

Continual atlas-based segmentation of prostate MRI

Amin Ranem¹ Camila González² Daniel Pinto dos Santos³ Andreas M. Bucher⁴
Ahmed E. Othman⁵ Anirban Mukhopadhyay¹

¹ Technical University of Darmstadt, ² Stanford University, ³ University of Cologne,
⁴ University of Frankfurt, ⁵ University Medical Center Mainz
amin.ranem@gris.tu-darmstadt.de

Abstract

Continual learning (CL) methods designed for natural image classification often fail to reach basic quality standards for medical image segmentation. Atlas-based segmentation, a well-established approach in medical imaging, incorporates domain knowledge on the region of interest, leading to semantically coherent predictions. This is especially promising for CL, as it allows us to leverage structural information and strike an optimal balance between model rigidity and plasticity over time. When combined with privacy-preserving prototypes, this process offers the advantages of rehearsal-based CL without compromising patient privacy. We propose **Atlas Replay**, an atlas-based segmentation approach that uses prototypes to generate high-quality segmentation masks through image registration that maintain consistency even as the training distribution changes. We explore how our proposed method performs compared to state-of-the-art CL methods in terms of knowledge transferability across seven publicly available prostate segmentation datasets. Prostate segmentation plays a vital role in diagnosing prostate cancer, however, it poses challenges due to substantial anatomical variations, benign structural differences in older age groups, and fluctuating acquisition parameters. Our results show that **Atlas Replay** is both robust and generalizes well to yet-unseen domains while being able to maintain knowledge, unlike end-to-end segmentation methods. Our code base is available under <https://github.com/MECLabTUDA/Atlas-Replay>.

1. Introduction

Continual learning (CL) plays a crucial role in safety-critical applications of Deep Learning (DL), particularly in

healthcare. In such domains, models must continually adapt to data drift over time while maintaining high performance on older data, even in cases where direct access to part of the data is restricted for privacy reasons. The objective of CL is to train a model that demonstrates high performance on sequentially arriving datasets, despite the constrained time-frame during which the datasets are accessible. Achieving this objective is challenging, as approaches tend to fall into one of the extremes: either suffering from *catastrophic forgetting* [19] by being too plastic, or *unable to learn* new tasks by being too rigid. Additionally, some methods exhibit linear growth in training time and resource requirements as the number of training tasks increases. CL models are further susceptible to domain shifts over time [15], increasing the amount of catastrophic forgetting even further. While various strategies have been proposed for CL, they often fail to perform well with medical data [9], resulting in segmentations that do not meet basic semantic standards. Therefore, striking a delicate balance between preserving previous knowledge and maintaining the necessary plasticity to learn new tasks is a key requirement when training in a continual fashion [8, 16].

As we have domain knowledge on the geometry of the organ to be segmented, we go back to the roots of medical image segmentation and contextualize atlas-based segmentation [26], as an alternative to end-to-end CL segmentation pipelines. In the context of CL, the access to structural information is key for achieving accurate and semantically coherent predictions. In atlas-based segmentation, registration is used to adapt the best-fitting labeled mask from an atlas (i.e. a pre-defined set of reference images) to a patient scan. The natural reliance of atlas-based approaches on geometrical aspects makes it ideally applicable for CL setups as the structural information can be leveraged to extract domain knowledge.

CL for end-to-end models can be generally classified

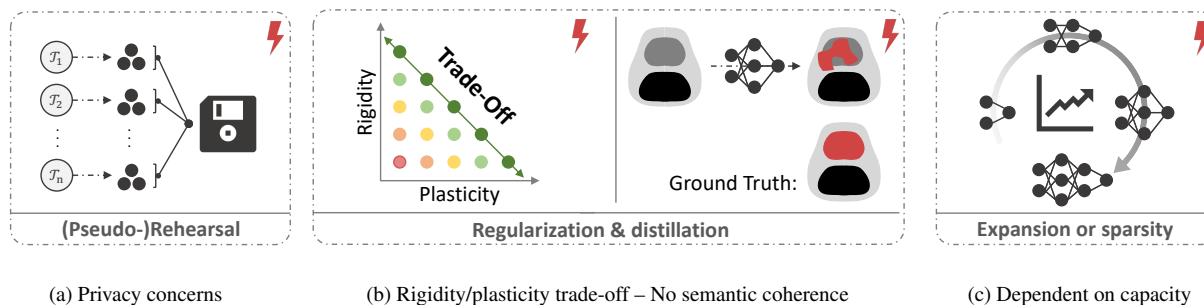


Figure 1. Demonstration of different problems CL methods face.

into one of the following three categories: (1) (pseudo-)rehearsal, (2) regularization/distillation, and (3) expansion. Figure 1 demonstrates the practical downsides of each strategy.

