

# Kaizen: Practical self-supervised continual learning with continual fine-tuning

Chi Ian Tang<sup>1,2</sup>, Lorena Qendro<sup>1</sup>, Dimitris Spathis<sup>1</sup>, Fahim Kawsar<sup>1</sup>,  
Cecilia Mascolo<sup>2</sup>, and Akhil Mathur<sup>1</sup>

<sup>1</sup>Nokia Bell Labs, UK

<sup>2</sup>University of Cambridge, UK

<sup>1</sup>{ian.tang, lorena.qendro, dimitrios.spathis, fahim.kawsar}@nokia-bell-labs.com

<sup>2</sup>{cit27, cm542}@cam.ac.uk

## Abstract

*Self-supervised learning (SSL) has shown remarkable performance in computer vision tasks when trained offline. However, in a Continual Learning (CL) scenario where new data is introduced progressively, models still suffer from catastrophic forgetting. Retraining a model from scratch to adapt to newly generated data is time-consuming and inefficient. Previous approaches suggested re-purposing self-supervised objectives with knowledge distillation to mitigate forgetting across tasks, assuming that labels from all tasks are available during fine-tuning. In this paper, we generalize self-supervised continual learning in a practical setting where available labels can be leveraged in any step of the SSL process. With an increasing number of continual tasks, this offers more flexibility in the pre-training and fine-tuning phases. With *Kaizen*<sup>1</sup>, we introduce a training architecture that is able to mitigate catastrophic forgetting for both the feature extractor and classifier with a carefully designed loss function. By using a set of comprehensive evaluation metrics reflecting different aspects of continual learning, we demonstrated that *Kaizen* significantly outperforms previous SSL models in competitive vision benchmarks, with up to 16.5% accuracy improvement on split CIFAR-100. *Kaizen* is able to balance the trade-off between knowledge retention and learning from new data with an end-to-end model, paving the way for practical deployment of continual learning systems.*

## 1. Introduction

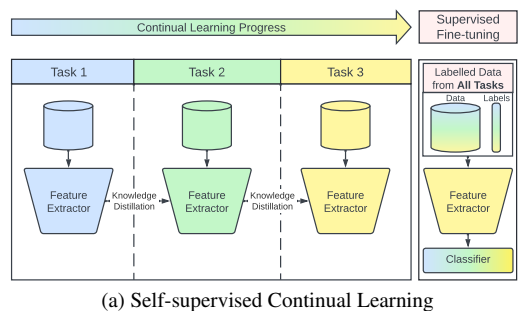
While traditional machine learning algorithms perform well on tasks they have been exposed to during training, they are unable to adapt to new concepts, classes, or data in-

troduced after training over time. This is particularly important in vision tasks, where deployed models such as those running on phones or CCTV cameras need to adapt to new information. Currently, model updates require, first, offline re-training on a central server and then distributing the updated model to the end users. This procedure greatly affects the usability and user experience since new functionalities such as recognizing new people in a photo album app [26], will not be included until a large amount of data has been collected for re-training. Given these limitations, a system that can adapt to new data and learn continually would be highly desirable.

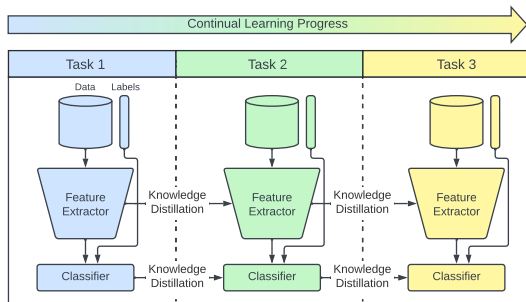
Continual learning (CL) methods aim to learn continually from an ever-changing stream of data, acquiring new knowledge while retaining performance on previous tasks [34]. The main difference to conventional deep learning is the data assumption: rather than having all training data available at once, data with different distributions (classes or domains) arrive over time. In practice, data can be only temporarily available for training because of storage, computing, and memory limitations as well as privacy concerns. Data availability comes also with the cost of labelling samples for supervised learning or fine-tuning, since only a limited amount of data can be annotated at a time.

