we describe suitability of pretraining methods that can be employed in our system.

Supervised learning with specific datasets. Models trained on a specific dataset by supervised learning can be used for our system. An example is a ResNet-50 model [22] trained on the ImageNet-1k classification task [14]. If input images are classified as the first class (e.g., tench) by the trained model, those input images are mapped to the feature maps so that they have a high similarity to the vector of weight parameters corresponding to the first class in the last fully connected layer. Therefore, we can employ the trained model as an image encoder by excluding the last fully connected layer. However, the trained model can overfit the classification setting of the specific dataset (e.g., the 1000 classes of ImageNet-1k).

Self-supervised contrastive learning. We can also employ self-supervised contrastive learning methods (e.g., SimCLR [9] and MoCo [21,10]), which learn image feature representations based on the similarity or dissimilarity of images between two or more views. Among the contrastive learning methods, CLIP [38] and ALIGN [25] jointly learn image and text feature representations of image and text encoders using a large number of image and text pairs collected from the internet. These methods do not require hand-crafted labeling or annotation, and the number of semantic labels is much larger than the number of labels on commonly used image datasets, such as the limited 1000 labels in ImageNet-1k.

Other self-supervised learning. Masked Auto Encoder [20] learns image feature representations through a task of reconstructing the original images from masked images. By fine-tuning on the labeled dataset, the encoders pretrained with Masked Auto Encoder have achieved high accuracy in the ImageNet-1k classification task. Because this pretraining method is label-independent, the pretrained model seems unlikely to overfit a specific classification setting (e.g., the 1000 classes of ImageNet-1k). However, the objective function of this pretraining method does not explicitly make the encoder model to map semantically similar images to the neighborhood in the latent space, so it is not suitable for similarity-based retrieval in our system.

Selecting pretraining methods. We conducted a preliminary experiment using three pretrained image encoder models: (1) a ResNet-50 model trained on the ImageNet-1k dataset with the labels by supervised learning, (2) a Vision Transformer [15] Base model with input patch sizes of 16×16 (ViT-B/16) trained by CLIP on 400 million image and text pairs collected from the internet, (3) a ViT-B/16 model trained by Masked Auto Encoder (MAE) on the ImageNet-1k dataset without the labels. We employed the three pretrained image encoders in our system, and evaluated the test accuracy on various image datasets (CIFAR-10, CIFAR-100 [29], STL-10 [11], and ImageNet-1k).

We summarize the test accuracy in Table 1. As discussed above, (1) the supervised learned model on ImageNet-1k achieves the best accuracy on ImageNet-1k, but not on the other datasets. This supervised learned model does not generalize well to unseen datasets. (3) the model pretrained by MAE shows poor performance since the objective of this pretraining method is not compatible with similarity-based retrieval in our system. On the other hand, (2) the model pretrained by CLIP achieves good accuracy on various datasets, indicating that the model is well generalized. Based on the result of this preliminary experiment, we employed image encoder models pretrained by CLIP in the experiments of Section 4. The exploration of better pretraining methods is our future work.







3.2 Knowledge Storing

To acquire knowledge, a pretrained image encoder model extracts feature maps from a support set, such as a training dataset of a user’s desired downstream task. Given a support set of n -labeled images $\{\mathbf{x}_{s,1}, \dots, \mathbf{x}_{s,n}\}$, the corresponding labels $\{y_1, \dots, y_n\}$, and the pretrained image encoder model $f(\cdot)$, the i -th support

Table 1. Accuracy comparison using different pretrained image encoder models.

	ImageNet-1k	CIFAR-10	CIFAR-100	STL-10
(1) Supervised learned on ImageNet-1k	74.9	85.9	64.7	96.7
(2) CLIP [38]	74.0	94.4	74.3	98.9
(3) MAE [20]	26.7	51.8	23.2	66.6

Correct classification

	Query	Top-5 similar support samples retrieved					Classified as tiger cat (correct)
Original images							
Distance	-	0.031	0.052	0.054	0.059	0.060	
Label	tiger cat (ground truth)	tiger cat	tiger cat	tiger cat	tiger cat	tiger cat	

Misclassification due to *incorrectly labeled samples







	Query	Top-5 similar support samples retrieved					Classified as tiger cat (wrong)
Original images							
Distance	-	0.062	0.068	0.076	0.080	0.080	
Label	tiger (ground truth)	tiger	tiger cat*	tiger cat*	tiger	tiger cat*	