The most successful CL methods follow some type of *replay* or *rehearsal*, which involves storing samples to interleave them during later training [33, 37]. Long-term storage of patient scans, however, violates data protection regulations [12]. Distillation methods such as *PLOP* [11] do not store data directly but rather distill knowledge from previously trained models. Yet recent work [14] shows that the additional computational burden associated with the pseudo-data generation hinders their use with high-dimensional medical images. Regularization methods [6, 19] on the other hand work by penalizing severe shifts from the previously-learned parameter space. These approaches have lower resource requirements, but disappointing performance across all tasks and merely allow a trade-off between rigidity and plasticity. Models with high plasticity increase the ability to learn new information, whereas rigid models maintain more knowledge from previous tasks and prevent catastrophic forgetting. Additionally, models trained end-to-end in a sequential manner with regularization or distillation approaches tend to generate predictions with no semantic coherence. Expansion techniques maintain stable performance across all tasks but grow the model size with the number of tasks [15, 17]. Other works [29, 31, 32] propose architectural modifications that are useful in certain settings but imply a high computational overhead. Instead of relying on end-to-end CL, we propose a modular atlas-based method for continual segmentation. *Atlas Replay* leverages structural information to maintain knowledge over time with the benefits of rehearsal while preserving patient privacy and avoiding model growing.

Atlas Replay generates *prototypes* built from a set of patient scans and the *VoxelMorph* registration framework [4] to perform registration between a specific patient scan and a prototype. A prototype is a combination of multiple patient scans – and corresponding segmentation masks – that disallows the direct identification of a subject. Prototypes are

registered to new images to generate segmentation masks for new patients during deployment.

We introduce an approach to build an atlas of prototypes, *irrespective of the anatomy*; and *propose a CL method* to perform atlas-based segmentation which maintains model plasticity while preserving previous knowledge and outperforms state-of-the-art (SOTA) end-to-end continual segmentation approaches. We can consider using a stored atlas to generate segmentation masks as a form of pseudo-rehearsal that maintains relevant information from previous examples without storing actual patient images.

The contributions of this work are three-fold. Our proposed approach:

- successfully builds privacy-preserved prototypes and therefore being more protective towards patient privacy compared to traditional replay-based methods,
- leverages prototypes and structural information to maintain knowledge over time,
- benefits from a rehearsal based approach and *VoxelMorph* to achieve stable performances for continual image segmentation.

To validate our method, we investigate the problem of prostate segmentation in T2-weighted Magnetic Resonance Images (MRI), which is an important step in the diagnosis and treatment of prostate cancer [35]. Prostates have relatively static shapes for which domain knowledge can be leveraged over time. Variations in imaging protocol, such as the diminishing use of endorectal coil over time [20], lead to domain shift. Such shifts clearly state the importance of CL from a clinical perspective.

With the introduction of *Atlas Replay*, we pave the path for integrating atlas-based methods into the realm of CL. We demonstrate that established conventional approaches like atlas-based segmentation, which have fallen into a certain neglect due to the current DL era, can be effectively utilized in dynamic clinical setups. This success can be attributed to the use of structural information, which showcases highly favorable outcomes in CL scenarios.

2. Methodology

Traditional end-to-end segmentation methods in CL typically struggle to strike a suitable balance between rigidity and plasticity, which presents a significant drawback in this field. Such models often face a trade-off where they either fail to acquire new knowledge by retaining information from previously seen data, or overly prioritize recent cases, leading to catastrophic forgetting and predictions lacking semantic coherence. Recognizing the importance of structural information for achieving a proper trade-off between rigidity and plasticity in CL setups, we turn to atlas-based segmentation as a solution. Most atlas-based segmentation methods are based on traditional (non-DL) techniques [1, 26]. We combine the advantages of DL with the ability to leverage domain knowledge of atlas-based segmentation for CL. An *atlas* Λ is a manually labeled set of patient scans [34]. Atlas-based segmentation is the *modular process* of registering an image from the atlas to a new scan to directly generate an accurate segmentation by transforming the respective mask in the same manner [26], thus leveraging structural information. This approach has been successfully applied to prostate cancer [3], heart regions [13], brain tissue [1] and aortic tissue [18] for MRI scans.

Fundamentals We start by introducing some key terminology: $\Omega \subset \mathbb{R}^3$ defines a 3D spatial domain. $\mathcal{T}_p \subset \Omega_{\mathcal{T}}$ is referred to as dataset p and consists of $m\{(f_m^i, f_m^s)\}$ pairs, where f^i is a patient scan and f^s the corresponding segmentation mask. Stage x in a continual setup defines the process of training the model on dataset x after it has been trained on all previous $\{1, \dots, x-1\}$ tasks. $\Omega_{\mathcal{T}}$ is a set of datasets and $\Omega_{\mathcal{P}}$, a set of prototypes. A prototype \mathcal{P}_k is a tuple $\mathcal{P}_k = (\mathcal{P}_k^i, \mathcal{P}_k^s)$, where \mathcal{P}_k^i is the scan and \mathcal{P}_k^s the corresponding segmentation mask. We define our set of prototypes $\Omega_{\mathcal{P}}$ as *privacy-preserving representations*.