Self-supervised Learning (SSL) techniques solve the data availability–annotation challenge even outperforming supervised methods in difficult vision tasks. SSL models learn meaningful representations that can be transferred to various downstream applications [1,8,9,16]. However, they usually operate under strict data assumptions where training is conducted on large aggregated datasets. These assumptions no longer hold when new unlabelled data is introduced over time and there is not a *single* snapshot of the dataset to train on. As a result, SSL models should adapt to the temporary availability of both unlabelled and labelled data. On the other hand, prior works in CL almost exclusively focus on supervised setups using ideas such as regulariza-

<sup>1</sup>The code for this work is available at  
<https://github.com/dr-bell/kaizen>



(a) Self-supervised Continual Learning



(b) Kaizen: Self-supervised Continual Learning with Continual Fine-tuning

Figure 1. **Self-supervised Continual Learning vs Continual Fine-tuning.** Existing approaches (a) wait until the end of the continual process to fine-tune, Kaizen (b) leverages distillation across both the feature extraction and fine-tuning steps for each task.

tion, replay, or parameter-isolation [21, 31, 32]. Our aim is to leverage both SSL and CL paradigms, therefore here we focus on Continual Self-Supervised Learning (CSSL), a still under-explored area. Some recent works [14, 24] have started looking at CSSL where models learn from a stream of unlabelled data, however, these methods focus on training a strong feature extractor assuming that labelled data from *all* tasks<sup>2</sup> are available for fine-tuning. We argue that this violates the core CL assumption where limited to no data from earlier tasks is available. A more practical setting would involve leveraging both unlabelled and labelled data *within each task’s training round*.

To achieve that, we introduce Kaizen<sup>3</sup>, a new CSSL architecture designed to work under a more practical data assumption, where end-to-end training is performed using the continuous arrival of unlabelled and labelled data. This architecture strikes a balance between (i) training the feature extractor and fine-tuning, and (ii) combating catastrophic forgetting and learning from new data by employing a novel

<sup>2</sup>A task in this context refers to a particular distribution of data, which can come from different sets of classes (for classification problems) or different domains.

<sup>3</sup>inspired by the East Asian concept of “continuous improvement”

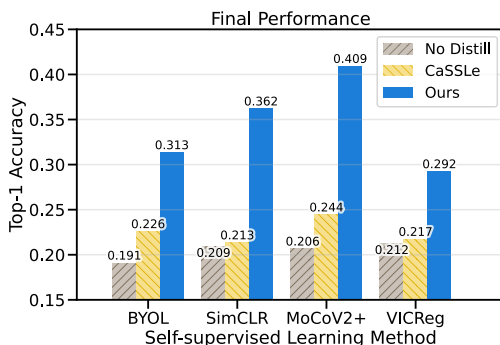



Figure 2. **Performance comparison.** Models trained using different self-supervised learning methods and knowledge distillation strategies on class-incremental CIFAR-100. The top figure shows the average performance across the entire continual learning process, while the bottom figure shows the performance in the final evaluation.

loss function that is specifically designed for these objectives.

We demonstrate that Kaizen outperforms previous works that focus on self-supervised continual learning in terms of their overall performance and the ability to retain knowledge after training on the final task, and more importantly after training on every task, by up to 14.4% in accuracy (in absolute terms) and 25.4% reduction in forgetting. Our main contributions are:

1. **Practical Continual Learning framework.** We propose Kaizen, a continual learning framework that can be deployed at any point during the continual learning process with a functional classifier. It leverages both unlabelled and labelled data in training the feature extractor and classifier instead of only using one type of data, in a carefully designed loss function and distillation mechanism, thus allowing higher flexibility in terms of storage requirements and in accommodating privacy concerns, which is important in real-world applications.
2. **Evaluation setup reflecting the real world.** We propose a novel evaluation setup with a range of evaluation

metrics that closely reflect how models perform in the real world, where there is a focus on the performance of the model over the entire continual learning process, instead of only the final model.

3. **Extensive empirical analysis.** We extensively evaluate `Kaizen` in various settings and vision benchmarks with state-of-the-art self-supervised learning techniques. We show that `Kaizen` is robust to catastrophic forgetting in classification tasks when encountering new data, while maintaining the quality of the feature extractor.

## 2. Related Work

**Continual Learning.** Traditional deep learning techniques suffer from catastrophic forgetting [21], a phenomenon where models tend to forget what has been previously learned by overriding it with new data. Continual learning methods address this issue by striking a balance between the stability (ability to retain knowledge) and plasticity (ability to learn new concepts) of deep learning models. Broadly, continual learning techniques can be categorized into three types [11]: replay-based [2, 7, 27, 28], where a small subset of previously seen data is retained and replayed to the model at every stage, regularisation-based [21, 23, 32, 35], in which regularisation objectives are used to guide the model to retain old knowledge, and parameter-isolation methods [30, 31], where separate regions/neurons of the models are dedicated to each task and remain frozen during further training. *Nevertheless, these techniques mostly focus on supervised learning setups, without dealing with label scarcity which is a common issue in real-world scenarios.*

**Self-supervised Learning (SSL).** SSL methods use unlabelled data with the goal of training more generalizable models. Recent techniques which involve bringing the representations of correlated views (e.g. transformed versions) of a single data point (e.g., image) closer to each other have proven powerful where the learned representations of the feature extractor achieve performance on par with those trained in a supervised manner [8, 16, 36]. These techniques include contrastive methods such as SimCLR [8], similarity maximization methods such as BYOL [16], momentum-based methods such as MoCo [9, 17], and reduction-based methods like VicReg [1]. However, in continual learning settings, these methods still suffer from catastrophic forgetting [11, 14]. Furthermore, they require substantial amounts of unlabelled data available *all at once* to work well, which is not practical in scenarios considering streams of unlabelled data.

