

CLAP4CLIP: Continual Learning with Probabilistic Finetuning for Vision-Language Models

Saurav Jha¹, Dong Gong^{1*}, Lina Yao^{2,1}

¹University of New South Wales, ²CSIRO's Data61

{saurav.jha, dong.gong}@unsw.edu.au; lina.yao@data61.csiro.au

Abstract

Continual learning (CL) aims to help deep neural networks to learn new knowledge while retaining what has been learned. Recently, pre-trained vision-language models such as CLIP [63], with powerful generalization ability, have been gaining traction as practical CL candidates. However, the domain mismatch between the pre-training and the downstream CL tasks calls for finetuning of the CLIP on the latter. The deterministic nature of the existing finetuning methods makes them overlook the many possible interactions across the modalities and deems them unsafe for high-risk CL tasks requiring reliable uncertainty estimation. To address these, our work proposes Continual LeArning with Probabilistic finetuning (CLAP). CLAP develops probabilistic modeling over task-specific modules with visual-guided text features, providing more reliable fine-tuning in CL. It further alleviates forgetting by exploiting the rich pre-trained knowledge of CLIP for weight initialization and distribution regularization of task-specific modules. Cooperating with the diverse range of existing prompting methods, CLAP can surpass the predominant deterministic finetuning approaches for CL with CLIP. Lastly, we study the superior uncertainty estimation abilities of CLAP for novel data detection and exemplar selection within CL setups. Our code is available at <https://github.com/srvCodes/clap4clip>.

1. Introduction

Learning in the real world involves dealing with the ever-changing distributions of task streams and their data [6, 26, 49, 85]. Given the constraints on resources and privacy, there is also no guarantee on re-training a network on all previously seen data [7]. Continual learning (CL) aims to learn from such data/task stream without *catastrophic forgetting* [6, 33] of past data/tasks. A challenging CL setup is

class-incremental learning, where new classes emerge with new tasks, and at test time, a model must infer from all seen classes without known task IDs [48].

Recent years have seen a surge of pre-trained multi-modal foundation models achieving state-of-the-art performances on several domains [1, 63, 84]. One such example for the vision-language (VL) domain is the Contrastive Language-Image Pre-training (CLIP) [63] model. CLIP comes with strong zero-shot generalization ability acquired by learning to match large-scale image-text pairs with contrastive loss [56]. However, to adapt well to downstream tasks, CLIP must be finetuned on the task-specific data [17, 94]. Considering both the need for continually finetuning pre-trained models to streaming tasks and their perks over training from scratch [70], our work investigates CL with CLIP.

An issue with the existing deterministic approaches [17, 94] is that these overlook the *uncertainties* arising from many possible interactions between visual and textual cues. For instance, on the textual side, while a good default hand-crafted prompt for a number of tasks is “A photo of a {class}”, there can be tasks where a more specific prompt can help improve the image-to-text coherence [63]. On the visual side, images from the same class have diverse range of backgrounds, poses, orientations, etc. Overlooking the uncertainties in image-text matching may cause overfitting to the downstream task data and forgetting of the generalizable knowledge [36]. This has particularly worse implications for CL where we seek to adapt CLIP on a stream of tasks. While some methods model the uncertainties through probabilistic finetuning [12, 47], these remain subpar at CL given: (a) their inaptness to existing prompt-based approaches [47], (b) their trading of in-domain performance for generalization [12].

Uncertainty-awareness can further be crucial for CL models deployed in mission-critical settings (healthcare, transport, etc.) as it can help calibrate predictions to reliably assess the models’ predictive confidences [27, 39]. Hence, to enhance the usage of pre-trained CLIP models for real-world CL tasks, we design a probabilistic finetuning

*D. Gong is the corresponding author. This work was partially supported by an ARC DECRA Fellowship (DE230101591) to D. Gong.

approach with **three properties**: A) **uncertainty-aware** modeling of cross-modal task cues for better generalization; B) compatibility with **existing prompt-based finetuning** methods [29, 70, 74, 94] to exploit their task-specific knowledge; C) leveraging the **rich pre-trained knowledge** of CLIP to further counter forgetting *without* extra training data.

To this end, we design a Variational Inference (VI) framework to model the distribution of visual-guided text features (rather than modeling the distribution of either modality alone), thus validating property #A. This helps our method implicitly address the biases in the image-text alignment [62, 95] (see Fig. 2b). Namely, we design a task-shared visual-guided attention module to encourage the alignment, and then perform specialized forward passes through lightweight task-specific inference modules to learn stochastic factors for CL tasks. To further counter the forgetting in these modules, we propose using the pre-trained language information of CLIP for weight initialization and task distribution regularization (property #C). Lastly, the plug-and-play nature of our probabilistic fine-tuning method makes it compatible with other model adaption/fine-tuning methods (such as [74]) and diverse types of prompts that are hand-crafted [63], uni-modal [94], multi-modal [29], or input-specific [74] (property #B). We backronymize our finetuning approach as **CLAP** – **Continual LeArning with Probabilistic** finetuning – for the pre-trained CLIP model.

Our experiments across several settings show that CLAP4CLIP enhances *prompt-based* finetuning for CLIP and surpasses the predominant deterministic finetuning methods for CL all while sharing a similar resource overhead. To further explore the out-of-the-box perks of CLAP’s probabilistic nature, we study its superior uncertainty quantification capabilities on a proposed *post-hoc* novel data detection setup and on exemplar selection for CL.

2. Related work

Continual Learning (CL). The existing CL literature is predominated by three categories of methods: (a) Regularization-based methods [2, 33, 44] alleviate forgetting by punishing changes to the parameters that are important to previous tasks; (b) Architecture-based approaches learn parameters that are specialized for individual tasks either by network expansion [14, 27] or by sub-network composition [28, 57]; (c) Rehearsal-based approaches [7, 65] rely on storing a fraction of the past task experiences in a memory to train with the current task. Nevertheless, each category has its own flaw – methods in (a) struggle to discriminate inter-task classes [41]; those in (b) often require task oracle during inference and can induce dramatic memory overhead for a larger number of tasks; those in (c) are

sensitive to the memory sizes besides being prone to overfitting on the memory samples [72]. Hence, practical CL calls for combining these. Our work leverages (a) via function-space regularization (Sec. 3.7), (b) via task-specific modules (Sec. 3.5), and (c) via herding-based replay [65] (Sec. 4).

Vision-Language Models (VLMs) finetuning. The powerful generalizability of pre-trained VLMs [1, 84] like the CLIP [63] has enabled their zero-shot applications to a range of downstream tasks, including CL [70]. In practice, their performance on downstream out-of-distribution data remains rather weak [60, 79]. For such cases, finetuning on task-specific data is a natural choice [79]. Instead of performing *full finetuning* on all parameters, some *parameter-efficient finetuning* methods learn a lightweight feature adapter module for textual and/or visual paths [17, 87]. Another line of parameter-efficient finetuning methods learns *soft prompts* which are a few continuous tokens serving as inputs to the frozen visual and/or textual encoder(s) to capture task-specific information [93, 94]. Existing works on CL with pre-trained CLIP have leveraged either [67, 74] or both [92] of these methods. However, such finetuning methods remain deterministic in nature. This imposes an explicit constraint on the modeling of the possible ways in which the visual and the textual semantics interact.

To address the aforesaid flaw, one could turn to adapting the existing probabilistic finetuning approaches to capture the cross-modal interactions in CL tasks. For instance, [47] learn the distribution of hand-crafted prompts while [12] leverage VI to model the distributions of soft prompts *conditioned* on the input images. Yet these methods are limited in their efficacies. While [47] is incompatible with conditional prompt learning [93], [12] lags in terms of in-domain performance and computational efficiency. Hence, our work aims to bridge these gaps for probabilistic finetuning all while adapting it for CL.

3. Methodology

3.1. Preliminaries

Continual Learning (CL). Class-incremental CL [48] aims to learn from a sequence of T tasks $[(C^1, D^1), (C^2, D^2), \dots, (C^T, D^T)]$. Each task $t \in [1, T]$ has its training data $D^t = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_{k^t}, y_{k^t})\}$, where \mathbf{x} and y are the input images and labels, respectively from the set of classes $C^t = \{c_1^t, c_2^t, \dots, c_{n^t}^t\}$. Following [24, 44, 48], we assume any two task-specific sets of classes to be disjoint: $C^i \cap C^j = \emptyset$. A neural network with parameters ϕ is then trained on task t using D^t to minimize the cross-entropy loss over C^t . At test time, the model is evaluated on all seen classes $\bigcup_{i=1}^t C^i$, where the past task predictions are prone to forgetting. As a solution, rehearsal-based methods [5, 80] replay past task samples from a memory \mathcal{M} during

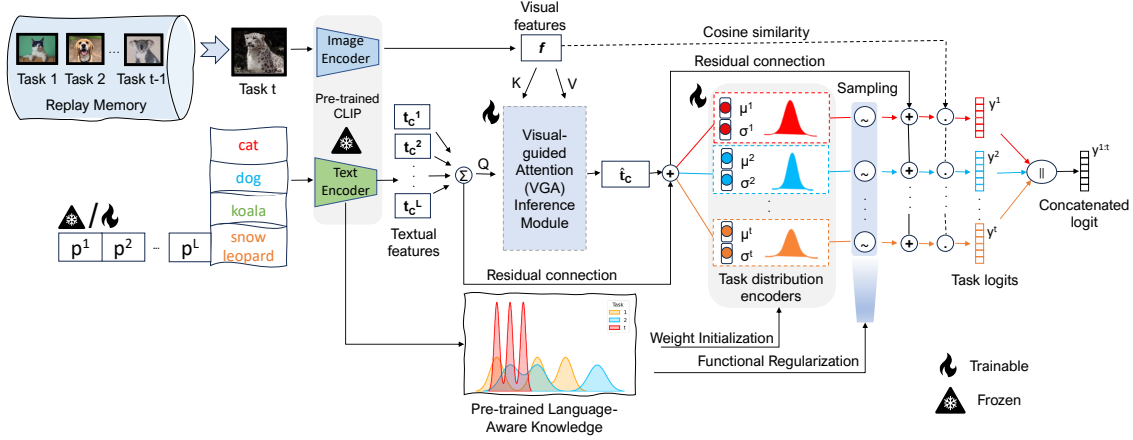


Figure 1. **CLAP4CLIP overview:** the visual-guided attention (VGA) inference module uses the text features as query (Q), and the visual features as keys (K) and values (V) to produce visual-guided text features. The task-specific text features are fed to their respective task distribution encoders (μ^t, σ^t) . The task distribution samples are then fused with the original task features prior to deriving the task logits y^t . All task logits are concatenated to produce the final prediction $y^{1:t}$.

training. We use herding [78] to maintain \mathcal{M} .

3.2. CLIP with prompt design

CLIP comprises an image encoder $f(\mathbf{x})$ acting on an image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ and a text encoder $g(\mathbf{t})$ acting on a word embedding vector $\mathbf{t} \in \mathbb{R}^{(L \times e)}$ derived from a text prompt $\mathbf{p} \in \mathbb{R}^{((L-1) \times e)}$. Here, H and W are the height and the width of the image, respectively, L is the text length, and e is the text embedding dimension. Both the encoders output a d -dimensional encoding, which are used in the prediction of the class y_i of the image as:

$$p(y_i|\mathbf{x}) = \frac{\exp(\langle f(\mathbf{x})^T, g(\mathbf{t}_i) \rangle / \tau)}{\sum_{c=1}^{|C^t|} \exp(\langle f(\mathbf{x})^T, g(\mathbf{t}_c) \rangle / \tau)}, \quad (1)$$

where τ is a learnable temperature parameter, $\langle \cdot, \cdot \rangle$ is the cosine similarity, and $\mathbf{t}_c = [\mathbf{p}, \mathbf{e}_c]$ is the result of adding a class-specific word embedding \mathbf{e}_c to the prompt \mathbf{p} . The textual encoding $g(\mathbf{t}_c)$'s for all classes are used as the weights of a linear classifier. In CL, $g(\mathbf{t}_c)$ encodes the class names seen until task t . Eq. (1) thus forms contrastive training criteria for the text and visual modalities, whose rich representation allows pre-trained CLIP to be used for zero-shot classification through *hard* prompt templates, i.e., $\mathbf{p}^c = \text{"A photo of a \{c^{\text{th}} \text{ class}\}"}$.

3.2.1 CLIP finetuning with soft prompts

To improve the CLIP performance on a downstream task t , prompt learning employs *soft* prompts as a set of learnable vector tokens $\mathbf{p} = \{\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^L\}$. CoOp [94] shares \mathbf{p} with all classes and averages them to use in t . MaPLe [29] learns multi-modal prompts by employing two such token sets \mathbf{p}_v and \mathbf{p}_t until the J -th layers of the vision and the text encoders of CLIP, respectively. AttriCLIP [74]

selects input-conditioned tokens $\{\{\mathbf{p}^j\}_{1 \leq j \leq L} | \mathbf{x}_k\}$. Learning \mathbf{p} (with frozen CLIP weights) thusly helps encode class/modality/instance-conditioned context for a given task.

3.2.2 CLIP finetuning with adapters

Adapter-based methods, e.g., CLIP-Adapter [17], learn a lightweight adapter module over the text and/or visual features of the frozen CLIP model. With a text adapter A_t , the updated text encoder features from Eq. (1) can now be rewritten (with a slight abuse of notation) as:

$$g(\mathbf{t}_i) = \alpha A_t(g(\mathbf{t}_i)) + \beta g(\mathbf{t}_i), \quad (2)$$

where α and β control the strength of the residual connection between the adapted and the pretrained features, e.g., $\beta = 1 - \alpha$ in [17]. Unlike soft prompts, however, the gradients of the adapter do not back-propagate through all the text encoder layers. This gives them a computational edge in training.

3.3. Probabilistic modeling and inference in CLAP4CLIP

In Sec. 3.3 - 3.7, we develop our CLIP-based probabilistic finetuning model with a variational inference (VI) framework, as shown in Fig. 1. We start by introducing a basic prediction-aware latent variable model for CLAP4CLIP in this section.