VoxelMorph [4] is a popular framework for medical image registration. The underlying architecture is a simple U-Net, to which additional convolutional layers are added to generate a deformation field ϕ . The network is trained by penalizing the difference between the warped moving image using ϕ and the target image. By using ϕ to alter the prototype segmentation mask, one can utilize the framework for atlas-based segmentation given a specific atlas Λ .

Prototype building We create four distinct prototypes based on the coil type used during acquisition. For each prototype \mathcal{P}_k , we extract r random samples from the associated dataset \mathcal{T}_p , a subset of $\Omega_{\mathcal{T}}$. In this study, we set $r = 7$, corresponding to the size of the smallest training set. To evaluate the performance between prototypes, we allocate three datasets from our data corpus to validate the inter-prototype performance in Section 4.4.

First, r random images are selected from the training dataset \mathcal{T}_p . The first image I represents the initial prototype \mathcal{P}_k . For each following image I^* , we refine \mathcal{P}_k through rigid alignment using *SimpleITK* [27, 38]. We then compute the average of \mathcal{P}_k and I^* to update the prototype, including the corresponding segmentation masks. The final \mathcal{P}_k is the prototype for dataset \mathcal{T}_p . Taking the average over multiple scans hinders the recovery of patient-sensitive information during storing or interleaving, as illustrated in Figure 2.

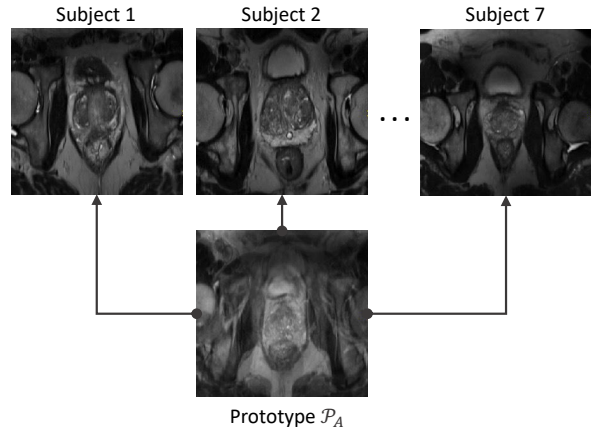


Figure 2. Illustration of prototype \mathcal{P}_A .

Figure 2 shows the difference between the final prototype \mathcal{P}_D and three out of seven subject scans that were used to build the prototype. The intensity distribution for every dataset is shown in Figure 3. We assess the effectiveness of our prototypes in maintaining privacy by performing a user study among senior radiologists, with more than 10 years of experience, in Section 4.3.

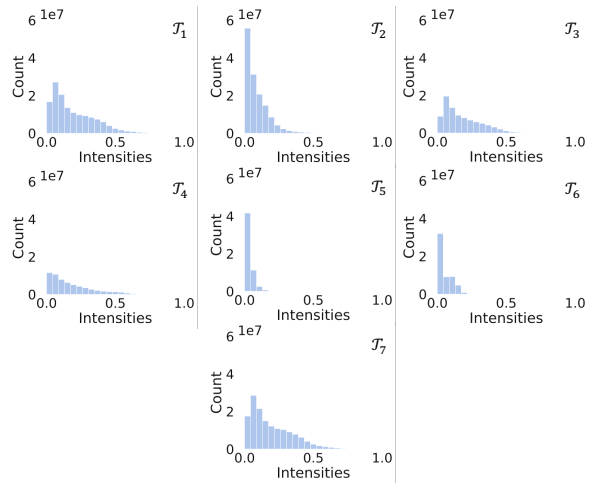


Figure 3. Intensity histograms for every dataset.

Prototype registration A patient scan f^i represents the fixed image that is registered to the best-fitting prototype \mathcal{P}_k – the moving image – using *VoxelMorph*. The deformation field ϕ from the network is then used to warp the prototypes’ segmentation mask: $\mathcal{P}_k^s \circ \phi$. As illustrated in Figure 4, our method only stores *prototypes* that preserve key information and therefore preserves patient privacy better than storing actual subjects.

Training continually using Atlas Replay Algorithm 1 demonstrates a continual training setup in simplified pseudo-code for our proposed Atlas Replay approach using the VoxelMorph (VxM) Framework.

