

Recall-Oriented Continual Learning with Generative Adversarial Meta-Model

Haneol Kang, Dong-Wan Choi*

Department of Computer Science and Engineering, Inha University, South Korea
haneol0415@gmail.com, dchoi@inha.ac.kr

Abstract

The *stability-plasticity dilemma* is a major challenge in continual learning, as it involves balancing the conflicting objectives of maintaining performance on previous tasks while learning new tasks. In this paper, we propose the *recall-oriented continual learning framework* to address this challenge. Inspired by the human brain’s ability to separate the mechanisms responsible for stability and plasticity, our framework consists of a two-level architecture where an inference network effectively acquires new knowledge and a generative network recalls past knowledge when necessary. In particular, to maximize the stability of past knowledge, we investigate the complexity of knowledge depending on different representations, and thereby introducing *generative adversarial meta-model* (GAMM) that incrementally learns task-specific parameters instead of input data samples of the task. Through our experiments, we show that our framework not only effectively learns new knowledge without any disruption but also achieves high stability of previous knowledge in both task-aware and task-agnostic learning scenarios. Our code is available at: <https://github.com/bigdata-inha/recall-oriented-cl-framework>.

Introduction

The ability to continuously acquire new knowledge without forgetting previously learned information is a major challenge in continual learning (CL) over deep neural networks. This challenge, often referred to as the *stability-plasticity dilemma* (Parisi et al. 2019), requires balancing two opposing objectives: preserving past knowledge (stability) while adapting to new tasks (plasticity). On the one hand, if we focus on stability by preventing the model’s state from undergoing significant changes, the model’s plasticity to acquire new information is impaired. On the other hand, making the model highly adaptable to incoming tasks leads to stability degradation on previous tasks, known as *catastrophic forgetting* (McCloskey and Cohen 1989). Despite considerable progress of recent studies in CL, the trade-off between learning new tasks and maintaining old knowledge remains somewhat inevitable. As shown in Figure 1, methods with high plasticity experience significant forgetting, while those with less forgetting lack the ability to adapt to new tasks.

*Corresponding Author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Trade-off between stability and plasticity in various CL baseline approaches on Split CIFAR-100.

In contrast, the human brain gracefully resolves the dilemma by separating yet complementing the parts responsible for stability and plasticity. More specifically, working memory (*a.k.a.* short-term memory) focuses on processing new information, whereas long-term memory is responsible for retaining and consolidating the information that is not currently used but may be needed in the future (McClelland, McNaughton, and O’Reilly 1995; O’Reilly et al. 2014). More interestingly, when we recall past knowledge, it is known that our brain creates something new in a generative manner from long-term memory, rather than simply retrieving something that explicitly exists in the brain (Schacter and Addis 2007).

Existing brain-inspired CL methods also consider a certain level of separation of new knowledge from previous knowledge, but mostly focus on the *rehearsal* process of human memory. This has led to the development of techniques like *memory replay* (Lopez-Paz and Ranzato 2017; Rebuffi et al. 2017) and *generative replay* (Kemker and Kanan 2018; Shin et al. 2017). Memory replay stores representative old samples in an extra buffer, while generative replay trains a generative model that can generate pseudo-samples of previous tasks. Then, in both techniques, previous samples or pseudo-samples are jointly trained with new data in a single neural network. However, due to limited memory and the complexity of underlying data distribution, neither approach can fully capture the entire knowledge of previous tasks, which results in less satisfactory performance.

In this paper, we propose the *recall-oriented CL frame-*

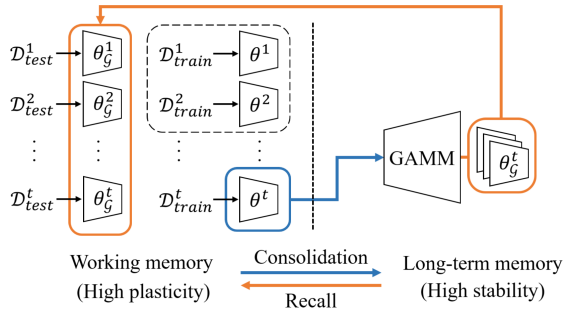


Figure 2: Separation of working memory and long-term memory in our framework.

work, where we divide learning process into two different networks corresponding to the human brain’s working memory and long-term memory, as illustrated in Figure 2. Firstly, a typical neural network takes the role of working memory, and freely learns new knowledge without any disruption to achieve pure plasticity. Secondly, inspired by the generative way of the human’s recall process (Schacter and Addis 2007), we introduce *generative adversarial meta-model* (GAMM) to be our counterpart of long-term memory. GAMM is a generative model that trains over model parameters rather than data samples, which thus can directly *recall* a previous task-specific model at inference time. The GAMM’s approach stems from our analysis on the complexity of knowledge based on different representations, assessed through two metrics: separability and volume. Our analysis reveals that the complexity of raw data is significantly higher than that of the corresponding learned parameters, implying that a trained model itself is a better form of knowledge to be accumulated in long-term memory than its training data distribution.

