

Anytime Continual Learning for Open Vocabulary Classification

Zhen Zhu[✉], Yiming Gong[✉], and Derek Hoiem[✉]

University of Illinois at Urbana-Champaign
{zhenzhu4,yimingg8,dhoiem}@illinois.edu

Abstract. We propose an approach for anytime continual learning (AnytimeCL) for open vocabulary image classification. The AnytimeCL problem aims to break away from batch training and rigid models by requiring that a system can predict any set of labels at any time and efficiently update and improve when receiving one or more training samples at any time. Despite the challenging goal, we achieve substantial improvements over recent methods. We propose a dynamic weighting between predictions of a partially fine-tuned model and a fixed open vocabulary model that enables continual improvement when training samples are available for a subset of a task’s labels. We also propose an attention-weighted PCA compression of training features that reduces storage and computation with little impact to model accuracy. Our methods are validated with experiments that test flexibility of learning and inference.

Keywords: Continual learning · Anytime learning · Open-vocabulary classification

1 Introduction

Continual learning aims to improve a system’s capability as it incrementally receives new data and labels, which gains increasing importance with the expanding scale and applications of visual learning. Continual learning for classification is traditionally approached with discrete label spaces. Adding labels or tasks over time changes the problem landscape, thereby increasing the difficulty of learning. The open vocabulary setting, however, frames classification as comparing continuous feature and label embeddings. Any images and textual labels can be embedded, so learning in this setting involves only improving on the existing problem definition, which may be more amenable to continuous improvement. Although open vocabulary models can predict over arbitrary label sets, even models like CLIP [37] that are trained on internet-scale data have unsatisfactory performance on many tasks [47, 48].

Our work aims to continually improve open vocabulary image classifiers as new labeled data is received. We call this “**anytime continual learning**” because the goal is to improve efficiently *any time* new examples are received and to maintain the ability to predict over arbitrary label sets at *any time*.

Recent work by Zhu et al. [62] extends CLIP by training linear classifiers and combining their predictions with those of label vectors from the original text embedding. To enable efficient and distributed learning, linear classifiers are trained for each partition of the feature space defined by online hierarchical clustering. This approach requires storing all training examples, but a new example is efficiently incorporated by training only on the data from the example’s nearest cluster. In experiments on their own benchmark designed for open vocabulary continual learning and existing benchmarks, Zhu et al. [62] outperform other recent approaches.

Our approach is to train online, fine-tuning the last transformer block while keeping the label embedding fixed. When a new training sample is received, we fill a class-balanced batch with stored samples and update the last transformer block in a single step. Our experiments show this performs better than retraining a classifier layer, and each new sample can be incorporated in milliseconds. Additionally, we introduce a modified loss, enabling the prediction of “none of the above” when the true label is absent from the candidate set, which improves overall performance.

Whenever a new training example arrives, our method updates an estimate of the tuned and original model’s accuracy for the given label. The tuned and original model’s predictions are then weighted in proportion to their expected accuracy for each label. This weighting accounts for the growing effectiveness of the tuned model, and greatly outperforms Zhu et al. [62]’s AIM weighting based on whether an example’s label is likely to have been seen in training.

The partial fine-tuning approach requires storing either full images that need to be re-encoded during training, or data-consumptive patch-level features at intermediate layers. We apply a per-image weighted PCA, which provides a 30x reduction of data with little impact on prediction accuracy. Such compression could be especially beneficial in a large-scale federated learning setting.

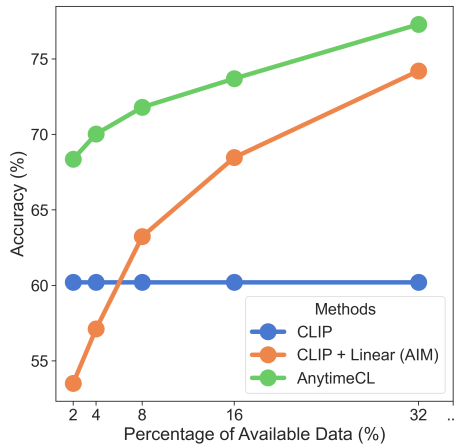


Fig. 1: Our AnytimeCL algorithm can be efficiently updated with each new example and continuously improve. By dynamically weighting predictions between a tunable model and a frozen open vocabulary model, our method can predict over any label set while gaining expertise. This figure shows our method outperforms previous SotA Zhu et al. [62] in every stage in the data incremental setting. Our method also outperforms in other settings like task-incremental and class-incremental.

In experiments on the open vocabulary continual learning benchmark proposed by Zhu et al. [62], our approach outperforms under all settings and stages, including data-incremental, class-incremental, and task-incremental learning, and zero shot prediction. Our ablations evaluate the effects of all key design parameters, including the partial fine-tuning, model prediction weighting, batch sampling, and the regularization loss term. We also show that our method can be applied to combine non-open vocabulary models like DINOv2 with CLIP to improve further while maintaining open vocabulary ability.

In summary, our **main contribution** is an open vocabulary continual learning method that can quickly incorporate new training examples received in any order and continually improve while maintaining open vocabulary performance. Our experiments demonstrate that our system’s accuracy and efficiency stem from multiple proposed innovations:

- Partial fine-tuning of features with a fixed label embedding;
- Online training with each batch composed of the new training sample and class-balanced stored samples;
- Online learning of per-label accuracy for effective combination of original and tuned model predictions;
- Loss modification to enable “none of the above” prediction, which also stabilizes open vocabulary training;
- Intermediate layer feature compression that reduces storage of training samples and improves speed without much loss to accuracy.

