

Fine-tuning Image Transformers using Learnable Memory

Mark Sandler Andrey Zhmoginov Max Vladymyrov Andrew Jackson
 Google Inc.
 {sandler, azhmogin, mxv, jacksona}@google.com

Abstract

In this paper we propose augmenting Vision Transformer models with learnable memory tokens. Our approach allows the model to adapt to new tasks, using few parameters, while optionally preserving its capabilities on previously learned tasks. At each layer we introduce a set of learnable embedding vectors that provide contextual information useful for specific datasets. We call these “memory tokens”. We show that augmenting a model with just a handful of such tokens per layer significantly improves accuracy when compared to conventional head-only fine-tuning, and performs only slightly below the significantly more expensive full fine-tuning. We then propose an attention-masking approach that enables extension to new downstream tasks, with a computation reuse. In this setup in addition to being parameters efficient, models can execute both old and new tasks as a part of single inference at a small incremental cost.

1. Introduction

Transformers [38], originally introduced for sequence problems, such as NLP and speech recognition, have been very successful in advancing state-of-the-art for vision tasks. “Vision transformers” most commonly have been trained on large amounts of data, either in supervised [21] or self-supervised [7] modes. The resulting model is then “fine-tuned” on a downstream task, either in whole or in part. This fine-tuning has been shown to help achieve high performance on a variety of downstream tasks, from classical tasks like image super-resolution [6] and object detection [3] to novel ones, such as Ising model simulations [19]. It has been observed that in order to reach the highest accuracy it is best to fine-tune the entire model on the target task [7,21]. However, this is often problematic, since Transformer models can have hundreds of millions or billions of parameters, making maintaining, transmitting and storing a full size model for each new task an expensive proposition. Additionally, full fine-tuning is often sensitive to the choice

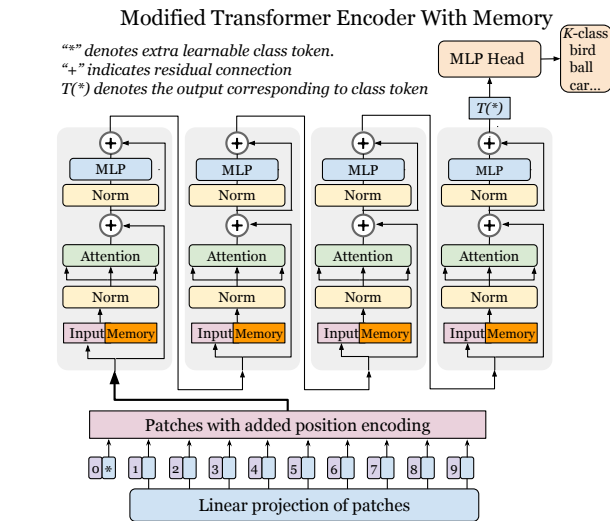


Figure 1. Structure of an encoder with learned memory. The input to each encoder layer is augmented with learned memory, which is used as an optional source of attention by the embeddings that propagate from previous layers. Memory tokens do not attend to other tokens. For detailed design of individual layer see Fig. 2.

of the learning rate [30]. In this paper we propose a method where we augment a pre-trained model with learnable tokens that are appended at *each* layer of the transformer block. These extra tokens can be intuitively thought of as a permanent memory that are used as a contextual reference to improve the final prediction. Another point of analogy is that these tokens can represent concepts at varying levels of abstractions: starting from tokens that run alongside pixel-based patches, and all the way to the tokens representing final concepts.

Our approach allows significant improvements to downstream task accuracy compared to head-only fine-tuning. Further, we believe that our particular architecture design is of independent interest as it provides a novel way of extending a trained model to perform new tasks while retain-