The major challenge in our framework is how to train GAMM in an efficient and scalable manner. One issue is the need of numerous versions of learned parameters for each task, as a typical generative model takes many samples to be trained. A possible approach can be training the task-specific model multiple times with different initialized parameters or different subsets of training data (Joseph and Balasubramanian 2020). However, this approach of redundant training is not only inefficient (Ratzlaff and Li 2019), but also less effective in acquiring the knowledge of new tasks, often requiring fine-tuning at inference time (Joseph and Balasubramanian 2020). Our strategy is to use *Bayesian neural network* (BNN) (Blundell et al. 2015; Maddox et al. 2019) as a task-specific model without redundant training, and GAMM takes a sufficient number of different models for the same task from the task-specific BNN. In terms of scalability, it is also important for GAMM not to use too large memory space to accumulate the learned knowledge. To this end, we observe that each individual task may not require a large capacity for plasticity, and thus we keep each task-specific model as lightweight as possible. This allows GAMM to incorporate many task-specific parameters within its limited capacity.

In our experiments, we show that our recall-oriented CL

framework nicely addresses the stability-plasticity dilemma with a small amount of memory. Compared to existing replay-based methods, our framework achieves the best performance while mostly consuming less memory space in both task-aware and task-agnostic CL scenarios. Also, our two-level architecture is observed to be highly effective at plasticity, and therefore outperforms the existing CL baseline approaches.

Related Works

Continual Learning for Less Forgetting. Existing methods of continual learning in neural networks can be grouped into three main categories: *regularization-based*, *architecture-based*, and *replay-based methods*. Regularization-based methods (Ahn et al. 2019; Ebrahimi et al. 2020a; Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017) more focus on stability by penalizing updates to important parameters, resulting in lower plasticity. On the other hand, architecture-based methods aim to adaptively utilize different parts of the network for each task by masking and attentions (Mallya and Lazebnik 2018; Serrà et al. 2018), or by dynamically expanding the network for incoming tasks (Ostapenko et al. 2019; Yoon et al. 2018). Replay-based methods attempt to mimic the human rehearsal mechanism by storing previous knowledge, typically in the form of exemplars (Lopez-Paz and Ranzato 2017; Rebuffi et al. 2017; Saha, Garg, and Roy 2021; Buzzega et al. 2020; Lyu et al. 2021) or as a generative model (Kemker and Kanan 2018; Shin et al. 2017; van de Ven, Siegelmann, and Tolias 2020). Both suffer from the stability-plasticity dilemma due to limited memory and complex data distributions.

Parameter Generation. A *parameter-generation* model (a.k.a. *meta-model*) is a type of model that is trained to generate parameters, rather than data samples. One popular framework is *hypernetwork* (Ha, Dai, and Le 2017), in which a target network is created by a deterministic hypernetwork and is not independently trained. Hypernetworks and their variants (Krueger et al. 2017; Ratzlaff and Li 2019) have been applied to various tasks including few-shot learning (Zhmoginov, Sandler, and Vladymyrov 2022), uncertainty estimation (Ratzlaff and Li 2019), and continual learning (von Oswald et al. 2020). Our parameter-generation scheme is different from hypernetworks in that we train a *generative adversarial network* (GAN) (Goodfellow et al. 2014) over multiple task-specific models, each separately trained. BNNs (Blundell et al. 2015; Maddox et al. 2019) can also be seen as parameter-generation models in that they treat network parameters as random variables and learn their distribution. Our framework trains GAMM by taking multiple parameters sampled from a single task-specific BNN. Although APD (Wang et al. 2018) also suggests training a GAN based on BNNs, they focus on how to replace MCMC-based BNNs requiring excessive training with a corresponding GAN, not considering any scenario where multiple tasks must be sequentially learned within a single GAN.

Continual Learning with Parameter Generation. A few studies have investigated continual learning using parameter-generation models. Most existing methods rely

on hypernetworks, where knowledge for each new task is not separately learned but directly consolidated into the meta-model (Henning et al. 2021; von Oswald et al. 2020). This one-phase training scheme seems excellent at retaining stability on previous tasks, achieving almost zero-forgetting, but it sacrifices plasticity as demonstrated by our experimental results. In addition, hypernetwork-based methods exhibit limited performance particularly when task identities are not given, likely due to the failure to adequately separate task-specific knowledge in the hypernetwork. Without using hypernetworks, our framework is similar to MERLIN (Joseph and Balasubramanian 2020) in that both employ non-deterministic meta-models, a *variational autoencoder* (VAE) (Kingma and Welling 2013) for MERLIN and a GAN for ours. However, MERLIN is designed for on-line learning scenarios and thus its meta-model is trained to generate rather small-sized networks (*e.g.*, an order of magnitude smaller than those of both hypernetworks and ours). This consequently requires fine-tuning at inference time with an extra memory buffer, as their goal is to generate well-initialized models. In contrast, our framework is capable of generating high-end models that are ready to use without fine-tuning.

Analysis on Representation Complexity

In a two-level CL approach, such as ours, where old and new knowledge are somehow separately maintained, a crucial question is what would be the best representation to preserve the learned knowledge. For example, in replay-based approaches, the previous knowledge is not only embedded in the inference model being trained but also separately stored in a buffer of exemplars or in a generative model, both of which however suffer from the stability-plasticity dilemma as mentioned earlier. To examine the above question, this section investigates how to quantify the complexity of knowledge depending on different representations, and introduces two metrics applicable to heterogeneous representations in different structures and dimensions. Finally, we specifically analyze and compare the complexity of raw data samples, feature vectors, and learned parameters.

Representation Complexity Metrics

In order to quantify the complexity of representations for a given training dataset, we take into account the following two perspectives, namely *volume* (\mathcal{V}) and *separability* (\mathcal{S}).