2 Related work

The goal of continual learning (CL) is to continually improve the performance when seeing more data, while online continual learning [1, 3, 8, 24, 34, 36] aims to achieve a good trade-off between performance and learning efficiency. Our problem setup and goals are distinguished by these important characteristics:

- *Flexibly learn*, receiving labeled examples in data-incremental (random order), class incremental (grouped by category), or task incremental (grouped by dataset);
- *Efficiently update* from a single example or a batch of examples;
- *Retain training data*, though potentially in compact, privacy-preserving forms;
- *Cumulatively improve* from new data without forgetting other classes or tasks;
- *Flexibly infer*, predicting over label sets defined at inference time, without rigid task boundaries.

Continual learning approaches can be broadly categorized into regularization [19, 23, 59], replay/rehearsal methods [24, 38, 42, 56], and parameter isolation or expansion [2, 27, 35, 36, 40, 41, 57, 60]. Regularization techniques generally impose constraints on the learning process to alleviate forgetting. Rehearsal methods involve storing and replaying past data samples during training [5, 35, 36, 42].

So far, when training data can be stored, simple replay methods tend to outperform others, whether limiting the number of stored examples [36] or training computation [35]. Parameter isolation methods maintain learning stability by fixing subsets of parameters [27, 41] or extending model with new parameters [40, 57, 60]. Recently, several works [16–18, 45, 51–53] adopt prompt tuning for continual learning, which can also fall into the parameter expansion category. These methods can be used for stage-wise incremental continual learning, but none demonstrates the efficacy on online incremental learning that is our focus. Prompt tuning provides an alternative to weight-based fine-tuning with less forgetting. We consider weight-tuning approaches here, but preliminary experiments indicate the same strategy is applicable to prompt tuning, with similar accuracy but slower training.

We now focus on those most relevant and influential to our work; see Wang et al. [50] for a comprehensive survey of continual learning.

2.1 CL for open-vocabulary classification

WiSE-FT [54] fine-tunes CLIP encoders on target tasks and averages fine-tuned weights with original weights for robustness to distribution shifts. PAINT [15] shares a similar idea as WiSE-FT but the weight mixing coefficient is found via a held-out set. With such approaches, generality degrades with each increment of continual learning. CLS-ER [4] exponentially averages model weights in different paces for its plastic and stable model to balance learning and forgetting. Closer to our work, ZSCL [61] fine-tunes CLIP encoders using a weight ensemble idea, similar to WiSE-FT, and applies a distillation loss with a large unsupervised dataset to reduce forgetting. Unlike these, we partially fine-tune CLIP with fixed label embeddings and use a dynamically weighted combination of the predictions from the tuned and original models, which enables improvement on target tasks without sacrificing the generality of the original model.

Another class of methods [13, 44] use attribute-based recognition [12, 22], continually learning attributes and attribute-class relations as a way to generalize from seen to unseen categories. We use CLIP, which provides zero-shot ability by training to match image-text pairs on a large training corpus, instead of the attribute-based approach. In reported results [13, 37, 44] and our own tests, CLIP has much higher zero-shot accuracy (e.g. for the SUN and aPY datasets) than the continual attribute-based methods achieve even after training, supporting the idea of building on vision-language models for open vocabulary image classification.

2.2 CL with constrained computation

Prabhu et al. [35] evaluate a wide variety of continual learning techniques under a setting of retaining data but limiting compute. Compared to all other techniques, they find that simple replay strategies, such as fine-tuning with uniformly random or class-balanced batches, are most effective in this setting. They also provide some evidence that partial-fine tuning of a well-initialized network may

be an efficient and effective approach. In this direction, we find that fine-tuning only the last transformer block of CLIP by performing one mini-batch update per incoming sample works surprisingly well.

2.3 Biologically inspired CL

The complimentary learning systems theory [31] (CLS theory) posits that humans achieve continual learning through the interplay of sparse retrieval-based and dense consolidated memory systems. This has inspired many works in continual learning, e.g. [4, 33, 62]. The most closely related is Zhu et al. [62], as detailed in the introduction. Many works are also inspired by the idea of wake/sleep phases of learning, where faster updates are required during the wake stage. Recent studies of mammalian memory and learning indicate that consolidation occurs in similar ways via replay during wakefulness and sleep [7, 43, 49]. Our online learning method mixes use of individual examples and consolidated networks for continuous improvement for continuous improvement, reflecting the idea of wakeful consolidation through replay.

2.4 Memory compression

Other work has compressed training data [9], or distilled training data into a network [58], or learned prompts [16–18, 45, 51–53] instead of maintaining original examples. AQM [9] uses an online VQ-VAE [29] to compress training samples to reduce storage for continual learning. We instead compress intermediate features to improve both storage and computation of partial fine-tuning.

3 Method

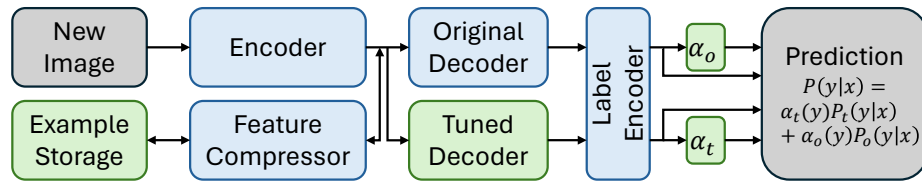


Fig. 2: Overview. On receiving a new training sample, the batch is completed with stored samples, the prediction is made using a label embedding, the tuned decoder and confidence weights (α_o , α_t) are updated in one step, and the new sample is stored. To save space and time, stored examples are encoded and compressed. In testing, the probability of each candidate label is determined by predictions from both decoders and their class-wise confidence weights. Our method enables constant-time updates from new samples while continually improving and maintaining open vocabulary performance. Green blocks are updated during training; blue blocks are not updated.

