

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Continual Hippocampus Segmentation with Transformers

Amin Ranem, Camila González, Anirban Mukhopadhyay GRIS, Technical University of Darmstadt Karolinenpl. 5, 64289 Darmstadt, Germany

{amin.ranem, camila.gonzalez, anirban.mukhopadhyay}@gris.informatik.tu-darmstadt.de

Abstract

In clinical settings, where acquisition conditions and patient populations change over time, continual learning is key for ensuring the safe use of deep neural networks. Yet most existing work focuses on convolutional architectures and image classification. Instead, radiologists prefer to work with segmentation models that outline specific regions-of-interest, for which Transformer-based architectures are gaining traction. The self-attention mechanism of Transformers could potentially mitigate catastrophic forgetting, opening the way for more robust medical image segmentation. In this work, we explore how recently-proposed Transformer mechanisms for semantic segmentation behave in sequential learning scenarios, and analyse how best to adapt continual learning strategies for this setting. Our evaluation on hippocampus segmentation shows that Transformer mechanisms mitigate catastrophic forgetting for medical image segmentation compared to purely convolutional architectures, and demonstrates that regularising ViT modules should be done with caution.

1. Introduction

While Continual Learning (CL) research focuses on image classification benchmarks, the importance of preventing forgetting for segmentation models in clinical settings is being noticed by both medical practitioners and regulatory entities [39]. Medical data such as Computer Tomography (CT) scans and Magnetic Resonance Images (MRIs) are particularly susceptible to domain shifts caused by changing acquisition parameters [4]. Yet privacy regulations set strict constraints on data availability which often make training a model sequentially the only option. As expected, this results in *Catastrophic Forgetting* for early tasks [13, 27, 30].

Transformer-based architectures [10, 38] are gaining traction in medical imaging for semantic segmentation [7, 17, 21, 36, 40]. Their ability to combine sequential input data with global knowledge by using self-attention [34, 43]

could potentially mitigate forgetting. Yet the application of Transformers in medical imaging is not focused on continual segmentation and the majority of CL research focuses on (1) convolutional-based architectures and (2) image classification [8, 16]. In this article, we analyse how susceptible Transformer-based segmentation architectures are to continual shifts in the data distribution, and explore whether self-attention mechanisms can be leveraged for continual segmentation.

The majority of recent work on semantic segmentation [1, 26, 28, 37, 42] consists of regularisation-based techniques in combination with convolutional architectures to reduce overfitting [35].

We propose a backbone model that combines state-ofthe-art medical segmentation with Transformers. Additionally, we examine whether popular regularisation-based methods improve the performance when used in combination with Transformers. We augment the state-of-the-art nnU-Net segmentation framework [19] with Transformer techniques that leverage the self-attention mechanism for medical image segmentation. More specifically, we propose an architecture where the Vision Transformer (ViT) is placed in between the encoding and decoding blocks of the U-Net, and the skip connections are used to build the input of the Transformer architecture (ViT U-Net).

| | Segmentation Performance | CL Performance | |
|-----------|-----------------------------|-------------------|--|
| nnU-Net | ++ | | |
| ViT | | + - | |
| ViT U-Net | + + | + + | |

Table 1: Comparison in terms of segmentation and CL performance. '+ +' means that the method is perfectly suited for this category, '+ -' that it is partially suitable '- -' that it is not.

Table 1 compares the state-of-the-art nnU-Net with the ViT and our proposed ViT U-Net in terms of segmentation

and CL performance. It shows that the proposed ViT U-Net is well suited for both categories.

We analyse the impact of freezing and regularising certain components within the context of Catastrophic Forgetting. The results show that freezing and regularising architecture components can have a positive influence on the amount of maintained knowledge if applied with caution.

We perform our evaluation on the problem of hippocampus segmentation, which is of utmost importance for diagnosing and selecting a promising treatment for neuropsychiatric disorders [3] yet very susceptible to distribution shifts [32]. Our results show that self-attention can be leveraged to maintain knowledge in a CL setup, whereas regularising the ViT attention module has a negative effect and should be applied with caution. To summarise, our contributions are two-fold:

- 1. we analyse how the problem of Catastrophic Forgetting manifests in Transformer vs. convolutional-based U-Net architectures, and
- we propose two variations of our ViT U-Net backbone model which leverages self-attention to mitigate Catastrophic Forgetting for medical image segmentation.

2. Methods

We first present our proposed ViT U-Net which is the composition of the well known nnU-Net [19] and ViT [10] into one self-adapting architecture, visualised in Figure 1. Secondly, we outline how we freeze certain architecture components and apply regularisation to assess the role of the ViT for CL. We define our problem setting as a classic *sequential* training setup, where a network is trained sequentially on a sequence of tasks while performing transfer learning. Finally, we briefly outline CL methods that we explore in combination with our proposed architecture.