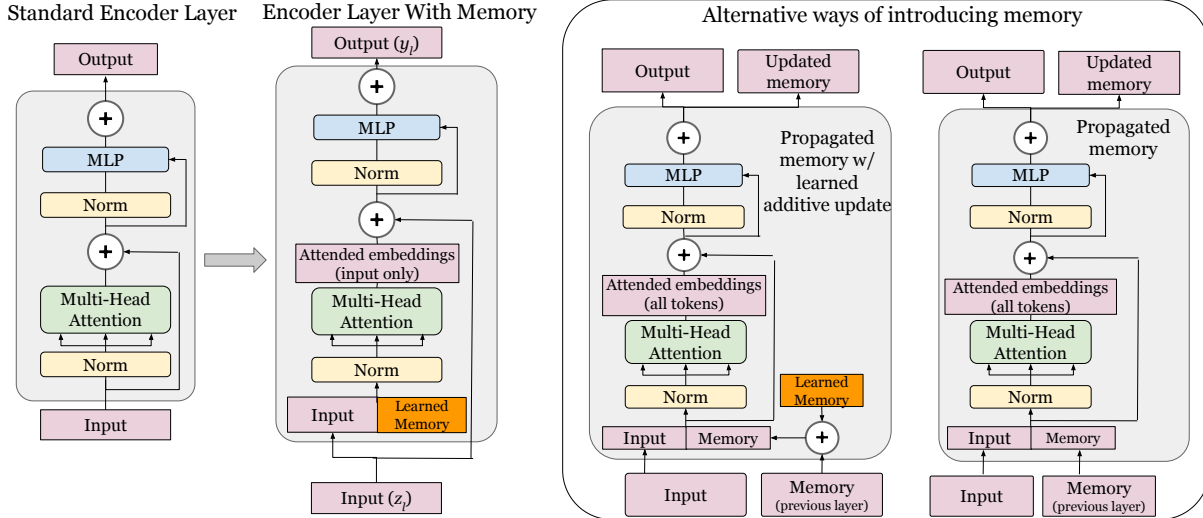


Figure 2. Detailed design of encoder layer with memory. The designs on the right are the alternative options that we explore in Sec. 4. The rightmost one is similar to the MemTransformer design in [5]. Here we use rounded boxes to indicate functions and square boxes to indicate input/output to highlight data propagation.

ing the ability of performing old tasks with only incremental compute cost. This is important in many applications, as the training data for different task often not available at the same time. Naively, by adding more memory one can easily extend a model to new tasks, however, if we try to run such model on an old task we would commonly observe 20%-40% percentage drop in accuracy. We show that with proper attention masking, additional cells can be introduced in a way that doesn't degrade the accuracy of the previous tasks in the presence of new memories, which is crucial for applications in life-long and continual learning. We show that the models can be both *extended*, where new functionality is added on top of existing ones, as well as *concatenated*, where the two (or more) independently trained models can be combined into a single network performing all tasks, with only incremental cost per task.

The paper is organized as follows. In Sec. 2 we describe the method of our memory model. Then in Sec. 2.2 we introduce our attention masking technique that independent sequential task training, without accuracy degradation. Then in Sec. 3 we provide an overview of related approaches to memory and fine-tuning approaches. Sec. 4 contains our experiments and ablation studies. The last Sec. 5 draws some early conclusions and identifies future areas of interest.

2. Memory model

To introduce memory, we build upon the Vision Transformer (ViT) encoder model [21], where an image $x \in \mathbb{R}^{h \times w \times c}$ is split into a grid of N patches, each of size $P \times P \times c$, which are flattened and fed through the multiple

layers of the transformer encoder. For this paper, we only consider classification models, and do not use a decoder. Following [21] we use learnable 1D position embeddings that are added to the input of the first layer. The input to the standard image transformer is defined as:

$$\mathbf{z}_0^{vit} := [x_{cls}, Ex_1, \dots, Ex_N] + E_{pos}$$

where E is a learnable linear transformation into the encoder embedding space and E_{pos} is a position embedding. Here x_{cls} is a special learnable token that is shared for all inputs and the corresponding output value is used as an embedding for the final classification as shown in Fig. 1. $x_1 \dots x_N$ are flattened image patches. The basic structure of the encoder layer is shown in Fig. 2. We refer to [21] for the detailed definition of each component.

Now, to add memory, we concatenate m learnable embeddings $E_{mem} \in \mathbb{R}^{m \times D}$, where D is the dimensionality of the input tokens, as an input to the layer:

$$\mathbf{z}_0^{mem} := [\mathbf{z}_0^{vit}; E_{mem}^0]$$

Thus our transformer receives $N + 1 + m$ tokens in total as an input. This input is then passed through a sequence of encoder-layers as in [21, 38]. The architecture of the individual layers exactly matches that of ViT, with a notable exception that we do not propagate updated memory: that is, the output of the self attention module is truncated to the first $N + 1$ tokens. Thus the output of the layer l , which we denote as y_l , has only $N + 1$ tokens. See the left part of Fig. 2 for a comparison between a standard layer block and our augmented version.