3.3.1 Prediction-aware variational modelling

We are interested in modeling the stochastic processes that generate the labels y for the inputs \mathbf{x} of a CL task t . To this end, we assume a prior distribution p_{χ} over the text features

$\{p_\chi(\mathbf{t}_c)\}_{c=1}^C$. In the feature space, we use the reparameterization trick [32] to represent $\mathbf{t}_c(\mathbf{p}) \sim p_\chi(\mathbf{t}_c(\mathbf{p}))$ as a linear combination of the deterministic text features $g(\mathbf{t}_c)$ and a stochastic latent variable $z \sim p_\chi$, which can be readily plugged into Eq. (1):

$$\mathbf{t}_c(\mathbf{p}) = g(\mathbf{t}_c(\mathbf{p})) + z \mid z \sim p_\chi, \quad (3a)$$

$$p(y_i|\mathbf{x}) = \int_{\chi} \frac{\exp(\langle f(\mathbf{x})^T, \mathbf{t}_c(\mathbf{p}) \rangle)}{\sum_{c=1}^{|C^t|} \exp(\langle f(\mathbf{x})^T, \mathbf{t}_c(\mathbf{p}) \rangle)} p(p_\chi(\mathbf{t}_c(\mathbf{p}))) d\chi, \quad (3b)$$

where $g(\cdot)$ is defined in Eq. (2). By modeling the generative processes behind the inputs, Eq. (3b) injects data-driven uncertainties into the predictions. It also offers us **further advantages** over other variational frameworks, such as [12]. First, as the latent variable z is used to infer the logits, it favors generalization by encoding *prediction-aware* information. Second, Eq. (3b) is *prompt-agnostic* as the text features could be derived from any existing soft [29, 74, 93] or hard [17, 70] prompt-based method. Third, defining the distribution over the encoder outputs rather than the prompts helps inherit the *computational efficiency* of adapters.

To deal with the intractability of the marginal likelihood in Eq. (3b), one could optimize for an evidence lower bound (ELBO) using a variational posterior q_ϕ that matches the prior p_χ based on the KL-divergence loss \mathbb{D}_{KL} :


$$\log p(y|\mathbf{x}) \geq \mathbb{E}_{q_\phi(z|\mathbf{t}_c)}[\log p(y|\mathbf{x}, z)] - \mathbb{D}_{\text{KL}}(q_\phi(z|\mathbf{t}_c) \| p_\chi). \quad (4)$$

Following standard practices [18, 32], q_ϕ can be modeled as a Gaussian distribution $\mathcal{N}(\mu(\mathbf{t}_c), \sigma(\mathbf{t}_c))$ whose mean μ and standard deviation σ are parameterized by linear adapter layers – we thus refer to $[\mu; \sigma]$ as probabilistic adapter. A simple choice for the static prior p_χ is $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Lastly, a finite set of Monte-Carlo (MC) samples $\{z_m\}_{m=1}^M$ can be drawn from the posterior by reparameterization [32].

However, finetuning adapters for CL is not straightforward. Namely, we have the overhead of search for the dataset-specific residual ratio α (see Eq. (2)) [17, 87]. This has particularly worse implications for a probabilistic adapter like ours, where a too large α can inject enough noise to corrupt the pre-trained representations to the point of catastrophic forgetting (see Fig. 2a). To preserve our property #B (Sec. 1) for backward compatibility with existing prompting techniques, we first seek to have *no additional overhead* of hyperparameter search for the residual ratio. Subsequently, we use $\alpha = \beta = 1$ through our work.

3.4. Cross-modal feature alignment

To investigate the source of catastrophic forgetting in finetuning methods, we study how CL effects the *cross-modal deviation* [53] between the text and the image features of CLIP. Here, we consider two basic CL models: the CoOp [94] and the CoOp with a CLIP-Adapter [17]. Then, for



(a) Effect of α on Accuracy and Backward Transfer (BWT).

(b) Avg. rotation angle [53] per step between image and text features.

Figure 2. **Need for Visual-guided Attention (VGA) inference module.** Fig. 2a: A simple adapter is inadequate at preventing catastrophic forgetting in CL – marked by high BWT scores; Fig. 2b: VGA module encourages cross-modal alignment among text and image features – marked by a decrease in average angle $\arccos\langle t, 1 \rangle$ between them – where otherwise the former deviates further with incremental training steps.

the base task ($t = 1$) test samples of CIFAR100, we compute the average of the Rotation Angle Matrix (RAM) [53] using the CL models’ frozen visual $f(\mathbf{x})$ and learnable textual $g(\mathbf{t}_c(\mathbf{p}))$ features at each incremental test step. Fig. 2b shows the deviation of the learned textual features from their (frozen) visual counterparts for the CoOp. This implies that the cross-modal retrieval performance of CLIP finetuned with learnable prompt deteriorates with incremental training. Moreover, a simple adapter (CoOp + Adapter) does not remedy the cross-modal deviation.

3.4.1 Modelling with visual-guided text features

To alleviate the cross-modal deviation due to incremental training, we consider explicit enriching of communication among the encoders’ features. To this end, we adopt a standard transformer-styled decoder block [71] as a visual-guided attention (VGA) module that refines the task-specific text features based on the visual features as context. While in a non-CL setup, the VGA module has been shown to guide the text features to attend to the informative per-pixel spatial features (obtained before global average-pooling in a ViT [13]), here we use the globally pooled visual features of the ViT for guiding the text features. This setting reflects our primary motivation for aligning the text features which we can later employ for variational modelling. Further, using globally pooled visual features favors the computational efficiency of our framework for large CL datasets where attending to per-pixel spatial features can incur much higher latency (App. 8.3).

Leveraging the task-specific text features $g(\mathbf{t}_c^t)$ as queries Q and the visual global context $f(\mathbf{x})$ as keys K and values V , the VGA performs context-context self-attention followed by context-query cross-attention. To eliminate the influence of the queries from other tasks, a *naive* strategy is to perform specialized VGA forward passes [14] with task-specific query. Instead, we replace several such costly passes with a single pass. To do so, we exploit the global nature of our context and mask out (set to $-\infty$) all inter-task connections in the query using a target mask. This ensures that only the task-specific text features in the query undergo

Effect of task-specific encoders on inter-class centroid distances




Figure 3. **Need for task-specific distribution encoders:** Cosine distance between the centroids of class-specific latent variables produced without (**left**) and with (**right**) the use of task-specific encoders on CIFAR100 (10 tasks, 10 classes per task). Centroids of class-specific latent variables are more separable on the right.

self-attention while the entire query cross-attends to the visual context:

$$\begin{aligned} O &= \{O^k\}_{k=1}^t \\ &= \{\hat{\mathbf{t}}_c^k\}_{k=1}^t \\ &= \text{VGA}(Q = \{\mathbf{t}_c^k\}_{k=1}^t, K = V = f(\mathbf{x})), \end{aligned} \quad (5)$$

where O is the set of task-specific visual-guided text features $\hat{\mathbf{t}}_c^t$. We fuse the original task-specific text features \mathbf{t}_c^t to derive the task embeddings $\tilde{\mathbf{t}}_c^t$:

$$\tilde{\mathbf{t}}_c^t = \hat{\mathbf{t}}_c^t + \mathbf{t}_c^t. \quad (6)$$

3.5. Task-specific distribution encoders

Instead of feeding the concatenated text encoding $\{\tilde{\mathbf{t}}_c^1, \tilde{\mathbf{t}}_c^2, \dots, \tilde{\mathbf{t}}_c^t\}$ to one shared adapter q_ϕ , we maintain task-specific encoders $\{q_\phi^1, q_\phi^2, \dots, q_\phi^t\}$; each comprising a probabilistic adapter to parameterize the t -th task-specific posterior distribution $\mathcal{N}(\mu^t, \sigma^t)$ over its task embeddings $\tilde{\mathbf{t}}_c^t$:

$$\{z_m^t\}_{m=1}^M \sim q_\phi^t(z|\tilde{\mathbf{t}}_c^t) = \mathcal{N}(\mu^t(\tilde{\mathbf{t}}_c^t), \sigma^t(\tilde{\mathbf{t}}_c^t)), \quad (7)$$

where $\{z_m^t\}_{m=1}^M$ are the M task-specific MC samples. Task-specific encoders help us sample latent variables that are more discriminative across tasks. This is depicted in Fig. 3 using the cosine distance between the embeddings of class-specific samples. With task-specific encoders (right), the cross-task class centroids are further apart than those with one global encoder (left).

During each incremental training step ($t > 1$), we keep the past task encoders frozen. Lastly, to reduce the forgetting of past task encoders, we follow other parameter-isolation techniques [5, 14, 82] to finetune on a class-balanced dataset of new data and rehearsal data \mathcal{M} at the end of each incremental training step ($t > 1$). We refer to this as memory consolidation training (see App. 8.2).

3.6. Alleviating forgetting with language-aware CLIP knowledge

Finetuning methods are prone to trading generalization for downstream task performances [93, 94]. On the contrary, a pre-trained CLIP with hand-crafted prompts has strong

generalizability. To alleviate forgetting (*i.e.*, losing generalization) in CL, we propose to leverage the rich language information captured by the text encoder with hand-crafted prompts. In the following, we assume $\{\mathbf{t}_y^{h,l} \in \mathbb{R}^d\}_{l=1}^L$ to be the features corresponding to the L hand-crafted textual prompts for the class $y \in C^t$.

3.6.1 Weight initialization

An informed weight initialization can help bridge the **stability gap** [37] in CL by stabilizing the learning for new task components over random initialization [20]. We thus leverage the t -th task text features $\{\mathbf{t}_y^{h,l}\}_{l=1}^L$ to initialize the weights $\mathbf{w}_t^\mu, \mathbf{w}_t^\sigma \in \mathbb{R}^{d \times d}$ of our t -th task mean and std. dev. modules. Let $\mathbf{s}_\mu, \mathbf{s}_\sigma \in \mathbb{R}^{|C^t| \times d}$ be the mean and the std. dev. of the L text features. We initialize \mathbf{w}_t^μ and \mathbf{w}_t^σ as:

$$\mathbf{w}_t^\mu = \frac{1}{d} \langle \mathbf{s}_\mu, \mathbf{s}_\mu^T \rangle, \mathbf{w}_t^\sigma = \frac{1}{d} \langle \mathbf{s}_\sigma, \mathbf{s}_\sigma^T \rangle. \quad (8)$$

3.6.2 Past-task distribution regularization

The functional spaces of the past task distributions are prone to forgetting in CL. Though replay helps alleviate the forgetting up to a certain degree, repeated training on the memory samples can lead to overfitting on these [43, 72]. To address this, previous works [27, 58] exploit functional priors for regularizing the visual space alongside memory replay. Here, we propose regularizing the past task distributions in the textual space by using the features $\{\mathbf{t}_y^{h,l}\}_{l=1}^L$ to distill the past task samples $\{\{z_m^i\}_{m=1}^M\}_{i=1}^{t-1}$. Namely, the probability of z_m^t belonging to class $y \in C^t$ is:

$$P_{\text{KD}}(y|z_m^t) = \frac{1}{M} \sum_{m=1}^M \frac{1}{L} \sum_{l=1}^L \frac{\exp(\langle \mathbf{t}_y^{h,l}, z_m^t \rangle)}{\sum_{c=1}^{|C^t|} \exp(\langle \mathbf{t}_c^{h,l}, z_m^t \rangle)}. \quad (9)$$

The resulting language-aware distillation loss using P_{KD} is thus given as:

$$\mathcal{L}_{\text{KD}} = - \sum_{t=1}^{T-1} \sum_{c=1}^{|C^t|} \log P_{\text{KD}}(c|z^t) y_c, \quad (10)$$

where \mathcal{L}_{KD} is a **data-free** (*i.e.*, no training samples required) text-to-text distribution regularizer that encourages the latent variables to stay close to the text features from hand-crafted prompts. Lastly, as \mathcal{L}_{KD} acts on the functional space of past tasks, this sets apart our setting from the non-CL setup of Bulat *et al.* [4] where the language-aware distillation loss regularizes the vector space.

3.7. Training objective and framework overview

3.7.1 Approximate ELBO

Unlike Eq. (4), the task-specific encoders q_ϕ^t now serve to approximate the intractable task-specific functional posteriors. The resulting ELBO (see App. 10 for derivation) with the static prior $p_\chi = \mathcal{N}(\mathbf{0}, \mathbf{I})$ is:

Algorithm 1: A forward CLAP4CLIP pass at test step t

```

Input      :  $\{\mathbf{t}^i\}_{i=1}^t$ : text features,  $f(x)$ : image features
Output    :  $\hat{y}^{1:t}$  (predictions for classes seen till task  $t$ )
1  $\{\mathbf{t}^i\}_{i=1}^t \leftarrow \text{VGA}(\{\mathbf{t}^i\}_{i=1}^t, f(x))$  // Eq. (5)
2 for  $i \leftarrow 1$ ;  $i \leq t$ ;  $i += 1$  do
3    $\tilde{\mathbf{t}}^i = \mathbf{t}^i + \mathbf{t}^i$  // Eq. (6)
4    $\mathcal{N}(\mu^i, \sigma^i) \leftarrow q_\phi^i(\tilde{\mathbf{t}}^i)$  // Sec. 3.5
5    $\hat{y}^i \leftarrow \emptyset$  // Null prediction set
6   for  $m \leftarrow 1$ ;  $m \leq M$ ;  $m += 1$  do
7      $z_m^i \sim \mathcal{N}(\mu^i, \sigma^i)$  // Sampling
8      $\tilde{\mathbf{t}}_m^i \leftarrow \tilde{\mathbf{t}}^i + z_m^i$  // Late fusion
9      $\hat{y}_m^i \leftarrow \langle f(\mathbf{x})^T, \tilde{\mathbf{t}}_m^i \rangle$  // Using Eq. (1)
10     $\hat{y}^i \leftarrow \hat{y}^i \cup \hat{y}_m^i$  // Set Union
11  $\hat{y}^{1:t} \leftarrow [\hat{y}^1, \dots, \hat{y}^t]$  // Concatenation

```

$$\log p(y^{1:T} | \mathbf{x}; \tilde{\mathbf{t}}_c) \geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi^t(z^t | \mathbf{x}; \tilde{\mathbf{t}}_c)} \left[\log p_\theta(y^t | z^t, \mathbf{x}; \tilde{\mathbf{t}}_c) \right] - \mathbb{D}_{\text{KL}}(q_\phi^t(z^t | \mathbf{x}; \tilde{\mathbf{t}}_c) \| p_\chi(z^t)) \right]. \quad (11)$$

3.7.2 Integrated objective

Denoting the loss weights by λ and γ , our total loss consists of the following three terms:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} - \lambda \mathbb{D}_{\text{KL}} + \gamma \mathcal{L}_{\text{KD}}, \quad (12)$$

where the cross-entropy \mathcal{L}_{CE} and the prior-matching \mathbb{D}_{KL} terms act on the outputs of all task encoders while the distribution regularization term \mathcal{L}_{KD} acts only on the past task encoders during finetuning. λ is set to 0.001. As the past task encoders are trainable only during the memory consolidation training stage, λ for these is set to 0 during training. γ is set to 15.