Volume. The volume is to measure the diversity of knowledge representations. Thus, a larger volume indicates the need for a larger capacity to learn diverse cases. To this end, we suggest measuring how widely representation points are distributed within each local group (*e.g.*, class), not across the entire representation space. This is based on the intuition that points belonging to different groups are easier to discriminate and thus simpler to learn when they are farther apart, which however undesirably lead to a wider distribution overall. Thus, in terms of the difficulty of learning, diversity needs to be measured locally, rather than globally across the entire space. In addition, the volume should not be dependent on dimensionality, and therefore dimension-

ality reduction is essential to compare representations with different dimensions. Based on these criteria, we define the volume (\mathcal{V}) as follows:

$$\mathcal{V} = \sum_{c=1}^C \det(\Sigma_c)^{1/2} = \sum_{c=1}^C \prod_{i=1}^d \lambda_i^c, \quad (1)$$

where C is the number of local groups (*e.g.*, number of classes), Σ_c is the covariance matrix for group $c \in [1, C]$, and λ_i^c is its i -th eigenvalue. Thus, \mathcal{V} is proportional to the square root of the determinant of its covariance matrix, which is the product of eigenvalues λ_i^c . As d is equally applied to representations with different dimensions, the volume is independent of dimensionality. It is also noteworthy that the more the local groups there are in the representation space, the larger the corresponding volume is, and thereby capturing the global diversity as well as the local diversity.

Separability. The separability refers to how well representations are separated across different groups (*e.g.*, classes) while being clustered within each group. Intuitively, well-separated representations lead to low entropy, making them easier to learn. Similar to Fisher’s discriminant ratio (Fisher 1936) and Davies-Bouldin Index (Davies and Bouldin 1979), we define the separability (\mathcal{S}) as the ratio of *between-group* variance to total variance as follows:

$$\mathcal{S} = \frac{\text{Var}[\mathbb{E}(x|c)]}{\text{Var}(x)} = \frac{\text{Var}[\mathbb{E}(x|c)]}{\text{Var}[\mathbb{E}(x|c)] + \mathbb{E}[\text{Var}(x|c)]}, \quad (2)$$

where the total variance $\text{Var}(x)$ is decomposed into *between-group* variance $\text{Var}[\mathbb{E}(x|c)]$ and *within-group* variance $\mathbb{E}[\text{Var}(x|c)]$. Thus, \mathcal{S} is minimized (*i.e.*, 0) when all the groups are not separated at all to the point that their means are exactly the same, making $\text{Var}[\mathbb{E}(x|c)] = 0$, whereas \mathcal{S} is maximized (*i.e.*, 1) when each group has the maximum cohesion with $\mathbb{E}[\text{Var}(x|c)] = 0$.

Analytic Results

Setting. We evaluate the complexity of three types of representations: input-level, feature-level, and parameter-level representations, using the MNIST (LeCun et al. 1998), SVHN (Netzer et al. 2011), and CIFAR-10 (Krizhevsky, Hinton et al. 2009) datasets, equally consisting of 10 classes. Input-level representations are training images themselves, while feature-level representations are their feature vectors from a model trained on the images. Finally, for the parameter-level representation, we train a Bayesian neural network (BNN) and take model samples from the BNN as many as the number of input images. Each sampled model is then split into the same number of chunks as the number of classes (*i.e.*, 10 for all three datasets), and these chunks and classes are regarded as local groups when computing the volume and separability. ResNet-32 (He et al. 2016) is commonly used for the model architecture throughout the analytic process.

Results. Figure 3 presents the separability and volume of three types of representations using different datasets, where the left y-axis represents values of $1 - \mathcal{S}$, while the right y-axis depicts the values of \mathcal{V} on a logarithmic scale. Thus,




Figure 3: Separability and volume of input images, feature vectors, and parameter chunks.

smaller bar heights indicate lower complexity in terms of separability and volume. As expected to some extent, input-level representations exhibit bars much longer than those of features and parameters. This implies that the representation of raw data is too complex either to be adequately covered by a few exemplars or a generative model. As a result, both memory replay and generative replay encounter challenges of preserving knowledge when dealing with complex datasets. While the feature-level representation demonstrates significantly reduced complexity compared to input images, it still proves to be more complex in terms of separability than parameter chunks. This observation is noteworthy since the dimensionality of parameter chunks is typically much higher than that of feature vectors.

In summary, despite their high dimensionality, learned parameters stand out as the most compact representation capable of efficiently capturing the underlying data distribution. Based on this analysis, our framework suggests the generative replay of learned parameters instead of using data samples or feature vectors.

Methodology

In this section, we first formulate the problem of continual learning (CL), and then present our proposed recall-oriented continual learning framework.

Problem Statement. In this paper, we follow the standard task-incremental learning (TIL) scenario, where a sequence of T tasks needs to be incrementally learned and each task t for $t \in [1, T]$ is associated with its training data \mathcal{D}_{train}^t and test data \mathcal{D}_{test}^t . The goal of TIL is to incrementally train a model that can make precise prediction for any task previously learned. In a typical TIL setting, we consider a task-aware scenario, where the task identity is explicitly given to the model at inference time as well as training time. Although our framework basically aims to deal with this task-aware version of TIL, it is also experimentally observed to be effective in a task-agnostic scenario without task-identifiers at inference time.

