

GENERATIVE KERNEL CONTINUAL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Kernel continual learning by [Derakhshani et al. \(2021\)](#) has recently emerged as a strong continual learner due to its non-parametric ability to tackle task interference and catastrophic forgetting. Unfortunately its success comes at the expense of an explicit memory to store samples from past tasks, which hampers scalability to continual learning settings with a large number of tasks. In this paper, we introduce generative kernel continual learning, which explores and exploits the synergies between generative models and kernels for continual learning. The generative model is able to produce representative samples for kernel learning, which removes the dependence on memory in kernel continual learning. Moreover, as we replay only on the generative model, we avoid task interference while being computationally more efficient compared to previous methods that need replay on the entire model. We further introduce a supervised contrastive regularization, which enables our model to generate even more discriminative samples for better kernel-based classification performance. We conduct extensive experiments on three widely-used continual learning benchmarks that demonstrate the abilities and benefits of our contributions. Most notably, on the challenging Split-CIFAR100 benchmark, with just a simple linear kernel we obtain the same accuracy as kernel continual learning with variational random features for one tenth of the memory, or a 10.1% accuracy gain for the same memory budget.

1 INTRODUCTION

Continual learning ([Ring, 1998](#); [Lopez-Paz & Ranzato, 2017](#); [Goodfellow et al., 2014](#)), also known as lifelong learning, strives to continually learn to solve a sequence of non-stationary tasks. The continual learner is required to accommodate new information, while maximally maintaining the knowledge acquired in previous tasks so as to be still able to complete those tasks. This is a challenge for contemporary deep neural networks as they suffer from catastrophic forgetting when learning over non-stationary data ([McCloskey & Cohen, 1989](#)).

Kernel continual learning by [Derakhshani et al. \(2021\)](#) recently emerged as a strong continual learner by combining the strengths of deep neural networks and kernel learners ([Schölkopf et al., 2001](#); [Schölkopf & Smola, 2002](#); [Rahimi & Recht, 2007](#)). Kernel continual learning deploys a non-parametric classifier based on kernel ridge regression, which systematically avoids task interference and offers an effective way to deal with catastrophic forgetting. Moreover, by sharing the feature extraction and kernel inference networks, useful knowledge is transferred across tasks. Kernel continual learning performs well in both the task-aware and domain incremental learning scenarios, while being simple and efficient in terms of architecture and training time.

Despite its appealing abilities, the success of kernel continual learning comes at the expense of an explicit memory that needs to maintain the data of all previously experienced tasks. It includes an episodic memory unit to store a subset of samples from the training data for each task, called the ‘coreset’, from which a classifier learns based on kernel ridge regression. In order to perform well, a large memory is required to construct a satisfactory kernel, which causes computation and storage overhead when learning along with a growing number of tasks. In addition, the coreset is constructed by drawing samples uniformly from existing classes in the same task. This uniform sampling strategy is unlikely to provide the most representative and discriminative samples per task, potentially hurting accuracy.




Figure 1: **Overview of generative kernel continual learning.** The variational auto-encoder is adopted as a generative model, which learns the data distribution of task t while doing replay over the data $\mathcal{R}_{<t}$ from previous tasks to avoid catastrophic forgetting. The trained decoder of the generative model is deployed to generate the coreset \mathcal{C}_t . The kernel \mathcal{K}_t for the current task t is constructed based on the generated \mathcal{C}_t and \mathcal{D}_t . The classifier for task t is constructed based on kernel ridge regression using \mathcal{K}_t . The encoder network q_ϕ^t , parameterized by ϕ , is shared and updated when training on the task sequence. The decoder network p_θ , parameterized with θ , adopts a gating mechanism that avoids task interference and minimizes forgetting. The kernel network f_γ^t , parameterized by γ , infers a task-specific kernel, which is shared among tasks and trained end-to-end. The shaded network q_ϕ^t is the trained encoder network from the generative model, which functions as a feature extractor to produce internal representations for kernel learning.

To alleviate the shortcomings of kernel continual learning while leveraging the merits, we introduce *generative kernel continual learning*, a memory-less variant of kernel continual learning that replaces the episodic memory unit with a generative model based on the conditional variational auto-encoder (van de Ven et al., 2020) that can learn task distributions sequentially. This change allows the kernel continual learning method to easily increase the coreset size and draw samples from each task’s data distribution based on data point likelihood. As a result, kernel approximation is enhanced, even for small sample sizes per class. We make three contributions in this paper:

1. We propose generative kernel continual learning by synergizing the strengths of generative models and kernels. It removes the dependence on an explicit memory while still being able to tackle catastrophic forgetting and task interference.
2. We introduce a supervised contrastive loss into the generative modelling, which increases the discriminability of the generated latent representations, further improving the model’s performance in terms of accuracy and forgetting.
3. We demonstrate generative kernel continual learning with the most simple linear kernel already outperforms the most advanced kernel continual learning variant with variational random features and sets a new state-of-the-art for both task-aware and domain incremental continual learning.

The schematic overview of the proposed generative kernel continual learning is depicted in Figure 1 and detailed next.

2 GENERATIVE KERNEL CONTINUAL LEARNING

A continual learning agent learns to solve a sequence of non-stationary tasks. Following the task setup in (Mirzadeh et al., 2020; Chaudhry et al., 2020; 2019b), we consider continual learning as a sequence of tasks $\{1, \dots, t, \dots, T\}$ that are proceeded one at a time, where T is the total number of tasks. Each task t is a classification problem with its own dataset $\mathcal{D}_t = (\mathbf{x}_i^t, \mathbf{y}_i^t)_{i=1}^{N_t}$ where \mathbf{x}_i^t is the i -th input vector, \mathbf{y}_i^t is its corresponding output target, and N_t refers to the number of input-output pairs regarding task t .