Our system receives training examples (x, y, \mathcal{Y}) one by one to update its models, in order to continually improve in prediction of y given (x, \mathcal{Y}) , where x is an input image, y is its target label, and \mathcal{Y} is its set of candidate labels.

Fig. 2 offers a system overview aimed at enhancing open vocabulary image classifiers by integrating a tuned model for learning target tasks with a frozen original model. The tuned model uses the same encoder as the original but incorporates a trainable decoder. For an image x , both the tuned model and the original model produce the probabilities for all candidate labels, denoted as $P_t(y|x)$ and $P_o(y|x)$. The final probability for the image is weighted through our Online Class-wise Weighting (OCW, Sec. 3.2):

$$P(y|x) = \alpha_o(y)P_t(y|x) + \alpha_t(y)P_o(y|x), \quad (1)$$

where $\alpha_o(y)$ and $\alpha_t(y)$ are the relative expected accuracies for the original and tuned models for label y . During training, new samples are encoded to intermediate features (feature vectors for patches plus a CLS token), optionally compressed (Sec. 3.3), and stored, for reuse in future steps.

3.1 Models

Original model. In our experiments, the original model is the publicly available CLIP [37] ViT model that has been trained contrastively on image-text pairs. In our experiments, the shared encoder is all but the last transformer block, and the original decoder is the last block. The CLIP model produces a probability for label y for image x given a set of candidate text labels \mathcal{Y} based on the dot product of image embedding e_x (CLS token) and text embedding e_y :

$$P_o(y|x) = \frac{\exp(100 \cdot \cos(e_x, e_y))}{\sum_{y_k \in \mathcal{Y}} \exp(100 \cdot \cos(e_x, e_{y_k}))}. \quad (2)$$

Tuned model. With each new sample, our goal is to efficiently update the tuned model to improve accuracy in received labels while maintaining general features that are effective for future learning. We tune only the last image transformer block while keeping the label embedding fixed, which helps the features to stay correlated with the text modality and reduces overfitting to received labels. Note that more blocks can be tuned, which could enable better performance but with increased computation. The tuned decoder is initialized with the weights of the original decoder.

Given a new sample, we form a batch of that sample and a class-balanced sampling of stored training samples. Specifically, we first determine the class count for selection: for a batch size of B and all seen labels \mathcal{Y}_t , we uniformly select $\min(B - 1, |\mathcal{Y}_t|)$ classes from \mathcal{Y}_t , where $|\cdot|$ indicates the length of a set. Then we uniformly sample an equal number of instances from each chosen class to form the batch. This class-balanced sampling ensures continual retention and improvement of prediction of training labels. Then, the tuned model is updated with one iteration based on a cross-entropy loss.

Additionally, we develop a regularization loss that helps with performance. The idea is that, if the true label is not in the label candidates, a low score should be predicted for every candidate label. We implement this with an “other” option within the candidate set, but since “other” does not have an appearance, we model it with just a learnable bias term. The combined loss for training the tuned model is therefore:

$$\mathcal{L}(x, y, \mathcal{Y}) = \mathcal{L}_{\text{ce}}(x, y, \mathcal{Y} \cup \text{other}) + \beta \mathcal{L}_{\text{ce}}(x, \text{other}, (\mathcal{Y} \cup \text{other}) \setminus y), \quad (3)$$

where β is a hyperparameter balancing the two loss terms, set to 0.1 based on validation experiments.

3.2 Online Class-wise Weighting

We estimate the probability of each candidate label by weighted voting of the tuned and original models (Eq. 1). We would like to assign more weight to the model that is more likely to be correct for a given label. There are two big problems: 1) training samples would provide a highly biased estimate of correctness for the tuned model; 2) maintaining a held out set or performing cross-validation would either reduce available training samples or be impractically slow. Our solution is to use each training sample, before update, to update our estimate of the likelihood of correctness for its label based on the tuned and original predictions.

Let $c_t(y)$ and $c_o(y)$ be the estimated accuracy of the tuned and original model for label y . We apply an exponential moving average (EMA) updating method to estimate them online, ensuring the estimations are reliable when evaluated at any time, concurring with our anytime continual learning goals. Assuming the EMA decay is set to η ($=0.99$ in our experiments), the estimated accuracy of the tuned model at the current step is:

$$c_t(y) = \eta \hat{c}_t(y) + (1 - \eta) \mathbb{1}[y_t(x) = y]. \quad (4)$$

Here, $\hat{c}_t(y)$ is the estimated accuracy of label y in the previous step; $y_t(x)$ denotes the predicted label of the tuned model for x . Since the exponential moving average depends on past values, we compute $c_t(y)$ as the average accuracy for the first $\lfloor \frac{1}{1-\eta} \rfloor$ samples. $c_o(y)$ is updated in the same way.

After getting $c_t(y)$ and $c_o(y)$, the weights of the two models are:

$$\alpha_t(y) = \frac{c_t(y)}{c_t(y) + c_o(y) + \epsilon}, \quad \alpha_o(y) = 1 - \alpha_t(y). \quad (5)$$

Here, ϵ is a very small number ($1e-8$) to prevent division by zero. For labels not seen by the tuned model, we set $\alpha_t(y) = 0$, so $\alpha_o(y) = 1$.