Fig. 2. Examples of query and the top-5 similar support samples retrieved. In these examples, validation and training data of ImageNet-1k are used for the query and the support set, respectively (the file names are also given in the examples). The top figure shows an example of **correct classification**, and the bottom figure shows one of **misclassification** due to *incorrectly labeled support samples. By reviewing the original images and labels, we can verify the validity of the classification results.

image is mapped to the d -dimensional latent space, and the extracted feature map $\mathbf{z}_{s,i}$ is obtained by $\mathbf{z}_{s,i} = f(\mathbf{x}_{s,i}), \mathbf{z} \in \mathbb{R}^d$.

Unlike fine-tuning, our system does not require iterative forward, backward and parameter-update operations, but it just requires one forward operation of each image in the support set for the feature extraction. Thus, the cost and effort are less than those of fine-tuning. Moreover, our system can avoid catastrophic forgetting without fine-tuning, even when the knowledge is iteratively updated in continual learning scenarios.

An extracted feature map is paired with the corresponding label such as $(\mathbf{z}_{s,i}, y_i)$, and registered in the external storage like a database. If the original images are also registered to the database, then they can be used to verify the validity of classification results (see Fig. 2 and Section 3.3 for details).

3.3 Inference

Our system classifies query images (i.e., inference) by referring to image feature maps and labels registered in the database. The pretrained encoder first extracts

a feature map \mathbf{z}_q of a query image \mathbf{x}_q as $\mathbf{z}_q = f(\mathbf{x}_q)$. Our system retrieves support feature maps that are similar to query one from the database by calculating the distance (D) based on cosine similarity between query and support feature maps as follows,

$$D(\mathbf{z}_q, \mathbf{z}_{s,i}) = 1 - \frac{\mathbf{z}_q \cdot \mathbf{z}_{s,i}}{\|\mathbf{z}_q\| \|\mathbf{z}_{s,i}\|}. \quad (1)$$

With $\arg \min_{i \in n} D(\mathbf{z}_q, \mathbf{z}_{s,i})$, we can retrieve the index of the nearest, that is, the most similar support feature map to the query one. Our system retrieves the index of top-k minimum distance, and the query image is classified based on the majority vote of the paired labels in order to mitigate the effect of outliers.

Fig. 2 shows examples of query and the retrieved samples from the support set. The original images in the support set are not required for classifying query images. However, the original images registered in the database can be used for evidence or for verification of the inference results. For instance, as shown in Fig. 2, not only the labels and distance values of retrieved feature maps, but also the original images, can be listed as inference logs. By reviewing the logs, we can verify that the retrieved samples contain inappropriate data (e.g., incorrectly labeled images). The bottom figure in Fig. 2 shows an example of the misclassification of a query image due to incorrectly labeled images in the support set (using a training dataset of ImageNet-1k). If incorrectly labeled images are obtained in the logs, then we can correct false knowledge in our system by fixing the incorrect labels or eliminating such samples from the database.

The above characteristics are not present in existing image classification systems. To correct false knowledge, existing systems need to fine-tune classifier models again after fixing incorrect labels in the support set. On the other hand, our system does not require fine-tuning, thus eliminating the cost and effort for fine-tuning, and can contribute to improving the explainability of AI by visualizing and reviewing referred images and labels when classifying input images.

4 Experiments

We implemented our image classification system using the PyTorch library [37]. The setup and results for each experiment are described in detail below.

4.1 Basic Performance Evaluation

Experimental setup. We evaluated the basic performance of our system using the CIFAR-10, CIFAR-100, STL-10, and ImageNet-1k datasets. We employed image encoder models pretrained by CLIP. We used ResNet-50 and 101 as CNN models and ViT-B/32, B/16, and L/14 as Vision Transformer models¹ for the image encoders. These pretrained models extracted feature maps from training images, and those extracted feature maps and labels were registered in the

¹ ViT-B/32, B/16, and L/14 correspond to Vision Transformer Base, Base, and Large model, with input patch sizes of 32×32 , 16×16 , and 14×14 , respectively.