Algorithm 1: Training using Atlas Replay

Input: Datasets to train on $\{\mathcal{T}_p\}_{p \leq |\Omega_{\mathcal{T}}|}$
Output: Trained model weights θ
 // Initialize \mathcal{M}_θ
 1 $\theta \leftarrow \text{initializeModel}()$
 // Build $k = 4$ prototypes
 2 $\{\mathcal{P}_k\} \leftarrow \text{buildPrototypes}(\{\mathcal{T}_p\}_{p \leq |\Omega_{\mathcal{T}}|}, k = 4)$
 // Select prototype for \mathcal{T}_1
 3 $\mathcal{P} \in \{\mathcal{P}_k\}$
 // Train with VxM
 4 $\theta \leftarrow \text{VxM}(\theta, \mathcal{T}_1, \mathcal{P})$
 5 **for** $i \leftarrow 2$ **to** $|\Omega_{\mathcal{T}}|$ **do**
 // Select prototype for \mathcal{T}_i
 6 $\mathcal{P} \in \{\mathcal{P}_k\}$
 // Train with VxM
 7 $\theta \leftarrow \text{VxM}(\theta, \mathcal{T}_i, \mathcal{P})$
 8 **end**

3. Experimental Setup

In this section, we briefly describe our corpus of publicly available datasets and report essential aspects of our experimental setup.

Data corpus Our prostate data corpus $\Omega_{\mathcal{T}}$ consists of seven publicly available datasets [2, 5, 21–24]. Table 1 shows the number of cases in each dataset and the random 80:20 train/validation split. The splits along with our code base is accessible under <https://github.com/MECLabTUDA/Atlas-Replay>.

Prototypes To give the reader a proper understanding of our set of prototypes $\Omega_{\mathcal{P}}$, these are illustrated in Figure 5.

Training setup The VoxelMorph framework [4] with default optimizer and learning rate is used for all registration experiments. To ensure a fair comparison between the

Dataset	Task ID	# Cases (train, val)	Vendor	Source
RUNMC	\mathcal{T}_1	30 – (24, 6)	Siemens	[5]
BMC	\mathcal{T}_2	30 – (24, 6)	Philips	[21]
HCRUDB	\mathcal{T}_3	19 – (15, 4)	Siemens	[21]
UCL	\mathcal{T}_4	13 – (10, 3)	Siemens	[22]
BIDMC	\mathcal{T}_5	12 – (9, 3)	GE	[22]
HK	\mathcal{T}_6	12 – (9, 3)	Siemens	
DecathProst	\mathcal{T}_7	32 – (25, 7)	Unknown	[2]

Table 1. Characteristics of our prostate data corpus; including the vendor of the acquisition device.

network’s performance, the same underlying U-Net architecture is used for end-to-end segmentation. Registration models are trained for 250 epochs using a composition of the Normalized Cross Correlation (\mathcal{L}_{NCC}), Cross-Entropy (CE, \mathcal{L}_{CE}), and a gradient smoothing loss for the deformation field (\mathcal{L}_{smooth}). The total registration loss is calculated in the following way:

$$\mathcal{L}_{reg} = \mathcal{L}_{NCC} + 2 \cdot \mathcal{L}_{CE} + \mathcal{L}_{smooth} \quad (1)$$

Based on our ablation results from Section 4.5, we use the double-weighted CE loss setup for all our VoxelMorph experiments. Segmentation networks are trained for 250 epochs with \mathcal{L}_{CE} only. All experiments were carried out on a GeForce RTX 3090 GPU (24 GB).

Metrics We report the mean Dice and standard deviation across the test images from all tasks as well as average backward (BWT) and forward (FWT) transferability [10, 25]. BWT indicates the amount of maintained knowledge on test samples $(f_m^i, f_m^s) \subset \mathcal{T}_j$ during training on different stages $\{\mathcal{T}_p\}_{p \leq |\Omega_{\mathcal{T}}|}; j < p$ over time. FWT on the other hand measures the impact of the current training stage $\{\mathcal{T}_p\}_{p \leq |\Omega_{\mathcal{T}}|}$ on test data $(f_m^i, f_m^s) \subset \mathcal{T}_j; j > p$ from an untrained stage.

Let \mathcal{T}_p be a specific task:

$$\begin{aligned} \text{BWT}(\mathcal{T}_p) = & \text{Dice}(\mathcal{M}_{[\mathcal{T}_1, \dots, \mathcal{T}_p, \dots, \mathcal{T}_n]}, \mathcal{T}_p) \\ & - \text{Dice}(\mathcal{M}_{[\mathcal{T}_1, \dots, \mathcal{T}_p]}, \mathcal{T}_p), \end{aligned} \quad (2)$$

where $\mathcal{M}_{[\mathcal{T}_1, \dots, \mathcal{T}_p]}$ is a network trained on stages $\{1, \dots, p\} \leq |\Omega_{\mathcal{T}}|$ and $\text{Dice}(\mathcal{M}_{[\mathcal{T}_1, \dots, \mathcal{T}_j]}, \mathcal{T}_p)$ indicates the Sørensen–Dice coefficient from a network trained on stages $\{1, \dots, j\}$ evaluated on dataset p . FWT is defined as:

$$\begin{aligned} \text{FWT}(\mathcal{T}_p) = & \text{Dice}(\mathcal{M}_{[\mathcal{T}_1, \dots, \mathcal{T}_{p-1}]}, \mathcal{T}_p) \\ & - \text{Dice}(\mathcal{M}_{[\mathcal{T}_p]}, \mathcal{T}_p). \end{aligned} \quad (3)$$