Proposed Framework. In order to achieve both high plasticity (ability to learn new knowledge) and high stability (ability to retain previous knowledge), we propose the recall-oriented CL framework that consists of three major steps, namely *knowledge acquisition*, *knowledge consolidation*, and *knowledge recall*. The overall process of the framework is illustrated in Figure 4. In the knowledge acquisition

step, knowledge of incoming task is acquired in the form of model parameters to be consolidated into the meta-model (*i.e.*, GMM). More specifically, we train a lightweight BNN, which corresponds to working memory, without any attempts to preserve the previous knowledge, which thus secures pure plasticity. The knowledge of working memory is then merged into GMM that corresponds to long-term memory in the knowledge consolidation step. Motivated by our analysis in the previous section, GMM is trained to generate learned parameters for the current task using the trained BNN, while trying not to forget model parameters for the past tasks through generative replay. Finally, at inference time, GMM *recalls* a *synthetic* model for each task, which is expected to achieve the performance comparable to that of the original task-specific model.

Knowledge Acquisition

The knowledge acquisition step aims to learn new knowledge by training a model specialized to each new task, and this task-specific knowledge is later consolidated into our long-term memory, GMM. In a few approaches with hypernetworks, the process of knowledge consolidation is simultaneously performed during learning each new task, with simple regularization on meta-model (Henning et al. 2021; von Oswald et al. 2020). Although these approaches have shown effectiveness at knowledge preservation in the meta-model, learning with regularization may lead to parameters less optimized for the current task, hindering plasticity. In order to ensure high plasticity, we propose a two-level approach of separating the process of learning new task from training on GMM.

This separated process of knowledge acquisition implies that we cannot feed data samples directly to GMM unlike hypernetworks being trained using data samples. Instead, we need multiple different models themselves trained for the same task, for GMM to use them as its training data. A straightforward approach to this end is training a task-specific model multiple times with differently initialized parameters (Lakshminarayanan, Pritzel, and Blundell 2017) or with different subsets of data samples (Joseph and Balasubramanian 2020), and collecting all those resulting models. However, it is prohibitively inefficient to train as many different models as necessary for GMM to effectively capture the underlying parameter distribution.

Task-Specific BNN. In order to address this issue, we propose using *Bayesian neural network* (BNN) as a task-specific model. Unlike a typical neural network that is trained to determine a single value for each weight, a BNN approximates the posterior distribution of weights $p(\theta|\mathcal{D})$, which allows us to sample each weight from the BNN (Blundell et al. 2015). Therefore, it suffices to train a single task-specific BNN $\theta_B^t \sim p(\theta|\mathcal{D}^t)$, instead of training multiple models for each t -th task. The training of θ_B^t relies on how well we represent the posterior distribution of weights. In many BNNs, the posterior distribution is approximated to be a Gaussian distribution. We also make this assumption, and hence we have: $\theta_B^t \sim \mathcal{N}(\mu_{\theta_B^t}, \Sigma_{\theta_B^t})$. Note that this Gaussian posterior distribution on model parameters is known to be

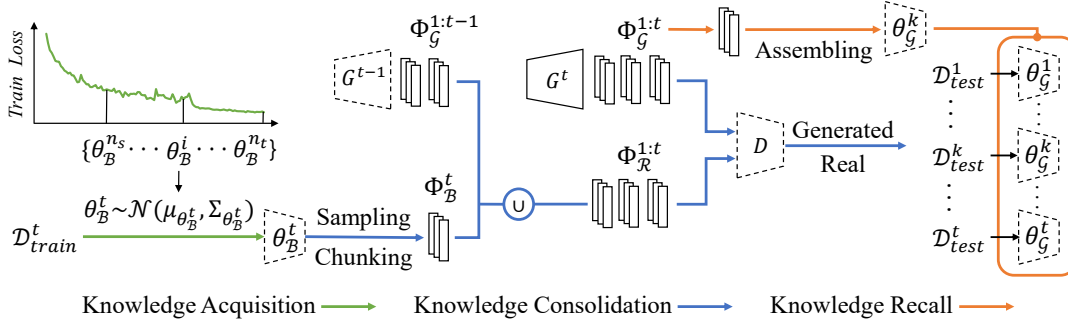


Figure 4: Illustration of three major steps in our recall-oriented continual learning framework.

effectively captured by the training history of a single neural network (Izmailov et al. 2018; Maddox et al. 2019), without the need for excessive training¹. Once θ_B^t is well-trained, we can take multiple parameter samples from the BNN and feed them to GAMM, enabling GAMM to effectively learn the parameter distribution of the model for the task t .

When training a task-specific BNN, we use a lightweight model rather than one with a large number of parameters. This is based on our intuition that each individual task may not require such a high capacity to be trained effectively. Also, using a lightweight model is essential for GAMM to reduce its required capacity.

Knowledge Consolidation

In the knowledge consolidation step, the acquired knowledge in working memory, which is learned by a task-specific BNN, is transferred to our long-term memory, GAMM. During this process, it is important to avoid forgetting previously learned knowledge that has been consolidated into GAMM. Thus, once plasticity is achieved in knowledge acquisition, GAMM is responsible for accumulating task-specific knowledge with high stability. As the knowledge is acquired in the form of parameters, our goal is to incrementally train GAMM in a way that it can re-generate task-specific parameters for all the tasks learned so far.

We design GAMM to be a conditional GAN, and thus we can apply generative replay to prevent GAMM from forgetting the parameter distributions of the past tasks. Unlike typical generative replay aiming at input-level generation, we replay parameters for the previous tasks, and train them with parameters of the current task. As analyzed in the previous section, the parameter-level representation is much more compact than the corresponding input-level representation, thereby achieving high stability with a limited capacity.

Initial Learning on GAMM. Let us now present in detail how to incrementally train GAMM by generative replay on parameter-level representations. For learning the first task with GAMM, we can adopt the standard way of training a GAN without replay. In typical GAN training, a generator G and a discriminator D are alternately trained with their conflicting objectives, namely generating samples close to the real ones and identifying those fake samples, respectively.