3.3 Storage efficiency and privacy

Partial tuning of our model requires either storing each image or storing the features (or tokens) that feed into the tuned portion of the model. Storing images

has disadvantages of lack of privacy and inefficiency in space and computation, due to need to re-encode in training. Storing features alleviates some of these problems, but still uses much memory or storage. Therefore, we investigate how to compress the training features, aiming to increase storage and computational efficiency while maintaining training effectiveness. The compressed features also provide more privacy than storing original images, though we do not investigate how well the images can be recovered from compressed features.

Well-trained networks learn data-efficient representations that are difficult to compress. Indeed, if we try to compress feature vectors with VQ-VAE [29] or PCA (principal component analysis) trained on a dataset, we are not able to achieve any meaningful compression without great loss in training performance. The features within *each image*, however, contain many redundancies. We, therefore, compute PCA vectors on the features in each image and store those vectors along with the coefficients of each feature vector. Further, not all tokens are equally important for prediction. Hence, we train a per-image attention-weighted PCA, weighted by attention between each token and the CLS token. Finally, we can compress further by storing min/max floating point values for each vector and for the coefficients and quantizing them to 8-bit or 16-bit unsigned integers. See the supplemental material for details. By storing only five PCA vectors and their coefficients this way, we can reduce the storage of fifty 768-dim tokens (7×7 patch tokens + CLS token) from 153K bytes to 5K bytes with less than 1% difference in prediction accuracy.

4 Experiments

Using the setup from [62], we sequentially receive training samples for a target task in a **data-incremental** (random ordering) or **class-incremental** (class-sorted ordering). Or, in a **task-incremental** setting, we receive all examples for each task sequentially. Within each set, we incorporate new examples **one by one** in an online fashion. After receiving each set, the model is evaluated on the entire task (including classes not yet seen in training) for data- or class-incremental, or on all tasks (including tasks not yet seen in training) for task-incremental.

Our main experiments use the same setup as Zhu et al. [62], which are designed to test flexible learning and flexible inference. A **target task** contains a subset of seen labels while a **novel task** contains none of these labels. Target tasks are CIFAR100 [21], SUN397 [55], FGVCAircraft [26], EuroSAT [14], OxfordIIITPets [32], StanfordCars [20], Food101 [6], and Flowers102 [28]. Novel tasks are ImageNet [39], UCF101 [46], and DTD [10]. Training examples are received for target tasks, but not for novel tasks. Under this setup, we have 226,080 training samples and 1,034 classes in total.

- **Task Incremental Learning** updates one model with each of the eight target tasks sequentially.
- **Class Incremental Learning** updates a per-task model with examples from a subset (one-fifth) of classes for a stage.

- **Data Incremental Learning** updates a per-task model with a random subset of the data in eight stages: 2%, 4%, 8%, 16%, 32%, 64%, 100%.

After each stage, accuracy is averaged across all tasks (including tasks/classes not yet seen).

Additionally, we provide a **union data incremental** scenario where we mix all target tasks together and union all target labels, mainly for hyperparameter selection and ablation experiments. **Flexible inference** is evaluated after the task-incremental learning for prediction over novel tasks and sets of candidate labels \mathcal{Y} : zero-shot, union of all target tasks/labels and zero-shot tasks/labels, and a mix of some target and zero-shot tasks/labels. See [62] for details. The union and mixed settings are most challenging because they require effective combined use of both the original and tuned models, without any task identifiers or indication whether an example’s label is in the tuned model’s domain.

Implementation details. We employ the ViT-B/32 model from CLIP as our network backbone. We use AdamW [25] as the optimizer with a weight decay of 0.05, and we adhere to CLIP’s standard preprocessing without additional data augmentations. For offline training, following [11], we set the learning rate at $6e-4$ and the batch size at 2048, employing a cosine annealing scheduler for the learning rate with a minimum rate of $1e-6$. For online training, we use a batch size of 32 and a learning rate of $9.375e = \frac{32 \times 6e-4}{2048}$, based on adjustments from the offline training hyperparameters. More method and training details are in the supplemental.

4.1 Main result

In this evaluation, we compare our method with the previous SotA Zhu et al. [62] (CLIP+LinProbe (AIM-Emb)) under the task-, class-, data-incremental scenarios, as well as flexible inference. From all plots in Fig. 3, our AnytimeCL method consistently outperforms CLIP+LinProbe (AIM). For task incremental, the improvement is mainly due to the partial finetuning with fixed label embeddings. To provide a more direct comparison with Zhu et al. [62], we train our approach only after receiving all data of a task (*e.g.*, 100% CIFAR100) and term this as AnytimeCL (Offline), which also shows steady improvements over Zhu et al. [62]. For early stages in class incremental and data incremental, the performance improvement mainly comes from the advantage of our online class-wise weighting method over AIM, which only considers whether the testing sample is likely to be from a known class.