2.1 KERNEL CONTINUAL LEARNING

Kernel continual learning (Derakhshani et al., 2021) deploys a non-parametric classifier based on kernel ridge regression, which does not require memory replay and systematically avoids task interference by adopting a task-specific classifier. Specifically, during training task t , dataset D_t is split into two disjoint subsets C_t and $D_t \setminus C_t$. As the model, a neural network h_θ is deployed that provides the feature vector $\psi(\mathbf{x}) = h_\theta(\mathbf{x}) \in \mathbb{R}^d$ as its output. $h_\theta(\cdot)$ is shared among all tasks, and θ denotes its parameters. To learn task t , a classifier f_c^t is built upon the feature vectors extracted from $h_\theta(\cdot)$. In the probabilistic perspective, the predictive prediction is denoted as follow:

$$\tilde{\mathbf{y}}' = f_c^t(h_\theta(\mathbf{x}')) = \text{softmax}\left(Y(\lambda I + \mathcal{K})^{-1}\tilde{K}\right), \quad (1)$$

where $\mathcal{K} = \psi(X)^\top \psi(X)$, $\tilde{K} = \psi(X)^\top \psi(\mathbf{x}')$, $\psi(X) \in \mathbb{R}^{d \times N_c}$ is the feature matrix that contains the feature vectors of the coreset samples, Y is the one-hot representation of the target of the coreset samples, and $\psi(\mathbf{x}')$ denotes the features extracted from a query sample of $D_t \setminus C_t$. Finally, h_θ is trained on the current task using the cross-entropy loss function. It is further supposed that the task boundary is known during training time and inference time, and C_t is accumulated in the memory unit \mathcal{M} to be used at inference time.

The performance of kernel continual learning is highly dependent on the quality of the coreset. In general, a coreset should satisfy two properties. It should be small as well as representative of the full dataset. Ideally, a model trained on the coreset performs as good as one trained on the full dataset (Borsos et al., 2020). For kernel continual learning, Derakhshani et al. (2021) confirm these facts and show a positive correlation exists between the coreset size and the continual learning model accuracy. That is, a larger, and as a result more representative, coreset leads to a better approximation of kernel terms \mathcal{K} and \tilde{K} in Eq. (1). Moreover, kernel continual learning also adopt variational random features to learn task-specific kernels, which also contributes greatly to the overall performance. In this paper, we work with predefined kernels without inducing learnable parameters in kernels.

2.2 GENERATIVE KERNEL LEARNING

In kernel continual learning, the coreset is constructed by selecting samples from the training set of each task (Derakhshani et al., 2021). However, this coreset selection mechanism has two complications. First, it relies on a uniform sampling strategy to draw samples from the task dataset D_t , with no guarantee that the selected samples are sufficiently representative of the data in each task. Second, the coreset size is usually bounded by memory storage constraints, which limits its capacity and leads to a drop in model performance as learning proceeds with an increasing number of tasks.

In order to alleviate the shortcomings while enjoying the merits of kernel continual learning, we propose to replace the fixed size memory unit with a generative model. The generative model removes the dependence on the memory unit in kernel continual learning by generating samples to construct kernels for each task. As an extra benefit, by using generative modelling, the model is allowed to flexible generate as many samples as needed, without being restricted by memory size in kernel continual learning. To be more specific, we resort the generative replay on internal representations (van de Ven et al., 2020) based on variational auto-encoders (Kingma & Welling, 2014), as it allows kernel continual learning to draw samples proportional to their likelihood values in the coreset selection phase.

As shown in Figure 2, similar to (von Oswald et al., 2020; Wortsman et al., 2020), the decoder applies a gating mechanism which allows the model to adjust its decoder output per each task or domain to minimize the task interference problem. More specifically, each layer consists of several gates in the decoder network. They are binary, randomly pre-initialized and fixed during training, and they are multiplied to the output of each layer as (van de Ven et al., 2020). Gates are defined according to the continual learning scenarios. For the task incremental learning, the number of gates in each layer is equivalent to the number of classes while for the domain incremental learning, it is equal to the number of domains in the given benchmark. Furthermore, in contrast to the vanilla variational auto-encoder (Kingma & Welling, 2014) that uses a factorized normal distribution as its prior distribution, we exploit a learnable mixture of Gaussian distributions, in which each component corresponds to one class. Treating the prior distribution in this manner helps capture the aggregated posterior objective better (Tomczak & Welling, 2018).

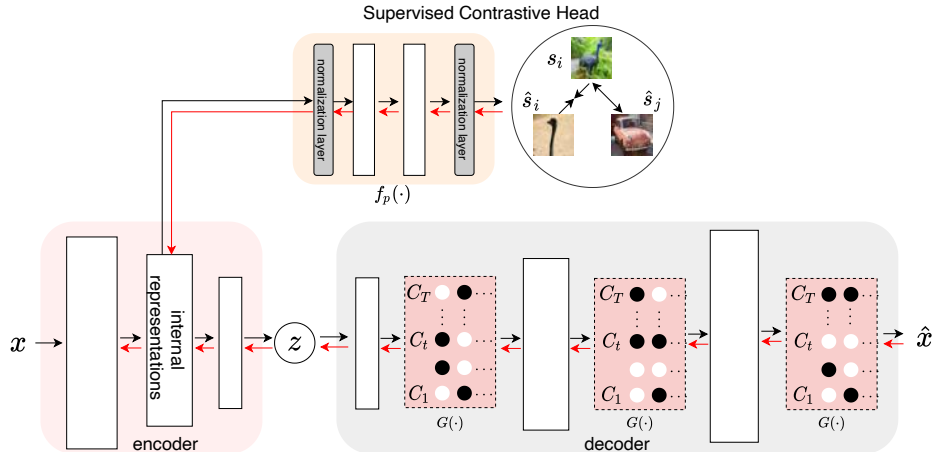


Figure 2: **Generative learning by a variational auto-encoder with supervised contrastive regularization.** The internal representations from the intermediate layers of the encoder are used to kernel learning. A projection network $f_p(\cdot)$ is applied to the internal representations, which are fed into the supervised contrastive loss function for discriminative sample generation. In the decoder network, as in (van de Ven et al., 2020), we adopt a gating function $G(\cdot)$ in red blocks to avoid task interference. Black and red arrows show the forward and backward path respectively.