**Continual Self-supervised Learning (CSSL).** CSSL approaches operate on unlabelled data while dealing with catastrophic forgetting. Earlier techniques focus either on self-supervised pre-training to later apply supervised

continual learning [3, 15] or extend the contrastive SSL paradigm to CL, too [6, 25]. However, these techniques have a narrow focus as they are tailored to specific SSL architectures. Few recent works [11, 14] started looking at general frameworks for unsupervised continual learning, where the model learns continually from a stream of unlabelled data using a combination of unsupervised learning techniques and knowledge retention mechanisms. Nevertheless, these methods only solve half of the problem in CL – they focus on training a strong feature extractor continually with unlabelled data and assume that all labelled data are stored and available for fine-tuning after the final CL step. *This violates the data assumption in continual learning where labelled data is only temporarily available, and the classifier should also learn continually.* `Kaizen` puts forward a general architecture in which both the feature extraction and the fine-tuning are performed continually from a stream of both large unlabelled and small labelled data.

## 3. Method

### 3.1. Background

**Continual Learning (CL).** In this work, we adopt the notion of *task-incremental* learning [11], in which the model sees the training data for one task at a time. We assume data  $(\mathcal{X}^{(t)}, \mathcal{Y}^{(t)})$ , which is randomly drawn from distributions  $P(\mathcal{X}^{(t)})$  and  $P(\mathcal{Y}^{(t)})$  for task  $t$ , is available for the model to learn from, with the goal to optimize for all seen tasks while having little to no access to data  $(\mathcal{X}^{(t')}, \mathcal{Y}^{(t')})$  from previous tasks  $t' < t$  [11]. Here we focus on the *class-incremental* learning [11, 19] setting where each task contains an exclusive subset of classes in a dataset, and thereby  $P(\mathcal{X}^{(i)}) \neq P(\mathcal{X}^{(j)})$  and  $P(\mathcal{Y}^{(i)}) \neq P(\mathcal{Y}^{(j)})$  if  $i \neq j$ , *without* the task label  $t$  provided to the model during inference or evaluation.

**SSL for visual representations.** Siamese representation learning or contrastive learning proposed in recent works [1, 4, 5, 8–10, 16, 17, 38] have demonstrated strong supervisory signal for a model to learn to extract distinctive features from large unlabelled datasets. The main idea of these methods involves passing two stochastically augmented views  $x_1$  and  $x_2$  of an input sample  $x$  through the same feature extractor (or one through the feature extractor and another through an exponentially updated momentum feature extractor in some works [16]), and optimizing using a loss function such that the extracted features are similar, with different mechanisms such as gradient stopping [10, 16] or a clustering task [5] to avoid collapse. The stochastic augmentations are designed such that the output images encode the semantics that the model learns to extract, in which they still contain almost the same meaning but the appearance might be different.

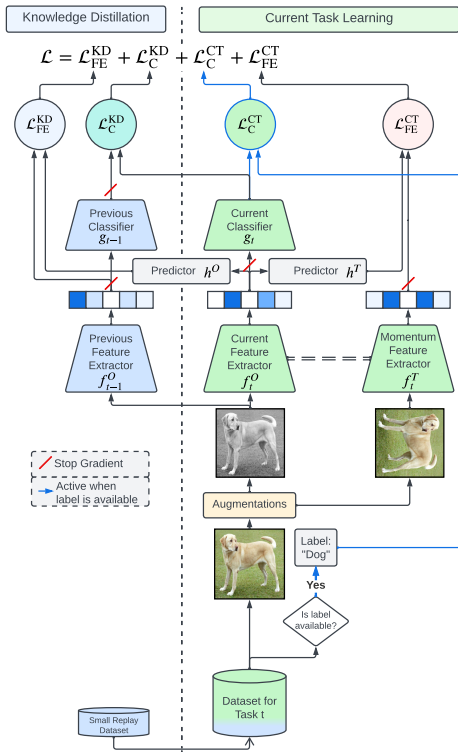


Figure 3. **Overview of the Kaizen framework.** Kaizen balances knowledge distillation and current-task learning in an end-to-end manner through a joint loss function. Available SSL methods can be used in training the feature extractors alongside knowledge distillation, while the classifiers are trained on both unlabelled and labelled data through knowledge distillation and fine-tuning.