4.2 Comparison under MTIL task incremental learning [61]

We compare the performance of AnytimeCL with Zhu et al. [62] and ZSCL [61] under the task incremental learning setting as introduced in ZSCL. Results are summarized in Tab. 1, using the Transfer, Last, and Average metrics from [61] for comparison. Improvements over CLIP are denoted in green under the Δ columns, while declines are marked in red. Our method consistently achieves

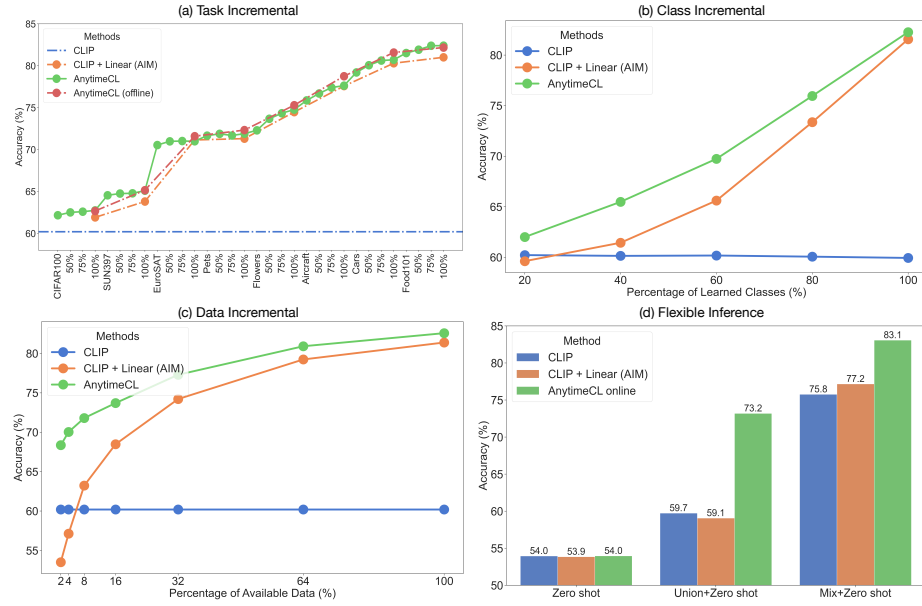


Fig. 3: Performance comparison of CLIP, CLIP + Linear (AIM), and AnytimeCL variants in task incremental (a); class incremental (b); data incremental (c); and flexible inference (d) settings. In (a), the online approach is represented by a solid line, while offline methods are depicted with dashed lines, assessed when the online algorithms receive 25%, 50%, 75%, and 100% data of a task, labeling tasks at the 25% point.

zero loss in Transfer in both experiments, indicating an absence of forgetting. On this benchmark, our method performs comparably to LinProbe (AIM) and outperforms other methods. This experimental setup does not fully capture our method’s key advantages of online training, ability for data and class incremental learning, and open vocabulary prediction. Also, in this setting, a simple weighting method is effective — to use the tuned model’s prediction for tasks with all candidate labels in the training set, or the original model’s prediction otherwise.

4.3 Training features compression

We test our per-image attention-weighted PCA compression for training features (Table 2). We compare the data size, average full time to process one 32-sample batch (including data loading, uncompressing, and forward/backward training passes), and final accuracy after completing fine-tuning. This test is on CIFAR100. Compared to processing the full image or full features, using our compressed features saves 30x the storage while achieving nearly the same accuracy. We were not able to achieve good accuracy when compressing with VQ-VAE, even when using large vocabularies, nor when using PCA vectors computed over the entire dataset. Computing PCA vectors per instance (over the 50 tokens) and weighting the tokens by their CLS-attention each improve the tuned accu-

Method	Transfer	Δ	Avg.	Δ	Last	Δ
CLIP	69.4	0.0	65.3	0.0	65.3	0.0
LwF [23]	56.9	-12.5	64.7	-0.6	74.6	+9.3
iCaRL [38]	50.4	-19.0	65.7	+0.4	80.1	+14.8
WiSE-FT [54]	52.3	-17.1	60.7	-4.6	77.7	+12.4
ZSCL [61]	68.1	-1.3	75.4	+10.1	83.6	+18.3
TreeProbe (AIM) (50k) [62]	69.3	-0.1	75.9	+10.6	85.5	+20.2
LinProbe (AIM) [62]	69.3	-0.1	77.1	+11.8	86.0	+20.7
AnytimeCL	69.4	0	77.0	+11.7	85.8	+20.5
AnytimeCL (Offline)	69.4	0	77.0	+11.7	86.2	+20.9

Table 1: Comparison of different methods on MTIL in Order I from ZSCL [61]. CLIP is our tested zero-shot performance. All other results are taken Zhu et al. [62] and ZSCL [61]’s papers. “Transfer” evaluates the model’s performance on zero-shot tasks; “Last” is the averaged accuracy on all target tasks after finishing the final task while “Avg.” computes the average task performance on all training stages.

racy. Quantizing the PCA vectors and coefficients further saves space at a small loss to accuracy. Using the compressed features also gives a nearly 2x speedup vs. using the full features (because the compressed features fit entirely in memory) and 6x vs. using the full image, though all times are fast. We provide more method details in the supplemental along with a comparison of using different per-instance PCA compression methods under the union data incremental scenario, to show how these compression methods influence the result at various timesteps.

Compression	KB/example	ms/batch	FT Accuracy
Full image	150.5	43.9	77.8
Full features	153.6	25.6	77.8
VQ	0.2	4.5	64.8
Dataset-wide PCA (200 components)	40.2	7.6	76.4
Per-instance PCA (5 components)	19.4	7.7	74.6
+ CLS-weight (5 components)	19.4	8.9	77.9
+ int-quantization (5 components)	5.3	13.9	77.5

Table 2: Comparison of different compression methods in terms of memory usage, processing speed, and test accuracy on CIFAR100. See supplementary material for details on how to get ms/batch and FT Accuracy. We use the same seed for all methods.