Training objective Training of generative kernel continual learning involves two successive steps. In the first step, we train the variational auto-encoder to be capable of generating samples regarding the current tasks t as well as samples from previous tasks ($\mathcal{R}_{<t}$). Then, in the second step, using the decoder of the trained generative model, we produce the coreset for the current task t . Next, we put both C_t and D_t through the encoder network q_ϕ and obtain their corresponding internal representations for learning kernels. These internal representations are decoupled from the model forward path and fed into the kernel network f_γ for construct classifiers based on kernel ridge regression. Finally, we train the classifiers on top of the features extracted from the kernel network as defined in Section 2.1.

During training the variational auto-encoder, we receive in each iteration the current task data $(\mathbf{x}_i, \mathbf{y}_i, \mathbf{t}_i)$ and the replayed data from previous tasks $(\mathbf{x}_j, \mathbf{y}_j, \mathbf{t}_j)$ where \mathbf{x} , \mathbf{y} and \mathbf{t} are the input vector, output target and the task identifier. Both input vectors \mathbf{x}_i and \mathbf{x}_j are fed in the encoder network and their corresponding posterior distributions $\mathbf{z}_i \sim q_\phi(\mathbf{z}|\mathbf{x}_i)$ and $\mathbf{z}_j \sim q_\phi(\mathbf{z}|\mathbf{x}_j)$ are estimated. Then, we reconstruct the two input vectors using the decoder network as $\mathbf{x}_i \sim p_\theta(\mathbf{x}|\mathbf{z}_i, \mathbf{y}_i, \mathbf{t}_i)$ and $\mathbf{x}_j \sim p_\theta(\mathbf{x}|\mathbf{z}_j, \mathbf{y}_j, \mathbf{t}_j)$. Finally, we minimize the following overall objective function per iteration:

$$\mathcal{L}(\theta, \phi) = \underbrace{\mathcal{L}(\theta, \phi; \mathbf{x}_i, \mathbf{y}_i, \mathbf{t}_i, \mathbf{z}_i)}_{\text{current task data}} + \underbrace{\mathcal{L}(\theta, \phi; \mathbf{x}_j, \mathbf{y}_j, \mathbf{t}_j, \mathbf{z}_j)}_{\text{replayed data}}, \quad (2)$$

where each term in Eq. (2) is the following evidence lower-bound objective function:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{y}, \mathbf{t}, \mathbf{z}) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{y}, \mathbf{t})] - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}|\mathbf{y})). \quad (3)$$

In Eq. (3), we define $p(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z}|\mu_y, \sigma_y I)$ as the prior distribution for class y , which is a factorized normal distribution and its parameters μ_y and σ_y are trainable.

Our formulation for generative kernel continual learning encourages the model to produce pseudo-samples from observed tasks and generate the coreset accordingly. However, its encoder network $q_\phi(\mathbf{z}|\mathbf{x})$, which is used as the feature extractor for kernel continual learning, would not necessarily be able to provide sufficiently discriminative samples. To alleviate this problem, we propose to exploit supervised contrastive learning (Khosla et al., 2020) in generative kernel continual learning to further improve the discrimination of the latent representation.

2.3 SUPERVISED CONTRASTIVE REGULARIZATION

In order to achieve better classification performance, the generated samples are required to be discriminative. To this end, we further introduce a supervised contrastive regularizer (Khosla et al.,

2020) into the optimization of the variational auto-encoders. We impose the supervised contrastive term on the international representations that are used for kernel learning. Generally, in contrastive learning methods (Chen et al., 2020; He et al., 2020; Chen & He, 2021), a data augmentation operation is applied over the current batch data to obtain two different copies of the input batch. Both copies are fed through an encoder network, and a normalized embedding vector is estimated. At training time, these normalized vectors are further forward propagated through a projection network that is thrown away at inference time.

We adopt a similar contrastive learning strategy. Given data of the current task $(\mathbf{x}_i, \mathbf{y}_i, \mathbf{t}_i)$ and the data to replay from previous tasks $(\mathbf{x}_j, \mathbf{y}_j, \mathbf{t}_j)$, we stack them together to obtain a new set: $\mathbf{x}=[\mathbf{x}_i; \mathbf{x}_j]$ with the associated target labels $\mathbf{y}=[\mathbf{y}_i; \mathbf{y}_j]$. We consider (\mathbf{x}, \mathbf{y}) as the current batch. Then, we put \mathbf{x} into the conditional variational auto-encoder as the data augmentation operation to generate a random augmentation of the input vector as $\hat{\mathbf{x}} = \text{Aug}(\mathbf{x}) = \text{Decoder}(\text{Encoder}(\mathbf{x}))$. As for the encoder network, we utilize $q_\phi(\mathbf{z}|\mathbf{x})$ that maps the \mathbf{x} and $\hat{\mathbf{x}}$ into an internal representation vector $r \in \mathbb{R}^k$. Then, this representation vector is normalized to the unit hypersphere to enhance the top-1 accuracy as done in (Khosla et al., 2020). In the next step, we forward the normalized r through a fully connected neural network, namely, the projection network $f_p(\cdot)$, and once more normalize the projection output $\mathbf{s} \in \mathbb{R}^v$. Finally, we compute the following supervised contrastive loss function and optimize it in conjunction with the main loss function introduced in Eq. (2):