FWT for the last model state as well as BWT for the first model state are not defined. Models with high plasticity are able to learn new knowledge and achieve higher FWT,

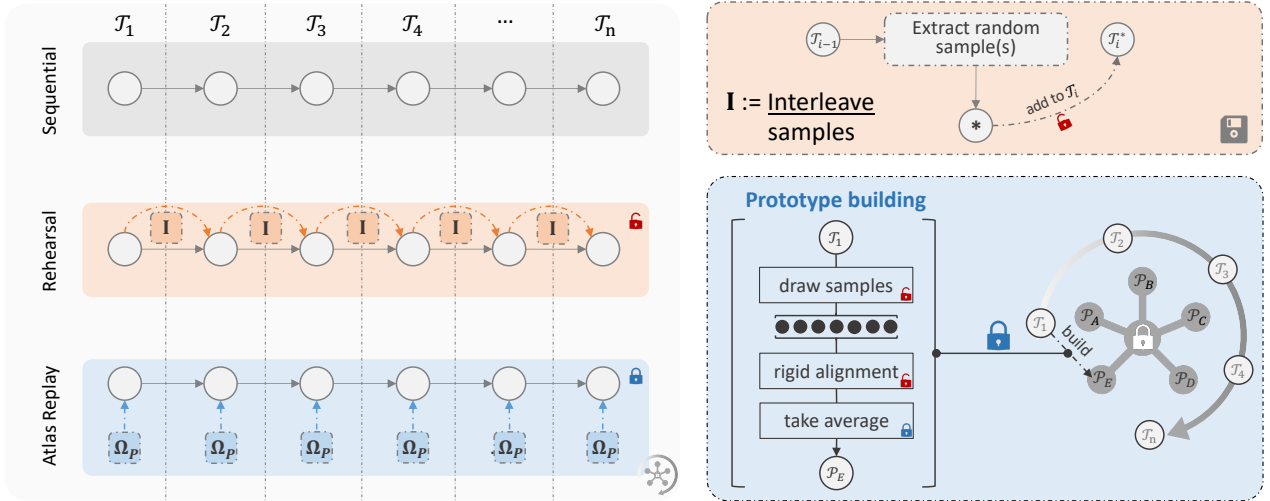


Figure 4. Comparison of sequential training on n tasks against traditional rehearsal and *Atlas Replay* in terms of privacy preservation; Ω_P represents the stored prototypes.




Figure 5. Prototypes and their overlapping segmentation mask. For every prototype, a random slice was selected.

whereas models that maintain most knowledge from previous tasks obtain a higher BWT.

To validate our user study, we report sensitivity, specificity, precision (positive predicted value), and the Matthews Correlation Coefficient (MCC) [7, 28, 36]. We calculate the MCC to provide an overall assessment of the classification performance with respect to both, true negative and true positive rates.

Baselines We compare Atlas Replay to end-to-end sequential training, the upper bound of simple replay training, which requires the storage of actual patient scans, and four popular CL methods: Elastic Weight Consolidation (EWC) [19], Riemannian Walk (RWalk) [6], Incremental Learning Techniques (ILT) [30] using distillation on the output (KD), intermediate (MSE), or both (KD, MSE) layers and Bias Correction (BiC) [37]. A hyperparameter search is conducted for EWC and RWalk, and the best settings are used of $\lambda = 2.2$ (EWC), $\alpha = 0.9$, and $\lambda = 1.7$ (RWalk). For ILT, we used the default distillation parameter $\lambda_D = 1$. Since BiC is a rehearsal-based method, from each task, seven samples were interleaved.

4. Results

We present a comprehensive evaluation of various aspects of our work. We start by analyzing the performance of Atlas Replay compared to sequential training, EWC, RWalk, BiC, and rehearsal training. We then proceed with a qualitative temporal evaluation in Section 4.2. Section 4.3 explores the effectiveness of privacy-preserving prototypes through a user study conducted with senior radiologists and computer scientists. We assess the inter-prototype performance during inference in 4.4. Furthermore, we conduct a loss ablation study in 4.5 and compare the U-Net’s end-to-end segmentation with our atlas-based segmentation approach in 4.6. The results are based on the hyperparameter search we conducted in Section 4.7. These evaluations provide multiple insights into the performance, robustness, and effectiveness of the proposed approach and prototype generation strategy.

4.1. Continual learning performance

In this section, we compare Atlas Replay to training a U-Net model sequentially alongside the continual learning methods EWC, RWalk, ILT and BiC. Further, we compare against the upper bound of rehearsal training (storing seven samples from each task).