### 3.2. Kaizen: Practical continual learner that balances self-supervised learning and fine-tuning

In order to continually learn from unlabelled data while maintaining a functional classifier, the Kaizen framework (as illustrated in Figure 3) is designed to balance different objectives including knowledge retention, self-supervised learning from unlabelled data, and supervised learning for classification. It consists of two main components: knowledge distillation and current task learning, one of each for the feature extractor and for the classifier.

**Current task learning.** This component trains the feature extractor and the classifier using the data for the current task (see the right half of Figure 3). In order to allow the model to learn from both unlabelled and labelled data, the feature extractor is trained on the current task using one of the self-supervised learning methods described above. Stochastic augmentation functions are applied to the input image which produces two different but correlated views of the image. One of these images is passed to the Current Fea-

ture Extractor  $f_t^O$ , while the other to the Momentum Feature Extractor  $f_t^T$  (or the same Current Feature Extractor depending on the SSL method), to obtain two different embeddings. The embedding from the Current Feature Extractor is passed to an additional, shallow neural network called the predictor  $h^T$ . The discrepancy between the embedding from the Momentum Feature Extractor and that from the predictor will then be reduced using the corresponding self-supervised loss function  $\mathcal{L}_{FE}^{CT}$ . The use of an additional predictor here is to allow the feature extractor to be flexible so that it does not have to produce the same embedding for the two augmented views, but should contain the same amount of information.

If the label for the image is available, we train the classifier in a supervised manner. The embedding obtained from the Current Feature Extractor will be fed to the classifier network  $g_t$  to obtain class probabilities after softmax activation. Categorical cross-entropy loss  $\mathcal{L}_C^{CT}$  is used for training the classifier. Note that the gradient updates do not back-propagate to the feature extractor. This is to allow the feature extractor to focus on extracting distinctive features from self-supervision, instead of specialising in the current classification task. This is important to ensure that the feature extractor remains general throughout the continual learning process, where the class distribution might shift from one task to another.

**Knowledge distillation.** Apart from the first task, the model is required to retain knowledge learned from previous tasks while learning from the new data. To achieve this, we make a frozen copy of the trained model ( $f_{t-1}^O$  and  $g_{t-1}$ ) from the previous task for knowledge distillation before we start training (Figure 3 (left)). For the feature extractor, we adopt a scheme similar to CaSSLe [14], where we mirror what we have done for the current task learning, by using an SSL method but with the Momentum Feature Extractor replaced by the feature extractor from the previous step  $f_{t-1}^O$ . The augmented image that was passed through the Online Feature Extractor  $f_t^O$  will also be fed to the Previous Feature Extractor. Similarly, the discrepancy between the embedding from the previous and the current feature extractor (after passing through a different predictor  $h^O$ ) is reduced using the corresponding self-supervised loss function  $\mathcal{L}_{FE}^{SSL}$ . The Previous Feature Extractor is frozen during training and gradients updates will not be applied to it.

Similar to the feature extractor, knowledge distillation is also performed for the classifier. The predictions from the Current Classifier  $g_t$  are made to be similar to the predictions (or soft labels) from the Previous Classifier ( $g_{t-1}$ ) using the categorical cross-entropy loss. Note that unlike in current task learning, knowledge distillation for the classifier is active regardless of the presence of a label given the input image.

**Memory replay.** One additional mechanism that our

method employs, is memory replay [20, 28, 29], where a small subset of actual samples or representative samples from previous tasks are kept and replayed during training. This has been shown to be crucial in maintaining model performance and combating catastrophic forgetting in task-label-free methods [11], where the task label is not provided to the model during inference.

**Overall Framework.** The overall loss function (Equation 1) consists of four components: the loss for knowledge distillation of the feature extractor  $\mathcal{L}_{FE}^{KD}$ , the loss for knowledge distillation of the classifier  $\mathcal{L}_C^{KD}$ , the cross-entropy loss of the classification task  $\mathcal{L}_C^{CT}$ , and the loss for self-supervised learning  $\mathcal{L}_{FE}^{CT}$ . These losses balance different learning objectives to ensure that the model, including the feature extractor and the classifier, is able to learn from new data while retaining knowledge from previous tasks.

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{FE}^{KD} + \mathcal{L}_C^{KD} + \mathcal{L}_C^{CT} + \mathcal{L}_{FE}^{CT} \\ &= \mathcal{L}_{SSL}(f_{t-1}^O(x_1), h^O(f_t^O(x_1))) + \\ &\quad \mathcal{L}_{CE}(g_{t-1}(f_{t-1}^O(x_1)), g_t(f_t^O(x_1))) + \\ &\quad \mathcal{L}_{CE}(g_t(f_t^O(x_1)), y) + \\ &\quad \mathcal{L}_{SSL}(f_t^T(x_2), h^T(f_t^O(x_2))) \end{aligned} \quad (1)$$

Here the  $\mathcal{L}_{CE}$  refers to the categorical cross-entropy loss, and  $\mathcal{L}_{SSL}$  refers to the self-supervised loss for the particular self-supervised learning method.