$$\mathcal{L}^s = \sum_{i \in I} \mathcal{L}_i^s, \quad (4)$$

$$\mathcal{L}_i^s = \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\mathbf{s}_i \cdot \mathbf{s}_p / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{s}_i \cdot \mathbf{s}_a / \tau)}, \quad (5)$$

where i is the index of all existing samples $\{1 \cdots 2N\}$ including samples in both \mathbf{x} and $\hat{\mathbf{x}}$, $A(i)=I \setminus \{i\}$, $P(i)=\{p \in A(i) : y_p=y_i\}$ is the set of all indices having the same class label as i , $|P(i)|$ is the set cardinality, and τ is the temperature parameter.

3 EXPERIMENTS

3.1 BENCHMARKS, METRICS & DETAILS

Three benchmarks We evaluate generative kernel continual learning for two well-established scenarios: *domain incremental learning* and *task incremental learning*. For the domain incremental learning scenario, we rely on PermutedMNIST (Kirkpatrick et al., 2017) and Rotated-MNIST (Mirzadeh et al., 2020), and for the task of incremental learning, we report on Split-CIFAR100 (Zenke et al., 2017). Each PermutedMNIST task is produced by a random permutation of the image pixels, such that this permutation remains fixed within the same task. Each Rotated-MNIST task is composed by randomly rotating the input images by a degree between 0 and 180, such that the rotation degree is the same within a task. SplitCIFAR100 is a continual variant of CIFAR100 where each task represents the data from 5 out of 100 random classes.

Two metrics We report the continual learning performance using average accuracy a_t and average forgetting F , following (Titsias et al., 2020; Mirzadeh et al., 2020; Chaudhry et al., 2020; Derakhshani et al., 2021). The average accuracy of a model when the training of task t is finished equals:

$$A_t = \frac{1}{t} \sum_{i=1}^t a_{t,i}, \quad (6)$$

where $a_{t,i}$ refers to the accuracy of the model on task i after being trained on task t . Average forgetting measures the drop in model performance between the highest accuracy and the last accuracy of each task when the continual learning model training is finished. It is calculated as:

$$F = \frac{1}{T-1} \sum_{i=1}^{T-1} \max_{1, \dots, T-1} (a_{t,i} - a_{T,i}). \quad (7)$$




Figure 3: **Generative coreset vs. uniform coreset** in kernel continual learning on SplitCIFAR100. Generative kernel continual learning needs about 10 times smaller coreset size for same accuracy, highlighting the expressiveness of the generated coreset.

Implementation details Generative kernel continual learning is implemented with a deep neural network composed of three networks: a variational auto-encoder, which itself has an encoder network $q_\phi(\mathbf{z}|\mathbf{x})$ and a decoder network $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{y}, \mathbf{t})$, a projection network $f_p(\cdot)$, and a fully-connected network f_γ for kernel learning. For PermutedMNIST and RotatedMNIST, the encoder $q_\phi(\mathbf{z}|\mathbf{x})$ consists of two fully connected layers, each having 2000 neurons. For SplitCIFAR100, similar to van de Ven et al. (2020), a pre-trained convolutional neural network precedes the encoder $q_\phi(\mathbf{z}|\mathbf{x})$. We pre-train this convolutional neural network on the classification task of CIFAR10 and only use its first five convolutional layers. For PermutedMNIST and RotatedMNIST, the pretrained network is an identity mapping function. The encoder $q_\phi(\mathbf{z}|\mathbf{x})$ consists of two fully connected layers, each of which has 256 neurons. For all benchmarks, the latent distribution of the variational auto-encoder consists of one fully connected layer of 100 neurons parameterizing the mean and log standard deviation. The decoder $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{y}, \mathbf{t})$ is exactly the transpose of the encoder $q_\phi(\mathbf{z}|\mathbf{x})$. With regard to projection network $f_p(\cdot)$, we use one fully-connected layer with hidden size of 196 neurons for PermutedMNIST and RotatedMNIST, and 256 neurons for SplitCIFAR100. Finally, f_γ includes two fully connected layers with 512 neurons, followed by a dropout layer for all benchmarks. We implement generative kernel continual learning in Pytorch and the code will be released.

In order to make fair comparisons with other models, we train f_γ with the same hyperparameters as (Derakhshani et al., 2021). Unless otherwise stated, all models in our experiments rely on generative kernel continual learning with supervised contrastive regularization, trained over 20 sequential tasks on 5 different random seeds. The coreset sizes for both training and inference are similar and all experiments exploit a simple linear kernel. We train the conditional generative model for 2000 iterations on PermutedMNIST and RotatedMNIST and 300 iterations on SplitCIFAR100 per task. Following (Derakhshani et al., 2021), the kernel network f_γ is trained for one epoch per task and the batch size is set to 10 (see appendix A.1 for more details).

3.2 RESULTS

Generative coreset vs. uniform coreset The number of samples per class as well as the quality of each sample in the coreset are considered to be two pivotal factors in obtaining good performance in kernel continual learning. To understand them better, we ablate the influence of the uniform coreset in kernel continual learning with the generative coreset in our approach using a varying number of samples per class in the coreset. As Figure 3 reveals, generative kernel continual learning performs better than kernel continual learning, independent of the number of samples per class. Interestingly, the difference between the average accuracy of the two methods increases in favor of our generative approach as the number of samples per class are decreased. Generative kernel continual learning only needs to generate 2 samples per class to achieve a similar accuracy as kernel continual learning with 20 samples. This shows that the generative model endows kernel continual learning with more representative and expressive coreset samples. The benefit increases when only a small number of samples per class can be stored. A smaller number of samples reduces the number of operations needed to compute kernels in equation 1, and as a result, it reduces the run time of generative kernel continual learning during both training and inference.