3.7.3 Algorithm overview

Algo. 1 outlines the pseudo-code of a forward pass of CLAP4CLIP at the t -th task test step. Here, a test image is to be classified among one of the seen classes $\{1, \dots, |C^t|\}$. Our method executes the computationally heavy VGA layers only once. The task-specific VGA outputs are passed through their respective task encoders. By limiting the quadratic complexity of the VGA module pass, our method induces minimal time overhead. Moreover, by only expanding the task encoder per task, our memory overhead is negligible compared to the large backbone of the pre-trained CLIP model (see Fig. 5b).

4. Experiments

Datasets. We evaluate our method on CIFAR100 [65, 85], ImageNet100 [24, 80], ImageNet-R [76], CUB200 [73], and VTAB [73]. CIFAR100 [34] and ImageNet100 [35] setups split their respective original datasets into 10 tasks with 10 classes each. ImageNet-R [23] and CUB200 split 200 classes into 10 tasks with 20 classes each. VTAB has

5 tasks with 10 classes each [91]. While CIFAR100, ImageNet100, and CUB200 are robust settings for evaluating CL methods in the face of large forgetting, ImageNet-R and VTAB are rather fair yet challenging settings for CL methods using pre-trained models as these might include test images in their pre-training set (see App. 8.1 for details).

Baselines. We compare CLAP4CLIP against several baselines and state-of-the-art finetuning methods. These include: (a) CLIP-based methods – Continual-CLIP [70], CoOp [94], CLIP-Adapter [17], AttriCLIP [74], MaPLE [29], and PROOF [92], (b) pre-trained vision-only methods – DualPrompt [76] and L2P [77], (c) the baseline CIL method – iCaRL [65]. For a fair comparison, we adhere to the experimental protocols of PROOF [92] throughout. We adopt ViT-B/16 with the pre-trained weights of OpenAI [63] as our backbone unless otherwise specified. As the upper bounds on performance, we use the CLAP4CLIP with single and task-specific encoders, trained on all tasks jointly (JOINT).

Variants. We integrate our method with four prompt-based approaches: Ours uses CLAP with hand-crafted prompt templates, CoOp + Ours with soft prompts [94], MaPLE + Ours uses multi-modal soft prompts [29], and AttriCLIP + Ours uses CLAP4CLIP with instance-conditioned soft prompts [74]. Lastly, Ours w/o Variational Inference (VI) is the deterministic variant of Ours (App. fig. 7).

Implementation and training. We train the model using SGD, with a batch size of 64, for 5 epochs, including 1 epoch of linear warmup. The initial learning rate (LR) is set to 1e-3 and decays with cosine annealing. At the end of each incremental task ($t > 1$), we perform memory consolidation training for 2 epochs, with an LR of 1e-4, on the class-balanced memory dataset (see App. 8.2).

Performance measure. We report the final accuracy after the last incremental step (Last) and the average of the accuracies after each step (Avg) [65].

4.1. Results

We report performances in Table 1 on all five datasets. Our method consistently achieves the best results among all the methods compared. Notably, on CIFAR100 and ImageNet100, our variants using the hand-crafted and multi-modal prompts outperform the others. On the challenging ImageNet-R setup with significant intra-class diversity, our method can better leverage the instance-conditioned prompt knowledge of AttriCLIP [74], which helps it outperform PROOF [92] by 1.46% in terms of average accuracy. On CUB200 and VTAB, sharing the prompt pool among all tasks gives CoOp [94] an edge over other baselines. Leveraging CoOp offers us the best results on these, while surpassing PROOF, which also builds upon CoOp with task-

Method	CIFAR100		ImageNet100		ImageNet-R		CUB200		VTAB	
	Avg \uparrow	Last \uparrow	Avg \uparrow	Last \uparrow	Avg \uparrow	Last \uparrow	Avg \uparrow	Last \uparrow	Avg \uparrow	Last \uparrow
Single-task JOINT	80.28		81.08		80.92		75.4		89.29	
Task-specific JOINT	82.9		83.55		83.07		85.72		94.6	
iCaRL [65]	72.93	57.6	68.62	59.5	66.34	43.71	82.39	75.1	53.38	41.6
L2P [77]	78.92	70.04	-	-	77.07	69.33	76.98	68.47	-	-
DualPrompt [76]	82.11	74.31	-	-	82.73	76.41	82.37	76.29	-	-
PROOF [92]	84.84	76.55	-	-	84.89	79.7	83.98	79.35	-	-
Continual-CLIP [70]	78.65	68.26	83.99	74.2	84.43	76.94	67.0	54.8	68.5	60.97
CoOp [94]	81.17	70.58	79.14	64.9	84.7	78.66	76.62	68.53	87.06	81.25
MaPLE [29]	82.74	74.52	79.23	64.06	85.28	79.71	73.38	64.43	83.91	81.81
AttriCLIP [74]	79.31	68.45	82.29	70.76	83.09	76.53	65.26	52.12	71.84	64.09
CLIP-Adapter [17]	78.75	68.32	84.13	73.96	84.49	78.1	67.41	54.49	68.23	61.02
Ours w/o VI	84.36	76.8	86.11	76.48	85.69	79.83	72.21	61.87	90.74	88.64
Ours	86.13	78.21	87.76	79.16	85.77	79.98	86.93	81.64	91.37	89.67
CoOp + Ours	85.71	77.4	86.8	78.18	85.32	79.52	86.99	81.95	92.51	91.28
MaPLE + Ours	86.06	78.48	87.47	79.02	86.25	80.56	81.53	74.24	90.97	88.83
AttriCLIP + Ours	78.06	67.59	87.37	79.3	86.35	80.6	83.71	79.01	74.84	71.12

Table 1. **Performance comparison** of different methods averaged over three runs. Best scores are in **bold**. Second best scores are in **blue**. The results for L2P, DualPrompt, and PROOF are taken from [92]. See App. Table 14 for std. dev. scores.




Figure 4. **Performance evolution of different methods.** The top-1 accuracy (%) is reported upon learning of each task (see App. fig. 8 for results on all datasets).

specific soft prompts. Lastly, Fig. 4 shows that our variants perform consistently better than their baseline prompting methods throughout the incremental steps.

Backward transfer (BwT) [46] is a crucial metric to monitor *forgetting*. Table 15 shows that in general, plugging CLAP4CLIP with prompt-based finetuning methods helps improve the BwT scores of the latter. It is worth noting that on the cross-dataset setting of VTAB [91], our variants are the only methods that effectively transfer the knowledge learned from incremental tasks to improve the performance on past tasks (*i.e.*, $BwT > 0$). This indicates that our probabilistic modeling strategy does not only counter forgetting but can also help bring anti-forgetting properties onto existing finetuning methods.

Transfer score [89] measures the extent of zero-shot transfer ability of finetuning methods. App. table 16 shows that our method consistently enhances the transfer scores of the underlying prompting algorithms to help solve the future tasks through increased positive knowledge transfer.

Expected calibration error (ECE) [52] measures the calibration of a CL model’s predictions to deem its fitness for high-risk applications [51]. App. table 17 compares the ECE of our variants and their respective underlying baselines at the last test step. In general, our variants help enhance (decrease) the ECE scores of the underlying prompting methods. This implies that even in the face of forgetting, CLAP retains more reliability in its predictive confidence.

4.2. Cross-Datasets Continual Learning (CDCL)

To simulate real-world settings with long sequence of tasks and large distribution shifts, the CDCL setting [74] trains a model sequentially on ImageNet100 and CIFAR100 (*i.e.*, on 20 tasks), and evaluates it jointly on these. For a fair comparison with [74], we adopt the ViT-L/14 as our CLIP backbone and set the train/test batch size to 32. All other settings remain the same as in Sec. 4.1. Table 2 reports the last task accuracy of different methods. While all our variants improve the CDCL performances of their respective baselines, combining ours with AttriCLIP [74] leads to the most gains. This further suggests that our framework can reliably leverage the diverse nature of learned prompts

Method	ImageNet100	CIFAR100 + ImageNet100
DualPrompt [76]	81.9	67.1
Continual-CLIP [70]	75.4	54.9
CoOp [94]	79.3	55.4
MaPLe [29]	84.81	76.2
AttriCLIP [74]	83.3	78.3
Ours	83.51	83.83
CoOp + Ours	82	82.63
MaPLe + Ours	82.97	83.6
AttriCLIP + Ours	84.14	84.56

Table 2. **Performance comparison on the CDCL setting [74].** All CLIP-based methods use the ViT-L/14 backbone.

to inherit their setting-specific advantages.

4.3. Ablation Studies

4.3.1 Influence of components.

We ablate the importance of different components of CLAP4CLIP in Table 3. On top of the base CLIP model, we first train a probabilistic encoder. Adding the VGA module and the memory consolidation training stage helps us achieve much stable performances while countering forgetting. We then apply task-specific encoders which makes the centroids of the class-specific latent variable more separable (see Fig. 3) thus improving the last task accuracy by 2.21%. Language-aware weight initialization and regularization help improve the last task accuracies by 0.78% and 0.23%, respectively. Weight initialization further helps us tackle the *stability gap* [20, 37] (see App. 8.5.4 for more ablations on language-aware components).

ID	Probabilistic Encoder	VGA	Consolidation training	Task-specific Encoders	Weight Initialization	Distribution Regularization	Avg	Last
#1	✓						22.78	11.49
#2	✓	✓					82.82	73.41
#3	✓	✓	✓				84.4	74.99
#4	✓	✓	✓	✓			85.7	77.2
#5	✓	✓	✓	✓	✓		86.01	77.98
#6	✓	✓	✓	✓	✓	✓	86.13	78.21

Table 3. **Ablations** of the key components of CLAP4CLIP on CIFAR100.

4.3.2 Effect of forgetting on task heads

We ablate the task head predictions over the test set of each task on CIFAR100 (see App. 8.5.1 for more details). Fig. 9a shows that the past task heads are relatively less accurate at identifying their respective test instances. However, their amount of forgetting do not necessarily correspond to the task order. This is further evident in Fig. 9b where at each step, the past task heads lag unevenly in performance. The uneven performances could be because of the role of task semantics behind forgetting [40, 64].

4.3.3 Probabilistic vs Deterministic inference

To understand our probabilistic inference modules further, we examine their performance against the deterministic

Prior type	Last ↑	Avg ↑	BWT ↑	ECE ↓	Runtime per iter. ↓
Static	78.21	86.13	-0.141	0.204	0.169
Data-driven	78.32	86.15	-0.115	0.216	0.172
Language-aware	78.38	86.22	-0.112	0.214	0.17

Table 4. **Performances of different priors** averaged over 3 runs on CIFAR100.

variant of ours (Ours w/o VI). Table 1 shows that our probabilistic inference module consistently outperforms its deterministic counterpart in terms of Avg and Last accuracy. This emphasizes the advantages of considering uncertainty in finetuning. We further introspect the effects of the number of layers for the VGA and task encoder modules in our framework in App. 8.5.2.

4.3.4 Sensitivity to the number of Monte Carlo (MC) samples

We vary the number of MC samples M from 1 to 50. In Fig. 5a, the accuracy is poorer in range $[1, 10]$, grows in range $[10, 20]$, and saturates thereafter. Hence, we set M to 20 for all our experiments.

4.3.5 More informed priors

To study the role of a more informed prior in our VI framework, we study three choices of priors to be used in the prior-matching term of eq. (11): (a) the static (standard normal) prior, (b) the language-aware prior using the distribution obtained from the task encoders using the hand-crafted prompts’ features $\{t_y^{h,l}\}_{l=1}^L$ (sec 3.6), (c) the data-driven prior using a randomly chosen subset of a training minibatch as the context set to condition the prior on (see App. 11 for more details). Table 4 shows that while (b) and (c) slightly improve over (a) in terms of accuracies and forgetting (BwT), these come at the cost of poorer model calibration and longer runtime per iteration.

4.3.6 Parameter analyses

The additional parameters in CLAP4CLIP come from the shared VGA module and the task-specific encoders. For a ViT-B/16 backbone of output dimension $d = 512$ on CIFAR100, the VGA module contains 4,204,032 parameters. The mean and the std. dev. layers for 10 tasks have $d \times d$ parameters each, *i.e.*, 524,2880 parameters. Hence, the CLAP4CLIP has 9.5 million extra parameters, which is negligible compared to the pre-trained CLIP with ≈ 150 million parameters. We report the parameter counts in fig. 5b.

4.3.7 Time analyses

We investigate the inference time per iteration for different methods. As shown in Table 18, our variants need more inference time than other finetuning methods for the performance gains. The increased time comes mainly from the VGA and from inferring the M latent variables.




Figure 5. **Ablations on CIFAR100 showing:** (a) performance trade-off with the number of MC samples M , (b) the number of trainable parameters in different finetuning methods.

5. Added utilities of probabilistic finetuning

We study the out-of-the-box utilities of CLAP4CLIP’s uncertainty quantification (UQ) capabilities. Our motivation for these is not to achieve state-of-the-art performance but to highlight the perks of probabilistic modeling in scenarios where the deterministic CL finetuning methods struggle.

5.0.1 Post-hoc novel data detection (PhNDD)

PhNDD uses a pre-trained classification model to identify novel data based on the output confidence [21, 83]. For CL, this can help discern the arrival of new tasks, expand the network, etc. To evaluate the PhNDD capabilities of models within a CL setup, we design a simple setting. Namely, at all but the last test step, we treat the test data from the past and the current tasks as *seen* while those from all future tasks as *novel*. We then use FPR95, AUROC [11], and AUPR [66] scores as our performance metrics (see App. 8.6 for more details). We average these metrics over all but last incremental test steps.

Method	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow
Continual-CLIP [70]	74.46	71.11	77.33
Ours w/o VI	82.29	78.88	68.83
Ours	82.21	79.54	68.72
CoOp [94]	80.15	77.62	66.8
+ Ours w/o VI	81.98	78.88	66.21
+ Ours	83.73	80.97	62.68

Table 5. PhNDD performances averaged over 3 runs on CIFAR100. Best scores for each variant are in **bold**.

To quantify the output confidence, we rely on the Energy score [45] given its aptness for pre-trained models. Table 5 compares the averaged PhNDD performances. Our probabilistic models enhance the PhNDD capabilities of their underlying prompting frameworks. Lastly, the inferior results of the deterministic (*i.e.*, w/o VI) versions of our models imply that considering uncertainty-awareness for PhNDD helps us maintain richer and more diverse features [88].