¹The details of $\mu_{\theta_B^t}$ and $\Sigma_{\theta_B^t}$ are presented in the Appendix.

In recent studies on GAN training, the state-of-the-art methods are mostly based on the WGAN framework (Arjovsky, Chintala, and Bottou 2017), which has the following objective function:

$$\min_G \max_D \mathbb{E}_{x \sim P_r} [D(x)] - \mathbb{E}_{\tilde{x} \sim P_g} [D(\tilde{x})], \quad (3)$$

where x represents real data with its distribution P_r and $\tilde{x} = G(z)$ represents generated data from the distribution P_g such that $z \sim p(z)$ s. GAMM is also trained by optimizing Eq. (3), not on real data samples ($x \sim P_r$) but on parameters sampled from the task-specific BNN (*i.e.*, $\theta_B \sim \mathcal{N}(\mu_{\theta_B}, \Sigma_{\theta_B})$) of the first task. In addition, rather than directly feeding entire models of θ_B to GAMM, we split θ_B into a set of equal-sized chunks $\Phi_B = \{\phi_B^1, \dots, \phi_B^m\}$ s.t. $m = \lceil |\theta_B| / |\phi_B^i| \rceil$, as used in a few studies (Ha, Dai, and Le 2017; Joseph and Balasubramanian 2020; von Oswald et al. 2020). Then, these sampled parameter chunks, each conditioned on its chunk-id, are fed into GAMM as training samples. This chunking technique reduces the output dimensionality of the meta-model to be the size of each chunk (*e.g.*, $|\phi_B^i| = 2000$) from the entire model size.

Incremental Learning on GAMM. Then, for the t -th task s.t. $t > 1$, the previous generator G^{t-1} is used to generate synthetic parameters of the previous tasks from 1 to $t-1$, and then G^t is jointly trained with the combined set of current and generated chunks. When GAMM learns the distribution of parameter chunks $\Phi_B^t \sim \mathcal{N}(\mu_{\Phi_B^t}, \Sigma_{\Phi_B^t})$ sampled from the BNN for task t , real chunks $\Phi_B^{1:t-1}$ from the previously learned BNNs are replaced by generated chunks $\Phi_G^{1:t-1} = \{G^{t-1}(z, c)\}$, where z is a latent vector sampled from $\mathcal{N}(0, 1)$ and c is a chunk-id randomly chosen from all chunk-ids of the previous tasks. As shown in Figure 4, the whole training dataset $\Phi_{\mathcal{R}}^{1:t}$, which corresponds to $x \sim P_r$ in Eq. (3), is the union of sampled chunks Φ_B^t and generated chunks $\Phi_G^{1:t-1}$, that is, $\Phi_{\mathcal{R}}^{1:t} = \Phi_B^t \cup \Phi_G^{1:t-1}$. Thus, GAMM at the t -th task is trained to generate $\Phi_G^{1:t}$ that are the chunks for all tasks learned so far with the following objective function:

$$\min_G \max_D \mathbb{E}_{\Phi_{\mathcal{R}}^{1:t} \sim P_r} [D(\Phi_{\mathcal{R}}^{1:t})] - \mathbb{E}_{\Phi_G^{1:t} \sim P_g} [D(\Phi_G^{1:t})]. \quad (4)$$

In the end of the knowledge consolidation step, G^t is only retained in GAMM for making inference as well as learning the next task, and all the other components are not stored.

Knowledge Recall

When the human brain recalls some information previously remembered, the corresponding knowledge is somehow generated from long-term memory (Schacter and Addis 2007), and brought into working memory for use. Similarly, whenever making inference for each task in our framework, the required knowledge, that is, the task-specific model itself, can be generated by GAMM. More precisely, given a task-identifier $k \in [1, T]$, GAMM produces a set of parameter chunks $\Phi_G^k = \{G(z, c^k)\}$, each of which is conditionally generated by a chunk-id c^k for task k . The generated chunks are then assembled into the corresponding task-specific model θ_G^k , as illustrated in Figure 4, which can be immediately used to make inference. This generated model is expected to output similar results to those of its original task-specific BNN, that is, $p(y|x, \theta_G^k) \approx p(y|x, \theta_B^k)$.

Task-Agnostic Inference. Although this paper focuses on the above TIL scenario using task-identifiers, the generation scheme of GAMM is also applicable to make prediction without task-identifiers at inference time. To this end, we can choose a model with the highest confidence out of task-specific models over all the tasks previously learned, and then use its output as our final prediction. We adopt the strategy of selecting a model with the smallest entropy \mathcal{H} (Henning et al. 2021), indicating the highest confidence, and finally return $p(y|x, \theta_G^{\hat{k}})$ s.t. $\hat{k} = \arg \min_{j \in [1, T]} \mathcal{H} [p(y|x, \theta_G^j)]$.

Experiments

This section evaluates our framework by comparing with the existing replay-based methods as well as the other CL baselines from the perspective of the stability-plasticity dilemma.