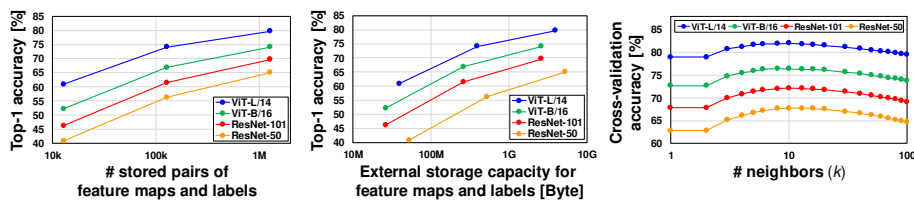


Fig. 3. **Left/center:** the relation between the top-1 accuracy on the ImageNet-1k dataset and the number of stored pairs of feature maps and labels/the external storage capacity for storing the pairs. **Right:** the relation between cross-validation accuracy and the number of neighbors (k) for kNN. For cross-validation, the 1.28M training images in ImageNet-1k are randomly split 9:1 into support and query sets.

database as support sets. With the support sets, we evaluated the test accuracy of our system using test datasets as query sets. As a baseline, we employed the test accuracy of the zero-shot CLIP classifier [38]. This classifier can be applied to those classification tasks without any fine-tuning like our system.

In addition to the accuracy evaluation, we measured the processing time for inference of the ImageNet-1k images. We used a NVIDIA A100 GPU and measured the processing time of our system and the zero-shot CLIP classifier as a baseline. In this measurement, both the parameters of the image encoder model and the feature maps of the support set were loaded into the GPU memory from the external storage prior to the inference. We then measured the processing time for transferring query images to the GPU, encoding the query images, and calculating the distance between the query and support feature maps. The batch size was set to 1000, and the average processing time per image was calculated.

Experimental results. The left graph in Fig. 3 shows the relation between the top-1 accuracy on the ImageNet-1k dataset and the number of stored pairs of feature maps and labels. Here, we randomly sample images from the training dataset, adjusting the number of the stored pairs to 12.8k, 128k, and 1.28M. We set the number of neighbors (k) for kNN to 10. The more stored pairs that there are in the database, the better the accuracy will be that our system can achieve. Moreover, the larger the model size, the better the performance of the feature extractor and also the accuracy.

The center graph in Fig. 3 shows the relation between the top-1 accuracy and the external storage capacity for storing feature maps and labels. In this experiment, the maximum storage capacity is 5.3 GB to store the feature maps and labels for 1.28M samples of ImageNet-1k. As shown in the center graph, the larger the capacity, the better the accuracy.

In the right graph of Fig. 3, we evaluated the relation between the accuracy and the number of neighbors of k by cross-validation. For cross-validation, the 1.28M training images are randomly split 9:1 into support and query sets. As

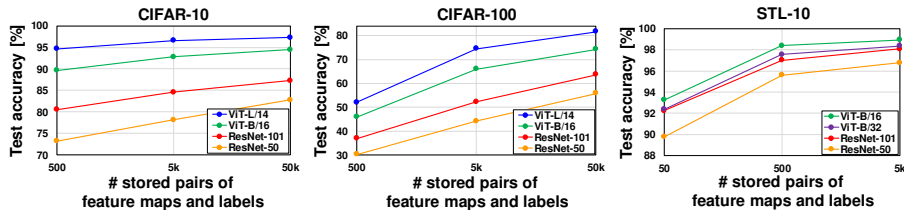


Fig. 4. Relation between the test accuracy and the number of stored pairs of feature maps and labels for various image datasets.

Table 2. Test accuracy evaluations for various image datasets. Baseline indicates the accuracy of the zero-shot CLIP classifier.

	CIFAR-10					CIFAR-100					STL-10					ImageNet-1k				
	ResNet		ViT			ResNet		ViT			ResNet		ViT			ResNet		ViT		
	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14
Baseline	75.6	81.0	91.3	91.6	96.2	41.6	49.0	65.1	68.7	77.9	94.3	96.7	97.2	98.2	99.3	59.6	62.2	63.2	68.6	75.3
Ours	82.8	87.3	92.7	94.4	97.3	55.7	63.6	71.5	74.3	81.7	96.8	98.1	98.4	98.9	99.6	65.0	69.7	67.0	74.0	79.7
Δ	+7.2	+6.3	+1.4	+2.8	+1.1	+14.1	+14.6	+6.4	+5.6	+3.8	+2.5	+1.4	+1.2	+0.7	+0.3	+5.4	+7.5	+3.8	+5.4	+4.4

shown in the right graph, the best accuracy is achieved when k is set to around 10. In the following experiments, k is set to 10 unless otherwise noted.