Figure 6 shows the mean Dice scores of the final networks evaluated across all seven datasets. EWC (6b) achieves more or less the same Dice performance across the tasks as the simple sequential setup (6a) showing slight shifts in BWT and FWT as seen in Table 2. RWalk, on the other hand, achieves a better BWT, meaning it maintains more knowledge. However, this comes at the cost of a lower mean Dice and FWT (6c). This indicates that even

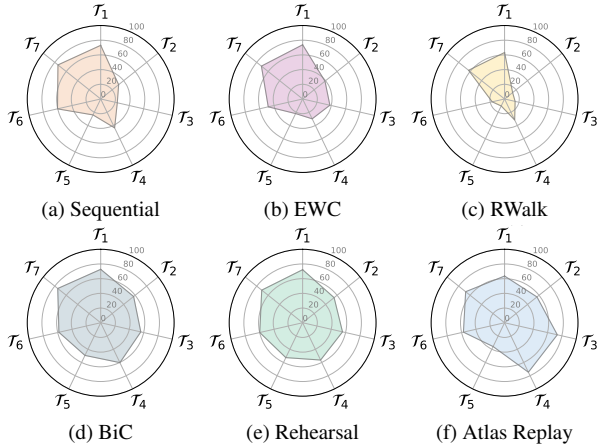


Figure 6. CL performance as Dice; the larger the area the better.

after tuning the hyperparameters, end-to-end state-of-the-art CL methods under-perform for medical image segmentation. BiC on the other hand preserves previous knowledge the best (BWT), but at the cost of plasticity loss (FWT), Table 2. Atlas Replay (6f) is the only method that maintains good performance over all seen tasks. BWT is higher than for the rehearsal upper bound, as indicated in 6d to 6f, without compromising the plasticity loss (FWT) as observed by BiC, Table 2. The different ILT versions on the other hand achieve slightly better performances than EWC.

Method	BWT \uparrow [%]	FWT \uparrow [%]
<i>Sequential</i>	-26.60 ± 17.88	-25.10 ± 17.44
<i>EWC</i>	-27.48 ± 15.00	-29.65 ± 18.32
<i>RWalk</i>	-22.79 ± 18.94	-51.48 ± 13.23
<i>ILT_{KD}</i>	-27.86 ± 11.85	-26.79 ± 21.32
<i>ILT_{MSE}</i>	-25.56 ± 14.90	-34.26 ± 18.94
<i>ILT_{KD, MSE}</i>	-16.59 ± 12.77	-34.58 ± 17.73
<i>BiC</i>	-0.56 ± 7.42	-29.01 ± 7.42
<i>Rehearsal</i>	-6.60 ± 7.02	-24.23 ± 15.58
<i>Atlas Replay</i>	-8.13 ± 7.68	-18.26 ± 15.22

Table 2. CL performance indicated by BWT and FWT along with standard deviation.

With the default U-Net specifications of VoxelMorph, the average mean segmentation Dice is about 61%. However, a significant improvement in our proposed method can be clearly observed.

4.2. Qualitative temporal evaluation

To analyze the robustness of our proposed method qualitatively, we visualize segmentation masks in Figure 7.

Atlas Replay consistently produces coherent segmentation masks, irrespective of the training stage. Sequential training and EWC on the other hand produce low-quality

segmentations until the network is trained on the particular stage 3 ($\{\mathcal{T}_p\}_{p \leq 3}$). In particular, the low performance on later tasks shows the effect of catastrophic forgetting, where the network adapts too strongly to the later training data. Atlas Replay is neither too rigid nor plastic, as it outputs robust predictions for data from both early and later training stages.

The image shows the impact of selecting the correct prototype during inference for the effectiveness of the registration method. Using a prototype that was built from a non-coil dataset (\mathcal{P}_A) to perform registration with a coil-acquired sample (\mathcal{T}_2) is expected to have lower performance as the geometric shape of the prostate changes significantly depending on the type of coil. Since the coil type is a central aspect of the examination that can be easily recognized, it is to be expected that the user applying the algorithm knows the type of coil that was used during acquisition and can select the correct prototype accordingly. If this were not known, a simple solution would be to register the desired sample across all prototypes and then use the prototype leading to the best registration performance in terms of NCC or MSE.

4.3. Effectiveness of privacy preserving prototypes

In order to evaluate the effectiveness of our privacy-preserving prototypes, we conducted a user study involving two senior radiologists with more than 10 years of experience (RAD) and two computer scientists (CS). The study aimed to assess the participants' ability to correctly select the used subject from a set of three patient scans, with only one scan being the correct match for the shown prototype. This selection process was performed five times for each prototype.

To measure the performance of the participants, we report four key metrics: sensitivity, specificity, precision (positive predicted value), and Matthews Correlation Coefficient (MCC). Sensitivity refers to the ability to correctly identify the true positive cases, i.e., correctly selecting the matching patient scan. Specificity measures the ability to correctly identify the true negative cases, meaning to correctly exclude the non-matching patient scans. Precision represents the proportion of correctly selected matching patient scans out of the total selected matching scans. MCC on the other hand provides an overall assessment of the classification performance, taking into account both true positive and true negative rates. Table 3 shows the average results of our user study for the radiologists and computer scientists including the standard deviation and random chance.