4.4 Ablations

In the following, we ablate several important factors that contribute to our method. Additional ablations and results are provided in the supplemental.

Weighting Strategies. Weighting is vital for our dual decoder approach. We compare several ways to compute the weights: 1) CLIP: namely $\alpha_t(y) = 0$ for

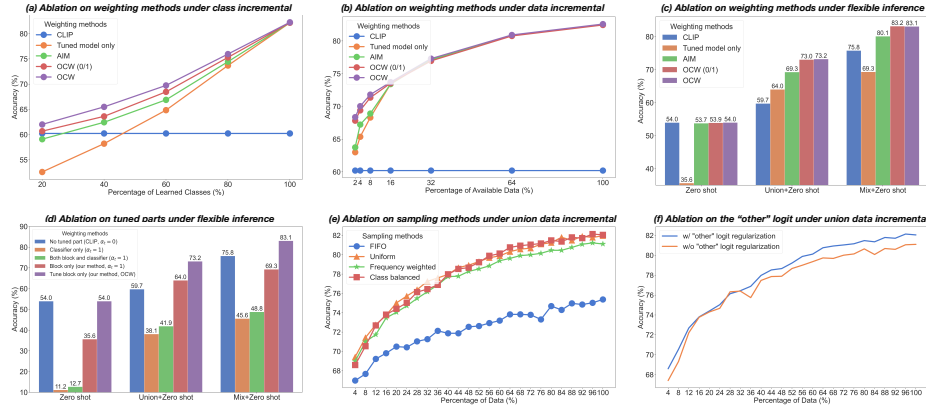


Fig. 4: Ablation results. Best viewed in color and zoomed-in.

any images; 2) Tuned model only: $\alpha_t(y) = 1$ for any images; 3) AIM [62]; 4) OCW (0/1): a variant of OCW where we round $\alpha_t(y)$ to 0 or 1 to use either the original or tuned model; 5) Our proposed OCW. We partially finetune the decoder with fixed label embeddings and combine the tuned model with the original model using different weighting strategies. The results are shown in Fig. 4 (a), (b) and (c), under data-, class-incremental and flexible inference. OCW performs the best of all methods in every setting and stage. Notably, OCW enables improvements even in the early stages when limited data is available for a class or task and better handles tasks that mix novel and target labels.

Tuned parts. Our proposed method tunes the last transformer block while keeping the label encoder fixed. In Fig. 4 (d), we compare this approach to alternatives of tuning only the label encoder, both the block and the label encoder, or neither of them, under the flexible inference test. When only using the tuned model for comparison ($\alpha_t = 1$), fine tuning only the last transformer best retains predictive ability for novel labels. Further analysis under the task incremental scenario is provided in the supplemental material.

Sampling methods. We compare different methods for sampling from the stored samples. FIFO cycles through samples in order of first appearance, “uniform” randomly draws samples for each batch, class-balanced (which we use in all other experiments) samples classes, and frequency weighted sampling (FWS) samples based on how many times a given sample has been batched in training. Class-balanced and uniform are similar in practice, and perform best (Fig. 4 (e)). This result is consistent with Prabhu et al. [35]’s finding that simple sampling methods outperform complex ones under a computational budget. See supplemental for more details.

The “other” logit regularization. The ablation experiments depicted in Fig. 4 (f) assess the impact of “other” logit regularization in the union data incremental scenario. The results demonstrate consistent enhancements when this regularization is applied, compared to its absence.

4.5 Scalability

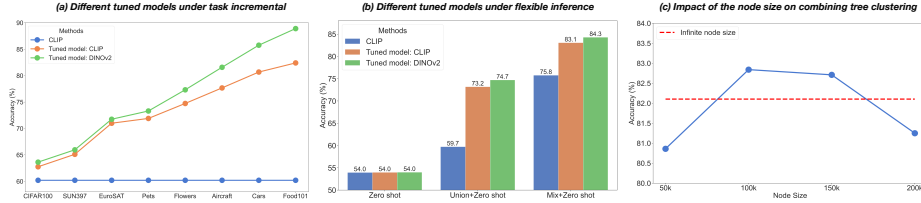


Fig. 5: (a) & (b): The “CLIP” method refers to the original model. For Tuned model: DINOv2, we use the ViT-B/14 checkpoint; (c): Infinite node size indicates only one tuned model regardless of the number of samples received.

DINOv2 as the Tuned Model. Using self-supervised DINOv2 [30] as our tuned model showcases our method’s adaptability beyond the original CLIP model [62]. Our approach allows for tuning the feature embedding space while keeping label embeddings fixed, enabling the incorporation of diverse encoders through an additional linear layer for dimension matching. This layer is trained in joint with the tuned decoder. As illustrated in Fig. 5 (a), replacing the tuned model with DINOv2 results in consistent performance improvements at every stage, with a notably steeper improvement curve in later stages. Furthermore, Fig. 5 (b) demonstrates that this modification preserves Zero-shot performance while significantly boosting Union+Zero-shot and Mix+Zero-shot outcomes, attributed to DINOv2’s enhanced performance on target tasks.

Tree clustering models. Zhu et al. [62] propose tree-based clustering as a way to limit the training time needed to incorporate new examples. Our method can already incorporate an example much faster in a way that does not depend on the number of training samples, but the clustering idea is still intriguing as a path to scalability in storage, memory, and model capacity. In Fig. 5 (b), we evaluate different node capacities and find that moderate-sized partitions may slightly increase performance. In these experiments, each node stores a copy of the last CLIP transformer block and updates it based on data that it receives.