## 4. Evaluation for Continual Learning

In this section, we first discuss the evaluation protocol and metrics adopted in this work and then present our experimental setup.

### 4.1. Self-supervised Continual Learning vs Continual Fine-tuning

Pure self-supervised continual learning focuses on training a strong feature extractor from data with changing distributions, and after learning from the final set of data, a classifier is trained with labelled data (see Figure 1a). Although the feature extractor is learning continually, this only solves half of the challenge in combating catastrophic forgetting. This is because if we want a classifier that is able to perform classification on all seen tasks, labelled data from all tasks must be retained or re-collected for training. In the evaluation, we show that if we limit the amount of labelled data from earlier tasks to a small subset, the performance suffers significantly.

We argue that in a more practical scenario, the classifier should be trained together with the feature extractor (see Figure 1b) because the labelled data is more likely to be available while training for a particular task, rather than after the feature extractor has been trained. This would require extra consideration with regard to knowledge distilla-

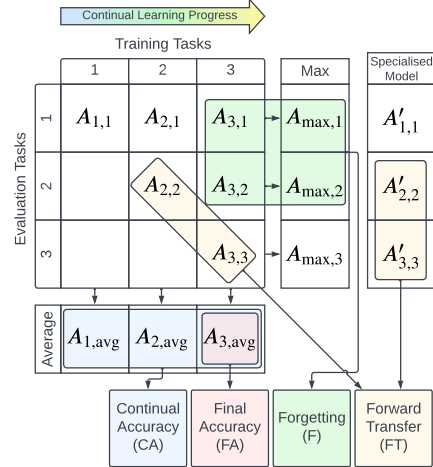


Figure 4. **Calculation of the evaluation metrics.** Illustration of our practical evaluation setup with regards to metrics and tasks used in each calculation.

tion for the classifier on top of the feature extractor. Our approach is designed to address this issue.

### 4.2. Evaluation metrics

In order to compare the performance of different methods fairly, we define the following evaluation metrics (see Figure 4) that reflect different aspects of a continual learning framework, taking into account metrics defined in previous works [11, 13, 20]. Here we denote the accuracy of a model on a particular task  $k$  (averaged across all classes in that task) after seeing the last sample from task  $t$  as  $A_{t,k}$ , and we assume that there are  $T$  tasks in total.

**Final Accuracy (FA)** defined as  $FA = \frac{1}{T} \sum_{i=1}^T A_{T,i}$  refers to the average model accuracy across all tasks, after it has been trained on the last task  $T$ .

**Continual Accuracy (CA)**, instead of only looking at the performance of the model at the final step, looks at the performance throughout the continual learning process, where a model might be deployed before the last task and be later re-trained. We define this to be the average accuracy of the model after being trained on each task  $CA = \frac{1}{T} \sum_{i=1}^T (\frac{1}{i} \sum_{j=1}^i A_{j,i})$ .

**Forgetting (F)** measures the extent of the performance a model has lost after training on new tasks. We calculate this by taking the difference between the highest performance of the model on a given task, and that of the model at the final step, excluding the final task:  $F = \frac{1}{T-1} \sum_{i=1}^{T-1} (A_{max,i} - A_{T,i})$ , where  $A_{max,i} = \max_{t \in \{1, \dots, T\}} A_{t,i}$ .

**Forward Transfer (FT)** refers to the use of previously acquired knowledge when learning new tasks. Here, we define it with model deployment considerations in mind: we take the difference between the performance of the model on

task  $k$  after seeing data from all tasks  $i \leq k$ , compared to a model which has only seen data of task  $k$ . We denote the accuracy of the model which is only trained on task  $k$  as  $A'_{k,k}$ . FT is thus calculated by:  $FT = \frac{1}{T-1} \sum_{i=2}^T (A_{i,i} - A'_{i,i})$ . A positive value would indicate that the training on previous tasks helped the model learn the new task. Note that this is a desired property that is hard to achieve.

This evaluation framework is set to facilitate comparisons of different properties of CL models in real-world applications and balance trade-offs between different properties such as forgetting and forward transfer. These metrics should provide guidance for which method to select depending on desired use cases.

### 4.3. Experimental setup

**Datasets.** We performed our evaluation using 2 datasets: CIFAR-100 [22], an object recognition dataset with 60,000 32x32 images evenly distributed across 100 classes, and ImageNet100 [33], also an object recognition dataset which is a 100-class subset of the original ImageNet dataset [12] and consists of 130,000 224x224 images spread across 1000 classes. In our experiments, we randomly split these datasets into tasks with an equal number of classes. For example, for 5 tasks, we have 20 classes in each task.