5.0.2 Exemplar selection

To study uncertainty-based exemplar selection, we first employ entropy of the output softmax distributions as our selection criteria [6]. Table 6 shows the efficacy of entropy-based rehearsal for our method, where other deterministic methods lag behind due to their inconsistent UQ capabilities. Next, we employ the energy [45] and the variance of

Method	Avg	Last
CoOp [94]	76.71	64.1
Clip-Adapter [17]	78.78	68.49
Ours w/o VI	84.44	76.55
Ours	85.18	77.92

Table 6. **Entropy-based exemplar selection results for different methods on CIFAR100.**

softmax distributions of the M predictions as our selection criterion, and contrast these against other criteria proposed in [6]. Figure 6 shows that variance-based exemplar selection outperforms random, and is only second to iCaRL [65] in terms of Last accuracy. Lastly, we note that deterministic finetuning methods are inherently limited at leveraging variance for exemplar selection.

6. Conclusion

In this paper, we propose CLAP4CLIP, a probabilistic finetuning method for learning task-specific distributions over visual-guided textual features. Our model shares the visual-guided text alignment module across all tasks while adding lightweight task-specific encoders to learn fine grained task distributions. Besides leading to little memory overhead, this architecture is compatible with several prompt-tuning based methods thus helping us inherit their respective perks on different CL settings. Our experiments show the superior results of CLAP4CLIP across several datasets and settings. We conclude with two out-of-the-box utilities of our method wherein existing continual learning methods lag: post-hoc novel data detection and uncertainty-based exemplar selection.

7. Future research directions

Few potential directions of research for CLAP4CLIP include the design of: (a) parameter-efficient adapters [25] for very large CL settings; (b) better regularization techniques to alleviate forgetting; and (c) more informed [9] yet computationally efficient priors for inference. Similarly, along the direction of alleviating forgetting and mitigating the stability gap [20, 37], it would be interesting to see how class-specific prompts generated by pre-trained Large Language Models (LLMs) can be exploited to obtain task-relevant language-aware CLIP knowledge while preserving the zero-shot transfer ability of the learned prompts (see App. table 19 for a preliminary investigation). Lastly, we consider applying CLAP4CLIP to more sophisticated Vision-Language tasks [68] as another possible direction for research (see App. table 20 for a preliminary investigation). We elaborate further on each of these directions in App. 9.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 1, 2
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, pages 139–154, 2018. 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 10
- [4] Adrian Bulat and Georgios Tzimiropoulos. Language-aware soft prompting for vision & language foundation models, 2023. 5
- [5] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 233–248, 2018. 2, 5, 1
- [6] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, pages 532–547, 2018. 1, 9
- [7] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 1, 2
- [8] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 1
- [9] Youngjae Cho, HeeSun Bae, Seungjae Shin, Yeo Dong Youn, Weonyoung Joo, and Il-Chul Moon. Make prompts adaptable: Bayesian modeling for vision-language prompt learning with data-dependent prior. *arXiv preprint arXiv:2401.06799*, 2024. 9
- [10] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. 1
- [11] Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240, 2006. 9, 7
- [12] Mohammad Mahdi Derakhshani, Enrique Sanchez, Adrian Bulat, Victor G Turrissi da Costa, Cees GM Snoek, Georgios Tzimiropoulos, and Brais Martinez. Bayesian prompt learning for image-language model generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15237–15246, 2023. 1, 2, 4
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 4
- [14] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9285–9295, 2022. 2, 4, 5, 1
- [15] Sebastian Farquhar and Yarin Gal. Towards robust evaluations of continual learning. *International Conference on Machine Learning (ICML) Workshop*, page 9, 2018. 1
- [16] Vincent Fortuin. Priors in bayesian deep learning: A review. *International Statistical Review*, 90(3):563–591, 2022. 13
- [17] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 1, 2, 3, 4, 6, 7, 9, 5
- [18] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018. 4, 13, 14
- [19] Andrew Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). 2006. 13
- [20] Md Yousuf Harun and Christopher Kanan. Overcoming the stability gap in continual learning. *arXiv preprint arXiv:2306.01904*, 2023. 5, 8, 9, 6
- [21] Jiangpeng He and Fengqing Zhu. Out-of-distribution detection in unsupervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3850–3855, 2022. 9, 7
- [22] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. 1
- [23] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 6
- [24] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 831–839, 2019. 2, 6, 1
- [25] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 9, 10

- [26] Saurav Jha, Martin Schiemer, Franco Zambonelli, and Juan Ye. Continual learning in sensor-based human activity recognition: An empirical benchmark analysis. *Information Sciences*, 575:1–21, 2021. [1](#)
- [27] Saurav Jha, Dong Gong, He Zhao, and Lina Yao. NPCL: Neural processes for uncertainty-aware continual learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [1](#), [2](#), [5](#), [13](#), [14](#)
- [28] Haeyong Kang, Rusty John Lloyd Mina, Sultan Rizky Hikmawan Madjid, Jaehong Yoon, Mark Hasegawa-Johnson, Sung Ju Hwang, and Chang D Yoo. Forget-free continual learning with winning subnetworks. In *International Conference on Machine Learning*, pages 10734–10750. PMLR, 2022. [2](#)
- [29] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [5](#)
- [30] Muhammad Uzair Khattak, Muhammad Ferjad Naeem, Muzammal Naseer, Luc Van Gool, and Federico Tombari. Learning to prompt with text only supervision for vision-language models. *arXiv preprint arXiv:2401.02418*, 2024. [10](#), [11](#)
- [31] Donggyun Kim, Seongwoong Cho, Wonkwang Lee, and Seunghoon Hong. Multi-task processes. In *International Conference on Learning Representations*, 2022. [13](#)
- [32] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [4](#)
- [33] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. [1](#), [2](#)
- [34] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [6](#)
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. [6](#)
- [36] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. *ICLR*, 2022. [1](#)
- [37] Matthias De Lange, Gido M van de Ven, and Tinne Tuytelaars. Continual evaluation for lifelong learning: Identifying the stability gap. In *The Eleventh International Conference on Learning Representations*, 2023. [5](#), [8](#), [9](#), [6](#)
- [38] Tuan Anh Le, Hyunjik Kim, Marta Garnelo, Dan Rosenbaum, Jonathan Schwarz, and Yee Whye Teh. Empirical evaluation of neural process objectives. In *NeurIPS workshop on Bayesian Deep Learning*, 2018. [13](#)
- [39] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62, 2022. [1](#)
- [40] Sebastian Lee, Sebastian Goldt, and Andrew Saxe. Continual learning in the teacher-student setup: Impact of task similarity. In *Proceedings of the 38th International Conference on Machine Learning*, pages 6109–6119. PMLR, 2021. [8](#)
- [41] Timothee LESORT and Andrei Stoian. Regularization shortcomings for continual learning, 2021. [2](#)
- [42] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020. [11](#)
- [43] Dongyue Li and Hongyang Zhang. Improved regularization and robustness for fine-tuning in neural networks. In *Advances in Neural Information Processing Systems*, 2021. [5](#)
- [44] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [2](#)
- [45] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in neural information processing systems*, 33:21464–21475, 2020. [9](#), [7](#)
- [46] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. [7](#)
- [47] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5206–5215, 2022. [1](#), [2](#)
- [48] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost Van De Weijer. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5513–5533, 2022. [1](#), [2](#)
- [49] Mark D McDonnell, Dong Gong, Amin Parvaneh, Ehsan Abbasnejad, and Anton van den Hengel. Ranpac: Random projections and pre-trained models for continual learning. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#)
- [50] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *The Eleventh International Conference on Learning Representations*, 2023. [10](#)
- [51] Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. *Advances in Neural Information Processing Systems*, 34:15682–15694, 2021. [7](#)
- [52] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*, 2015. [7](#)
- [53] Zixuan Ni, Longhui Wei, Siliang Tang, Yueting Zhuang, and Qi Tian. Continual vision-language representation learning with off-diagonal information, 2023. [4](#)
- [54] M-E Nilsback and Andrew Zisserman. A visual vocabulary for flower classification. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR’06)*, pages 1447–1454. IEEE, 2006. [1](#)

- [55] Bernt Øksendal and Bernt Øksendal. *Stochastic differential equations*. Springer, 2003. 13
- [56] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 1
- [57] Oleksiy Ostapenko, Pau Rodriguez, Massimo Caccia, and Laurent Charlin. Continual learning via local module composition. *Advances in Neural Information Processing Systems*, 34:30298–30312, 2021. 2
- [58] Pingbo Pan, Siddharth Swaroop, Alexander Immer, Runa Eschenhagen, Richard Turner, and Mohammad Emtiyaz E Khan. Continual deep learning by functional regularisation of memorable past. In *Advances in Neural Information Processing Systems*, pages 4453–4464. Curran Associates, Inc., 2020. 5
- [59] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3498–3505, 2012. 1
- [60] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Kenji Kawaguchi, Hanxiao Liu, Adams Wei Yu, Jiahui Yu, Yi-Ting Chen, Minh-Thang Luong, Yonghui Wu, et al. Combined scaling for zero-shot transfer learning. *Neurocomputing*, page 126658, 2023. 2
- [61] Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15691–15701, 2023. 10
- [62] Longtian Qiu, Renrui Zhang, Ziyu Guo, Ziyao Zeng, Yafeng Li, and Guangnan Zhang. Vt-clip: Enhancing vision-language models with visual-guided texts. *arXiv preprint arXiv:2112.02399*, 2021. 2
- [63] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2, 6
- [64] Vinay V Ramasesh, Ethan Dyer, and Maithra Raghu. Anatomy of catastrophic forgetting: Hidden representations and task semantics. *ICML workshop on Continual Learning*, 2020. 8
- [65] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 2, 6, 7, 9, 1
- [66] Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3): e0118432, 2015. 9, 7
- [67] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11909–11919, 2023. 2
- [68] Tejas Srinivasan, Ting-Yun Chang, Leticia Pinto Alva, Georgios Chochlakis, Mohammad Rostami, and Jesse Thomason. Climb: A continual learning benchmark for vision-and-language tasks. *Advances in Neural Information Processing Systems*, 35:29440–29453, 2022. 9, 11
- [69] Tejas Srinivasan, Furong Jia, Mohammad Rostami, and Jesse Thomason. I2i: Initializing adapters with improvised knowledge. *arXiv preprint arXiv:2304.02168*, 2023. 11
- [70] Vishal G. Thengane, Salman A. Khan, Munawar Hayat, and Fahad Shahbaz Khan. Clip model is an efficient continual learner. *ArXiv*, abs/2210.03114, 2022. 1, 2, 4, 6, 7, 8, 9, 5
- [71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4, 5
- [72] Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9385–9394, 2021. 2, 5
- [73] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [74] Runqi Wang, Xiaoyue Duan, Guoliang Kang, Jianzhuang Liu, Shaohui Lin, Songcen Xu, Jinhua Lü, and Baohang Zhang. Attriclip: A non-incremental learner for incremental knowledge learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3654–3663, 2023. 2, 3, 4, 6, 7, 8, 1, 5
- [75] Zifeng Wang, Zheng Zhan, Yifan Gong, Geng Yuan, Wei Niu, Tong Jian, Bin Ren, Stratis Ioannidis, Yanzhi Wang, and Jennifer Dy. Sparcl: Sparse continual learning on the edge. *Advances in Neural Information Processing Systems*, 35:20366–20380, 2022. 10
- [76] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *European Conference on Computer Vision*, pages 631–648. Springer, 2022. 6, 7, 8
- [77] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 6, 7
- [78] Max Welling. Herding dynamical weights to learn. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1121–1128, 2009. 3, 1
- [79] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. 2

- [80] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 374–382, 2019. 2, 6
- [81] Qingsen Yan, Dong Gong, Yuhang Liu, Anton van den Hengel, and Javen Qinfeng Shi. Learning bayesian sparse networks with full experience replay for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 109–118, 2022. 10
- [82] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023, 2021. 5, 1
- [83] Jingkan Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021. 9
- [84] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *Transactions on Machine Learning Research*, 2022. 1, 2
- [85] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International conference on machine learning*, pages 3987–3995. PMLR, 2017. 1, 6
- [86] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867*, 2019. 1
- [87] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European Conference on Computer Vision*, pages 493–510. Springer, 2022. 2, 4
- [88] Xinwei Zhang, Jianwen Jiang, Yutong Feng, Zhi-Fan Wu, Xibin Zhao, Hai Wan, Mingqian Tang, Rong Jin, and Yue Gao. Grow and merge: A unified framework for continuous categories discovery. *Advances in Neural Information Processing Systems*, 35:27455–27468, 2022. 9
- [89] Z. Zheng, M. Ma, K. Wang, Z. Qin, X. Yue, and Y. You. Preventing zero-shot transfer degradation in continual learning of vision-language models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19068–19079, Los Alamitos, CA, USA, 2023. IEEE Computer Society. 7, 4
- [90] Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Deep class-incremental learning: A survey. *arXiv preprint arXiv:2302.03648*, 2023. 1
- [91] Da-Wei Zhou, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Revisiting class-incremental learning with pre-trained models: Generalizability and adaptivity are all you need. *arXiv preprint arXiv:2303.07338*, 2023. 6, 7, 1
- [92] Da-Wei Zhou, Yuanhan Zhang, Jingyi Ning, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Learning without forgetting for vision-language models. *arXiv preprint arXiv:2305.19270*, 2023. 2, 6, 7, 1
- [93] Kaiyang Zhou, Jingkan Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. 2, 4, 5
- [94] Kaiyang Zhou, Jingkan Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 1, 2, 3, 4, 5, 6, 7, 8, 9
- [95] Beier Zhu, Yulei Niu, Saeil Lee, Minhoe Hur, and Hanwang Zhang. Debaised fine-tuning for vision-language models by prompt regularization. *arXiv preprint arXiv:2301.12429*, 2023. 2

CLAP4CLIP: Continual Learning with Probabilistic Finetuning for Vision-Language Models

Supplementary Material

8. Experiments and Benchmarks

8.1. Datasets

We evaluate our method on five datasets, the details of which are reported in Table 7. Following Zhou *et al.* [92], we shuffle the order of training classes for all but the VTAB dataset with the random seed 1993. While the original VTAB [86] includes 19 evaluation tasks from three categories (*natural*, *specialized*, and *structured*) and their respective sub-domains, we rely on the five datasets cross-domain class-incremental subset proposed in SimpleCIL [91]. The five datasets (used in the same streaming order) include Resisc45 [8], DTD [10], Pets [59], EuroSAT [22], and Flowers [54]. To make the classes emerge from domain to domain, we do not shuffle the class order for VTAB.