2.1. ViT U-Net

With the recent publication of the ViT architecture by Dosovitskiy *et al.* [10], new interesting possibilities are opened to leverage the self-attention mechanism known from Transformers for Computer Vision tasks. Within this context, the ViT can potentially be used for medical image segmentation to **increase the amount of maintained knowledge in a CL setup** as the attention mechanism is designed to access and build upon previously seen states [34, 43]. For this purpose, we propose two different versions of the ViT U-Net, a composition of the self-adapting nnU-Net framework and the ViT architecture, where the ViT is placed in between the encoding and decoding blocks of the U-Net (see Figure 1).

The *U-Net* architecture proposed by Ronneberger *et al.* [31] consists of encoding and decoding blocks – *indicated in the shape of a U* – which gives the architecture its



Figure 1: Composition of the nnU-Net and ViT, our proposed ViT U-Net V2. E indicates the encoding and D the decoding blocks of the nnU-Net.

prominent name. Between these sections, skip connections ensure that spatial information is maintained.

The nnU-Net is a dynamic framework that performs relevant pre- and post processing steps and adapts the architecture and training configuration of U-Net models based on characteristics of the training data. The framework is stateof-the-art for several medical segmentation challenges [20, Figure 1].

As mentioned by Chen *et al.* [7], the use of intermediate convolutional feature maps as input of the ViT yields better results as opposed to raw input images. For this purpose, the nnU-Net skip connections are used as the input of the ViT. In addition, we solely focus on skip connections in order to keep the amount of resource allocation as low as possible. Particularly in medical imaging this plays a key role in the applicability of the models due to the large dimensionality of CT and MRI scans.

We present two versions of our proposed architecture where we augment an nnUNet with ViT modules:

- The high-level version (V1) considers only high-level features, *i.e.* the very first skip connection is used to build the ViT dimensions and thus used as input for the ViT.
- The all-level version (V2) on the other hand considers both high- and low-level features by combining the first and last skip connection using transposed convolutional layers. Based on the findings of Hua *et al.* [18], those extracted features are combined through a channel-wise addition as shown in Figure 1 symbolised with the ⊕ symbol.

Recent work by Lee *et al.* [24] introduced Shifted Patch Tokenization (SPT) and Locality Self-Attention (LSA) to enable an efficient training on small size datasets when using ViT. We also explore this strategy in our ablation study.

2.2. Role of ViT for Continual Learning

To analyse the role of the ViT in the ViT U-Net, we conduct two different but similarly composed analyses.

In our first analysis we consider freezing different components of the ViT U-Net. This helps to understand the purpose and benefit of using Transformer-based architectures for medical image segmentation. Models are trained in a sequential setup, where we freeze the specific components after training on the first dataset, as illustrated in Figure 2. With this method, we gain insights on which components are most relevant for learning new patterns and which ones are used to maintain previous knowledge.

Following the same structure, we apply regularisation on specific architecture components during the training process. Among other strategies, we use the well known Elastic Weight Consolidation [22]. With the help of this analysis we want to learn how the regularisation loss affects the process of maintaining knowledge as well as the impact on Transformer-based architectures. To indicate that regularisation is applied on a specific architectural part <arch>, we use the \mathcal{E} : <arch> notation. When it is applied on the whole network, a simple * is used instead.



(a) ViT U-Net architecture with frozen U-Net.



(b) ViT U-Net architecture with frozen ViT.

Figure 2: ViT U-Net V2 architecture with frozen parts represented in grey. E indicates the encoding and D the decoding blocks of the nnU-Net.

2.3. Continual Learning methods

We explore the following continual learning methods. As a replay-based strategy, we use simple *rehearsal* which interleaves 25% of the data randomly selected from every previous task the network has been trained on.

In terms of regularisation-based methods, we explore the popular *Elastic Weight Consolidation (EWC)* [22], which makes use of the Fisher Information Matrix \mathcal{F} to ensure that important parameter values are not heavily modified. We also consider *Riemannian Walk (RWalk)* [6], which combines EWC with Path Integral to measure parameter importance, but calculates \mathcal{F} in an online fashion throughout the training process instead of after training with each task.

Modeling the Background (MiB) [5] is a method specifically developed for incremental learning combining knowledge distillation and a modified Cross Entropy Loss that should tackle the *Background Shift* phenomenon in semantic segmentation. The *Pseudo-labeling and Local Pod* (*PLOP*) [11] falls into the same category as MiB, however it combines a multi-scale spatial distillation loss with pseudolabeling to increase the amount of maintained knowledge. Based on this method, we consider the *Pooled Outputs Distillation (POD)* method as well. Unlike the POD-Net presented by Douillard *et al.* [12], we follow the same principle as PLOP but only apply the spatial POD embedding as distillation loss while pseudo-labeling is omitted. We only focus on the spatial POD embedding to understand the importance of pseudo-labeling for medical imaging.



(a) Replay-based approach.(b) Regularisation (red) as in RWalk and EWC [6, 22].



(c) Knowledge Distillation as in MiB, PLOP and POD [5, 11, 12].

Figure 3: Existing CL methods.

Figure 3 demonstrates the different existing CL methods used in this work. In general, replay-based methods inter-

leave data from previous tasks as shown in Figure 3a. As mentioned previously, regularisation-based methods ensure that certain parameters are not modified or only to a certain degree as done by Kirkpatrick *et al.* [22] or Chaudhry *et al.* [6]. Within this context, Figure 3b shows the general idea behind regularising different layers/parameters during training. Knowledge distillation is commonly used as well, *e.g.* in the MiB [5], PLOP [11] or POD [12] method. Figure 3c visually demonstrated the concept behind it.