Table 3 clearly shows that correctly selecting the used subject from a set of three samples, along with the prototype, proved to be very challenging for senior radiologists as well as computer scientists. The results illustrate that possessing technical knowledge about the process of proto-




Figure 7. Temporal analysis for sequential, EWC, and Atlas Replay using Case 3, Slice 19 from \mathcal{T}_3 .

Group	Sensitivity \uparrow [%]	Specificity \uparrow [%]	Precision \uparrow [%]	MCC \uparrow [%]
random chance	–	–	33.33 ($\frac{1}{3}$)	–
RAD	47.50 \pm 7.50	74.00 \pm 4.00	47.50 \pm 7.50	21.50 \pm 14.50
CS	45.00 \pm 10.00	73.00 \pm 5.00	45.00 \pm 10.00	18.00 \pm 15.00

Table 3. Results of our user study evaluation on privacy-preserving prototypes compared to random chance. The evaluation included participation from both radiologists and computer scientists.

type building has minimal impact on selecting the correct samples as the classification rates are very similar to the ones from the radiologists. Such a difficulty in correctly identifying the matching patient scan demonstrates the effectiveness of our privacy-preserving prototype building approach. The average precision achieved by the radiologists was 47.50%, indicating the complexity of the task and the privacy preservation capabilities of our prototypes.

4.4. Inter-prototype performance




Figure 8. Combination of scatter and box plot showing performance distribution and change in Dice for every task based on the used prototype during inference with the final model trained sequentially on all datasets.

To assess the inter-prototype performance, we trained se-

quentially on all seven tasks using Atlas Replay and utilized the final model to analyze the influence of prototypes during inference (Figure 8). Specifically, for each sample from the test sets we performed registration using all four prototypes.

Figure 8 illustrates the distribution of Dice for every task based on the utilized prototype during inference with the sequentially trained model. This visualization effectively shows how the choice of prototype during the registration process influences the performance observed in inference across different tasks.

4.5. Loss ablation study

To determine the optimal setup for atlas-based segmentation, we conduct an ablation study where we modify the traditional registration loss proposed in the VoxelMorph paper [4]. This study allows us to identify the most effective VoxelMorph configuration. The networks are trained on the joint dataset $\bigcup_{p=1}^{|\Omega_{\mathcal{T}}|} \mathcal{T}_p$ and the evaluation is performed across all seven datasets $\{\mathcal{T}_p\}_{p \leq |\Omega_{\mathcal{T}}|}$.

Table 4 shows that weighting the Cross-Entropy loss twice leads to the best segmentation performance across all tasks. A network trained with no segmentation loss term, i.e. only for registration, achieves a limited contribution in terms of segmentation performance. Given this insight, we use the double weighted Cross-Entropy loss setup for all our VoxelMorph related experiments as shown in Equation 1.

4.6. End-to-end segmentation vs. atlas-based segmentation

To better assess the generalizability of the models, we train U-Net and Atlas Replay networks for every dataset $\mathcal{T}_p \subset \Omega_{\mathcal{T}}$ and validated them across all datasets.

Figure 9 illustrates the difference between an end-to-end segmentation approach (U-Net) and Atlas Replay in terms of generalizability and segmentation performance. Using

Ablation ($\mathcal{L}_{reg}, \mathbf{1}$)	Dice $\uparrow \pm \sigma \downarrow$ [%]						
	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7
$\mathcal{L}_{NCC} + 0 \cdot \mathcal{L}_{CE} + \mathcal{L}_{smooth}$	37.03 \pm 17.05	62.63 \pm 1.97	66.56 \pm 6.06	69.78 \pm 2.08	58.48 \pm 6.94	63.73 \pm 8.74	38.84 \pm 11.86
$\mathcal{L}_{NCC} + 1 \cdot \mathcal{L}_{CE} + \mathcal{L}_{smooth}$	56.01 \pm 15.05	68.95 \pm 3.46	70.83 \pm 3.36	73.80 \pm 1.27	60.18 \pm 8.94	63.28 \pm 11.71	57.89 \pm 10.83
$\mathcal{L}_{NCC} + 2 \cdot \mathcal{L}_{CE} + \mathcal{L}_{smooth}$	70.55 \pm 6.32	73.29 \pm 3.65	71.39 \pm 1.94	79.84 \pm 3.49	60.07 \pm 4.88	65.36 \pm 8.89	69.47 \pm 6.70
\mathcal{L}_{CE} (U-Net joint)	79.62 \pm 2.50	77.01 \pm 11.10	78.95 \pm 6.61	79.72 \pm 4.13	83.49 \pm 1.97	82.45 \pm 1.98	80.34 \pm 1.94

Table 4. VoxelMorph ablations trained on joint prostate data with different loss variations (top) and U-Net joint results (bottom).