8.1.1 Exemplar selection

Following [24, 65], we employ the herding algorithm [78] to choose the exemplars for our main experiments. Following the previous works [90, 92], we rely on two typical methods to populate the memory:

1. **Fixed memory budget** maintains a static memory \mathcal{M} with K instances. Upon having seen $|\mathcal{Y}_b|$ number of classes after an incremental training stage, the model selects $\frac{K}{|\mathcal{Y}_b|}$ exemplars per class.
2. **Expandable exemplar set** dynamically expands the memory \mathcal{M} with the arrival of more incremental tasks. After each incremental training stage, the model here stores $|\mathcal{Y}_b| \times k_c$ exemplars, where k_c is the number of exemplars per class.

For CIFAR100, ImageNet100 and VTAB, given their lesser number of classes, we employ the first policy, and keep a total of 2,000, 1,000, and 1,000 exemplars, respectively. This amounts to the respective sub-totals of 20 and 10 exemplars per class after the last incremental stage. We choose these sizes for a straightforward comparison with the existing works, *i.e.*, PROOF [92] for CIFAR100 and AttriCLIP [74] for ImageNet100. For VTAB, the chosen memory size reflects the fact that we have only 1,796 training instances in total (see Table 7). For ImageNet-R and CUB200 with 200 classes each, we adopt the second policy and store 20 exemplars per class.

8.2. Hyperparameter selection and tuning

To the end goal of obtaining task-agnostic hyperparameters [15], we tuned our hyperparameters using a validation set

comprising 10% of the CIFAR-100 training dataset. Similar to [14], performing the hyperparameter search only on the CIFAR100 setup helps us avoid optimizing for the number of tasks while generalizing across all our other setups. Table 8 shows the candidate values for the hyperparameter grid search and their best chosen values. Tables 9, 10, 11, and 12 report the last task accuracy scores (Last) corresponding to the hyperparameter search for the number of training epochs, the number of finetuning epochs, the coefficient γ and the coefficient λ , respectively. Fig. 5a in the main paper reports the accuracy and the runtimes for the different number of MC samples M . *We will release the full source code upon the acceptance of our paper.*

Training for memory consolidation. To alleviate the forgetting of past tasks, we finetune on the class-balanced dataset of new data and rehearsal data \mathcal{M} at the end of each incremental training step ($t > 1$) [14, 24]. Following other well-established parameter-isolation CL algorithms [5, 14, 82], we freeze the past task encoders during the normal training. This helps us avoid knowledge interference from the dominant new task training samples. During the memory consolidation training stage, we optimize all the task encoders while freezing the task-shared VGA module parameters.

8.3. Latency comparison for VT-CLIP styled VGA vs Ours

We compare the performance of VT-CLIP styled VGA with Ours. To align the text features with the image features, the former uses per-pixel spatial features obtained from the ViT prior to global pooling while we use the globally pooled features. Table 13 shows that VT-CLIP styled VGA achieves similar accuracy as ours while incurring $\approx 6\times$ higher inference time.

8.4. Results

8.4.1 Performance evolution

To complement the results in Table 1, Fig. 8 compares the accuracy of different methods at each evaluation step across all datasets. Our major conclusions are briefed as follows. A) The base task performance of CLAP4CLIP (ours) is consistently higher than other methods including the state-of-the-art PROOF [92]. This suggests that our probabilistic finetuning framework is effective for general downstream tasks in a non-incremental setting. B) For the CL settings in Table 1 where either of the CLAP4CLIP variants achieve the best performances, their performance

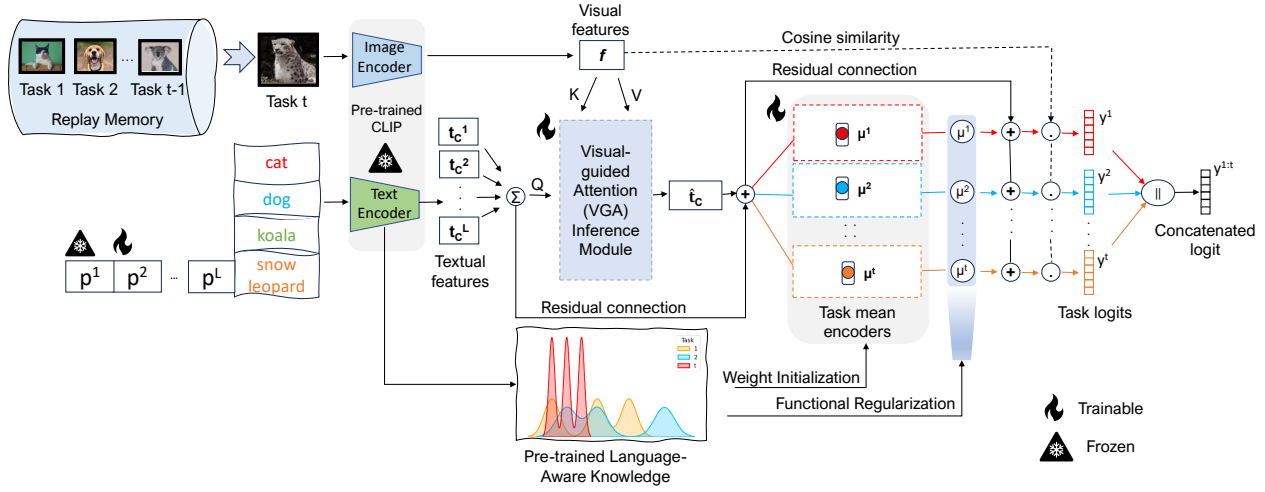


Figure 7. **Illustration of the deterministic variant of Ours (Ours w/o VI in Table 1)**: the task-specific text features are fed to their respective task encoders, consisting of only the mean μ layer each. There is no sampling involved and the task mean outputs are fused directly with the original task features prior to deriving the task logits y^t . All task logits are concatenated to produce the final prediction $y^{1:t}$.

Dataset	# training instances	# testing instances	# Classes	# Tasks	Link
CIFAR100	50,000	10,000	100	10	URL
ImageNet100	130,000	5,000	100	10	URL
ImageNet-R	24,000	6,000	200	10	URL
CUB200	9,430	2,358	200	10	URL
VTAB	1,796	8,619	50	5	URL

Table 7. **Benchmark datasets** and their details.

Hyperparameter	Range	Chosen value
Learning rate	$5e-3, 1e-3, 5e-4$	$1e-3$
Epochs	3, 5, 7	5
Warmup epochs	0.5, 1, 1.5	1
Finetuning epochs	1, 2, 3, 4	2
γ	1, 5, 10, 15, 20, 25	15
λ	0.0001, 0.001, 0.01, 0.1	0.001
M	1, 5, 10, 15, 20, 25, 30, 50	20

Table 8. **Hyperparameter tuning**: we run a gridsearch on the CIFAR100 setup with a validation set comprising 10% of the training set. The chosen values are reused across all other setups.

Epochs	3	5	7
Last	77.32	78.21	78.18

Table 9. Accuracy vs. Training epochs

Finetuning ep.	1	2	3	4
Last	77.65	78.21	78.2	78.18

Table 10. Accuracy vs. Finetuning epochs

γ	1	5	10	15	20	25
Last	78.04	77.94	78.1	78.21	77.96	77.14

Table 11. Accuracy vs. weight “ γ ” for \mathcal{L}_{KD}

λ	0.0001	0.001	0.01	0.1
Last	78.16	78.21	77.99	77.4

Table 12. Accuracy vs. weight “ λ ” for \mathbb{D}_{KL}

curves also consistently retain superior results across all evaluation steps. This validates the effectiveness of our method at tackling forgetting. C) Similar to Zhou *et al.* [92],

we notice that CLAP4CLIP achieves a significant performance improvement over vision-only methods (L2P and DualPrompt). This indicates the merits of considering text and visual cues together for continual learning.

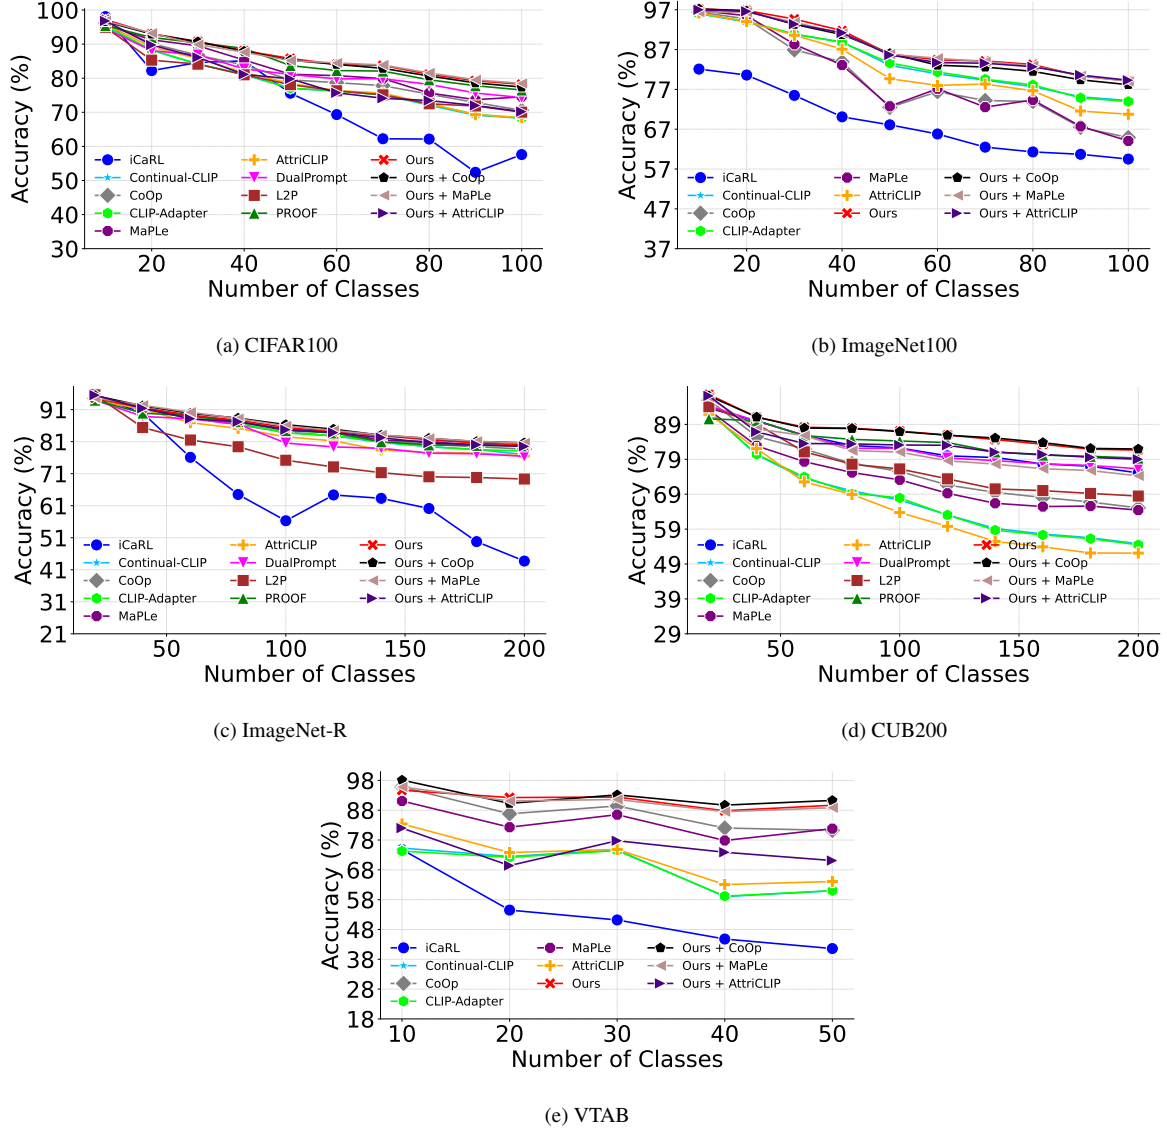


Figure 8. **Performance evolution of different methods.** The top-1 accuracy (%) is reported upon learning of each task.

Method	Avg.	Last	Inference time (s)
VT-CLIP styled VGA	86.54	77.98	0.94
Ours	86.13	78.21	0.16

Table 13. Performance comparison of VT-CLIP styled VGA with Ours on CIFAR-100.

8.5. Ablation studies

8.5.1 Task head selection interpretation and results.

To elaborate on Sec. 4.3, Fig. 9a reports the accuracy of test set samples corresponding to the task encoder heads at the end of incremental training on the last task. Here, the first row is to be interpreted as follows: 77% of test sam-

ples belonging to test set of the first task (test set ID 1) were *correctly* allocated to task head 1, 1.4% of test samples belonging to test set of the first task (test set ID 1) were *incorrectly* allocated to task head 2, and so on. Visualizing the task head selection results for the last task evaluation helps us uncover the amount of forgetting among the individual task heads at the end of the incremental training.

Fig. 9b compares the evolution of the task head selection accuracy across the incremental test steps. Here, at the first test step, we have only one task head and thus the task head selection accuracy is 100%. At the second test step, we have the test samples from two seen tasks as well as two available task heads. Out of all test samples of task 1, the reported 94.5% were correctly classified into the task head 1 while

Method	CIFAR100	ImageNet100	ImageNet-R	CUB	VTAB
Continual-CLIP [70]	1.416	2.175	1.98	2.087	0.614
+Ours	1.39	2.19	1.86	2.06	0.443
CoOp [94]	1.57	2.47	1.95	1.99	0.54
+Ours	1.533	2.074	2.011	1.885	0.516
MaPLe [29]	1.3	2.052	2.16	1.803	0.49
+Ours	1.36	1.956	1.84	1.62	0.407
AttriCLIP [74]	1.781	2.54	2.37	2.419	0.996
+Ours	1.677	2.019	2.388	2.410	0.98

Table 14. **Standard deviation** (std. dev.) scores comparison for Avg. accuracy scores of Table 1 between our variants and their corresponding baseline prompt-based finetuning methods over three runs. In general, our std. dev. scores are comparable to or lower than the corresponding baseline methods and are thus statistically significant.

Method	CIFAR100	ImageNet100	ImageNet-R	CUB	VTAB
Continual-CLIP [70]	-0.086	-0.091	-0.066	-0.124	-0.041
+Ours	-0.106	-0.117	-0.107	-0.117	0.012
CoOp [94]	-0.257	-0.338	-0.12	-0.162	-0.007
+Ours	-0.129	-0.139	-0.112	-0.106	0.011
MaPLe [29]	-0.209	-0.352	-0.1	-0.145	0.037
+Ours	-0.105	-0.112	-0.093	-0.102	0.005
AttriCLIP [74]	-0.128	-0.152	-0.082	-0.151	-0.099
+Ours	-0.143	-0.1	-0.092	-0.037	0.041

Table 15. **Backward Transfer** (BwT) scores \uparrow comparison between our variants and their corresponding baseline prompt-based finetuning methods averaged over three runs. Best scores across each pair is highlighted in **bold**.