Varying coreset size at inference time Due to the generative nature of our proposal, the coreset size no longer needs to be the same during training and inference time, as in kernel continual learning. We run an ablation to show how average accuracy of our model changes by fixing the coreset size during training and varying the coreset size during inference. To do so, we train generative




Figure 4: **Varying coreset size at inference time.** Average accuracy over 20 sequential tasks on SplitCIFAR100 when the coreset size differs between training (legend) and inference time. Increasing the coreset size during inference increases average accuracy considerably.

kernel continual learning with three different coreset sizes where we generate 1, 2 and 5 samples per class as depicted in Figure 4. For each coreset, we evaluate the trained model on seven different coreset sizes at inference time. As expected, model accuracy decreases when the coreset size at inference time is smaller than the size used during training time. That is to say, as we decrease the coreset size during inference below the coreset size during training time, the kernel network f_γ receives less information than it expects to observe, and consequently it does poorly on the kernel approximation. However, enlarging the coreset size at inference time, beyond the size used during training, always leads to an improvement. This behaviour shows the potency of kernel network f_γ to utilize the extra information to enhance the model accuracy. Hence, we suggest to train generative kernel continual learning with a smaller coreset size to reduce the training time and evaluate the same model on a larger coreset size for better model accuracy.

Influence of supervised contrastive regularization We introduce a supervised contrastive regularizer in generative kernel continual learning to enhance the discriminative power of the encoder’s latent representation. To measure its benefit, we ablate generative kernel continual with and without the supervised contrastive regularization term in left Table 1, where we report each experiment on five different random seeds. As expected, adding the supervised contrastive regularizer leads to a consistent improvement in accuracy on all three benchmarks. In right Table 1, we further experiment with temperature parameter τ in Eq. (4) on SplitCIFAR100 over 1 random seed, where hyperparameter τ adjusts the compactness and concentration of positive and negative samples in contrastive learning. It shows that our method is hardly sensitive to this hyperparameter and we consider fine-tuning optional.

Ability to generate multiple kernel types To determine how the kernel type changes the generative kernel continual learning accuracy, we train the proposed method with three different kernel types: linear, radial basis function and polynomial kernels, on all three benchmarks. To provide fair comparisons, we keep all hyperparameters fixed, and only change the kernel type. Table 2 shows the results. Generative kernel continual learning provides promising results, independent of the kernel type. As could be expected, the best kernel varies per benchmark. The linear kernel and radial basis function kernels perform better for RotatedMNIST and PermutedMNIST, while for SplitCIFAR100 the polynomial kernel provides best average accuracy. By and large, a simple linear kernel is a good baseline, but the kernel type could be taken into account during hyperparameter optimization when accuracy is critical.

Comparison with state-of-the-art To compare our proposed method against the state-of-the-art, we report average accuracy and average forgetting of generative kernel continual learning with su-

Table 1: **Influence of supervised contrastive regularization** on generative kernel continual learning over PermutedMNIST, RotatedMNIST and SplitCIFAR100. Left: the contrastive regularizer further enhances continual learning average accuracy. Right: optimizing temperature parameter τ in Eq. (4) results in a small gain on SplitCIFAR100.

	RotatedMNIST	PermutedMNIST	SplitCIFAR100	SplitCIFAR100					
Without regularizer	79.80	88.23	71.87	Temperature (τ)	0.02	0.04	0.08	0.1	1
With regularizer	82.48	89.23	72.79	Average Accuracy	73.70	74.25	73.70	73.78	74.10

Table 2: **Ability to generate multiple kernel types.** Results are promising, independent of kernel type.

	RotatedMNIST	PermutedMNIST	SplitCIFAR100
Linear	82.48	89.23	72.79
Polynomial	82.08	89.23	75.33
Radial Basis Function	81.42	89.59	72.68

Table 3: **Comparison with state-of-the-art.** Results for other methods are adopted from Derakhshani et al. (2021). Column *unit* indicates whether methods exploit a memory unit \mathcal{M} or a generative model \mathcal{G} . Generative kernel continual learning yields competitive results on PermutedMNIST and RotatedMNIST and outperforms alternatives on SplitCIFAR100 by a large margin.

Method	Unit	Permuted MNIST		Rotated MNIST		Split CIFAR100	
		Accuracy	Forgetting	Accuracy	Forgetting	Accuracy	Forgetting
Naive-SGD (Mirzadeh et al., 2020)	\times	44.4 \pm 2.46	0.53 \pm 0.03	46.3 \pm 1.37	0.52 \pm 0.01	40.4 \pm 2.83	0.31 \pm 0.02
EWC (Kirkpatrick et al., 2017)	\times	70.7 \pm 1.74	0.23 \pm 0.01	48.5 \pm 1.24	0.48 \pm 0.01	42.7 \pm 1.89	0.28 \pm 0.03
AGEM (Chaudhry et al., 2019a)	\mathcal{M}	65.7 \pm 0.51	0.29 \pm 0.01	55.3 \pm 1.47	0.42 \pm 0.01	50.7 \pm 2.32	0.19 \pm 0.04
ER-Reservoir (Chaudhry et al., 2019b)	\mathcal{M}	72.4 \pm 0.42	0.16 \pm 0.01	69.2 \pm 1.10	0.21 \pm 0.01	46.9 \pm 0.76	0.21 \pm 0.03
Stable SGD (Mirzadeh et al., 2020)	\times	80.1 \pm 0.51	0.09 \pm 0.01	70.8 \pm 0.78	0.10 \pm 0.02	59.9 \pm 1.81	0.08 \pm 0.01
Kernel Continual Learning (Derakhshani et al., 2021)	\mathcal{M}	85.5 \pm 0.78	0.02\pm0.00	81.8 \pm 0.60	0.01 \pm 0.00	62.7 \pm 0.89	0.06 \pm 0.01
<i>Generative Kernel Continual Learning</i>	\mathcal{G}	89.2\pm0.44	0.08 \pm 0.00	82.4\pm1.33	0.01\pm0.01	72.8\pm0.68	0.04\pm0.00