Q \ K	INP	CLS	C_1	C_2	...	C_k	M_1	M_2	...	M_k
INP	✓	✓	-	-	...	-	-	-	...	-
CLS	✓	✓	-	-	...	-	-	-	...	-
C_1	✓	✓	✓	-	...	-	✓	-	...	-
C_2	✓	✓	★	✓	...	-	★	✓	...	-
C_k	✓	✓	★	★	...	✓	★	★	...	✓

Table 1. Attention mask for model extension and concatenation. Here ✓ indicates that corresponding token type Q (query) attends to token type K (key), - that it does not attend, and ★ indicates that attention is used for model extension, but not for concatenation. For brevity, we denoted CLS- K as C_k and MEM- K as M_k and omitted memory rows since they do not attend to other tokens. See Sec. 2.2 for details.

All consequent layers receive input that follows the same pattern:

$$\mathbf{z}_l^{mem} = [\mathbf{y}_{l-1}; E_{mem}^l]$$

where \mathbf{y}_{l-1} is the truncated output of the previous layer.

In our experiments we explored several different variants of memory models, including ones where memory itself actively attends and propagates to the next layer, however we found that such modifications generally hurt performance. We report some of those results in the ablation study in Sec. 4.3.

2.1. Fine-tuning with full attention

The main use of memory that we explore in this paper, is applying it to fine-tuning existing models. To do this we introduce randomly-initialized memory tokens as described above and perform gradient descent to learn the memory token, the classifier head, and the class token x_{cls} . While this method gives excellent results, the resulting model can no longer be used to solve its original task, due to change in hidden activations. In the next section we introduce *Attention Masking* that allows to lift this constraint.

2.2. Attention masking for computation reuse

After fine-tuning the class token of a transformer model on a new dataset, or adding memory, the network performance on the original task generally degrades. This might not be a problem if we are only interested in solving the new task. However it is often desirable to be able to execute both original and new tasks as part of the same inference, with only incremental cost per each new task. A common way to approach this is via multi-task learning where all datasets are present at training time and we learn a universal model with a shared backbone and per-task head solves all tasks. However, it is not always possible, as in practice data is often owned by different entities, and can not be learned within the same environment or at the same time. In this section we describe an extension of our model where

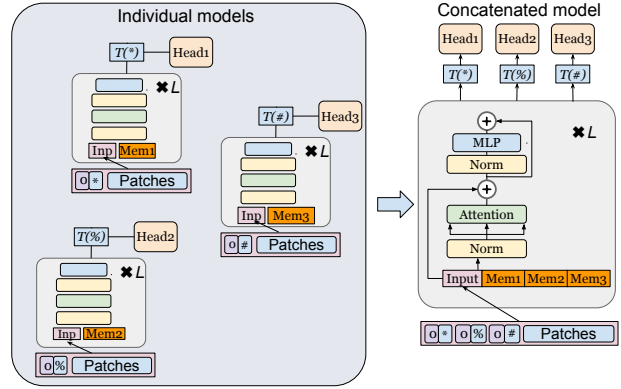


Figure 3. Concatenating multiple models together. The class tokens are concatenated at the input level, while individual models’ memory are concatenated at each layer.

we introduce memory a new dataset class token x_{cls}^{new} and a per-task head in such a way that the model output for the original dataset token x_{cls} is preserved, i.e., identical to the output obtained for this dataset token in the original model. As a result, we can solve multiple tasks for the same input, effectively reusing *both* parameters *and* computation. We achieve this by using an attention mask that prevents the original x_{cls} and input patches from attending to the memory tokens and x_{cls}^{new} that we introduce for the new dataset. A head for each task is attached to the corresponding dataset token. (see Tab. 1).

Model concatenation vs. model extension. Attention masking can be used to extend memory as more and more tasks are added. For model extension each new dataset \mathcal{D}_k is trained by fine-tuning its individual memory and the dataset class token $x_{cls}^{(k)}$ and the attention is constrained to explicitly prevent older datasets from attending to tokens added by following tasks (see Tab. 1). Alternatively, if we have two independently trained subtasks¹, we can *concatenate* these models into one as shown in Fig. 3. In this case the attention mask for the class token $x_{cls}^{(k)}$ in the concatenated model will be chosen to only attend to x_{cls} and input patches, and not to other task tokens or memory, thus allowing us to concatenate memory and tokens for models trained independently. Interestingly, as we discuss in Sec. 4.4 we did not see much improvement when extending a model, despite it theoretically being more powerful, as opposed to concatenating them. However we believe this remains an interesting direction to explore in future work.

¹trained with attention masking preventing x_{cls} and input patches from attending to x_{cls}^{new}

3. Related work

Transformers have been the subject of extensive research in the last few years, first in the NLP community [38]. Vision transformers were first introduced in [21, 35], where they achieved impressive results on a series of complex classification tasks. We refer to [20, 27] for an in-depth survey of recent literature.