Method	CIFAR100	ImageNet100	ImageNet-R	CUB	VTAB
Continual-CLIP [70]	65.34	53.13	61.67	59.55	65.13
+Ours	65.47	53.07	64.05	58.11	66.91
CoOp [94]	64.09	52.6	60.93	62.11	69.38
+Ours	66.2	55.09	63.44	58.6	74.1
MaPLe [29]	68.22	57.04	66.56	61.6	71.51
+Ours	76.17	62.33	70.03	67.8	78.29
AttriCLIP [74]	61.45	50.4	56.41	57.04	61.59
+Ours	61.87	50.56	58.03	57.95	64.3


Table 16. **Transfer** scores [89] \uparrow comparison between our variants and their corresponding baseline prompt-based finetuning methods averaged over three runs. Best scores across each pair is highlighted in **bold**.

Method	CIFAR100	ImageNet100	ImageNet-R	CUB	VTAB
Continual-CLIP [70]	0.288	0.238	0.206	0.208	0.186
+Ours	0.216	0.207	0.201	0.203	0.165
CoOp [94]	0.245	0.3	0.191	0.21	0.191
+Ours	0.224	0.217	0.207	0.204	0.136
MaPLe [29]	0.168	0.243	0.149	0.195	0.195
+Ours	0.214	0.208	0.146	0.184	0.159
AttriCLIP [74]	0.256	0.256	0.205	0.209	0.191
+Ours	0.304	0.205	0.19	0.198	0.304


Table 17. **Expected Calibration Error** (ECE) scores \downarrow (computed over 15 bins) comparison between our variants and their corresponding baseline prompt-based finetuning methods averaged over three runs. Best scores across each pair is highlighted in **bold**.

the rest 5.5% were incorrectly classified into the task head 2. Similarly, for test samples belonging to task 2, 3.1% were incorrectly classified into the task head 1 while the reported 96.9% were correctly classified into the task head

2, and so on. Hence, by studying the task head selection per incremental step, we can investigate the trend of forgetting among the individual task heads.



(a) Last step task head selection accuracies



(b) Per step task head selection accuracies

Figure 9. **Task head selection accuracies** reported on CIFAR-100 upon: (a) evaluation on the last step, (b) evaluation on each incremental step.


8.5.2 Effect of inference module types.

To further investigate the effects of inference modules on performances, we vary the number of layers for the VGA module (sec. 3.4) and for the task-specific encoders (sec. 3.5). Fig. 10a reports the results of varying the number of Transformer Decoder layers [71] in the VGA module. As the number of layers grow, the average accuracy (Avg) increases while the last task accuracy (Last) decreases. This indicates that while a larger number of layers in the VGA module lead to an increase in the initial tasks’ performances, these are amenable to larger forgetting on later incremental steps.


In Fig. 10b, we report the performances for varying number of MLP layers in the mean and the standard deviation heads of the task distribution encoders. Unlike the VGA module, here we observe a consistent trend of decreasing last and average task accuracy with the increase in the number of layers. This clearly indicates the superiority of using a single-layered task distribution encoder.

8.5.3 Inference time for different finetuning methods.

Table 18 investigates the inference time per iteration for different methods. Among the compared prompt-based methods, the inference time for AttriCLIP [74] is notably the highest. This is because it relies on selecting test instance-conditioned prompt tokens from a pool of prompt



(a)



(b)

Figure 10. **Ablation studies on CIFAR100 showing:** (a) the variation of accuracy with the number of Transformer decoder layers in the VGA module, (b) the variation of accuracy with the number of linear layers in the task-specific mean and standard deviation encoders.

tokens. The instance-specific prompts are fed to the text encoder which further outputs an equivalent number of instance-specific text features to be used in the derivation of logits through eq. 1. These operations increase the inference time of AttriCLIP beyond our proposed variants of CLAP4CLIP with hand-crafted prompts (Ours), class-conditioned prompts (CoOp + Ours), and multi-modal prompts (MaPLe + Ours) where the latter three outperform AttriCLIP significantly across all our settings.

Method	Inference time (s)
Continual-CLIP [70]	0.017
CoOp [94]	0.018
MaPLe [29]	0.035
AttriCLIP [74]	0.257
CLIP-Adapter [17]	0.019
Ours	0.163
CoOp + Ours	0.182
MaPLe + Ours	0.064
AttriCLIP + Ours	0.299

Table 18. Average inference time for different finetuning methods on CIFAR100.

8.5.4 Influence of language-aware knowledge components on training dynamics.

Continuing our ablations from sec. 4.3, here we visualize the effects of using language-aware pre-trained knowledge, *i.e.*, weight initialization and task distribution regularization on the training dynamics of our model. For thorough analyses, we consider four variants of our model: (a) Ours uses both weight initialization and task distribution regularization, (b) Ours without weight initialization, (c) Ours without task distribution regularization, and (d) Ours without either of the language-aware components.

Does language-aware weight initialization help alleviate stability gap [37]? To answer this, we first investigate the evolution of the training loss during the initial training stages of each incremental task. Figure 11 shows the loss \mathcal{L} (eq. (12)) during the initial 100 training iterations of each task. We observe that our proposed weight initialization technique leads to lower training losses for the scenarios with or without task distribution regularization, *i.e.*, in general, **red values** < **green values** and **blue values** < **orange values**. Following Harun *et al.* [20], our observations support the hypothesis that larger loss values lead to the **stability gap** [37] for CL, and that an informed weight initialization method can help tackle it by reducing the initial training loss.

To further verify the benefit of our proposed weight initialization strategy for reducing the stability gap, we ablate the accuracy evolution of the first task test samples during the early training stages of each task. Figures 12 and 13 contrast these for CIFAR100. In general, our proposed weight initialization strategy helps mitigate the drop in accuracy during the initial training phases. On average, the first task accuracy upon the first iteration of training across all tasks remains 78.12 without weight initialization and grows to 79.5 with weight initialization, *i.e.* a gain of 1.38 percentage points.

How does language-aware knowledge help learning of task distributions in general? To understand the effect of language-aware knowledge on task distribution learning, we next investigate the evolution of the means and standard deviations learned by the past and the new task heads throughout the training iterations. To this end, Fig. 14 and Fig. 15 report the training iterations against the L2 norm of means and standard deviations for the past task heads (at each incremental training step) and the new task heads (at each training step). We observe two consistent trends regarding the evolution of distributions of the past and the new task heads. First, the proposed initialization of weights helps stabilize the learning of the means and standard deviations with (**red** against **green**) or without (**blue** against **orange**) regularizing the task distributions. Second, regularizing the task distributions increases the L2 norms of the learned mean and the standard deviation as these now have

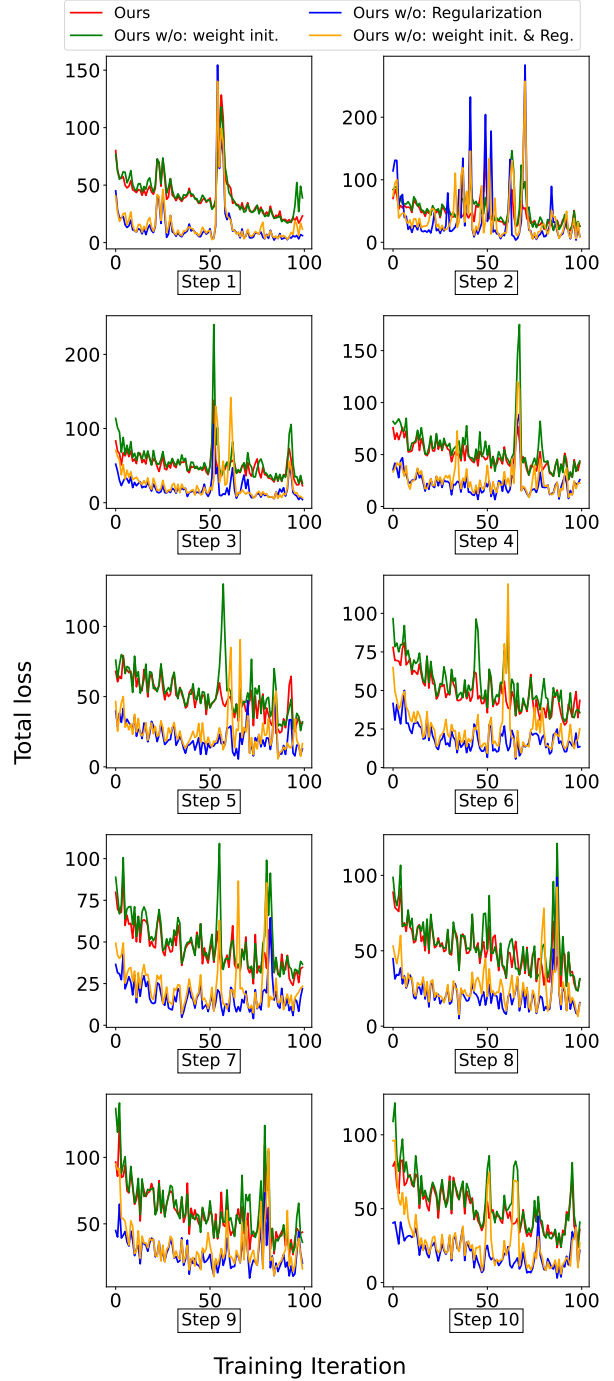


Figure 11. Evolution of the loss value \mathcal{L} during the first 100 training iterations of each task on CIFAR100. Training with our proposed weight initialization strategy consistently leads to lower training losses thus bridging the **stability gap** [37] in CL.

to encode more information to mimic the distributions of the hand-crafted text features.

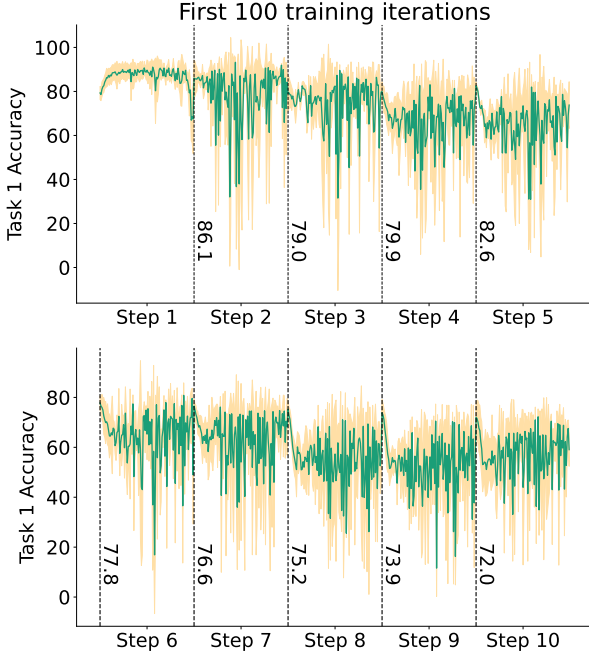


Figure 12. Test accuracy **without weight initialization** on the first task for the initial 100 iterations of incremental training on all ten tasks of CIFAR100. The green lines are the means over three different runs, the orange shades denote ± 1 standard error of the mean. The labels to the vertical bars denote the accuracy values for the first iteration of training on each task.

8.6. Post-hoc novel data detection

Our post-hoc novel data detection (PhNDD) setting aims to evaluate the continual learning methods at identifying novel data on the fly. To do so, we design an evaluation setup that uses no additional data resource other than that provided by the dataset-specific CL setting. Starting from the first test step, we treat the test data of the future tasks as *novel* while those of the seen tasks (including the most recently trained one) as *seen*. Since the last test step of a CL dataset has no future tasks, we exclude this step for our PhNDD evaluation, *i.e.*, we carry our PhNDD evaluation of CL models starting from the first until the penultimate test step.

Following other standard practices [21, 45], we use the Energy scores [45] of the outputs for each test sample as a measure of the model’s confidence score. The samples assigned with a confidence score below the pre-defined confidence threshold are classified as novel. By assuming the seen data as the positive class and the novel data as the negative class, we can obtain a series of true positives rate (TPR) and false positive rate (FPR) by varying the confidence thresholds. One of our PhNDD evaluation metrics – the FPR95 then measures the FPR when the TPR is 0.95. As such, a lower FPR95 score indicates better PhNDD performance. Our other two PhNDD performance metrics include