Fig. 4 shows the relation between the test accuracy and the number of stored pairs of feature maps and labels on the other datasets. For all datasets, our system improves the test accuracy as the number of stored pairs increases and as the model size becomes larger.

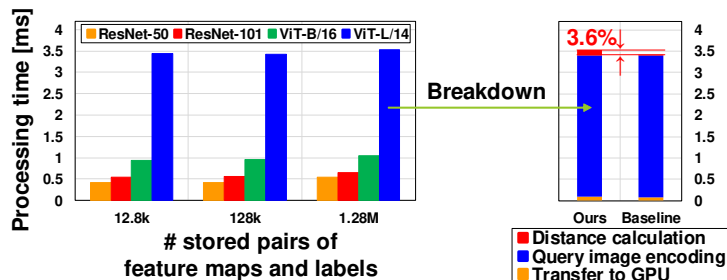
Table 2 summarizes the test accuracy of our system. Our system can improve accuracy from the baseline (pretrained but not fine-tuned models) on the various image datasets by storing image feature maps and labels of training datasets.

Based on the experimental results in Figs. 3 and 4, we utilize a data augmentation technique to increase the number of stored pairs of feature maps and labels. Data augmentation is typically used to prevent models from overfitting during training and to improve the generalization performance. In particular, we use only horizontal flipping to augment support images ($2\times$), and set the number of neighbors of k to 20. Table 3 summarizes the accuracy with and without data augmentation. As shown in the results for CIFAR-100 and ImageNet-1k in Table 3, our system can improve accuracy by doubling the number of pairs of stored feature maps and labels. In the results for CIFAR-10 and STL-10, data augmentation did not always improve accuracy, but the accuracy was already high enough without data augmentation (the accuracy is higher than 94%).

Fig. 5 shows the processing time for inference of the ImageNet-1k images that was measured on the NVIDIA A100 GPU. The left graph shows the processing time of our system when storing 12.8k, 128k, and 1.28M pairs of feature maps and labels with different image encoder models. In this experiment, the maximum required memory capacity is 5.3 GB when using 1.28M pairs of feature maps and labels, and all the feature maps can be fully loaded into GPU memory. By exploiting massive parallel computations of GPU, we execute a linear search

Table 3. Accuracy comparison with and without data augmentation (DA) on the support sets.

	CIFAR-10					CIFAR-100					STL-10					ImageNet-1k				
	ResNet		ViT			ResNet		ViT			ResNet		ViT			ResNet		ViT		
	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14	50	101	B/32	B/16	L/14
w/o DA	82.8	87.3	92.7	94.4	97.3	55.7	63.6	71.5	74.3	81.7	96.8	98.1	98.4	98.9	99.6	65.0	69.7	69.0	74.0	79.7
w/ DA	83.6	87.8	93.0	94.3	97.3	56.8	63.9	71.8	74.7	81.7	97.0	97.9	98.3	99.1	99.6	65.4	70.2	69.2	74.3	79.8
Δ	+0.8	+0.5	+0.3	-0.1	0.0	+1.1	+0.3	+0.3	+0.4	0.0	+0.2	-0.2	-0.1	+0.2	0.0	+0.4	+0.5	+0.2	+0.3	+0.1

**Fig. 5.** Processing time for inference of the ImageNet-1k images as measured on an NVIDIA A100 GPU (left) and the breakdown comparison for ViT-L/14 with 1.28M stored pairs of feature maps and labels (right). The baseline is the processing time of the zero-shot CLIP classifier.

based on cosine distance between a query and all the feature maps calculated in Eq. (1). As shown in the left graph, as the number of stored pairs increases, the processing time becomes slightly longer (e.g., 3% in the case of ViT-L/14). The right graph in Fig. 5 is the breakdown of the processing time for ViT-L/14 with 1.28M stored pairs. For comparison, the right graph also includes the breakdown of the zero-shot CLIP classifier as a baseline. The processing time for query image encoding is dominant, but that for the distance calculation is short enough (about 0.1 ms). The overhead from the baseline for retrieving similar feature maps from 1.28M stored ones is as small as 3.6%.

4.2 Continual Learning

Experimental setup. We evaluated the applicability of our system to continual learning, especially in task incremental and class incremental settings.