There have been a few attempts to couple a learning network with an *external* immutable memory via a learned differentiable module for encoding and retrieval of the information [13, 33, 41]. Subsequent work focused on representing memory using a more explicit data structure, such as stacks [18]. More recently, in the context of transformers for language modeling, several papers [10, 12, 14, 25, 39] propose to augment a model with access to fixed knowledge base.

On the other hand, *internal* episodic memory in deep neural networks has been subject of active research since early on, often in the context of recurrent networks. Hopfield networks [16] were one of the first attempts to store information within the layers of a neural network. In a recent years, there has been proposed a series of improvements to the energy function of such networks to increase its capacity [24], help protect against adversarial attacks [23] or as a replacement for attention in transformer networks [32].

More generally, different variations of internal recurrent memory networks are actively used in many sequential models, such as in GRU [9] or LSTM [15] architectures. For LSTM, the recurrent cells could be “enabled” to preserve or forget a cumulative state as a document was scanned sequentially. Transformers represent a detour where instead of keeping the intermediate state as sequential data is scanned, they introduce a global attention mechanism that enables individual tokens to do random access to the full window of data. As a result, the memory is limited to the size of the window of data, which scales quadratically with respect to the attention window. Several papers attempted to reduce this cost by either introducing sparsity [8], low-rank approximations [40] or splitting the attention matrix into hierarchical modules connected locally [4]. In a similar vein, BigBird [45] and ETC [1] models introduce special global tokens that are not part of the input and serve as connector tokens for locally connected segments. More recent work expanded on the use of memory to preserve connections in long sequences to avoid the high computational cost of the full attention matrix [2, 42]. The main difference from our paper is that the memory considered in earlier works was episodic – it was preserved only for the duration of the sequence. In the present work the memory is *learned* as part of the training process and it does not represent information about any individual samples, but rather is a collective summary representing the entire dataset.

The closest approach to our work is a recent paper by

Burtsev *et al.* [5] and that considered a special case of memory where individual memory tokens were added at the first layer alongside the data and were treated as regular inputs and propagated using the same multi-head attention. A similar line of inquiry was also considered in Li *et al.* [26], where they considered single task-specific token. All these work was applied to NLP tasks. In contrast, we learn memory that is passed alongside inputs, as a source of attention, *at every layer* but it does *not* itself attend to any tokens, nor does it propagate to deeper layers. We show that this improves the performance when compared to only using learned tokens at the input-level layer. Another line of inquiry that uses a similar approach is a recent work on data-efficient transformer training by Touvron *et al.* [35]. There, authors train an additional class token for distillation purposes only. Using a dedicated token for a new task can be seen as a special memory based approach. This was further refined in the follow-up work by the same authors in [36], where they moved the class token into later layers of the architecture. The class token is then used as a read-out token. This can be seen as a hybrid of propagated-memory and finetuning the class token approach.

An alternative way of viewing the present work can be seen as modifying the individual layers during fine-tuning to redirect their attention in novel ways. Such approaches were explored in [17, 29, 31]. There, a set of special layers called adapters are introduced for each of the fine-tuned tasks. During fine-tuning, only those layers and/or normalization parameters [30, 44] are updated. Such approaches are more complex because they introduce additional structures and in contrast with our approach, they do not allow non-destructive fine-tuning.

4. Experiments

4.1. Datasets and training setup

For all our experiments, with the exception of Tab. 3 we use the ViT-B/32 base transformer model from [21]. We use the pre-trained model released in [21], which was trained on Imagenet-21K [11]. This model has about 80M parameters. We use a cosine learning rate schedule [28] and followed the setup of fine-tuning in [21], to generate our baselines. We used batch size 512, and all fine-tuning runs were run for 20000 steps. We note that shorter runs generally reached slightly worse results but preserved the relative relationship between different fine-tuning setups. We used SGD with Momentum with gradient clipping, though we did not see much benefit with or without it. We used a 5-step linear rate warmup.