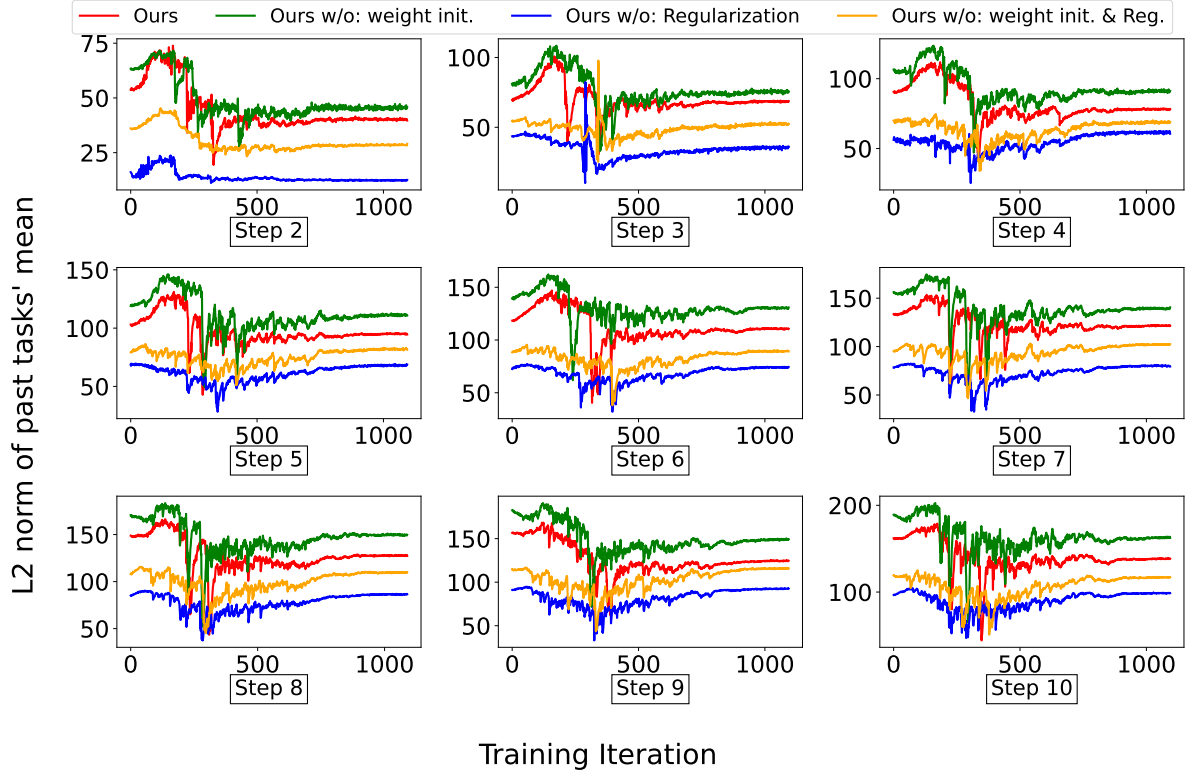


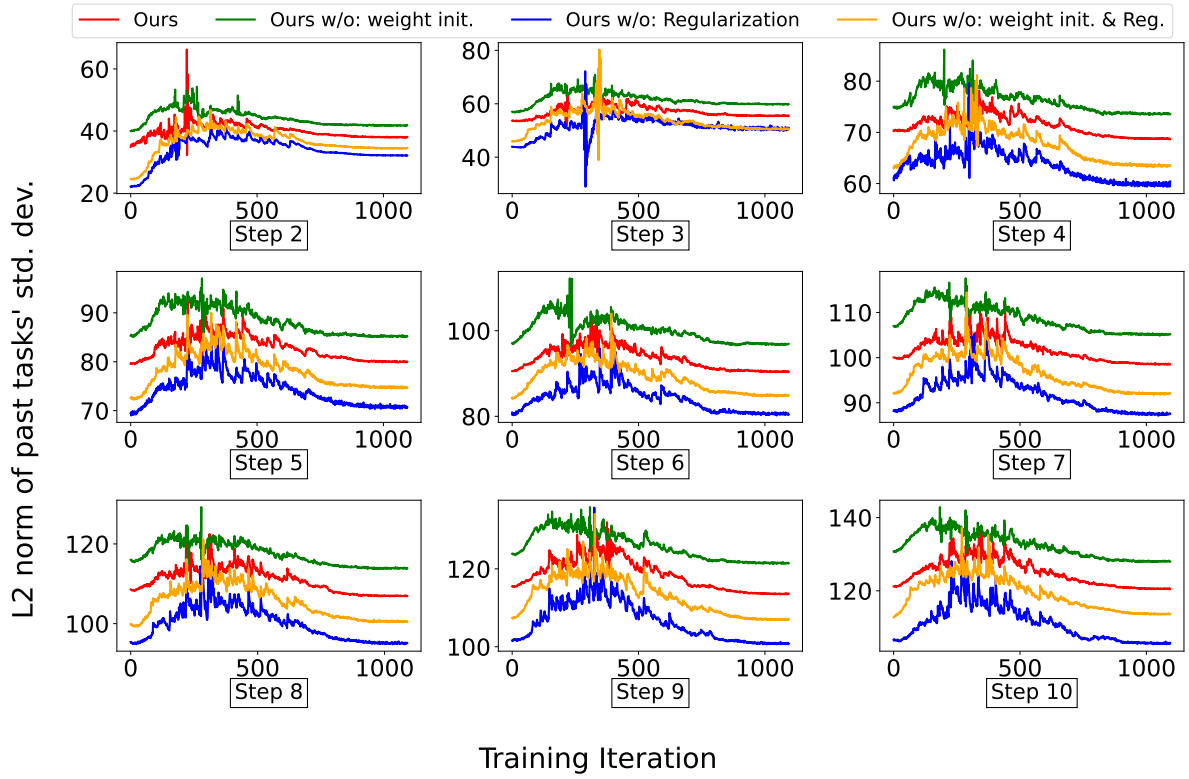
Figure 13. Test accuracy **with weight initialization** on the first task for the initial 100 iterations of incremental training on all ten tasks of CIFAR100. The green lines are the means over three different runs, the orange shades denote ± 1 standard error of the mean. The labels to the vertical bars denote the accuracy values for the first iteration of training on each task.

the the area under receiver operating characteristic curve (AUROC [11]) calculated based on FPR and TPR, and the precision-recall curve (AUPR [66]). Higher values of AUROC and AUPR indicate better PhNDD performance.

In Table 5 in the main paper, we report the PhNDD metrics averaged over all the evaluated steps. Here, in Fig. 16, we show the evolution of these metrics with each evaluation starting from the first test step until the penultimate test step of CIFAR100. We observe that the zero-shot Continual-CLIP [70] has the poorest PhNDD performances (highest FPR95, and least AUROC and AUPR scores) across all steps given that it has not been finetuned on the downstream CL tasks. Among the finetuned methods, the CoOp [94] exhibits the poorest performances across all tasks. Among the variants of our method, combining CoOp with ours (CoOp + Ours) achieves the best PhNDD performances across all tasks. Furthermore, the deterministic versions: Ours w/o VI and CoOp + Ours (w/o VI) remain sub-optimal to their respective probabilistic variants, *i.e.*, Ours and CoOp + Ours. The latter results validate the added perks of our probabilistic modeling framework for post-hoc novel data detection.

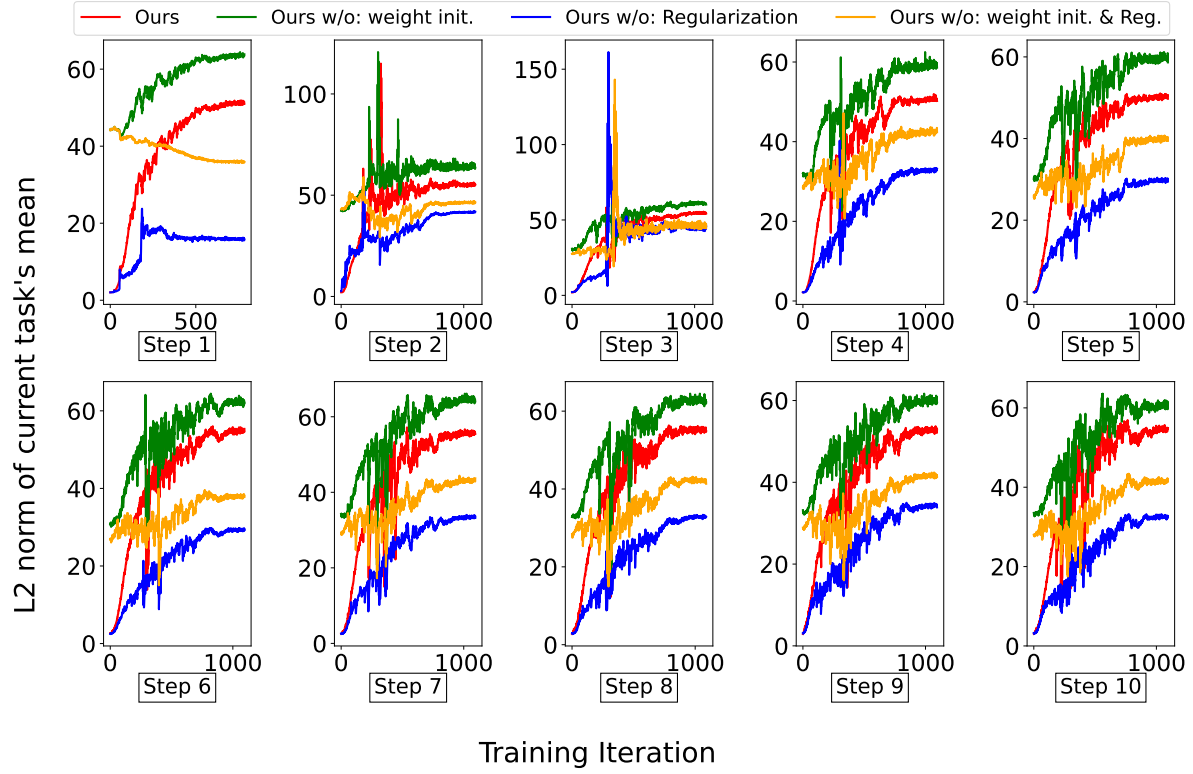


(a) L2 norm of mean of past task heads

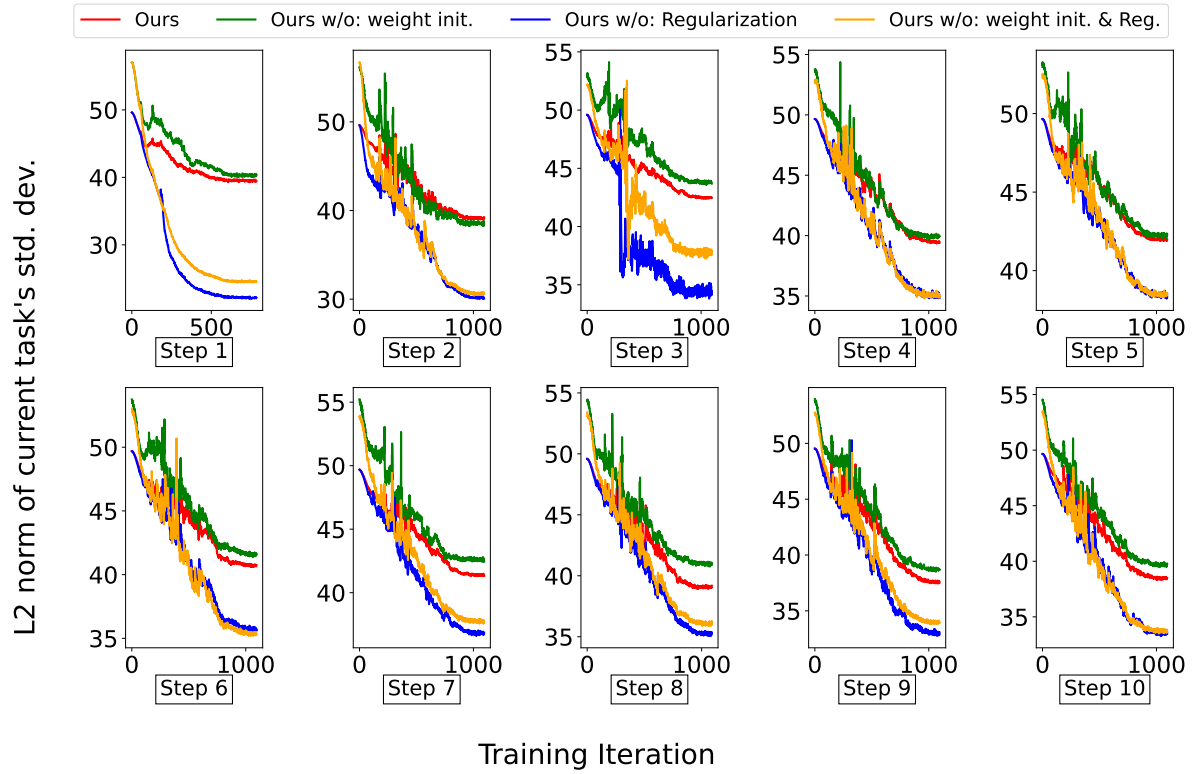


(b) L2 norm of standard deviation of past task heads

Figure 14. Evolution of mean and standard deviation of past task encoders with training iterations.



(a) L2 norm of mean of current task heads



(b) L2 norm of standard deviation of current task heads

Figure 15. Evolution of mean and standard deviation of task encoders (recorded at the step where they were first introduced) with training iterations.

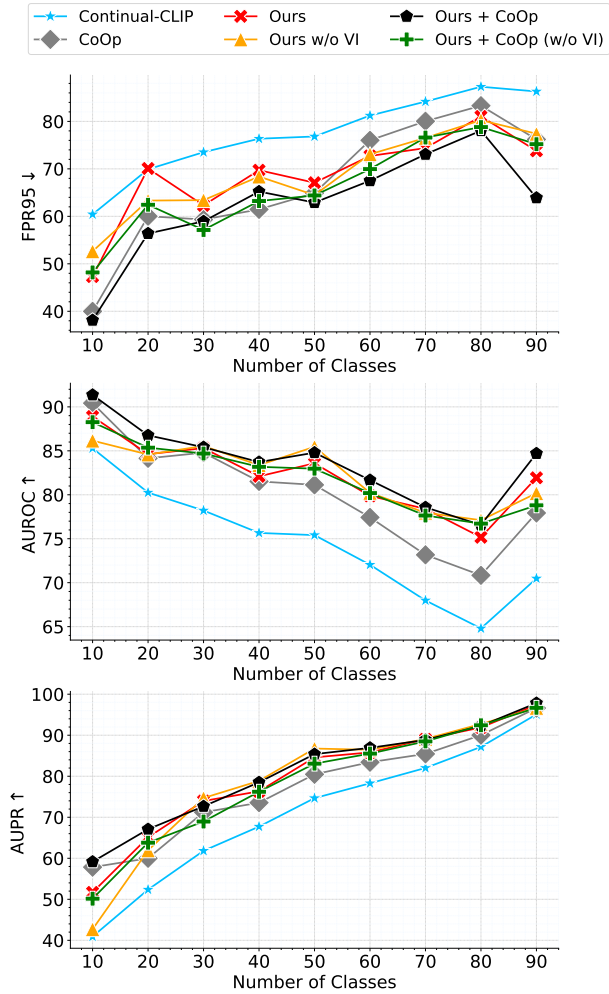


Figure 16. **Performance comparisons for post-hoc novel data detection** averaged over 3 runs on CIFAR100: FPR95 (left), AUROC (middle), and AUPR (right). The evaluations are carried over all but the last incremental test step.

9. Limitations and further research directions

Parameter overhead. For each incoming task, CLAP4CLIP initializes a new head consisting of $d \times d$ parameters where d is the output dimension of the CLIP model’s encoder. For a very large number of real-world CL tasks, the number of finetunable parameters for CLAP4CLIP may thus become comparable to or larger than that of the pre-trained CLIP model’s ≈ 150 million parameters. For example, using a ViT-B/16 encoder with $d = 512$ brings an overhead of $\approx 525,000$ new parameters with each incoming task. After having seen

≈ 300 new tasks, the number of CLAP parameters to be finetuned thus amount to ≈ 158 million, which is larger than the frozen CLIP itself, and thus defeats the purpose of finetuning at the first place. One solid future direction to use CLAP4CLIP for very large real-world CL settings could thus be introducing more strict parameter-efficiency measures [75, 81] and/or learning probabilistic adapters with low-rank weights [25].

Design choices. Other future directions for improving CLAP4CLIP could include the use of better regularization techniques to further prevent forgetting (see Table 15 for the current forgetting in CLAP), and the search for more informed yet computationally efficient priors (see Table 4 for the computational overhead attached with more informed priors).