Task incremental learning is a setting where the number of tasks increases step by step. We used a Split CIFAR-100 dataset [7,52] in which the CIFAR-100 dataset is split into 20 disjoint subsets. Each subset consists of five randomly sampled classes without duplication from a total of 100 classes. Additionally, we applied ImageNet-100 [24,49,16] and ImageNet-1k datasets to the incremental learning setting. ImageNet-100 consists of 100 classes randomly sampled from 1000 classes in the ImageNet-1k dataset. We split the ImageNet-100 and ImageNet-1k datasets into 10 disjoint subsets, each containing 10 and 100 classes,

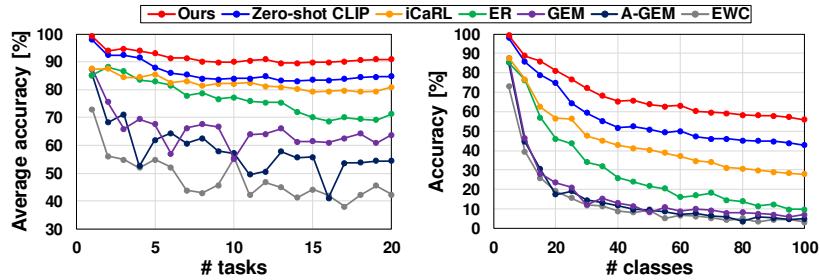


Fig. 6. Accuracy curves on the Split CIFAR-100 dataset in the incremental learning settings (left: task incremental, right: class incremental).

respectively. We did not balance the number of samples in each subset. We used each subset in Split CIFAR-100, Split ImageNet-100, and Split ImageNet-1k as one task for a 5-class, 10-class, and 100-class classification, respectively.

Class incremental learning is a setting in which the number of classes is continually added. In our experiment, we increased the number of classes by 5, 10, and 100 by sequentially adding the classes in each subset of Split CIFAR-100, Split ImageNet-100, and Split ImageNet-1k, respectively.

For each setting, we used ResNet-50, ViT-B/32, and ViT-L/14 models for the image encoder, and employed the pretrained models by CLIP as Section 4.1. We evaluated the accuracy of our system, the zero-shot CLIP classifier, and the conventional methods proposed for the incremental learning settings (iCaRL, ER, GEM, A-GEM, and EWC [39,40,35,7,27]). We applied the same pretrained models to the initial values for all the methods. In iCaRL, ER, GEM, and A-GEM, a part of training data learned in the past is stored in memory, and when new data is added, the model parameters are updated along with the stored data. EWC does not store any training data learned in the past. The detailed conditions such as the hyperparameter settings are described in the supplementary materials.

Experimental results. The left graph in Fig. 6 shows the average accuracy on the Split CIFAR-100 dataset for task incremental learning. The average accuracy is calculated by averaging the test accuracy evaluated in each task. In the conventional methods (except the zero-shot CLIP classifier), as the number of tasks increases, the average accuracy over tasks gradually drops because the models forget the knowledge for older tasks. On the other hand, our system stores extracted image feature maps and labels, so it does not require fine-tuning to acquire new knowledge. Therefore, our system avoids catastrophic forgetting and keeps the average accuracy over 90%, even when the number of tasks exceeds 15. The right graph in Fig. 6 shows the accuracy for class incremental learning. The accuracy curve of our system is higher than those of the conventional methods, even as the number of classes increases.

Table 4. Accuracy evaluations in task and class incremental learning settings.

ResNet-50 on Split CIFAR-100			ViT-B/32 on Split ImageNet-100		
Method	task	class	Method	task	class
EWC [27]	42.2	13.5	Zero-shot CLIP [38]	92.0	80.2
A-GEM [7]	54.5	15.8	Ours	94.3	85.1
GEM [35]	63.7	17.4			
ER [40]	71.3	29.4	ViT-L/14 on Split ImageNet-1k		
iCaRL [39]	80.9	43.9	Method	task	class
Zero-shot CLIP [38]	84.8	56.6	Zero-shot CLIP [38]	93.0	82.3
Ours	90.8	67.9	Ours	94.2	85.5

Table 4 summarizes the average accuracy on the Split CIFAR-100, Split ImageNet-100, and Split ImageNet-1k datasets. The accuracy of task incremental learning is the average of the accuracy evaluated on each task after all tasks have been added. The accuracy of class incremental learning is the average of the each accuracy from the first to the last class setting, which is introduced in Ref. [39] (e.g., the average of the accuracies when the number of classes is from 5 to 100 in the right figure of Fig. 6). Our system achieves better accuracy than the conventional methods, which indicates that our system is effective in incremental learning settings.