pervised contrastive regularization for 20 sequential tasks on 5 different random seeds over three well-established continual learning benchmarks in Table 3. To provide a fair comparison with Derakhshani et al. (2021), the coreset size at both inference and training time is similar and equal to 20 for all benchmarks. Moreover, we exploit the simple linear kernel for all benchmarks. The temperature hyperparameter τ is 0.08 for PermutedMNIST and RotatedMNIST and is 0.04 for SplitCIFAR100. As shown, generative kernel continual learning outperforms all methods, setting a new state-of-the-art on PermutedMNIST, RotatedMNIST and SplitCIFAR100. In Figure 5, we visualize and compare our proposed model with alternative methods in terms of running average accuracy for 20 sequential tasks. It can be seen that generative kernel continual learning performs consistently better than other methods on PermutedMNIST, RotatedMNIST and, especially, on SplitCIFAR100. We attribute the 10.1% accuracy improvement on SplitCIFAR100 to the quality of the coreset samples provided by our generative model. For SplitCIFAR100, we also train generative kernel continual learning with a polynomial kernel, a coreset size of 5 samples per class during training, and a coreset size of 80 samples per class during inference. We keep all other hyperparameters fixed. In this setting, we obtain an average accuracy of **76.5 \pm 0.35** and an average forgetting of **0.03 \pm 0.00**.

4 RELATED WORK

Following Lange et al. (2019), we divide continual learning methodologies into three different categories. We have regularisation-based methods (Kirkpatrick et al., 2017; Aljundi et al., 2018; Lee et al., 2017; Zenke et al., 2017; Kolouri et al., 2019) that regularize the neural network parameters to not change drastically from those learned on previous tasks. This goal is accomplished by estimating a penalty term for each parameter of the network using the Fisher information matrix Kirkpatrick et al. (2017), the gradient magnitude of each parameter Aljundi et al. (2018), or by sequential Bayesian inference Nguyen et al. (2018). Recently, Kapoor et al. (2021) proposes to exploit a variational auto-regressive Gaussian processes to improve the posterior distribution of sequential Bayesian inference methods due to the sequential nature of continual learning data. In general, these methods encourage the continual learning model to prioritize preserving old knowledge rather than absorbing new information from new tasks. Hence, in contrast to our proposal, these methods fail to scale-up for longer task sequences.

In the second category, we have replay/rehearsal based methods that attempt to simultaneously re-train the continual learning model over the previous tasks data and the current task data to avoid catastrophic forgetting (Lopez-Paz & Ranzato, 2017; Riemer et al., 2019; Rios & Itti, 2018; Shin et al., 2017; Zhang et al., 2019; Rebuffi et al., 2017; Chaudhry et al., 2019b;a). Obtaining samples or knowledge from previous tasks is usually performed in two different ways: (i) adding a generative model to the continual learning model and producing samples of earlier tasks Shin et al. (2017), (ii) augmenting a memory unit to the continual learning model and accumulating a small subset of raw input data Chaudhry et al. (2019b) or gradient parameters Chaudhry et al. (2019a). Recently,




Figure 5: **Comparison with state-of-the-art** over 20 consecutive tasks, in terms of average accuracy. Our model consistently outperforms alternatives on all three benchmarks, especially on the more challenging SplitCIFAR100.

Saha et al. (2021) propose a memory-based method where a new task is learned by taking gradient steps in the orthogonal direction to the gradient subspaces marked as crucial for previous past tasks. The method employs SVD after learning each task to find the crucial subspaces and stores them in a memory. In our proposed method, rather than using a memory unit, we augment the kernel continual learning model by Derakhshani et al. (2021) with a conditional generative model to enable generation of samples for each task.

The third category covers architecture-based methods (Rusu et al., 2016; Yoon et al., 2018; Jerfel et al., 2019; Li et al., 2019; Wortsman et al., 2020; Mallya & Lazebnik, 2018). These methods aim to directly minimize the task interference problem by either pruning and model expansion Mallya & Lazebnik (2018), or allocating a set of new parameters when observing a new task Rusu et al. (2016), or by partitioning a neural network into several sub-networks using gating mechanism Wortsman et al. (2020). Recently, Kumar et al. (2021) presents a Bayesian framework to learn the structure of deep neural networks by unifying the variational Bayes based regularization and architecture based methods. This method supports knowledge transfer among tasks by overlapping sparse subsets of weights learned by different tasks. Using an expectation maximization method, Lee et al. (2021) introduces a transfer mechanism that selectively chooses the transfer architecture configuration for each task. For each task, this would allow the method to dynamically select which layers to transfer and which to keep as task-specific. Veniat et al. (2021) introduce a compositional neural architecture for continual learning, where each module in the network represents an atomic skill and can be combined to solve a certain task. Similarly, our generative kernel continual learning adopts a gating mechanism in the decoder of the conditional variational auto-encoder to minimize task interference.