Following the standard practice we used standard inception preprocessing [34] for all datasets except CIFAR-100, where we used random clipping. For the latter we only used random flips. The memory was initialized using

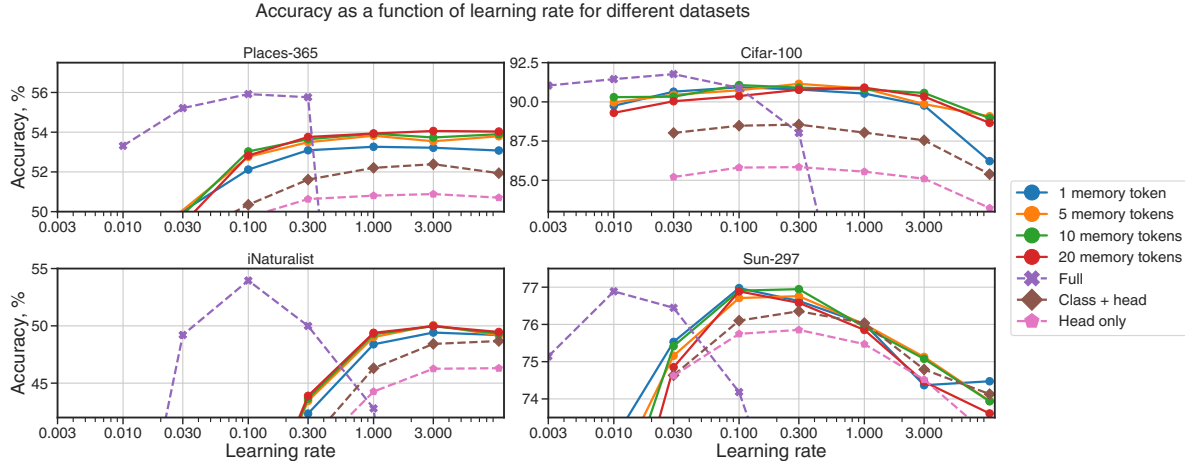


Figure 4. Performance of memory-augmented transformers for different datasets and comparison with baseline fine-tuning methods.

	Full	Head Only	Head + Class	1 cells	5 cells	10 cells	20 cells
Sun-297	76.9	75.9	76.4	77.0	76.8	76.9	76.9
iNaturalist	54.0	46.3	48.7	49.4	50.1	50.0	50.0
Cifar-100	91.8	85.8	88.6	90.9	91.0	91.1	90.9
Places-365	55.9	50.9	52.4	53.3	53.8	53.9	54.1

Table 2. Accuracy for different datasets using the optimal learning rate for each fine-tuning regime.

	Full	Head	Head + CLS	5 cells
ViT-B/32	55.9	50.9	52.4	53.8
ViT-B/16	58.0	52.3	53.8	55.4
ViT-L/32	56.6	52.4	49.7	54.8

Table 3. Accuracy for Places-365 dataset using optimal learning rate for different ViT architectures.

	Accuracy	Params (ViT-B/16)
Memory	50.1	768 * 12 * m \approx 46K
Adapters [17]	49.9	3078 * 5 * 4 * m \approx 737K
Combined	50.3	46K + 737K \approx 783K

Table 4. Performance for i-Naturalist for adapters [17] with bottleneck of size 5, vs learnable memory with 5 cells. (ours). Parameter counts exclude the size of the head which is the same for all methods

$\mathcal{N}(0, 0.02)$, following the initialization of other variables in Vision Transformers in [21]. We evaluate the performance of our approach on the following datasets CIFAR-100 [22], i-Naturalist [37], Places-365 [46] and SUN-397 [43]. In all cases we used 95% of the training split for training, the remaining 5% as a dataset for early stopping and hyperparameter selection, and we used the validation split to report final accuracy.

4.1.1 Baseline fine-tuning experiments

We compare our methods with the following fine-tuning modes.

Full fine-tuning the entire model is fine-tuned. This is the most expensive and thus often the least practical regime if multiple downstream tasks are present, because each task requires a full model. It is also prone to overfitting and is sensitive to the learning rate. For example, we found that often the optimal learning rate for other regimes resulted in full fine-tuning producing dramatically sub-optimal results, as can be seen in Fig. 4

Head-only fine-tuning as might be expected, only tunes the head of the classifier. This approach essentially fully reuses the embedding learned by a pretrained model. One important advantage of this mode is that it allows reuse of both the *parameters* and the *compute*, because the computation of embeddings is shared across tasks.

Head + Class token where we fine-tune the head and the class token. Note that this regime, while using only marginally more parameters, does not allow any significant computation reuse, since an input tokens attention (and embedding) will change as they attend to a new class-token. We note that this method can be seen as a precursor to memory training, since tuning the class token allows us to learn a representation of a dataset.