3. Experimental Setup

We start by briefly presenting our dataset corpus. Afterwards, we introduce our hardware settings, training and parameter selection process.

3.1. Hippocampus data corpus

Our data corpus contains a total of three hippocampus segmentation datasets, each consisting of T1-weighted MRI scans. The *Harmonized Hippocampal Protocol* dataset [2], which we refer to as *HarP*, contains healthy subjects and patients with Alzheimer's disease. The second dataset, *Dryad* [23], contains 50 cases of healthy patients. Lastly, the hippocampus data from the Medical Decathlon Challenge [33], henceforth referred to as *DecathHip*, comprises cases of healthy and schizophrenia patients. As the datasets are annotated with either one or two labels which are not consistent throughout the corpus, we align label characteristics by joining them into one single *hippocampus* class indicating both the posterior and anterior regions of the hippocampus. Table 2 gives an overview of the different dataset characteristics, among others the mean resolution and voxel spacing.

| Dataset | # Cases (train, val) | Resolution | Spacing | Source |
|-----------|-------------------------|------------|------------------|--------|
| HarP | 270 - (216, 54) | [48 64 64] | [1.00 1.00 1.00] | [2] |
| Dryad | 50 - (40, 10) | [48 64 64] | [1.00 1.00 1.00] | [23] |
| DecathHip | 260 - (208, 52) | [36 50 35] | [1.00 1.00 1.00] | [33] |

Table 2: Characteristics of our hippocampus data corpus; including number of cases, resolution and voxel spacing.

Closely looking at the table shows that all three datasets have the same voxel spacing for all dimensions and nearly the same resolution. DecathHip has on average a smaller resolution than HarP or Dryad. Figure 4 shows samples from our data corpus with corresponding Ground Truth (GT) segmentations.

3.2. Hardware settings

All experiments are conducted on a server system with 256GB DDR4 SDRAM, 2 Intel Xeon Silver 4210 CPUs and 8 NVIDIA Tesla T4 (16 GB) GPUs. To avoid any run time



Figure 4: Samples from the hippocampus data with corresponding GT segmentations – [2, 23, 33].

interference from multiple experiments on the same GPU, only one experiment is run on a single GPU at a time.

3.3. Training Setup

If not otherwise specified, every model is trained in a two-dimensional slice-by-slice model using the previously introduced hippocampus data corpus after being preprocessed by the nnU-Net framework. For further details on the pre-processing, the reader is referred to the original publication of the nnU-Net [19].

A random 80:20 data split is used for training and testing. Models are trained sequentially always following the same order: Harp \rightarrow Drayd \rightarrow DecathHip. All networks are trained for 250 epochs on every dataset from the corpus, *i.e.* 750 epochs in total. Further, the models are only trained once.

Whenever a setting is evaluated, the final model after training on the entire corpus is used to extract predictions. We use the Sørensen–Dice (Dice) coefficient with corresponding standard deviation (σ) between patients as a measure of model performance.

3.4. Hyperparameter Tuning

To select the best hyperparameters, a total of three different settings are tested for every CL method after training 125 epochs per task. To this end, a random 80:20 split from the original training split is used. In case a method expects multiple hyperparameters, some parameters are fixed to a specific – most commonly the default – value. Only those parameters are fixed where the outcome is somewhat predictable and the authors introduced a specific default value in their publication, *e.g.* the α used to calculate the Fisher values in RWalk. Hyperparameters such as λ in EWC have a huge impact on the overall performance and thus are never fixed but rather tuned.

Table 3 summarises the parameter settings considered

| Method | Hippocampus | | | |
|---|--|----------------------------------|---|--|
| | Fixed params Tuned params | | Best params | |
| $\begin{array}{l} EWC_{nnU-Net}, \ \mathcal{E}: \ \ast \\ EWC_{ViT}, \ \mathcal{E}: \ \ast \\ EWC_{ViT}, \ \mathcal{E}: \ nnU-Net \\ EWC_{ViT}, \ \mathcal{E}: \ viT \end{array}$ | _ | $\lambda = \{0.01, 0.20, 0.50\}$ | $\lambda = 0.01$ $\lambda = 0.01$ $\lambda = 0.01$ $\lambda = 0.20$ | |
| RWalk _{nnU-Net} RWalk _{ViT} | $\label{eq:alpha} \begin{split} \alpha &= 0.9, \\ \text{update} &= 10 \end{split}$ | $\lambda = \{0.80, 2.10, 3.30\}$ | $\begin{array}{l} \lambda = 2.10 \\ \lambda = 3.30 \end{array}$ | |
| MiB _{nnU-Net} MiB _{ViT} | $\alpha = 0.9$ | $\lambda = \{0.10, 1.00, 2.50\}$ | $\begin{array}{l} \lambda = 1.00 \\ \lambda = 0.10 \end{array}$ | |
| POD _{nnU-Net} POD _{ViT} | scales $= 3$ | $\lambda = \{0.01, 0.10, 0.20\}$ | $\begin{array}{l} \lambda = 0.01 \\ \lambda = 0.10 \end{array}$ | |
| PLOP _{nnU-Net} PLOP _{ViT} | scales $= 3$ | $\lambda = \{0.01, 0.10, 0.20\}$ | $\begin{aligned} \lambda &= 0.10\\ \lambda &= 0.10 \end{aligned}$ | |