4.3 Correcting false knowledge

Experimental setup. Datasets used as support sets can contain incorrectly labeled images. In this section, we observe the impact of incorrectly labeled images on classification results and accuracy, and evaluate the effect of eliminating them. According to Refs. [36,1], the ImageNet-1k dataset contains a small number of incorrectly labeled images (at least 6%). Ref. [1] reassesses the label of each image in ImageNet-1k and releases an ImageNet-ReaL dataset that eliminates the incorrectly labeled images from the original ImageNet-1k dataset. We compared the accuracy before and after eliminating the incorrectly labeled images from the training dataset, namely, the support set.

Experimental results. Fig. 7 shows an example of samples retrieved with the ViT-L/14 model before and after eliminating incorrectly labeled images from the support set. Before the elimination, a query image is misclassified as a ruffed grouse (correctly as a partridge) due to the retrieval of feature maps of incorrectly labeled images. After the elimination, the incorrectly labeled samples are no longer included in the top-10 retrieved samples, and the query image is correctly classified as a partridge. These verification results of Fig. 7 can be visualized with the support images (like Fig. 2 in Section 3.2). In this experiment, when all the support images in the ImageNet-1k dataset are stored for the verification, the required storage capacity is 152GB.

Table 5 shows the accuracy on the ImageNet-ReaL dataset before and after the elimination. Our system can correct false knowledge and improve accuracy by simply eliminating feature maps of incorrectly labeled images. Unlike exist-

Before			After		
Query image	Label (ground truth)	File name	Query image	Label (ground truth)	File name
	partridge	ILSVRC2012_val_00019750.JPEG		partridge	ILSVRC2012_val_00019750.JPEG
Retrieved support samples			Retrieved support samples		
Distance	Label	File name	Distance	Label	File name
0.000	ruffed grouse*	n01797886_6586.JPEG*	0.062	partridge	n01807496_4822.JPEG
0.062	ruffed grouse*	n01797886_15633.JPEG*	0.064	partridge	n01807496_5569.JPEG
0.062	partridge	n01807496_4822.JPEG	0.066	partridge	n01807496_18631.JPEG
0.064	ruffed grouse*	n01797886_6603.JPEG*	0.066	partridge	n01807496_28384.JPEG
0.064	partridge	n01807496_5569.JPEG	0.066	partridge	n01807496_15316.JPEG
0.066	partridge	n01807496_18631.JPEG	0.069	partridge	n01807496_2534.JPEG
0.066	partridge	n01807496_28384.JPEG	0.07	partridge	n01807496_6715.JPEG
0.066	partridge	n01807496_15316.JPEG	0.074	partridge	n01807496_9868.JPEG
0.068	ruffed grouse*	n01797886_11301.JPEG*	0.076	partridge	n01807496_10464.JPEG
0.069	ruffed grouse*	n01797886_2986.JPEG*	0.076	partridge	n01807496_18094.JPEG
Classified as a ruffed grouse (wrong)			Classified as a partridge (correct)		

Fig. 7. Examples of query and top-10 similar support samples retrieved from ImageNet-1k before and after eliminating the incorrectly labeled samples from the support set. Before the elimination, the query image is **misclassified as a ruffed grouse** due to *incorrectly labeled samples. After the elimination, the query image is **correctly classified as a partridge** in all samples.

Table 5. **Left:** validation accuracy on the ImageNet-Real dataset before and after eliminating incorrectly labeled images from the support sets. **Right:** the details for the support and query sets.

	ResNet		ViT		
	50	101	B/32	B/16	L/14
Before	71.2	75.7	75.0	79.3	83.9
After	71.4	76.1	75.4	79.6	84.0
Δ	+0.2	+0.4	+0.4	+0.3	+0.1

	Support set	Query set
Before	1,281,167 images in ImageNet-1k	46,837 images in ImageNet-Real
After	1,148,659 images in ImageNet-Real	46,837 images in ImageNet-Real

ing image classification systems, our system does not need to fine-tune model parameters again on the modified dataset.