LLM-generated class descriptions as language-aware knowledge. In Sec. 3.6, we proposed using the text features from hand-crafted prompts as language-aware CLIP knowledge to help alleviate forgetting. However, hand-crafted prompts require manual labelling of data which is not always practical. Hence, several recent works [30, 50, 61] have opted to mining Large Language Models (LLMs) for efficiently obtaining the class-specific descriptions. To study the feasibility of alleviating forgetting using such LLM-generated class descriptions, we leverage the diverse prompts from CuPL [61] obtained using the GPT-3 [3] model. Our preliminary investigation suggests that the hand-crafted prompts have an upper hand over GPT-3 based prompts for CLAP4CLIP performance (see Table 19). This could be because of the broad range of knowledge encoded in the GPT-generated prompts – which at times are irrelevant for the test images.

Prompt type	Last \uparrow	Avg \uparrow	BWT \uparrow	ECE \downarrow	Runtime per iter. \downarrow
Hand-crafted	78.21	86.13	-0.141	0.204	0.169
GPT-3	77.76	85.7	-0.099	0.219	0.151

Table 19. **Performance comparison on CIFAR100** using hand-crafted vs. LLM-generated prompts for encoding language-aware CLIP knowledge. The results reported are averages over 3 runs. Best results across the metrics are highlighted in **bold**.

Based on the above finding, we suggest leveraging task-relevant LLM-generated descriptions as language-aware knowledge to be another promising future research direction. Lastly, it is worth noting that a number of existing methods that rely on LLM-generated prompts are limited in their transferable knowledge across unseen classes and datasets [50, 61] (e.g., any new class at test-time would require mining the LLM descriptions in advance). On the contrary, our proposed weight initialization and task distribution regularization strategies provide a natural frame-

work for LLM-generated prompts to be used alongside arbitrary learnable prompts (*e.g.* replacing $\mathbf{t}_{y'}^{h,l}$ in eq. (9)). This compliments the idea of LLM-based text-only supervision frameworks [30] that seek to enrich *zero-shot transfer* of prompts to new classes by extracting rich contextual information from LLM data.¹

Compatibility with Vision-Language datasets The tasks we have covered so far in the paper are based solely on Vision datasets. To further demonstrate that our method is compatible with more sophisticated vision-language datasets, we here consider using a toy Visual Question Answering (VQAv2) task from the CLiMB dataset [68]. The CLiMB dataset hosts a number of tasks/settings to evaluate multi-modal and low-shot transfer abilities of CL algorithms. However, given the intricacies of these tasks (visual question answering, reasoning, etc.), we leave a full in-depth engagement with CLiMB [68] as a separate future direction for research.²

To show the aptness of our method for the dataset’s tasks, we carry out preliminary experiments on the single-task learning setting [69] of the VQAv2 subset of CLiMB. Following [69], we rely on the BART model [42] for text generation here. Table 20 shows that our method surpasses the Continual-CLIP by 9.29 percentage points on the VQAv2 task, thus showing that ours enhances the zero-shot generalization capability of CLIP.

Model	VQAv2 task score
Continual-CLIP	57.42
Ours	66.71

Table 20. Single-task learning performance on the VQAv2 subset of the CLiMB dataset.

¹Given that new classes might emerge at test time for which we do not have the LLM-generated descriptions, it is important that the learned prompts preserve their zero-shot generalization ability.

²The CLiMB dataset [68] was introduced as an independent CL benchmark with a number of tasks (Visual Question Answering/Reasoning/Entailment) and training settings including low-shot and unimodal learning. Existing works [69] that study CLiMB thus rely solely on it *and not on additional datasets* for evaluations.

10. Derivation of ELBO for the static prior.

We seek to maximize the likelihood $p(y^{1:T})$ for all observed labels $y^{1:T}$. To derive the predictions, our framework uses the visual-aligned text features $\tilde{\mathbf{t}}_c^{1:T}$ and the image inputs \mathbf{x} (see eq. (1)). Our evidence is thus $p(y^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T})$ for which we derive the lower bound (ELBO). In the following, we denote the prior network as $p_\theta(z^t)$ for which the true posterior is $p_\theta(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)$. We approximate the true posterior using the variational posterior $q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)$. Our derivation ends up with the reconstruction term $p_\theta(y^t|z^t, \mathbf{x}; \tilde{\mathbf{t}}_c^t)$ that can be seen as a deterministic function converting a given latent vector z^t and an input image \mathbf{x} into an observation y^t . For our CLIP-based variational framework, this deterministic function is the cosine similarity operation followed by the softmax application (Eq. (3b)).

$$\begin{aligned}
& \log p_\theta(y^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T}) && \text{(Log-likelihood of evidence)} \\
&= \log p_\theta(y^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T}) \int q_\phi(z^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T}) dz^{1:T} && (\because \int q_\phi(z^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T}) dz^{1:T} = 1) \\
&= \int q_\phi(z^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T}) (\log p_\theta(y^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T})) dz^{1:T} && \text{(Bring evidence into integral)} \\
&= \mathbb{E}_{q_\phi(z^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T})} [\log p_\theta(y^{1:T}|\mathbf{x}; \tilde{\mathbf{t}}_c^{1:T})] && \text{(Definition of Expectation)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} [\log p_\theta(y^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)] \right] && \text{(Rewrite using sum)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t, z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{p_\theta(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && \text{(Re-introduce } z^t \text{ by Chain rule of probability)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t, z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t) q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{p_\theta(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t) q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && \text{(Multiply by } 1 = \frac{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}) \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t, z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] + \mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{p_\theta(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && \text{(Split the expectation)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t, z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] + \mathbb{D}_{\text{KL}}(q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t) \| p_\theta(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)) \right] && \text{(Definition of KL divergence)} \\
&\geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t, z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && (\because \text{KL divergence} \geq 0) \\
&\geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log \frac{p_\theta(y^t|z^t, \mathbf{x}; \tilde{\mathbf{t}}_c^t) p_\chi(z^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && \text{(Chain rule of probability)} \\
&\geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log p_\theta(y^t|z^t, \mathbf{x}; \tilde{\mathbf{t}}_c^t) \right] + \mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\frac{p_\chi(z^t)}{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \right] \right] && \text{(Split the Expectation)} \\
&\geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t)} \left[\log p_\theta(y^t|z^t, \mathbf{x}; \tilde{\mathbf{t}}_c^t) \right] - \mathbb{D}_{\text{KL}}(q_\phi(z^t|\mathbf{x}; \tilde{\mathbf{t}}_c^t) \| p_\chi(z^t)) \right] && \text{(Definition of KL divergence)}
\end{aligned}$$

11. Data-driven prior

The choice of prior is pivotal to a Bayesian inference workflow like ours [16]. While a standard Gaussian prior $p_\chi = \mathcal{N}(0, I)$ adapts well to a range of settings, it is (seemingly) uninformative regarding the nature of a given task [19]. With the end goal of deriving more informative priors, we thus seek to replace p_χ with task data-dependent prior distributions p^t , wherever applicable.

To this end, we first note that the outputs of the VGA module remain **invariant** not only to the order of the input text features (due to self-attention) but also to the order of the contextual image features (due to cross-attention). The latter invariance implies that the joint task-specific distribution learned by the encoder q_ϕ^t (conditioned on the VGA outputs $\hat{\mathbf{t}}_c^t$ from eq. 5) is preserved if we were to permute the elements of the task-specific visual context set. More formally, this observation helps guarantee the (finite) *exchangeability* and the *consistency* properties of a stochastic process [55].

Motivated by the above, we treat the t -th task image features \mathbf{x}^t as the target set \mathcal{T}^t and employ a randomly chosen subset of it as our context set \mathcal{C}^t to align the t -th task text features and to condition our prior p^t on:

$$\begin{aligned}\hat{\mathbf{t}}_c^t &= \text{VGA}(Q = \mathbf{t}_c^t, K = V = \mathcal{C}^t), \\ p^t &= q_\phi^t(\tilde{\mathbf{t}}_c^t) = (\mu^t(\tilde{\mathbf{t}}_c^t), \sigma^t(\tilde{\mathbf{t}}_c^t))\end{aligned}\quad (14)$$

where $\tilde{\mathbf{t}}^t$ is the fused task-specific text feature following eq. (6). The task-specific prior p^t thus endows our training framework with a resemblance to the neural process (NP) architectures [18, 27, 31]. Following NPs, we use the same encoder q_ϕ^t to parameterize the conditional prior and the variational posterior. This results in the following approximate ELBO (see App. 10 for the ELBO derivation):

$$\begin{aligned}\log p(\mathbf{y}^{1:T} | \mathbf{x}, \mathcal{C}^{1:T}) &\geq \\ \sum_{t=1}^T &\left[\mathbb{E}_{q_\phi(z^t | \mathbf{x})} \left[\log p(\mathbf{y}^t | z^t, \mathbf{x}_{\mathcal{T}^t}, \mathcal{C}^t) \right] \right. \\ &\left. - \mathbb{D}_{\text{KL}}(q_\phi(z^t | \mathcal{T}^t) \| q_\phi(z^t | \mathcal{C}^t)) \right]\end{aligned}\quad (15)$$

where in practice, the entire set of t -th images in a training minibatch form the target set \mathcal{T}^t and a randomly chosen subset of the targets make up the context \mathcal{C}^t [38]. Note that *unlike* NPs, our framework does not entirely rely on data-driven priors. Namely, while training on a CL task t , the past-task encoders are frozen and we have ample t -th task data points to condition the prior on. We thus resort to optimizing the ELBO (15) during training. On the other hand, during finetuning, we have limited task-specific data points to condition our context on. As such, we empirically found that switching to the static prior yields better results and thus resort to optimizing the ELBO (11) during finetuning.

11.1. Effect of the context size on data-driven prior.

Table 4 in the main paper compares the results of using a data-driven prior against the uniform normal prior and the language-aware prior (see Sec. 3.6), where the latter is driven from the pre-trained text encoder using hand-crafted prompts. We observe that data-driven prior leads to minor accuracy improvements over the standard normal prior but falls narrowly behind the language-aware prior. Here, we study the influence of the batch size of the context set selected at random to derive our prior from.

Table 21 shows the last task accuracy with varying context sizes and a fixed target batch size of 64. We find that a smaller context set size hurts the performance of the model to the extent of falling behind the standard normal prior. Given that the context sets are the sampled subsets of the training (target) minibatches, a much smaller context set can lead to the increase in the prior matching loss values. We find that the context set batch size of 40 performs the best, and thus use this to ablate the prior-dependent performances in the main paper.

	4	8	16	32	40	50	Static prior
Accuracy	76.1	76.99	77.41	78.03	78.32	77.35	78.21

Table 21. **Influence of the context set size** used to derive the data-driven prior on CIFAR100.

11.2. Derivation of ELBO for the data-driven prior.

Similar to App. 10, we start with the log-likelihood of the evidence which now involves conditioning on an additional context set $\mathcal{C}^{1:T}$. The t -th task context set is used to condition our prior network $p_\theta(z^t|\mathcal{C}^t)$. Following the standard practices of other data-driven prior frameworks [18, 27], we introduce parameter-sharing between our conditional prior and variational posterior networks. This allows us to replace our prior network with the variational posterior network $q_\phi(z^t|\mathcal{T}^t)$, where \mathcal{T} is the target set for task t .

$$\begin{aligned}
& \log p_\theta(Y_{\mathcal{T}}^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T}) && \text{(Log-likelihood of evidence)} \\
&= \log p_\theta(Y_{\mathcal{T}}^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T}) \int q_\phi(z^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T}) dz^{1:T} && (\because \int q_\phi(z^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T}) dz^{1:T} = 1) \\
&= \int q_\phi(z^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T}) (\log p_\theta(Y_{\mathcal{T}}^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T})) dz^{1:T} && \text{(Bring evidence into integral)} \\
&= \mathbb{E}_{q_\phi(z^{1:T} | \mathcal{T}^{1:T})} [\log p_\theta(Y_{\mathcal{T}}^{1:T} | \mathbf{x}_{\mathcal{T}}^{1:T}, \mathcal{C}^{1:T})] && \text{(By definition)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} [\log p_\theta(Y_{\mathcal{T}}^t | \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t)] \right] && \text{(Rewrite using sum)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log \frac{p_\theta(Y_{\mathcal{T}}^t, z^t | \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t)}{p_\theta(z^t | \mathbf{x}_{\mathcal{T}}^t, Y_{\mathcal{T}}^t, \mathcal{C}^t)} \right] \right] && \text{(Re-introduce } z^t \text{ by Chain rule of probability)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log \frac{p_\theta(Y_{\mathcal{T}}^t | z^t, \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t) p_\theta(z^t | \mathcal{C}^t)}{p_\theta(z^t | \mathcal{T}^t)} \right] \right] && \text{(By Chain rule of probability; } \mathcal{C} \subset \mathcal{T}) \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log \frac{p_\theta(Y_{\mathcal{T}}^t | z^t, \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t) p_\theta(z^t | \mathcal{C}^t) q_\phi(z^t | \mathcal{T}^t)}{p_\theta(z^t | \mathcal{T}^t) q_\phi(z^t | \mathcal{T}^t)} \right] \right] && \text{(Equivalent fraction)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log p_\theta(Y_{\mathcal{T}}^t | z^t, \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t) \right] \right. \\
&\quad \left. + \mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log \frac{p_\theta(z^t | \mathcal{C}^t)}{q_\phi(z^t | \mathcal{T}^t)} \right] + \mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log \frac{q_\phi(z^t | \mathcal{T}^t)}{p_\theta(z^t | \mathcal{T}^t)} \right] \right] && \text{(Split the expectation)} \\
&= \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log p_\theta(Y_{\mathcal{T}}^t | z^t, \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t) \right] \right. \\
&\quad \left. - \mathbb{D}_{\text{KL}}(q_\phi(z^t | \mathcal{T}^t) \| p_\theta(z^t | \mathcal{C}^t)) + \mathbb{D}_{\text{KL}}(q_\phi(z^t | \mathcal{T}^t) \| p_\theta(z^t | \mathcal{T}^t)) \right] && \text{(By definition of KL divergence)} \\
&\geq \sum_{t=1}^T \left[\mathbb{E}_{q_\phi(z^t | \mathcal{T}^t)} \left[\log p_\theta(Y_{\mathcal{T}}^t | z^t, \mathbf{x}_{\mathcal{T}}^t, \mathcal{C}^t) \right] - \mathbb{D}_{\text{KL}}(q_\phi(z^t | \mathcal{T}^t) \| p_\theta(z^t | \mathcal{C}^t)) \right] && (\because \text{KL divergence} \geq 0)
\end{aligned}$$