Table 3: Overview of parameter setting considered for each CL method.

for every method, where we regard the setting as best which achieves the highest mean Dice over all intermediate networks evaluated on all tasks. In case two parameter settings result in the same Dice, the one with the lowest σ is used. This measure indirectly takes into account both the amount of forgetting and the Forward Transfer.

4. Experimental Results

We first conduct an ablation study to select the best performing ViT U-Net variation. In a second step, we analyse the relevance of the ViT when trained in a simple sequential setup, while freezing different components of the network. In a third set of experiments we further look into the behaviour and impact of different architecture components when applying the EWC regularisation term. Last but not least, we present the performance results of all considered CL methods in terms of backward and forward transfer.

Our experiments provide us with multiple valuable insights. Firstly, we give a clear understanding of the interplay between the U-Net and ViT within the architecture as well as their influence on each other. We also thoroughly demonstrate the importance and role of the self-attention mechanism in a CL setup. Furthermore, the relevance of the different architecture components is analysed in order to comprehend which are key for maintaining performance on previous tasks and which can be safely tuned for acquiring new knowledge. Within the context of a CL setup, this indicates what network modules can be safely regularised.

4.1. Metrics

Backward Transfer (BWT) indicates how much knowledge is maintained during training on a sequence of tasks and Forward Transfer (FWT) is used to demonstrate how a model state trained on $\{1, \ldots, i-1\}$ performs on unseen, yet to train on dataset *i*.

We define BWT for a specific task \mathfrak{T}_i as

$$\begin{aligned} \text{BWT}\left(\mathfrak{T}_{i}\right) &= \text{Dice}\left(\mathcal{M}_{\left[\mathfrak{T}_{1},\ldots,\mathfrak{T}_{i},\ldots,\mathfrak{T}_{n}\right]},\mathfrak{T}_{i}\right) \\ &- \text{Dice}\left(\mathcal{M}_{\left[\mathfrak{T}_{1},\ldots,\mathfrak{T}_{i}\right]},\mathfrak{T}_{i}\right), \end{aligned} \tag{1}$$

where $\mathcal{M}_{[\mathfrak{T}_1,\ldots,\mathfrak{T}_k]}$ is any network trained on data $\{1,\ldots,k\}$ and $\operatorname{Dice}(\mathcal{M}_{[\mathfrak{T}_1,\ldots,\mathfrak{T}_j]},\mathfrak{T}_i)$ indicates the performance from a network trained on data $\{1,\ldots,j\}$ evaluated on dataset *i*.

FWT for \mathfrak{T}_i is correspondingly defined as

$$FWT(\mathfrak{T}_{i}) = Dice\left(\mathcal{M}_{[\mathfrak{T}_{1},...,\mathfrak{T}_{i-1}]},\mathfrak{T}_{i}\right) - Dice\left(\mathcal{M}_{[\mathfrak{T}_{i}]},\mathfrak{T}_{i}\right), \qquad (2)$$

based on [9, 25]. As Lopez-Paz and Ranzato mention in their work [25], BWT for the first model state as well as FWT for the last model state are not defined.

Throughout this section, the Dice is used in multiple variations, where *Dice mean* indicates the mean Dice of the final network evaluated on all three tasks, *Dice first* indicates the Dice of the final network evaluated on the first task and *Dice last* on the third task.

4.2. Ablation study

Shifted Patch Tokenization (SPT) and Locality Self-Attention (LSA) are ViT modifications to enable an efficient training of the ViT on small size datasets, introduced by Lee *et al.* [24]. Combining our proposed ViT U-Net versions with the ViT modifications as suggested by Lee *et al.* [24], we end up with multiple ViT U-Net variations. To select the best-performing architecture, we first conduct an ablation study by analysing all variations for V1 and V2, *i.e.* without using the proposed SPT or LSA method, using either one of those ViT adaptations or using both proposed modifications. We provide an overview of the results using the final model in Table 4.

| Ablation | | Dice $\uparrow \pm \sigma \downarrow [\%]$ | | |
|--------------|------------|--|------------------|------------------|
| | | HarP | Dryad | DecathHip |
| ViT U-Net V1 | unmodified | 4.08 ± 1.27 | 3.82 ± 0.76 | 89.63 ± 0.54 |
| | SPT | 4.79 ± 1.35 | 2.57 ± 0.70 | 89.71 ± 0.55 |
| | LSA | 3.63 ± 1.24 | 5.49 ± 0.52 | 89.76 ± 0.56 |
| | SPT + LSA | 3.86 ± 1.42 | 6.35 ± 0.38 | 89.68 ± 0.57 |
| ViT U-Net V2 | unmodified | 2.77 ± 0.89 | 3.70 ± 0.51 | 89.66 ± 0.57 |
| | SPT | 4.82 ± 1.83 | 16.19 ± 0.94 | 89.69 ± 0.55 |
| | LSA | 2.98 ± 1.05 | 6.19 ± 0.43 | 89.75 ± 0.53 |
| | SPT + LSA | 3.88 ± 1.09 | 4.94 ± 0.41 | 89.67 ± 0.61 |