5 Discussion

Our AnytimeCL approach shows promise of breaking free of staged train and deploy processes, with many benefits to applications that constantly or sporadically receive new data and have customizable or evolving label spaces.

The *anytime* has three main parts: (1) being able to do a task even when training data do not fully cover the label space for that task; (2) being able to incorporate a new training example quickly; and (3) not losing the benefit of older training data. Our key to the first part is using complimentary models. For that we are motivated by CLS theory [31] and more specifically adopt the

framework of Zhu et al. [62], improving it with our prediction weighting method and by partially fine-tuning the encoder. Our key to the second part is that we do not need to incorporate one million examples quickly, just one. This reframes the computationally budgeted CL problem of Prabhu et al. [35], who aim to incorporate a new large batch of examples in time that allows visiting only a portion of the existing data. For some settings, our goal is more practical, as it enables an opportunistic approach that interleaves pure online learning during busy times and offline learning when time allows. While humans are often held up as excellent continual learners, they also learn one thing fast but many things slowly over time and make use of restful wake and sleep to continue learning. The effectiveness of our online learning is also due, in part, to leveraging a strong vision-language model, just recently possible. The key to the third goal is simple — do not throw out the data. If one cares about privacy and storage, then finding ways of compressing and preserving privacy, like our PCA-based approach, may be better than trying to retain the benefits of a large dataset without one.

Potential directions for future work:

- Beyond classification: Our partial fine-tuning approach is applicable to other tasks, such as semantic segmentation, visual question answering, and object detection. Our weighting method is applicable to open vocabulary detection and segmentation.
- Multi-model inference: Extending our approach to multiple models, including task-specific models, would further extend flexible learning and model re-use opportunities.
- Scalability: We expect that tree-based data clustering and training one model per cluster, as in [62], provides a good mechanism for scalability. Larger scale experiments, involving dozens of datasets and many millions of examples, are needed to fully test this.
- Federated learning: Combining tree-based clustering and feature compression, we can encode training examples in the client or central server and transmit to a node that stores and updates the tunable block and its data. This enables fully distributed training on inexpensive nodes.

6 Conclusion

We propose an effective approach for anytime continual learning of open vocabulary image classification. The main innovation is dynamic class-sensitive weighting for combining predictions of open-vocabulary and example-based tuned models, and we also find that partial fine-tuning with sufficiently small learning rates enables a surprisingly effective online learner. We offer a per-image attention-weighted PCA approach to compress features, with benefits to storage, computation, and privacy. Our experiments show that our approach is very promising, and further exploration is merited to more fully understand its effective scope and limitations.

Acknowledgement

This work is supported in part by ONR award N00014-21-1-2705, ONR award N00014-23-1-2383, and U.S. DARPA ECOLE Program No. #HR00112390060. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, ONR, or the U.S. Government.

References

1. Aljundi, R., Caccia, L., Belilovsky, E., Caccia, M., Lin, M., Charlin, L., Tuytelaars, T.: Online continual learning with maximally interfered retrieval. *CoRR* (2019)
2. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. In: *CVPR* (2017)
3. Aljundi, R., Lin, M., Goujaud, B., Bengio, Y.: Gradient based sample selection for online continual learning. In: *NeurIPS* (2019)
4. Arani, E., Sarfraz, F., Zonooz, B.: Learning fast, learning slow: A general continual learning method based on complementary learning system. In: *ICLR* (2022)
5. Bang, J., Kim, H., Yoo, Y., Ha, J., Choi, J.: Rainbow memory: Continual learning with a memory of diverse samples. In: *CVPR* (2021)
6. Bossard, L., Guillaumin, M., Gool, L.V.: Food-101 - mining discriminative components with random forests. In: *ECCV* (2014)
7. Brodt, S., Inostroza, M., Niethard, N., Born, J.: Sleep—a brain-state serving systems memory consolidation. *Neuron* (2023)
8. Caccia, L., Aljundi, R., Asadi, N., Tuytelaars, T., Pineau, J., Belilovsky, E.: New insights on reducing abrupt representation change in online continual learning. In: *ICLR* (2022)
9. Caccia, L., Belilovsky, E., Caccia, M., Pineau, J.: Online learned continual compression with adaptive quantization modules. In: *ICML* (2020)
10. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: *CVPR* (2014)
11. Dong, X., Bao, J., Zhang, T., Chen, D., Gu, S., Zhang, W., Yuan, L., Chen, D., Wen, F., Yu, N.: CLIP itself is a strong fine-tuner: Achieving 85.7% and 88.0% top-1 accuracy with vit-b and vit-l on imagenet. *CoRR* (2022)
12. Farhadi, A., Endres, I., Hoiem, D., Forsyth, D.: Describing objects by their attributes. In: *CVPR. IEEE* (2009)
13. Gautam, C., Parameswaran, S., Mishra, A., Sundaram, S.: Online lifelong generalized zero-shot learning. *Neural networks* (2021)
14. Helber, P., Bischke, B., Dengel, A., Borth, D.: 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* (2019)
15. Ilharco, G., Wortsman, M., Gadre, S.Y., Song, S., Hajishirzi, H., Kornblith, S., Farhadi, A., Schmidt, L.: Patching open-vocabulary models by interpolating weights. In: *NeurIPS* (2022)
16. Jin, W., Cheng, Y., Shen, Y., Chen, W., Ren, X.: A good prompt is worth millions of parameters: Low-resource prompt-based learning for vision-language models. In: *Proc. ACL* (2022)
17. Khattak, M.U., Rasheed, H.A., Maaz, M., Khan, S.H., Khan, F.S.: Maple: Multi-modal prompt learning. In: *CVPR* (2023)