5 CONCLUSION

In this paper, we introduce generative kernel continual learning, a memory-less variant of kernel continual learning that replaces the episodic memory with a generative model based on variational auto-encoder. We further introduce supervised contrastive regularization, which enables our model to generate even more discriminative samples for better classification performance. We conduct extensive experiments on three benchmarks for continual learning. Our experiments highlight the effectiveness of generative kernel continual learning. First, it is shown that synergizing the strengths of generative models and kernels leads to remove the dependence on an explicit memory while being able to tackle catastrophic forgetting and task interference. Second, it is demonstrated that adding a supervised contrastive loss into the generative modelling increases the discriminability of the generated latent representations, improving the model’s performance in terms of accuracy and forgetting. Moreover, we show that our generative kernel continual learning already achieves state-of-the-art performance on all benchmarks with a simple linear kernel.

ETHICS STATEMENT

Being able to adapt to non-stationary data distributions and continuously changing environments, our method has potential inherent impact in the applications that often encounter dynamic environments in practice, e.g., medical imaging, astronomical imaging, and autonomous driving. Accordingly, our method would also potentially face some negative social impacts accompanying with applications, e.g., lack of fairness with the model trained by incomplete data, legal compliance, and the privacy of patients in medical imaging.

REPRODUCIBILITY STATEMENT

We refer to Section 3.1 for detailed information on benchmarks, metrics, and the implementation of generative kernel continual learning in terms of architecture and training. In addition, we refer to Appendix A.1 for the list of all hyperparameters used to train generative kernel continual learning. We will further open-source all code, scripts to reproduce the experiments with the exact hyperparameters, and scripts to calculate the evaluation results at: <https://github.com/<redacted>/<redacted>>.

REFERENCES

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *European Conference on Computer Vision*, 2018.
- Zalán Borsos, Mojmír Mutný, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. *Advances in Neural Information Processing Systems*, 2020.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with A-GEM. In *International Conference on Learning Representations*, 2019a.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K. Dokania, Philip H. S. Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. In *Advances in Neural Information Processing Systems*, 2019b.
- Arslan Chaudhry, Naeemullah Khan, Puneet K Dokania, and Philip HS Torr. Continual learning in low-rank orthogonal subspaces. In *Advances in Neural Information Processing System*, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 2020.
- Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- Mohammad Mahdi Derakhshani, Xiantong Zhen, Ling Shao, and Cees Snoek. Kernel continual learning. In *International Conference on Machine Learning*, 2021.
- Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradientbased neural networks. In *International Conference on Learning Representations*, 2014.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Ghassen Jerfel, Erin Grant, Thomas L. Griffiths, and Katherine A. Heller. Reconciling meta-learning and continual learning with online mixtures of tasks. In *Advances in Neural Information Processing Systems*, 2019.
- Sanyam Kapoor, Theofanis Karaletsos, and Thang D Bui. Variational auto-regressive gaussian processes for continual learning. In *International Conference on Machine Learning*, 2021.

- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 2020.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.
- Soheil Kolouri, Nicholas Ketz, Xinyun Zou, Jeffrey Krichmar, and Praveen Pilly. Attention-based structural-plasticity. *arXiv preprint arXiv:1903.06070*, 2019.
- Abhishek Kumar, Sunabha Chatterjee, and Piyush Rai. Bayesian structural adaptation for continual learning. In *International Conference on Machine Learning*, 2021.
- Matthias Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ale Leonardis, Gregory G. Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv preprint arXiv:1909.08383*, 2019.
- Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. In *Advances in Neural Information Processing Systems*, 2017.
- Seungwon Lee, Sima Behpour, and Eric Eaton. Sharing less is more: Lifelong learning in deep networks with selective layer transfer. In *International Conference on Machine Learning*, 2021.
- Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, 2019.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017.
- Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Academic Press*, 1989.
- Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh. Understanding the role of training regimes in continual learning. In *Advances in Neural Information Processing Systems*, 2020.
- Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.
- Sylvestre-Alvise Rebuffi, Alexander I Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *International Conference on Learning Representations*, 2019.
- Mark B Ring. Child: A first step towards continual learning. *Learning to learn*, 1998.
- Amanda Rios and Laurent Itti. Closed-loop gan for continual learning. In *International Joint Conference on Artificial Intelligence*, 2018.

- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. In *Advances in Neural Information Processing Systems*, 2016.
- Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. In *International Conference on Learning Representations*, 2021.
- Bernhard Schölkopf and Alex J Smola. *Learning with kernels*. MIT Press, 2002.
- Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola. A generalized representer theorem. In *International Conference on Computational Learning Theory*, 2001.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, 2017.
- Michalis K Titsias, Jonathan Schwarz, Alexander G de G Matthews, Razvan Pascanu, and Yee Whye Teh. Functional regularisation for continual learning using gaussian processes. In *International Conference on Learning Representations*, 2020.
- Jakub Tomczak and Max Welling. VAE with a VampPrior. In *International Conference on Artificial Intelligence and Statistics*, 2018.
- Gido M van de Ven, Hava T Siegelmann, and Andreas S Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 2020.
- Tom Veniat, Ludovic Denoyer, and MarcAurelio Ranzato. Efficient continual learning with modular networks and task-driven priors. In *International Conference on Learning Representations*, 2021.
- Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. *International Conference on Learning Representations*, 2020.
- Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. In *Advances in Neural Information Processing Systems*, 2020.
- Jaehong Yoon, Eunho Yang, Jungtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations*, 2018.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017.
- Mengmi Zhang, Tao Wang, Joo Hwee Lim, and Jiashi Feng. Prototype reminding for continual learning. *arXiv preprint arXiv:1905.09447*, 2019.

A APPENDIX

In section A.1, we report all hyperparameters used to reproduce the results in Table 3 and Figure 5. In section A.2, we study the performance of generative kernel continual learning with and without kernel network f_γ . Finally, In section A.3, we visualize the internal representation of conditional variational auto-encoder, which is used as the input for the kernel network f_γ , with and without contrastive regularization term.

A.1 HYPERPARAMETERS

In Table 4 and 5, we report all hyperparameters used to generate results in Table 3 and Figure 5 in the main paper.