**SSL methods.** Since we do not modify the self-supervised learning component from its original formulation [1, 8, 9, 16] and the CaSSLe adaptation [14], our method is compatible with existing self-supervised learning methods. We selected the contrastive-based methods SimCLR [8] and MoCoV2+ [9, 17], the asymmetric-model-based method BYOL [16] and the cross-correlation-based method VICReg [1] as the backbone SSL method for training. Our aim is to investigate whether our proposed architecture generalizes across different self-supervised learning methods and identify limitations.

**Hyperparameters.** We build upon the Pytorch implementation introduced by [14]. For each task, we train the model for 500 epochs for CIFAR-100 with 5 tasks, 250 epochs for CIFAR-100 with 20 tasks, and 200 epochs for ImageNet100. The batch size is 256 for CIFAR-100 and 128 for ImageNet100. We keep the amount of data replayed to the model during continual learning (for our method) and during classifier training (for other baselines) to be 1% of the original data as this offered the best tradeoff between accuracy and amount of labelled data used. ResNet18 [18] is used as the architecture for the feature extractor, while the classifier consists of one fully-connected layer of 1000 units and another one with the number of output classes, which is 100. LARS [37] is used for large batch training. A weighting factor of 2 is used for the knowledge distillation loss for the classifier ( $\mathcal{L}_C^{KD}$  in Equation 1). We keep all other hyperparameters unmodified. In the supplementary material, we provide a more detailed discussion of the hyperparameters

and the source code for reproducibility.

**Baselines.** We compare our framework against the state-of-the-art CSSL pipeline, CaSSLe [14], and the *No distill* setup, where no extra measures are taken to mitigate catastrophic forgetting, and the entire model is fine-tuned from task to task. Since the baseline methods are pure self-supervised pre-training methods, the classifier is trained after the feature extractor is fully trained using the task data. To ensure a fair comparison, these classifiers are trained with memory replay enabled: the same subset of data that our model is trained on, is also available for these classifiers to train on.

## 5. Results

### 5.1. Performance comparison against CSSL

As discussed in section 4.3, we focus on the comparison of our *Kaizen* against CaSSLe and the *No distill* setup. Figure 2 compares the Continual Accuracy and Final Accuracy of these three methods with different SSL models for the feature extractor, on the CIFAR-100 dataset split into 5 tasks of equal sizes. *Kaizen* outperforms the other two baselines irrespective of the evaluation metric or the SSL model, achieving the highest continual accuracy at 0.570 and final accuracy at 0.409 using MoCoV2+, outperforming CaSSLe by 0.138 and 0.165 respectively. It is important to note that the data availability is kept the same across all methods, where at each task, every method has access to data of the current task and 1% of replay data from previous tasks. Our method outperforms the rest by incorporating knowledge distillation and fine-tuning into the pipeline, instead of training the classifier training in the end. This validates our hypothesis and shows that *Kaizen* is overall effective in retaining knowledge from previous tasks, by performing well at both the final step and throughout the continual learning process. It is interesting to note that MoCoV2+ is the most effective of the four SSL methods; this could potentially be attributed to the use of the Memory Buffer of representations, which acts as a knowledge retention mechanism that enables gradual learning.

### 5.2. Performance variation across time

Here we examine the performance of different methods at different stages of the continual learning process. Figures 5 and 6 show the average accuracy over seen tasks of *Kaizen* and the baselines after training on each task, using CIFAR-100 and ImageNet-100 (5 tasks). We did not perform an evaluation on ImageNet-100 for the *No distill* pipeline because it consistently underperforms CaSSLe. We find that in general, all methods show lower performance after training on more tasks, partially due to the fact that the classification problem becomes more difficult as the number of classes increases, as well as catastrophic for-

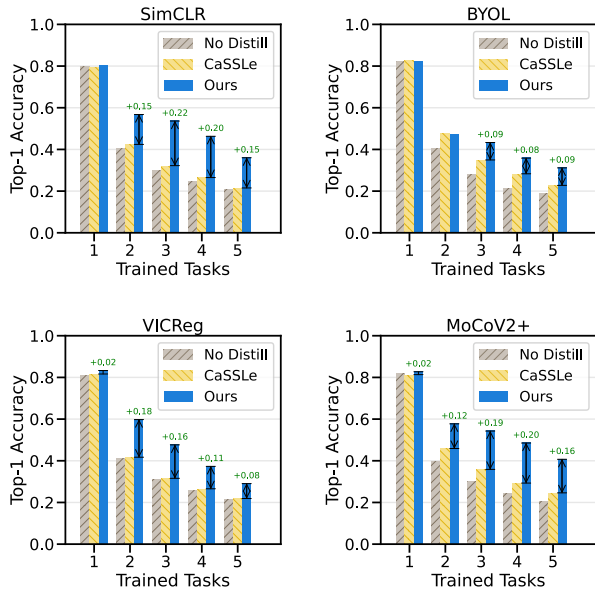


Figure 5. Average performance over tasks on CIFAR-100. Comparison between Kaizen and baselines using 4 SSL algorithms and 5 tasks. Our model consistently outperforms baselines and is robust to forgetting in later tasks.