Experimental Settings for CL Benchmarks

Performance Metrics. We evaluate CL performance using three metrics: *average accuracy* (ACC), *backward transfer* (BWT) and *learning accuracy* (LA) as follows: $ACC = \frac{1}{T} \sum_{i=1}^T R_{T,i}$, $BWT = \frac{1}{T-1} \sum_{i=1}^{T-1} R_{T,i} - R_{i,i}$, $LA = \frac{1}{T} \sum_{i=1}^T R_{i,i}$, where T is the total number of tasks and $R_{i,j}$ is the classification accuracy on the j -th task after learning the i -th task. ACC is the average accuracy of all the tasks in the end. BWT measures how much we forget the past knowledge, where the more negative values, the more forgetting occurs. LA is to measure the model’s plasticity (Riemer et al. 2019), by taking the average accuracy for each task right after it has been learned. Our final goal is to achieve highest ACC with near-zero BWT yet large LA values.

Datasets. We consider the datasets commonly used in CL benchmarks, namely Split CIFAR-10 and Split CIFAR-100 for comparison with replay-based methods, and PMNIST, Split CIFAR-100 and 5-Datasets (Ebrahimi et al. 2020b) for comparison with the other CL baselines. Split CIFAR-10 is constructed by dividing the original CIFAR-10 dataset into 5 tasks, each with 2 classes. Similarly, Split CIFAR-100 consists of 10 tasks, each with 10 classes. PMNIST is a variant of MNIST, where each task has a different random permutation on image pixels. 5-Datasets consists of five differ-

ent datasets: CIFAR-10, MNIST, SVHN, FMNIST (Xiao, Rasul, and Vollgraf 2017), and notMNIST (Bulatov 2011), each of which is learned as a single task.

Baselines. We test various replay-based methods, which are classified into two categories: memory replay (Mem) and generative replay (Gen). Also, their replaying targets can be in one of the following three levels: input-level (I), feature-level (F), and parameter-level (P). Both A-GEM (Chaudhry et al. 2019a) and ER (Chaudhry et al. 2019b) replay input-level representations with a fairly large memory buffer of 100 exemplars per class (van de Ven, Tuytelaars, and Tolias 2022). DGR (Shin et al. 2017) and BI-R (van de Ven, Siegelmann, and Tolias 2020) belong to generative replay with input-level and feature-level representations, respectively. HNET (von Oswald et al. 2020) and PR-BBB (Henning et al. 2021) replay parameter-level representations via hypernetworks, which are the state-of-the-arts in replay-based methods. We consider the other types of baselines; regularization-based EWC (Kirkpatrick et al. 2017), architecture-based HAT (Serrà et al. 2018), and memory-based GPM (Saha, Garg, and Roy 2021).

Model Architectures. We use a simple GAN architecture for GAMM, where the generator G consists of 100 input units followed by two fully-connected (FC) layers of 200 units and the discriminator D has one FC layer of 256 units followed by a binary output head. The output dimensionality of G varies depending on the chunk size. For inference models to make prediction, GAMM generates lightweight neural networks, which are reduced versions of the original architectures (e.g., 0.15 times the original size of ResNet-32) that have been commonly used in the literature. Those original architectures include: (1) ResNet-32 for Split CIFAR-10 (Henning et al. 2021), (2) ResNet-18 (Henning et al. 2021) and 5-layer AlexNet (Saha, Garg, and Roy 2021) for Split CIFAR-100, (3) a FC network with two layers of 100 units (Saha, Garg, and Roy 2021) for PMNIST, and (4) reduced ResNet-18 (Saha, Garg, and Roy 2021) for 5-Datasets. Some baselines utilize their own architectures, such as BI-R (van de Ven, Siegelmann, and Tolias 2020) with an autoencoder-based model, and HNET (von Oswald et al. 2020) and PR-BBB (Henning et al. 2021) with hypernetwork-based meta-models (Henning et al. 2021). More experimental details are presented in the Appendix.

CL Performance Comparison and Discussion

Comparison with Replay-Based Methods. We compare the performance of replay-based methods in both task-aware and task-agnostic settings. Table 1 summarizes the experimental results. Our framework with GAMM outperforms all the baselines, consistently showing the highest ACC that indicates the best stability-plasticity balance. In task-aware settings, GAMM achieves almost zero-forgetting with its near-zero BWT values, while it also shows the excellent plasticity with high LA values. The performance gap between GAMM and the compared methods becomes larger in task-agnostic scenarios, which shows the highest effectiveness of GAMM even without task-identifiers at inference time. As our expectation in analysis on representation

Method	Strategy	Split CIFAR-10					Split CIFAR-100				
		Task-aware		Task-agnostic			Task-aware		Task-agnostic		
		ACC	BWT	LA	ACC	Memory	ACC	BWT	LA	ACC	Memory
A-GEM	Mem (I)	85.97	-11.96	95.54	20.15	≈ 1.2 M	56.36	-35.15	87.99	9.80	≈ 18.9 M
ER	Mem (I)	88.57	-6.22	93.55	50.32	≈ 1.2 M	80.85	-5.52	85.82	31.82	≈ 18.9 M
DGR	Gen (I)	70.38	-17.31	84.23	13.83	14.3 M	16.55	-48.12	59.86	12.44	25.0 M
BI-R	Gen (F)	94.99	-0.18	95.14	55.19	13.3 M	82.33	-0.64	82.90	22.83	13.3 M
HNET	Gen (P)	94.53	-0.45	94.89	56.67	1.0 M	86.50	-0.26	86.73	41.63	5.0 M
PR-BBB	Gen (P)	95.30	-0.28	95.52	58.94	20.2 M	84.78	-0.05	84.82	41.66	14.7 M
GAMM	Gen (P)	96.73	-0.13	96.84	78.49	0.9 M	87.19	<u>-0.22</u>	<u>87.39</u>	44.75	5.1 M

Table 1: Performance of replay-based methods, where we report the mean of ACC, BWT, and LA over five runs with different task orders. The approximation symbol (\approx) represents the memory usage considering both parameters and exemplars.