Table 4: Performance of the ViT U-Net ablations trained on the hippocampus data corpus without applying any CL mechanism. The *unmodified* rows indicate the unmodified ViT as opposed to the ViT variations (V1 or V2) and their combination (using any combination between SPT and LSA). We highlight the best method bold as well as the highest performance per task. The results show that the ViT U-Net V2 using only SPT and the base ViT architecture with 12 layers and self-attention heads performs best when trained in a sequential setting on the hippocampus data corpus. It achieves the highest performances in most cases and already maintains a significant larger amount of knowledge compared to all other ablations. Therefore, this particular setting is used to conduct all further experiments.

4.3. Impact of ViT in Sequential Training settings

To directly visualise if the self-attention mechanism helps to maintain more knowledge, the ViT U-Net architecture is used in two additional experiments. In the first, the U-Net is frozen after training on the first task – a setting that we name *Frozen U-Net* –. In a second experiment, we freeze the ViT architecture instead after the first task, which we refer to as *Frozen ViT*. Figure 5 shows the performance results of our four variations. For plotting purposes, we normalize the BWT and FWT values from the [-1; 1]range into a positive [0; 2] interval by adding 1.0 to all values.

Comparing the two sequential training setups, a performance increase can easily be seen between the traditional nnU-Net (5a) and ViT U-Net (5b) architecture, as the amount of BWT is higher. Closely analysing the Frozen U-Net experiment (5c) one can deduct that knowledge from previous tasks is maintained the best as the BWT achieves the highest values. However, as the nnU-Net is frozen, the performance on the last task is very low and indicates that the nnU-Net components are crucial to learn new knowledge. The ViT architecture does not perform that good for medical image segmentation on its own. In comparison, the Frozen ViT variation (5d) achieves very high results on the last task, however the performance on the first task is significantly lower indicating that the longer the ViT architecture is frozen, the more knowledge is forgotten. These findings confirm our expectations that the self-attention mechanism of the ViT can indeed be leveraged to maintain knowledge when trained on a sequence of tasks for medical image segmentation.

4.4. Impact of ViT during regularisation

In order to shed even more light on the role of the ViT in the network, we conduct four EWC-related experiments to demonstrate the influence of the architecture components and the effect of regularisation applied on different layers. Closely analysing Figure 6, one can see that the EWC on a plain nnU-Net (6a) already achieves very good results, as the enclosed area is large. However, using the ViT U-



Figure 5: Comparison of the network performances when freezing different parts of the network. No CL mechanism was used during training. The enclosed area is defined by mean Dice over all heads, the first and last Dice value as well as mean BWT and FWT over all tasks in [%]. The larger the area, the better the performance of the method.

Net instead (6b), improves the overall performance. For instance, when FWT is increased, the Dice on the last task is enhanced as well, indicating that the ViT U-Net architecture maintains the amount of knowledge from previous tasks while achieving a higher Dice on the last trained task. Looking at the parameter settings from Table 3, the ViT U-Net uses the same λ setting as the plain nnU-Net which clearly indicates that the EWC regularisation term has a softer impact on the ViT U-Net architecture, *i.e.* decreasing the rigidity, as the learning on new tasks is fairly improved.

Performing regularisation only on the ViT architecture (6d) leads to a much higher FWT value as well as higher performances on the last task as opposed to applying EWC only on the nnU-Net architecture (6c). However, only regularising the nnU-Net components leads to a significantly larger BWT as well as a better performance on the first task. Those findings indicate that the regularisation on the ViT architecture interferes with the self-attention mechanism and thus decreases the amount of maintained knowledge. Furthermore, this analysis shows that the composition of the nnU-Net with the ViT architecture creates some sort of interplay, *i.e.* convolutions and self-attention are complementary as mentioned by Park and Kim in their recent work [29].



Figure 6: Comparison of the network performances when applying EWC on different parts of the network. Enclosed area is defined by mean Dice over all heads, the first and last Dice value as well as mean BWT and FWT over all tasks in [%]. The larger the area, the better the performance of the method.

4.5. Performance of CL methods

All CL methods are deployed on the state-of-the-art nnU-Net to ensure a fair comparison. Table 5 shows an overview over all conducted experiments while reporting BWT and FWT and highlighting the differences in performance.

Analysing Table 5, it is straightforward that the sequential, POD and PLOP training methods perform the worst. In fact, the POD and PLOP methods even perform worse than the sequential setting under certain circumstances. The rehearsal method embodies our upper bound and EWC achieves a performance which is more or less in the same range in terms of BWT and FWT. The Dice on the last task, however, is – *due to the rigidity of the regularisation* – approximately 10% worse compared to the results of applying rehearsal.