Figure 6. Average performance over tasks on ImageNet-100. Comparison between Kaizen and CaSSLe for 5 tasks. Our model outperforms in SimCLR and the first tasks in BYOL.

getting. Echoing the results of the previous section, our method maintains a higher level of accuracy overall, even though all methods start from similar performance on the first task. The difference in performance is more evident in CIFAR-100 (Figure 5) than in ImageNet-100 (Figure 6), which could be attributed to the larger data size and more available labelled data.

### 5.3. Longer continual learning scenarios

In certain scenarios, a model could be required to continually learn from data with changing distribution over a long time. Figure 7 displays the performance of Kaizen against CaSSLe on CIFAR-100 when it is split into 20 tasks




Figure 7. Average performance over 20 tasks on CIFAR-100. Our model achieves higher accuracy for the first 10 tasks.




Figure 8. Detailed breakdown of performance over tasks on CIFAR-100. Fine-grained accuracy for every additional task across Kaizen and CaSSLe, with a fixed SSL backbone.

(5 classes each) instead of 5 tasks only. We can see that Kaizen is able to outperform the state-of-the-art in the first 10 tasks, but reaches a similar level of performance or even lower in later tasks. We believe that these results demonstrate the trade-off between knowledge retention and learning new tasks: our method trades performance on new tasks for combating catastrophic forgetting. In continual learning with many tasks, this might not be the best strategy because the number of new tasks is usually higher than the number of existing classes. Lower performance on any particular new task might lead to lower performance down the road.

### 5.4. Per-task performance breakdown

To support our intuition that our method prioritises knowledge retention over learning from new tasks, we investigate the per-task performance at each step. Figure 8 shows the performance of the model on each individual task throughout the continual learning process, using CIFAR-100 with 5 tasks. We find that Kaizen is able to forget acquired knowledge more gracefully [11] than CaSSLe, where the performance on previous tasks gracefully degrades over time. Without a knowledge distillation scheme, CaSSLe, on the other hand, is able to perform on the new task better than Kaizen, but the performance on previous tasks drops significantly in just one step. This shows that our method strikes a balance between performance on new tasks and knowledge retention. The supplementary mate-

SSL	Baseline	FA $\uparrow$	CA $\uparrow$	F $\downarrow$	FT $\uparrow$
BYOL	No Distill	0.191	0.381	0.794	0.003
	CaSSLe	0.226	0.431	0.758	<b>0.010</b>
	<b>Kaizen</b>	<b>0.313</b>	<b>0.481</b>	<b>0.610</b>	-0.028
SimCLR	No Distill	0.209	0.391	0.755	<u>0.021</u>
	CaSSLe	0.213	0.403	0.747	0.019
	<b>Kaizen</b>	<b>0.362</b>	<b>0.547</b>	<b>0.451</b>	-0.170
MoCoV2+	No Distill	0.206	0.393	0.760	<b>-0.010</b>
	CaSSLe	0.244	0.432	0.709	-0.012
	<b>Kaizen</b>	<b>0.409</b>	<b>0.570</b>	<b>0.396</b>	-0.210
VICReg	No Distill	0.212	0.400	0.750	<b>0.004</b>
	CaSSLe	0.217	0.405	0.739	-0.001
	<b>Kaizen</b>	<b>0.292</b>	<b>0.516</b>	<b>0.580</b>	-0.110

Table 1. **SSL & baseline performance comparison.** Evaluation of *Kaizen* compared to four SSL and two continual baselines across four metrics on CIFAR-100. Our model with the MoCoV2+ backbone outperforms in most metrics. Best performance across SSL methods is bold, and across metrics is underlined. (FA: Final Accuracy, CA: Continual Accuracy, F: Forgetting, FT: Forward Transfer)

rial contains additional results with ImageNet-100 and other SSL techniques. Trade-off along this spectrum is left for future work.

## 5.5. Comprehensive evaluation of continual learning and SSL methods

Table 1 presents *Kaizen* compared to other baselines on the four considered SSL paradigms. These results on CIFAR-100 show that *Kaizen* improves the performance of all SSL methods as represented by the final and continual accuracy. As expected, the absence of distillation in No Distill definitely hurts the accuracy of the overall continual learning framework as it is unable to maintain knowledge with each additional task. This can also be seen in the Forgetting metric where No Distill is on average 0.255 higher than *Kaizen* and 0.026 higher than CaSSLe. *Kaizen* always outperforms CaSSLe across the FA, CA, and F metrics achieving the best performance with MoCoV2+. On the other hand, as shown in the results above, CaSSLe and *No distill*, which prioritise learning from new data, show a slight positive transfer in some cases, while our model suffers from negative forward transfer.