Methods	Permuted MNIST		Split CIFAR-100		5-Datasets	
	ACC	BWT	ACC	BWT	ACC	BWT
EWC	$92.01 \pm 0.56^\dagger$	-0.03 $\pm 0.00^\dagger$	71.17 ± 0.46	-2.23 ± 0.47	$88.64 \pm 0.26^\dagger$	$-0.04 \pm 0.01^\dagger$
HAT	-	-	73.25 ± 0.30	-2.79 ± 0.43	$91.32 \pm 0.18^\dagger$	-0.01 $\pm 0.00^\dagger$
A-GEM	$83.56 \pm 0.16^\dagger$	$-0.14 \pm 0.00^\dagger$	$63.98 \pm 1.22^\dagger$	$-0.15 \pm 0.02^\dagger$	$84.04 \pm 0.33^\dagger$	$-0.12 \pm 0.01^\dagger$
ER	$87.24 \pm 0.53^\dagger$	$-0.11 \pm 0.01^\dagger$	$71.73 \pm 0.63^\dagger$	-0.06 $\pm 0.01^\dagger$	$88.31 \pm 0.22^\dagger$	$-0.04 \pm 0.00^\dagger$
GPM	$93.91 \pm 0.16^\dagger$	-0.03 $\pm 0.00^\dagger$	72.45 ± 0.54	-0.43 ± 0.43	$91.22 \pm 0.20^\dagger$	-0.01 $\pm 0.00^\dagger$
GAMM	96.40 ± 0.02	-0.68 ± 0.12	74.90 ± 0.94	-0.18 ± 0.08	93.68 ± 0.20	-0.31 ± 0.14

Table 2: Performance of the proposed framework and baselines. Mean and standard deviation of ACC and BWT over five runs with different task orders are reported. The values with \dagger are from (Saha, Garg, and Roy 2021).

complexity, the methods of replaying either exemplars (A-GEM and ER) or pseudo-samples (DGR) are observed to suffer from severe forgetting, showing most negative BWT values. This confirms that replaying input-level representations is indeed not effective at knowledge preservation due to their high complexity. BI-R using feature-level replay has higher stability with better BWT values, but shows limited LA values due to the stability-plasticity dilemma that is partially unresolved by replaying feature vectors. HNET and PR-BBB also exhibit better BWT compared to input-level replay methods, but their plasticity is worse than GAMM in terms of LA. This is likely due to their one-phase training scheme within a hypernetwork, where the process of learning new knowledge can be hindered by the process of retaining past knowledge. This one-phase training seems to make worse effect on the performance in task-agnostic settings. In contrast, GAMM manages to achieve the best performance not only in task-aware scenarios but also in task-agnostic scenarios by separating the mechanisms of learning new knowledge from preserving old knowledge.

Memory Usage. Table 1 also shows the amount of memory required for each replay-based model to maintain the learned knowledge. The memory usage is measured by the total number of parameters (or values) in the main model, the generative model, or the stored exemplars. GAMM shows better efficiency in memory usage, compared to the memory replay and generative replay methods, and uses a similar amount of memory to what HNET takes, while ensuring al-

ways the best performance.

Comparison with CL Baselines. We compare our framework with the other baselines with respect to ACC and BWT. Table 2 summarizes the experimental results, where PM-NIST is a task-agnostic setting and the others follows a task-aware setting. Our framework outperforms all the baselines with clear margins in ACC. Most methods achieve BWT close to 0 while their ACC is not as good as that of our framework, which implies that their techniques to alleviate forgetting have impaired plasticity of models. In contrast, our method achieves the highest ACC thanks to its high plasticity via the separation of working and long-term memory.

Conclusion

In this paper, we proposed a novel framework to resolve the stability-plasticity dilemma in continual learning, where the task-specific knowledge is separately learned, and then consolidated into the meta-model. Based on our analysis on representation complexity, we found that the parameter-level representation is a proper form of knowledge to be maintained, and thereby proposing GAMM that can re-generate task-specific models themselves to make prediction at inference time. Our experimental results confirm that the proposed framework is capable of achieving the best balance between stability and plasticity without a large amount of memory consumption, outperforming the existing methods in both task-aware and task-agnostic scenarios.

Acknowledgments

This work was supported in part by Institute of Information & communications Technology Planning & Evaluation (IITP) grants funded by the Korea government (MSIT) (No.2022-0-00448, Deep Total Recall: Continual Learning for Human-Like Recall of Artificial Neural Networks, No.RS-2022-00155915, Artificial Intelligence Convergence Innovation Human Resources Development (Inha University)), in part by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) (No.2021R1F1A1060160, No.2022R1A4A3029480), and in part by INHA UNIVERSITY Research Grant.