Surprisingly, RWalk performs much worse than EWC although it uses the same EWC regularisation as part of the overall regularisation method. This indicates that either the parameter importance scores worsens the regularisation or the online calculation of \mathcal{F} does not represent a good alternative for medical imaging as opposed to the original calculation of \mathcal{F} in EWC. On the other hand, the MiB method

achieves – *excluding EWC and rehearsal* – the best performances when only used with the nnU-Net.

| CL. Method | BWT↑ [%] | | FW | FWT↑ [%] | |
|-------------------------------|------------------|--------------------|-------------|------------------|--|
| | HarP | Dryad | Dryad | DecathHip | |
| Sequential _{nnU-Net} | -83.47 | $-85.09 \\ -74.94$ | -0.22 | -0.08 | |
| Sequential _{ViT} | - 81.03 | | 0.31 | 0.02 | |
| POD _{nnU-Net} | - 82.80 | -86.84 | 0.53 | $0.19 \\ 0.19$ | |
| POD _{ViT} | -85.02 | - 85.16 | 0.93 | | |
| PLOP _{nnU-Net} | -81.47 | -82.64 | 0.37 | $-1.02 \\ -0.92$ | |
| PLOP _{ViT} | -83.29 | -84.87 | 0.23 | | |
| RWalk _{nnU-Net} | -81.34 | - 83.12 | -0.34 | -0.05 | |
| RWalk _{ViT} | -82.22 | -87.78 | 0.43 | 0.01 | |
| MiB _{nnU-Net} | -82.61 | - 65.86 | -0.15 | $-0.89 \\ -0.17$ | |
| MiB _{ViT} | -83.74 | -83.32 | 0.30 | | |
| EWC _{nnU-Net} | $-2.22 \\ -1.76$ | -2.87 | -3.65 | -18.57 | |
| EWC _{ViT} | | -3.95 | -2.39 | -11.41 | |
| Rehearsal _{nnU-Net} | - 0 .23 | -0.54 | 0.14 | -0.02 | |
| Rehearsal _{ViT} | -0.31 | -0.59 | 0.44 | 0.20 | |

Table 5: nnU-Net and ViT U-Net performance results for CL methods in form of BWT and FWT as defined in Equations 1 and 2.

Only in some cases does the ViT U-Net actually decrease the amount of Catastrophic Forgetting. This undermines our previous finding, that regularising the ViT has a negative influence on the self-attention mechanism which might be an indicator for why the BWT is decreased in most of the cases. If regularisation is not introduced on the ViT architecture as shown in the sequential training – Table 4 and Figure 5 –, the results are significantly increased. In most of the cases however, the ViT U-Net achieves improved values for FWT.

Finally, Figure 7 shows an example prediction on a random HarP scan using ViT U-Nets trained only on HarP and Dryad in a sequential and EWC setting to showcase the ViT's attention and correspondingly influence on the network's prediction. For this purpose, the attention map from the last ViT head is used, *i.e.* head 12 as we only use the base ViT architecture in the scope of this manuscript. The Figure also shows the predictions of the networks to allow a comparison with the GT segmentation.

The prediction of the sequential training method (7d) shows that a certain amount of forgetting has already taken place even though the network has not even started on task three yet. Comparing EWC with sequential training, one can certainly confirm that the EWC predictions match the GT (7a) quite well. Focusing on the attention heat maps (7b) and corresponding overlay with the original image (7c), one can easily see that the attention of the sequential network is off, *i.e.* the attention is not even close to the area that should actually be segmented. This showcases that a

significant amount of forgetting has taken place throughout the training on the second task.

Looking at the results from the MiB network, one can hardly find the network's prediction as it is a very small area. Interestingly, the main attention is located at the right top corner of the image, whereas parts of the attention are located over the GT area. This visually shows that the ViT has an influence on the prediction based on its focus within the image. If the focus is on a different area, the segmentation result is not good either, however this does not mean that the attention has a direct influence on the network's prediction as shown in the sequential setup. This shows that the nnU-Net component still has a larger influence on the prediction's shape and location within the image.

Shifting to the network trained in an EWC setting, the attention of the last ViT head focuses adequately on the hippocampus area where the GT is defined. However, the main focus is still a little off and not in the middle of the actual area that should be segmented. Nevertheless, the area of focus certainly captures the GT shown in Figure 7a.



Figure 7: Analysis of the self-attention from the last ViT head using the intermediate networks trained on HarP and Dryad in a sequential (top row), EWC (middle row) and MiB (bottom row) setting using a random HarP image.

5. Limitations

Despite several insightful takeaways, our study has certain limitations. The first is our reduced scope, as we only focus on domain incremental learning and one anatomy for our analysis due to limited computational resources. It would be very interesting to know the differences and results when applied in a class incremental setup instead. As the ViT U-Net is a large network with many parameters due to the self-adapting process, the amount of allocated resources are increased as compared to the traditional nnU-Net. Replacing convolutional layers with convolutional block attention modules for example might have a positive influence on the overall performance while reducing the amount of parameters [41]. Investigating different optimisations to lighten the amount of resource allocation and with it the overall ViT U-Net architecture is an interesting topic on its own.