## 5.6. Ablation on replay dataset size

Replay was shown to be one of the crucial components in combating catastrophic forgetting for methods that do not rely on task labels [11]. In this evaluation, we vary the amount of data being replayed to *Kaizen* throughout the training process ranging from 0%, 1% (default setting in all experiments), 5% to 10% on CIFAR-100 using MoCoV2+ as the SSL method, and display these results in Figure 9.




Figure 9. **Ablation analysis for different amounts of replay data.** Performance of our model on CIFAR100 with MoCoV2+ backbone with varying amounts of replay data (1% is the default replay in the rest of the paper). A higher amount of replay data leads to better overall performance, less forgetting, but lower forward transfer.

As expected, by replaying more data, the model achieves higher continual accuracy as well as higher final accuracy and lower forgetting. However, we cannot see the full picture without looking at the change in forward transfer. Here, we observe that by replaying more data, the model prioritises knowledge retention even more, and forward transfer suffers as a result. This trade-off is important to be taken into account when designing a continual learning system, in addition to storage and privacy concerns. It is important to note that *Kaizen* works reasonably well without any replay (0% case), which would satisfy a strict data assumption when replay is not practical.

## 6. Conclusion

In order to deploy continual learning systems in real-world scenarios, we have to address many challenges that are present in existing works. Towards this goal, here we introduced *Kaizen*, a continual learning method that addresses shortcomings of self-supervised and supervised continual learning methods, in which a classifier is continually trained with the possibility of deployment at any point during the training process. We carefully design the learning objective to balance feature extractor and classifier training, as well as knowledge retention and learning from new data. Through extensive evaluation and a comprehensive set of evaluation metrics, we have demonstrated that *Kaizen* is able to balance the trade-off between knowledge retention and learning from new data more gracefully compared to the state-of-the-art, and achieves higher performance overall.

## 7. Acknowledgements

This work is partially supported by Nokia Bell Labs through their donation for the Centre of Mobile, Wearable Systems and Augmented Intelligence to the University of Cambridge.



## References

- [1] Adrien Bardes, Jean Ponce, and Yann Lecun. Vi-creg: Variance-invariance-covariance regularization for self-supervised learning. In *ICLR 2022-10th International Conference on Learning Representations*, 2022. [1](#), [3](#), [6](#), [11](#)
- [2] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020. [3](#)
- [3] Lucas Caccia and Joelle Pineau. Special: Self-supervised pretraining for continual learning. In *Continual Semi-Supervised Learning: First International Workshop, CSSL 2021, Virtual Event, August 19–20, 2021, Revised Selected Papers*, pages 91–103. Springer, 2022. [3](#)
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020. [3](#)
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [3](#)
- [6] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 9516–9525, 2021. [3](#)
- [7] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with agem. *arXiv preprint arXiv:1812.00420*, 2018. [3](#)
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [1](#), [3](#), [6](#), [11](#)
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [1](#), [3](#), [6](#), [11](#)
- [10] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021. [3](#)
- [11] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. [3](#), [5](#), [7](#), [8](#)
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [6](#)
- [13] Natalia Díaz-Rodríguez, Vincenzo Lomonaco, David Filliat, and Davide Maltoni. Don’t forget, there is more than forgetting: new metrics for continual learning. In *Workshop on Continual Learning, NeurIPS 2018 (Neural Information Processing Systems)*, 2018. [5](#)
- [14] Enrico Fini, Victor G Turrissi da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022. [2](#), [3](#), [4](#), [6](#), [11](#)
- [15] Jhair Gallardo, Tyler L Hayes, and Christopher Kanan. Self-supervised training enhances online continual learning. *arXiv preprint arXiv:2103.14010*, 2021. [3](#)
- [16] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [1](#), [3](#), [6](#), [11](#)
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [3](#), [6](#)
- [18] 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](#)
- [19] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. [3](#)
- [20] David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [5](#)
- [21] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. [2](#), [3](#)
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [6](#)
- [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [3](#)
- [24] Zhiwei Lin, Yongtao Wang, and Hongxiang Lin. Continual contrastive learning for image classification. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2022. [2](#)
- [25] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Rethinking the representational continuity: Towards unsupervised continual learning. *arXiv preprint arXiv:2110.06976*, 2021. [3](#)
- [26] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021. [1](#)
- [27] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic

- plasticity driven framework for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11321–11329, 2019. 3
- [28] 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. 3, 5
- [29] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. 5
- [30] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 3
- [31] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018. 2, 3
- [32] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [33] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020. 6
- [34] Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 1
- [35] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. 3
- [36] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 3
- [37] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. 6
- [38] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 3