References

- Ahn, H.; Cha, S.; Lee, D.; and Moon, T. 2019. Uncertainty-based Continual Learning with Adaptive Regularization. In *Advances in Neural Information Processing Systems*.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, 214–223. PMLR.
- Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; and Wierstra, D. 2015. Weight uncertainty in neural network. In *International conference on machine learning*, 1613–1622. PMLR.
- Bulatov, Y. 2011. Notmnist dataset. Google (Books/OCR). Technical report, Tech. Rep.[Online]. Available: <http://yaroslavvb.blogspot.it/2011/09...>
- Buzzega, P.; Boschini, M.; Porrello, A.; Abati, D.; and Calderara, S. 2020. Dark Experience for General Continual Learning: a Strong, Simple Baseline. In *Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*.
- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2019a. Efficient Lifelong Learning with A-GEM. In *International Conference on Learning Representations*.
- Chaudhry, A.; Rohrbach, M.; Elhoseiny, M.; Ajanthan, T.; Dokania, P. K.; Torr, P. H.; and Ranzato, M. 2019b. Continual Learning with Tiny Episodic Memories. *arXiv preprint arXiv:1902.10486, 2019*.
- Davies, D. L.; and Bouldin, D. W. 1979. A Cluster Separation Measure. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(2): 224–227.
- Ebrahimi, S.; Elhoseiny, M.; Darrell, T.; and Rohrbach, M. 2020a. Uncertainty-guided Continual Learning with Bayesian Neural Networks. In *International Conference on Learning Representations*.
- Ebrahimi, S.; Meier, F.; Calandra, R.; Darrell, T.; and Rohrbach, M. 2020b. Adversarial continual learning. In *The European Conference on Computer Vision*, 386–402.
- Fisher, R. A. 1936. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2): 179–188.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27, 2672–2680.
- Ha, D.; Dai, A. M.; and Le, Q. V. 2017. HyperNetworks. In *International Conference on Learning Representations*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, 770–778.
- Henning, C.; Cervera, M. R.; D’Angelo, F.; von Oswald, J.; Traber, R.; Ehret, B.; Kobayashi, S.; Grewe, B. F.; and Sacramento, J. 2021. Posterior Meta-Replay for Continual Learning. In *Advances in Neural Information Processing Systems*.
- Izmailov, P.; Podoprikin, D.; Garipov, T.; Vetrov, D. P.; and Wilson, A. G. 2018. Averaging Weights Leads to Wider Optima and Better Generalization. In *Uncertainty in Artificial Intelligence*.
- Joseph, K. J.; and Balasubramanian, V. N. 2020. Meta-Consolidation for Continual Learning. In *Advances in Neural Information Processing Systems*.
- Kemker, R.; and Kanan, C. 2018. FearNet: Brain-Inspired Model for Incremental Learning. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Krueger, D.; Huang, C.-W.; Islam, R.; Turner, R.; Lacoste, A.; and Courville, A. 2017. Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *Advances in Neural Information Processing Systems*.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11): 2278–2324.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Lyu, F.; Wang, S.; Feng, W.; Ye, Z.; Hu, F.; and Wang, S. 2021. Multi-Domain Multi-Task Rehearsal for Lifelong Learning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Maddox, W. J.; Izmailov, P.; Garipov, T.; Vetrov, D. P.; and Wilson, A. G. 2019. A Simple Baseline for Bayesian Uncertainty in Deep Learning. In *Advances in Neural Information Processing Systems*.
- Mallya, A.; and Lazebnik, S. 2018. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 7765–7773.
- McClelland, J. L.; McNaughton, B. L.; and O’Reilly, R. C. 1995. Why there are complementary learning systems in

- the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3): 419.
- McCloskey, M.; and Cohen, N. J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, 109–165. Elsevier.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning.
- Ostapenko, O.; Puscas, M.; Klein, T.; Jahnichen, P.; and Nabi, M. 2019. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11321–11329.
- O’Reilly, R. C.; Bhattacharyya, R.; Howard, M. D.; and Ketz, N. 2014. Complementary learning systems. *Cognitive science*, 38(6): 1229–1248.
- Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54–71.
- Ratzlaff, N.; and Li, F. 2019. HyperGAN: A Generative Model for Diverse, Performant Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning*, 5361–5369.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Riemer, M.; Cases, I.; Ajemian, R.; Liu, M.; Rish, I.; Tu, Y.; and Tesauro, G. 2019. Learning to Learn without Forgetting by Maximizing Transfer and Minimizing Interference. In *International Conference on Learning Representations*.
- Saha, G.; Garg, I.; and Roy, K. 2021. Gradient Projection Memory for Continual Learning. In *International Conference on Learning Representations*.
- Schacter, D. L.; and Addis, D. R. 2007. The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481): 773–786.
- Serrà, J.; Suris, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming Catastrophic Forgetting with Hard Attention to the Task. In Dy, J. G.; and Krause, A., eds., *Proceedings of the International Conference on Machine Learning*.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual Learning with Deep Generative Replay. In *Advances in Neural Information Processing Systems*, 2990–2999.
- van de Ven, G. M.; Siegelmann, H. T.; and Tolias, A. S. 2020. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1): 1–14.
- van de Ven, G. M.; Tuytelaars, T.; and Tolias, A. S. 2022. Three types of incremental learning. *Nature Machine Intelligence*, 1–13.
- von Oswald, J.; Henning, C.; Sacramento, J.; and Grewe, B. F. 2020. Continual learning with hypernetworks. In *International Conference on Learning Representations*.
- Wang, K.; Vicol, P.; Lucas, J.; Gu, L.; Grosse, R. B.; and Zemel, R. S. 2018. Adversarial Distillation of Bayesian Neural Network Posteriors. In *Proceedings of the 35th International Conference on Machine Learning*, 5177–5186.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2018. Life-long Learning with Dynamically Expandable Networks. In *International Conference on Learning Representations*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual Learning Through Synaptic Intelligence. In *Proceedings of the International Conference on Machine Learning*.
- Zhmoginov, A.; Sandler, M.; and Vladymyrov, M. 2022. HyperTransformer: Model Generation for Supervised and Semi-Supervised Few-Shot Learning. In *International Conference on Machine Learning*, 27075–27098.