Another rather intriguing aspect for future work is the importance and role of Layer Norms within the ViT, as recent work has shown that Batch-Norm Layers play an important role in terms of Catastrophic Forgetting [14, 15]. Within this scope, it would be interesting to see the outcome when replacing the ViT Layer Norms with Batch-Norm Layers, how the architecture responds if the ViT Norm Layers are frozen or even regularised using EWC or any suitable knowledge distillation.

6. Conclusion

As recent publications show, Transformers are taking a bigger role in the fields of medical imaging and semantic segmentation. With the ViT U-Net we successfully incorporate the ViT into the well-known and commonly used nnU-Net framework. Our ablation study shows that the primary task of image segmentation performance is not jeopardised by the combination of both architectures.

Additionally, we demonstrate that the proposed model leverages the self-attention of Transformers to maintain knowledge and thus mitigates Catastrophic Forgetting for medical imaging compared to the purely convolutional nnU-Net.

However, regularising the ViT component over time has a negative effect in terms of maintaining knowledge, as regularisation interferes with the self-attention mechanism. Instead, discouraging the amount of change for convolutional layers is preferable for maintaining previous knowledge without compromising plasticity.

Given our promising findings, we hope that other researchers are inspired to conduct further research within the scope of Transformers in CL and medical image segmentation.

7. Reproducibility

All datasets used in this manuscript are publicly available under the corresponding citations. The random data split and trained networks along with instructions on how to run all experiments will be provided upon acceptance. The code including all implementations is publicly available under https://github.com/MECLabTUDA/ Lifelong-nnUNet.

References

- Chaitanya Baweja, Ben Glocker, and Konstantinos Kamnitsas. Towards continual learning in medical imaging. arXiv preprint arXiv:1811.02496, 2018.
- [2] Marina Boccardi, Martina Bocchetta, Félix C Morency, D Louis Collins, Masami Nishikawa, Rossana Ganzola, Michel J Grothe, Dominik Wolf, Alberto Redolfi, Michela Pievani, et al. Training labels for hippocampal segmentation based on the eadc-adni harmonized hippocampal protocol. *Alzheimer's & Dementia*, 11(2):175–183, 2015. 4
- [3] Diedre Carmo, Bruna Silva, Clarissa Yasuda, Letícia Rittner, Roberto Lotufo, Alzheimer's Disease Neuroimaging Initiative, et al. Hippocampus segmentation on epilepsy and alzheimer's disease studies with multiple convolutional neural networks. *Heliyon*, 7(2):e06226, 2021. 2
- [4] Daniel C Castro, Ian Walker, and Ben Glocker. Causality matters in medical imaging. *Nature Communications*, 11(1):1–10, 2020. 1
- [5] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9233–9242, 2020. 3, 4
- [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. 3, 4
- [7] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306, 2021. 1, 2
- [8] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1
- [9] Natalia Díaz-Rodríguez, Vincenzo Lomonaco, David Filliat, and Davide Maltoni. Don't forget, there is more than forgetting: new metrics for continual learning. *arXiv preprint arXiv:1810.13166*, 2018. 5
- [10] 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. *ICLR*, 2021. 1, 2
- [11] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Plop: Learning without forgetting for continual semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4040–4050, 2021. 3, 4
- [12] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *European Con-*

ference on Computer Vision, pages 86–102. Springer, 2020. 3, 4

- [13] Camila Gonzalez, Georgios Sakas, and Anirban Mukhopadhyay. What is wrong with continual learning in medical image segmentation? *arXiv preprint arXiv:2010.11008*, 2020.
 1
- [14] Sharut Gupta, Praveer Singh, Ken Chang, Mehak Aggarwal, Nishanth Arun, Liangqiong Qu, Katharina Hoebel, Jay Patel, Mishka Gidwani, Ashwin Vaswani, et al. The unreasonable effectiveness of batch-norm statistics in addressing catastrophic forgetting across medical institutions. *arXiv preprint arXiv:2011.08096*, 2020. 8
- [15] Sharut Gupta, Praveer Singh, Ken Chang, Liangqiong Qu, Mehak Aggarwal, Nishanth Arun, Ashwin Vaswani, Shruti Raghavan, Vibha Agarwal, Mishka Gidwani, et al. Addressing catastrophic forgetting for medical domain expansion. arXiv preprint arXiv:2103.13511, 2021. 8
- [16] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028– 1040, 2020. 1
- [17] Ali Hatamizadeh, Dong Yang, Holger Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. arXiv preprint arXiv:2103.10504, 2021. 1
- [18] Yuansheng Hua, Lichao Mou, and Xiao Xiang Zhu. Lahnet: A convolutional neural network fusing low-and high-level features for aerial scene classification. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 4728–4731. IEEE, 2018. 2
- [19] Fabian Isensee, Jens Petersen, Andre Klein, David Zimmerer, Paul F Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian Wirkert, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. arXiv preprint arXiv:1809.10486, 2018. 1, 2, 4
- [20] Fabian Isensee, Jens Petersen, Simon AA Kohl, Paul F Jäger, and Klaus H Maier-Hein. nnu-net: Breaking the spell on successful medical image segmentation. arXiv preprint arXiv:1904.08128, 1:1–8, 2019. 2
- [21] Davood Karimi, Serge Vasylechko, and Ali Gholipour. Convolution-free medical image segmentation using transformers. arXiv preprint arXiv:2102.13645, 2021. 1
- [22] 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. 3, 4
- [23] Jessie Kulaga-Yoskovitz, Boris C Bernhardt, Seok-Jun Hong, Tommaso Mansi, Kevin E Liang, Andre JW Van Der Kouwe, Jonathan Smallwood, Andrea Bernasconi, and Neda Bernasconi. Multi-contrast submillimetric 3 tesla hippocampal subfield segmentation protocol and dataset. *Scientific Data*, 2(1):1–9, 2015. 4
- [24] Seung Hoon Lee, Seunghyun Lee, and Byung Cheol Song. Vision transformer for small-size datasets. *arXiv preprint arXiv:2112.13492*, 2021. 2, 5