Table 4: Hyperparameters used to train the kernel network f_γ .

	RotatedMNIST	PermutedMNIST	SplitCIFAR100
Batch size	10	10	10
Learning rate (LR)	0.1	0.1	0.3
LR decay factor	0.8	0.8	0.95
Momentum	0.8	0.8	0.4
Dropout	0.1	0.1	0.0
Coreset size	20	20	20
Kerenl type	Linear	Linear	Linear
Optimizer	SGD	SGD	SGD

Table 5: Hyperparameters used to train the conditionl variational auto-encoder.

	RotatedMNIST	PermutedMNIST	SplitCIFAR100
Learning Rate	0.001	0.001	0.001
Batch size	512	512	512
Replay size	512	512	512
Number of iteration	2000	2000	300
Optimizer	Adam	Adam	Adam
Temperature (τ)	0.08	0.08	0.04

A.2 INFLUENCE OF KERNEL NETWORK

In this section, we study the performance of generative kernel continual learning with and without kernel network f_γ . This network maps the internal representation of conditional variational auto-encoder to another new space, and upon this new space, we construct each task’s non-parametric classifier. For fair comparison, in both experiments, we exploit same hyperparameters and architecture as provided in Tables 4 and 5. Moreover, we train both models for 20 sequential tasks on SplitCIFAR100 benchmark, and report the average accuracy over 5 different random seeds. Results are presented in Table 6. As it is shown, augmenting generative kernel continual learning with kernel network f_γ enhances the model performance by 12%.

Table 6: **Influence of kernel network f_γ .** Average accuracy of our proposed method with and without kernel network over 20 sequential tasks for five different random seeds on SplitCIFAR100. Incorporating kernel network into the conditional variational auto-encoder increases average accuracy considerably.

	SplitCIFAR100
with kernel network	72.79
without kernel network	59.68

A.3 FEATURE SPACE VISUALIZATION

To highlight the benefit of the supervised contrastive regularization term in enhancing the discriminability of the generated internal representations, we further visualize the internal representation of the encoder network q_ϕ on the generated coreset of task 1 in Figure 6 where we train the generative kernel continual learning model over 20 sequential tasks. In this figure, the first row shows the scenario where we train generative kernel continual learning without supervised contrastive regularization term while the second row is the case where we exploit supervised contrastive regularization. Comparing these two scenarios shows that the regularization term leads to improve the discriminability of internal representation of the conditional variational auto-encoder. To explore more, we perform similar experiments over the test dataset in Figure 7. Same conclusion as coreset is inferred.




Figure 6: **Coreset internal representation.** In this figure, we visualize the internal representation of the encoder network of conditional variational auto-encoder of the first task when we observe 20 tasks with (w/) and without (w/o) the contrastive regularization term. As shown, the regularization term allows our proposed method to obtain better concentrated features.

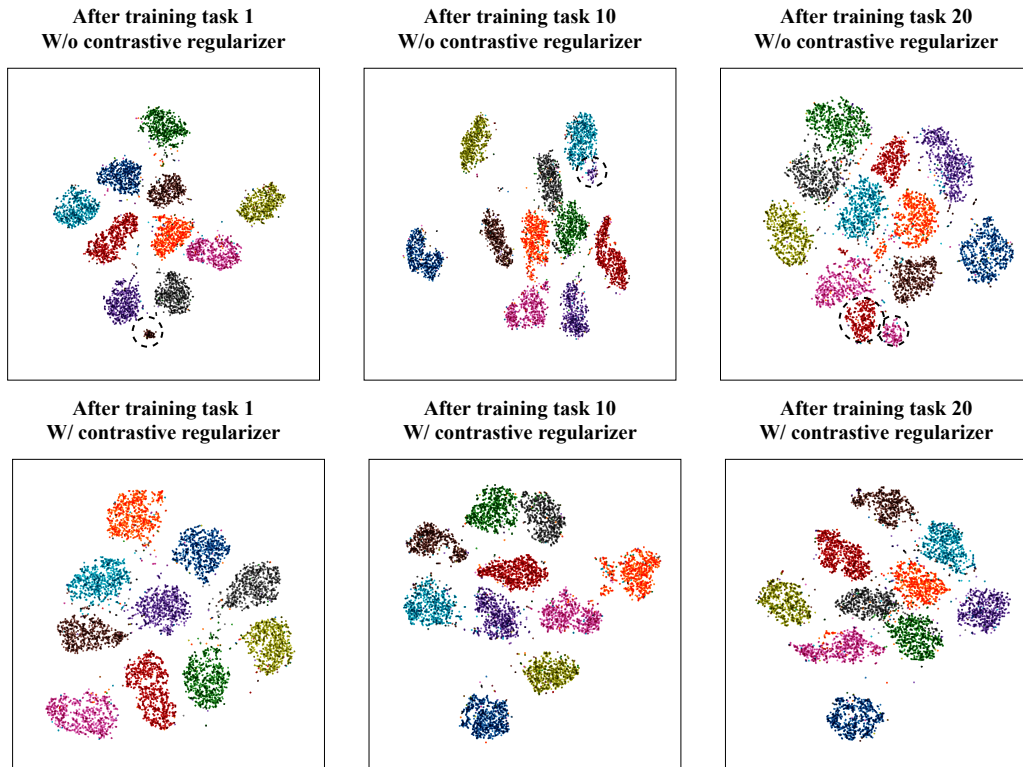


Figure 7: **Test dataset internal representation.** In this figure, we visualize the internal representation of the encoder network of conditional variational auto-encoder of the first task when we observe 20 tasks with (w/) and without (w/o) the contrastive regularization term. As shown, the regularization term allows our proposed method to obtain better concentrated features.