	Full	Head only	Class	Mem (5)
Params	80M	$768 \times k$	768	$768 \times 12 \times 5$
Flops	$\approx 1\text{G}$	0	0	$\approx 25\text{M}$

Table 5. Incremental cost on number of parameters and FLOPS for ViT-32-B for each fine-tuning regime. Class and head fine-tuning do not incur any extra computation cost as the architecture stays unchanged. Here k denotes number of classes, and we assume 5 memory tokens per layer.

4.1.2 Memory fine-tuning

Our main experiments use the following two regimes.

Memory + head + class token We refer to this as “memory” fine-tuning. This is our method. We conduct experiments on architectures containing 1, 2, 5, 10 and 20 cells. We did not see improvement when using more than 20 cells per layer. In fact most of the benefit is realized (see Tab. 2), when we use 5 cells for all our datasets, thus we conjecture that using 5 cells provides good universal size that can be used without any additional hyper-parameter search.

In Fig. 4 we show performance of learned memory vs other fine-tuning methods for different learning rate. As can be seen, especially fine-tuning performs best as much lower learning rates. Thus in all our experiments we used the best learning rate for each regime.

Memory with attention mask This is the regime which combines the benefit from *head-only* fine-tuning of allowing computation reuse, while achieving accuracy above *Head + Class* token, and slightly below the full-attention fine-tuning. The advantage of this regime is that models can be independently fine-tuned on multiple tasks, and combined in a single model that reuses most of the compute.

The number of variables used by our different approaches is shown in Tab. 5.

The rest of the section is organized as split into three parts. In Sec. 4.2 we discuss our results on standard transfer learning tasks, and in Sec. 4.4 we discuss the results on attention-masking based fine-tuning.

4.2. Performance of memory for transfer learning

In Fig. 4 we show our results for fine-tuning with memory versus the baseline regimes as a function of learning rate. Remarkably, full fine-tuning is much more sensitive to the learning rate. The performance dropped dramatically for a large learning rate, while memory and head-only fine-tuning achieve better performance with higher learning rates. As can be seen from the same figure, models with memory outperform all other fine-tuning regimes.

In Tab. 2 we show the comparison where we picked the best learning rate for each fine-tuning regime and for different numbers of memory tokens. As can be seen, for

most datasets the highest performance is achieved when the number of memory tokens is between 5 and 20. We hypothesize that such a small difference is caused by the increased potential for over-fitting as more memory is introduced. In the supplemental materials we show that training accuracy increases steadily as the number of memory tokens increases. For Cifar-100 and Sun-297, the performance of memory-fine-tuning nearly reaches the performance of full fine-tuning, while using $1/100$ of the number of variables.

4.3. Ablation Study

We now turn our attention to exploring which components of our approach are important, as well as showing some alternative approaches. First of all, we consider introducing memory at the first layer only and propagating through the entire transformer as shown in Fig. 2. This can be as seen as introducing additional class tokens as a source of attention for other tokens in the inputs. This approach is similar to that introduced in [5] for NLP tasks. While we see small improvement, as seen in Fig. 6, it is minimal compared to improvements from using full memory. A variant of this approach that can be seen as hybrid of propagating and learned memory, is when the memory is both propagated and then learned memory is added. This method worked better, but overall it appeared that using learnable memory works best without propagation.

Another question we consider is which layers are important for memory introduction. For that we introduce k layers of memory starting from the earliest layers (close to input), and compare it against memory introduced at deepest layers. As can be seen in Fig. 5 the performance gain is mostly gradual. It appears that adding memory to the initial layers has less performance benefit than adding them to the last 4 layers, but it appears that all layers benefit from adding memory.

4.4. Attention masking

In this section we study attention masking for model fine-tuning as discussed in Sec. 2.2. First we verify that performance degrades if we simply concatenate memory and class tokens fine-tuned independently without attention masking. We start by independently fine-tuning 4 separate models with dataset-specific class tokens and memory on SUN-397, Places-365, i-Naturalist and CIFAR-100. All four models are fine-tuned from the same pre-trained ViT-B/32 model. Starting with a model trained on i-Naturalist, we then progressively add class tokens and memory from models trained on Places-365, CIFAR-100 and SUN-397 in that order. As a result, the original i-Naturalist accuracy evaluated using a corresponding class token degrades from 50.0 to 28.1 (after adding Places-365), 27.8 (after adding CIFAR-100) and finally to 17.3 (after adding SUN-397). In other words, naive concatenation of class tokens and mem-

Performance when different layers have memory




Figure 5. Introducing memory in different layers. All results are for Places-365 dataset. The left chart shows the impact of serially introducing memory in the first 0, 4, 8 and 12 layers, respectively the left chart shows the effect of introducing memory in the deepest layers first. The right-most chart shows performance for best optimal learning rate for each count k .