5 Conclusions

In this paper, we investigated a system that stores knowledge for image classification not in model parameters but in external high-capacity storage and refers to this storage like a database when classifying input images. Our system can increase knowledge by adding image feature maps and their corresponding labels to the database. Without fine-tuning the model parameters, our system avoids catastrophic forgetting in continual learning scenarios and achieves better accuracy than the conventional methods. By reviewing the neighborhood samples retrieved by the kNN classifier, we can verify the classification results. If incorrectly labeled samples are included in the classification results, then our system can correct the false knowledge and improve accuracy without fine-tuning by simply eliminating the incorrectly labeled samples from the database. In future work, we will work on scaling up our system with high-capacity storage by utilizing a large-scale fast ANN (approximate nearest neighbor) search.

References

1. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., van den Oord, A.: Are we done with ImageNet? arXiv [abs/2006.07159](#) (2020)
2. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N.S., Chen, A.S., Creel, K., Davis, J.Q., Demszky, D., Donahue, C., Doumbouya, M., Durmus, E., Ermon, S., Etchemendy, J., Ethayarajh, K., Fei-Fei, L., Finn, C., Gale, T., Gillespie, L., Goel, K., Goodman, N.D., Grossman, S., Guha, N., Hashimoto, T., Henderson, P., Hewitt, J., Ho, D.E., Hong, J., Hsu, K., Huang, J., Icard, T., Jain, S., Jurafsky, D., Kalluri, P., Karamcheti, S., Keeling, G., Khani, F., Khattab, O., Koh, P.W., Krass, M.S., Krishna, R., Kuditipudi, R., et al.: On the Opportunities and Risks of Foundation Models. arXiv [abs/2108.07258](#) (2021)
3. Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., van den Driessche, G., Lespiau, J., Damoc, B., Clark, A., de Las Casas, D., Guy, A., Menick, J., Ring, R., Hennigan, T., Huang, S., Maggiore, L., Jones, C., Cassirer, A., Brock, A., Paganini, M., Irving, G., Vinyals, O., Osindero, S., Simonyan, K., Rae, J.W., Elsen, E., Sifre, L.: Improving language models by retrieving from trillions of tokens. arXiv [abs/2112.04426](#) (2021)
4. Buzzega, P., Boschini, M., Porrello, A., Calderara, S.: Rethinking Experience Replay: a Bag of Tricks for Continual Learning. In: ICPR (2021)
5. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging Properties in Self-Supervised Vision Transformers. arXiv [abs/2104.14294](#) (2021)
6. Chaudhry, A., Gordo, A., Dokania, P.K., Torr, P.H.S., Lopez-Paz, D.: Using hindsight to Anchor Past Knowledge in Continual Learning. In: AAAI (2021)
7. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient Lifelong Learning with A-GEM. In: ICLR (2019)
8. Chefer, H., Gur, S., Wolf, L.: Transformer Interpretability Beyond Attention Visualization. In: CVPR (2021)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A Simple Framework for Contrastive Learning of Visual Representations. arXiv [abs/2002.05709](#) (2020)
10. Chen, X., Fan, H., Girshick, R.B., He, K.: Improved Baselines with Momentum Contrastive Learning. arXiv [abs/2003.04297](#) (2020)
11. Coates, A., Ng, A., Lee, H.: An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In: AISTATS (2011)
12. Cunningham, P., Delany, S.J.: k-Nearest Neighbour Classifiers: 2nd Edition (with Python examples). arXiv [abs/2004.04523](#) (2020)
13. Dai, Z., Liu, H., Le, Q.V., Tan, M.: CoAtNet: Marrying Convolution and Attention for All Data Sizes. arXiv [abs/2106.04803](#) (2021)
14. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: CVPR (2009)
15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: ICLR (2021)
16. Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E.: PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning. In: ECCV (2020)

17. Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: With a Little Help From My Friends: Nearest-Neighbor Contrastive Learning of Visual Representations. In: ICCV (2021)
18. French, R.M.: Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* **3** (1999)
19. Hassani, A., Walton, S., Shah, N., Abuduweili, A., Li, J., Shi, H.: Escaping the Big Data Paradigm with Compact Transformers. arXiv [abs/2104.05704](https://arxiv.org/abs/2104.05704) (2021)
20. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.B.: Masked Autoencoders Are Scalable Vision Learners. arXiv [abs/2111.06377](https://arxiv.org/abs/2111.06377) (2021)
21. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum Contrast for Unsupervised Visual Representation Learning. In: CVPR (2020)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: CVPR (2016)
23. Hoi, S.C., Sahoo, D., Lu, J., Zhao, P.: Online learning: A comprehensive survey. *Neurocomputing* **459** (2021)
24. Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a Unified Classifier Incrementally via Rebalancing. In: CVPR (2019)
25. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In: ICML (2021)
26. Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., Yih, W.t.: Dense Passage Retrieval for Open-Domain Question Answering. In: EMNLP (2020)
27. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114** (2017)
28. Kossen, J., Band, N., Lyle, C., Gomez, A., Rainforth, T., Gal, Y.: Self-Attention Between Datapoints: Going Beyond Individual Input-Output Pairs in Deep Learning. In: NeurIPS (2021)
29. Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep. (2009)
30. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet Classification with Deep Convolutional Neural Networks. In: NIPS (2012)
31. Lange, M.D., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G.G., Tuytelaars, T.: Continual learning: A comparative study on how to defy forgetting in classification tasks. arXiv [abs/1909.08383](https://arxiv.org/abs/1909.08383) (2019)
32. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86** (1998)
33. Li, C., Yang, J., Zhang, P., Gao, M., Xiao, B., Dai, X., Yuan, L., Gao, J.: Efficient Self-supervised Vision Transformers for Representation Learning. In: ICLR (2022)
34. Long, A., Yin, W., Ajanthan, T., Nguyen, V., Purkait, P., Garg, R., Blair, A., Shen, C., van den Hengel, A.: Retrieval Augmented Classification for Long-Tail Visual Recognition. arXiv [abs/2202.11233](https://arxiv.org/abs/2202.11233) (2022)
35. Lopez-Paz, D., Ranzato, M.: Gradient Episodic Memory for Continual Learning. In: NIPS (2017)
36. Northcutt, C.G., Athalye, A., Mueller, J.: Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks. In: NeurIPS Datasets and Benchmarks Track (2021)

37. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic Differentiation in PyTorch. In: NIPS 2017 Workshop on Autodiff (2017)
38. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning Transferable Visual Models From Natural Language Supervision. In: ICML (2021)
39. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: iCaRL: Incremental Classifier and Representation Learning. In: CVPR (2017)
40. Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., Wayne, G.: Experience Replay for Continual Learning. In: NeurIPS (2019)
41. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization. In: ICCV (2017)
42. Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. In: ICLR (2015)
43. Stock, P., Cisse, M.: ConvNets and ImageNet Beyond Accuracy: Understanding Mistakes and Uncovering Biases. In: ECCV (2018)
44. Talmor, A., Herzig, J., Lourie, N., Berant, J.: CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In: NAACL-HLT (2019)
45. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is All you Need. In: NIPS (2017)
46. Wang, W., Xie, E., Li, X., Fan, D., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction without Convolutions. arXiv [abs/2102.12122](#) (2021)
47. Wang, X., Liu, Z., Yu, S.X.: Unsupervised Feature Learning by Cross-Level Instance-Group Discrimination. In: CVPR (2021)
48. Wu, D., Wang, Y.: Adversarial Neuron Pruning Purifies Backdoored Deep Models. In: NeurIPS (2021)
49. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Fu, Y.: Large Scale Incremental Learning. In: CVPR (2019)
50. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised Feature Learning via Non-parametric Instance Discrimination. In: CVPR (2018)
51. Xu, Y., Zhu, C., Wang, S., Sun, S., Cheng, H., Liu, X., Gao, J., He, P., Zeng, M., Huang, X.: Human Parity on CommonsenseQA: Augmenting Self-Attention with External Attention. arXiv [abs/2112.03254](#) (2021)
52. Zenke, F., Poole, B., Ganguli, S.: Continual Learning Through Synaptic Intelligence. In: ICML (2017)
53. Zhong, J., Ding, G., Guo, Y., Han, J., Wang, B.: Where to Prune: Using LSTM to Guide End-to-end Pruning. In: IJCAI (2018)
54. Zhou, B., Khosla, A., Lapedriza, À., Oliva, A., Torralba, A.: Learning Deep Features for Discriminative Localization. In: CVPR (2016)
55. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: Image BERT Pre-training with Online Tokenizer. In: ICLR (2022)
56. Zhou, Y., Zhang, Y., Wang, Y., Tian, Q.: Accelerate CNN via Recursive Bayesian Pruning. In: ICCV (2019)