- [25] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. Advances in neural information processing systems, 30, 2017. 5
- [26] Patrick McClure, Charles Y Zheng, Jakub Kaczmarzyk, John Rogers-Lee, Satra Ghosh, Dylan Nielson, Peter A Bandettini, and Francisco Pereira. Distributed weight consolidation: a brain segmentation case study. *NeurIPS*, 31, 2018. 1
- [27] Marius Memmel, Camila Gonzalez, and Anirban Mukhopadhyay. Adversarial continual learning for multi-domain hippocampal segmentation. In *Domain Adaptation and Representation Transfer, and Affordable Healthcare and AI for Resource Diverse Global Health,* pages 35–45. Springer, 2021. 1
- [28] Sinan Özgün, Anne-Marie Rickmann, Abhijit Guha Roy, and Christian Wachinger. Importance driven continual learning for segmentation across domains. In *International Workshop* on Machine Learning in Medical Imaging, pages 423–433. Springer, 2020. 1
- [29] Namuk Park and Songkuk Kim. How do vision transformers work? arXiv preprint arXiv:2202.06709, 2022. 6
- [30] Oleg S Pianykh, Georg Langs, Marc Dewey, Dieter R Enzmann, Christian J Herold, Stefan O Schoenberg, and James A Brink. Continuous learning ai in radiology: implementation principles and early applications. *Radiology*, 297(1):6–14, 2020. 1
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2
- [32] Antoine Sanner, Camila Gonzalez, and Anirban Mukhopadhyay. How reliable are out-of-distribution generalization methods for medical image segmentation? In DAGM German Conference on Pattern Recognition, pages 604–617. Springer, 2021. 2
- [33] Amber L. Simpson, Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram van Ginneken, Annette Kopp-Schneider, Bennett A. Landman, Geert Litjens, Bjoern H. Menze, Olaf Ronneberger, Ronald M. Summers, Patrick Bilic, Patrick Ferdinand Christ, Richard K. G. Do, Marc Gollub, Jennifer Golia-Pernicka, Stephan Heckers, William R. Jarnagin, Maureen McHugo, Sandy Napel, Eugene Vorontsov, Lena Maier-Hein, and M. Jorge Cardoso. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *CoRR*, abs/1902.09063, 2019. 4
- [34] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7262–7272, 2021. 1, 2
- [35] Jonas Teuwen and Nikita Moriakov. Chapter 20 convolutional neural networks. In S. Kevin Zhou, Daniel Rueckert, and Gabor Fichtinger, editors, *Handbook of Medical Image Computing and Computer Assisted Intervention*, The Elsevier and MICCAI Society Book Series, pages 481–501. Academic Press, 2020. 1
- [36] Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical transformer: Gated axial-

attention for medical image segmentation. *arXiv preprint arXiv:2102.10662*, 2021. 1

- [37] Karin van Garderen, Sebastian van der Voort, Fatih Incekara, Marion Smits, and Stefan Klein. Towards continuous learning for glioma segmentation with elastic weight consolidation. In *International Conference on Medical Imaging with Deep Learning–Extended Abstract Track*, 2019. 1
- [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 *Advances in neural information processing systems*, pages 5998–6008, 2017. 1
- [39] Kerstin N Vokinger, Stefan Feuerriegel, and Aaron S Kesselheim. Continual learning in medical devices: Fda's action plan and beyond. *The Lancet Digital Health*, 3(6):e337– e338, 2021. 1
- [40] Wenxuan Wang, Chen Chen, Meng Ding, Hong Yu, Sen Zha, and Jiangyun Li. Transbts: Multimodal brain tumor segmentation using transformer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 109–119. Springer, 2021. 1
- [41] Yimin Yang and Siamak Mehrkanoon. Aa-transunet: Attention augmented transunet for nowcasting tasks. arXiv preprint arXiv:2202.04996, 2022. 8
- [42] Jingyang Zhang, Ran Gu, Guotai Wang, and Lixu Gu. Comprehensive importance-based selective regularization for continual segmentation across multiple sites. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 389–399. Springer, 2021. 1
- [43] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 1, 2