Fine-tuning	Comp. Reuse	CIFAR-100	Places-365	SUN-397	i-Naturalist
<i>Class, no memory</i>	N	88.6	52.4	76.4	48.7
<i>Class, 10 memory tokens, full attention</i>	N	91.1	53.9	76.9	50.0
<i>Head-only, no memory</i>	Y	85.8	50.9	75.9	46.3
<i>Class, masking, no memory</i>	Y	87.2	52.5	75.8	49.0
<i>Class, 10 memory tokens, masking</i>	Y	88.7	53.3	76.4	49.5
<i>Full attention, 4 models combined</i>	-	71.5	47.6	67.8	17.3

Table 6. Test accuracies of models fine-tuned from a pre-trained ViT-B/32 model on different datasets. All models included an additional separate class token. We consider models trained with full attention and models trained with attention-masking. The models with masking can be combined with each other without any accuracy degradation, while also retaining ability to perform original Imagenet-21K predictions as part of the same computation. To demonstrate *interference* of models trained without attention masking, in a separate experiment (last row), we concatenated all models fine-tuned with full attention and memory, and then evaluated different dataset heads.




Figure 6. Performance of different memory architectures. All results are for Places-365 dataset. “Full memory” is our approach with learnable memory at each layer. “Propagated 1st” propagates memory introduced at the first layer, as described in [5]. “Propagated/added” propagates memory from previous layer and adds new at each layer. See Fig. 2 and Sec. 4.3 for more detail.

ory pre-trained without attention masking, results in destructive interference between the elements of the sequence, hurting the performance of individual model heads.

Alternatively, we could instead fine-tune individual models with attention masking and concatenate the individual

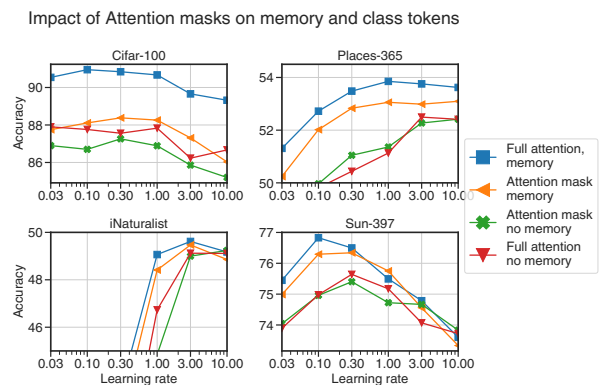


Figure 7. Performance of fine-tuning with attention mask. The top line (■) shows the baseline with full attention. The second (◀) shows attention-mask with memory. The last two are class+head fine-tune with full attention (▼) and with attention mask (×).

class tokens and memory tokens with a mask that would prevent them from interfering with each other (see Sec. 2.2), retaining the original accuracy of individual models. Fig-

ure 6 shows test accuracy for models fine-tuned on different datasets with attention masking. Our results, summarized in Tab. 6, show attention mask produces models that are less accurate compared to their full-attention counterparts, but there is still an advantage of using additional memory tokens compared to models without them and this approach to training allows us to assemble multiple models into a single model sharing most of the necessary computation while still retaining Imagenet-21K head predictions.

Finally, we conducted experiments where we *extended* the original ViT-B/32 model by adding dataset-specific class tokens and memory for each new consecutive dataset. In this setup, as discussed in Sec. 2.2, the attention mask restricted old class tokens and memory from attending to tokens added later for new datasets, but new tokens could attend to all existing tokens. As in all our experiments, we started with ViT-B/32 model and fine-tuned it on two datasets in one of two ways, either: (a) fine-tuned it first on SUN-397 and then on Places-365, or (b) fine-tuned in on Places-365 and then on SUN-397. We did not see any improvement from the case when SUN-397 and Places-365 were trained independently and concatenated post-training. We hypothesize that in order to see advantage in model extension, the intermediate datasets need to provide a significant amount of relevant samples compared to both the pre-training dataset and the new final dataset.

5. Conclusions, Limitations and Future work

We proposed a novel method of incorporating memory into transformer models. Our results show that memory improves the performance of a fine-tuned model significantly and allows training models for multi-task scenarios without compromising on performance. The memory tokens are shown to be an important part of the attention model. We believe they could be a critical ingredient for lifetime and continuous learning. There are several directions that are subject to future exploration. There are some limitations in our work that will be a subject of future work.