18. Khattak, M.U., Wasim, S.T., Naseer, M., Khan, S., Yang, M., Khan, F.S.: Self-regulating prompts: Foundational model adaptation without forgetting. In: ICCV (2023)
19. Kirkpatrick, J., Pascanu, R., Rabinowitz, N.C., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. CoRR (2016)
20. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13) (2013)
21. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Master's thesis, Department of Computer Science, University of Toronto (2009)
22. Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-based classification for zero-shot visual object categorization. IEEE TPAMI (2014)
23. Li, Z., Hoiem, D.: Learning without forgetting. In: ECCV (2016)
24. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. In: NeurIPS (2017)
25. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
26. Maji, S., Rahtu, E., Kannala, J., Blaschko, M.B., Vedaldi, A.: Fine-grained visual classification of aircraft. CoRR (2013)
27. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: CVPR (2018)
28. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing (2008)
29. van den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. In: NeurIPS (2017)
30. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P., Li, S., Misra, I., Rabbat, M.G., Sharma, V., Synnaeve, G., Xu, H., Jégou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision. CoRR (2023)
31. O'Reilly, R.C., Bhattacharyya, R., Howard, M.D., Ketz, N.: Complementary learning systems. Cogn. Sci. (2014)
32. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.V.: Cats and dogs. In: CVPR (2012)
33. Pham, Q., Liu, C., Hoi, S.: Dualnet: Continual learning, fast and slow. In: NeurIPS (2021)
34. Prabhu, A., Cai, Z., Dokania, P.K., Torr, P.H.S., Koltun, V., Sener, O.: Online continual learning without the storage constraint. CoRR (2023)
35. Prabhu, A., Hammoud, H.A.A.K., Dokania, P.K., Torr, P.H.S., Lim, S., Ghanem, B., Bibi, A.: Computationally budgeted continual learning: What does matter? CVPR (2023)
36. Prabhu, A., Torr, P., Dokania, P.: Gdumb: A simple approach that questions our progress in continual learning. In: ECCV (2020)
37. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021)
38. Rebuffi, S., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: CVPR (2017)

39. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *IJCV* (2015)
40. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. *arXiv:1606.04671* (2016)
41. Serrà, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: *ICML* (2018)
42. Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. In: *NeurIPS* (2017)
43. Siegel, J.M.: Memory consolidation is similar in waking and sleep. *Curr Sleep Med Rep* (2021)
44. Skorokhodov, I., Elhoseiny, M.: Class normalization for (continual)? generalized zero-shot learning. In: *ICLR* (2021)
45. Smith, J.S., Karlinsky, L., Gutta, V., Cascante-Bonilla, P., Kim, D., Arbelle, A., Panda, R., Feris, R., Kira, Z.: Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In: *CVPR* (2023)
46. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR* (2012)
47. Tong, S., Jones, E., Steinhardt, J.: Mass-producing failures of multimodal systems with language models. In: *NeurIPS* (2023)
48. Tong, S., Liu, Z., Zhai, Y., Ma, Y., LeCun, Y., Xie, S.: Eyes wide shut? exploring the visual shortcomings of multimodal llms. *CoRR* (2024)
49. Wamsley, E.J.: Offline memory consolidation during waking rest. *Nature Reviews Psychology* (2022)
50. Wang, L., Zhang, X., Su, H., Zhu, J.: A comprehensive survey of continual learning: Theory, method and application. *CoRR* (2023)
51. Wang, Y., Huang, Z., Hong, X.: S-prompts learning with pre-trained transformers: An occam’s razor for domain incremental learning. In: *NeurIPS* (2022)
52. Wang, Z., Zhang, Z., Lee, C., Zhang, H., Sun, R., Ren, X., Su, G., Perot, V., Dy, J.G., Pfister, T.: Learning to prompt for continual learning. In: *CVPR* (2022)
53. Wang, Z., Zhang, Z., adn Ruoxi Sun, S.E., Zhang, H., Ren, C.Y.L.X., Su, G., Perot, V., Dy, J., Pfister, T.: Dualprompt: Complementary prompting for rehearsal-free continual learning. In: *ECCV* (2022)
54. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., Schmidt, L.: Robust fine-tuning of zero-shot models. In: *CVPR* (2022)
55. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: *CVPR* (2010)
56. Yan, S., Xie, J., He, X.: DER: dynamically expandable representation for class incremental learning. In: *CVPR* (2021)
57. Yoon, J., Yang, E., Lee, J., Hwang, S.J.: Lifelong learning with dynamically expandable networks. In: *ICLR* (2018)
58. Yu, R., Liu, S., Wang, X.: Dataset distillation: A comprehensive review (2023)
59. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: *ICML*. pp. 3987–3995 (2017)
60. Zhang, J.O., Sax, A., Zamir, A., Guibas, L.J., Malik, J.: Side-tuning: A baseline for network adaptation via additive side networks. In: *ECCV* (2020)
61. Zheng, Z., Ma, M., Wang, K., Qin, Z., Yue, X., You, Y.: Preventing zero-shot transfer degradation in continual learning of vision-language models. *CoRR* (2023)

62. Zhu, Z., Lyu, W., Xiao, Y., Hoiem, D.: Continual learning in open-vocabulary classification with complementary memory systems. CoRR (2023)