Scaling memory. We observed that as the number of memory tokens increases, there appears to be little benefit from using more than 10 tokens. One potential direction is to explore limiting the attention to the Top-K memory tokens. This would reduce the amount of background noise from irrelevant tokens.

Combining episodic memory with learnable memory. Episodic memory provides a way to memorize very long sequences, without remembering everything. This type of memory is typically used during inference and not learned directly. Unifying permanent and episodic memory within the same framework is an interesting extension.

Incremental Model extension. We showed that multiple attention masking models can be trained independently

and concatenated into one model with efficient computation reuse while outperforming the baseline. However, showing that models trained *sequentially* can benefit from intermediate memory accumulation is important for curriculum learning and is a subject of future work.

Novel domains and architectures An important next step is to introduce memory into other transformer models, and to apply our method to different domains, such as NLP.

References

- [1] Joshua Ainslie, Santiago Ontañón, Chris Alberti, Philip Pham, Anirudh Ravula, and Sumit Sanghai. ETC: encoding long and structured data in transformers. *arXiv e-prints*, pages arXiv–2004, 2020.
- [2] Anonymous. Memorizing transformers. *Under review for ICLR*, 2021.
- [3] Josh Beal, Eric Kim, Eric Tzeng, Dong Huk Park, Andrew Zhai, and Dmitry Kislyuk. Toward transformer-based object detection. *arXiv preprint arXiv:2012.09958*, 2020.
- [4] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [5] Mikhail S. Burtsev and Grigory V. Sapunov. Memory transformer. *CoRR*, abs/2006.11527, 2020.
- [6] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12299–12310, 2021.
- [7] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020.
- [8] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [10] Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, and William Cohen. Mention memory: incorporating textual knowledge into transformers through entity mention attention. *arXiv preprint arXiv:2110.06176*, 2021.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [12] Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. Entities as experts: Sparse memory access with entity supervision. *arXiv preprint arXiv:2004.07202*, 2020.
- [13] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

- [14] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. Retrieval augmented language model pre-training. In *International Conference on Machine Learning*, pages 3929–3938. PMLR, 2020.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [16] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [17] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019.
- [18] Armand Joulin and Tomas Mikolov. Inferring algorithmic patterns with stack-augmented recurrent nets. *Advances in neural information processing systems*, 28:190–198, 2015.
- [19] Onur Kara, Arijit Sehanobish, and Hector H Corzo. Fine-tuning vision transformers for the prediction of state variables in ising models. *arXiv preprint arXiv:2109.13925*, 2021.
- [20] Salman H. Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *CoRR*, abs/2101.01169, 2021.
- [21] Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, Thomas Unterthiner, and Xiaohua Zhai. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [22] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [23] Dmitry Krotov and John Hopfield. Dense associative memory is robust to adversarial inputs. *Neural computation*, 30(12):3151–3167, 2018.
- [24] Dmitry Krotov and John J Hopfield. Dense associative memory for pattern recognition. *Advances in neural information processing systems*, 29:1172–1180, 2016.
- [25] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *arXiv preprint arXiv:2005.11401*, 2020.
- [26] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *CoRR*, abs/2101.00190, 2021.
- [27] Yang Liu, Yao Zhang, Yixin Wang, Feng Hou, Jin Yuan, Jiang Tian, Yang Zhang, Zhongchao Shi, Jianping Fan, and Zhiqiang He. A survey of visual transformers, 2021.
- [28] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [29] Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489*, 2021.
- [30] Pramod Kaushik Mudrakarta, Mark Sandler, Andrey Zhmoginov, and Andrew Howard. K for the price of 1. parameter efficient multi-task and transfer learning. In *International Conference on Learning Representations*, 2019.
- [31] Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. AdapterHub: A framework for adapting transformers. *arXiv preprint arXiv:2007.07779*, 2020.
- [32] Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020.
- [33] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850. PMLR, 2016.
- [34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [35] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [36] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*, 2021.
- [37] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [39] Pat Verga, Haitian Sun, Livio Baldini Soares, and William W Cohen. Facts as experts: Adaptable and interpretable neural memory over symbolic knowledge. *arXiv preprint arXiv:2007.00849*, 2020.
- [40] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [41] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.
- [42] Qingyang Wu, Zhenzhong Lan, Jing Gu, and Zhou Yu. Memformer: The memory-augmented transformer, 2020.
- [43] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492, June 2010.

- [44] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018.
- [45] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. In *NeurIPS*, 2020.
- [46] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2017.