Figure 9. Scatter plot showing the performance distribution in Dice when trained on one dataset and evaluated across all seven datasets from the prostate corpus $\Omega_{\mathcal{T}}$; trend line is based on mean Dice performance across all evaluated samples.

the proposed registration-based approach, we obtain robust results regardless of what dataset is used for training. The U-Net performance on the other hand is clearly dependent on the training dataset, as indicated by the strong distribution shifts in the figure, i.e. lack of generalizability. \mathcal{T}_5 , for instance, shows a significant performance deterioration. Besides an increase in generalizability, the performance of the registration-based method is also higher, as shown by the mean Dice trend line, which is consistently above that of the U-Net. These results show the increased versatility and robustness of Atlas Replay even before continual learning is performed.

4.7. Hyperparameter search

For every used end-to-end CL method, a hyperparameter search was performed using four different settings. The setup with the highest mean performance in terms of Dice, BWT, and FWT together was selected for our experiments. Table 5 shows the results for each hyperparameter setting. The highest values for a method and the best parameter settings are marked in bold.

The results clearly show the addressed trade-off problem from Figure 1b. Depending on how the hyperparameter(s) are set, the trade-off between plasticity and knowledge preservation varies. The network either performs well in terms of maintaining knowledge over time – increased λ for EWC – or the results are very similar no matter how the hyperparameter is set – λ for RWalk.

Method	Prostate				
	Fixed params	Tuned param	Dice \uparrow [%]	BWT \uparrow [%]	FWT \uparrow [%]
EWC	-	$\lambda = 0.4$	42.44 \pm 20.67	-28.47 \pm 12.35	-28.87 \pm 19.48
		$\lambda = 1.1$	44.83 \pm 20.86	-29.28 \pm 14.82	-27.47 \pm 16.29
		$\lambda = 1.7$	44.03 \pm 20.49	-30.96 \pm 15.33	-30.59 \pm 16.66
		$\lambda = 2.2$	46.06 \pm 18.46	-27.48 \pm 15.00	-29.65 \pm 18.32
RWalk	$\alpha = 0.9,$ update = 20	$\lambda = 0.4$	29.94 \pm 23.59	-29.01 \pm 12.84	-46.69 \pm 13.29
		$\lambda = 1.1$	27.80 \pm 20.07	-25.91 \pm 17.21	-50.06 \pm 6.757
		$\lambda = 1.7$	30.58 \pm 21.38	-22.79 \pm 18.94	-51.48 \pm 13.23
		$\lambda = 2.2$	30.48 \pm 23.26	-21.43 \pm 10.50	-54.31 \pm 11.07

Table 5. Results of the hyperparameter search considered for the two used end-to-end CL methods; mean Dice, BWT and FWT over all tasks including standard deviation [in %]; highest values and best parameter setup are marked in bold.

5. Conclusion

We introduce *Atlas Replay*, a robust atlas-based segmentation technique for continuous training in clinical environments with data drift. We evaluate our approach on seven prostate segmentation scenarios and show that it outperforms state-of-the-art continual learning methods by maintaining knowledge from early stages without compromising model plasticity. *Atlas Replay* is more generalizable than U-Net models even when trained statically with only data from one site, with a performance difference of 15% on average. We additionally introduced a prototype-building method with initial privacy preservation from a human perspective that can be leveraged to maintain structural information over time resulting in a proper balance between rigidity and plasticity for CL setups. Future work should address more stringent privacy preservation for prototypes and a self-growing atlas technique in combination with proper prototype adjustments over time to further push the potential of registration for segmentation in terms of CL. By releasing our code base along with instructions and pre-trained networks, we hope to amplify and inspire CL research beyond end-to-end models that incorporates atlas-based segmentation for medical settings.

6. Acknowledgements

This work was (partially) supported by RACOON (NUM), under BMBF grant number 01KX2021 and EVA-KI, under BMG grant number ZMVII- 2520DAT03A.

References

- [1] Paul Aljabar, Rolf A Heckemann, Alexander Hamers, Joseph V Hajnal, and Daniel Rueckert. Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *Neuroimage*, 46(3):726–738, 2009. 3
- [2] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, Bram van Ginneken, et al. The medical segmentation decathlon. *arXiv preprint arXiv:2106.05735*, 2021. 4
- [3] Takahiro Aoyama, Hidetoshi Shimizu, Tomoki Kitagawa, Kazushi Yokoi, Yutaro Koide, Hiroyuki Tachibana, Kojiro Suzuki, and Takeshi Kodaira. Comparison of atlas-based auto-segmentation accuracy for radiotherapy in prostate cancer. *Physics and Imaging in Radiation Oncology*, 19:126–130, 2021. 3
- [4] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019. 2, 3, 4, 7
- [5] Nicholas Bloch, Anant Madabhushi, Henkjan Huisman, et al. Nci-isbi 2013 challenge: Automated segmentation of prostate structures. *Medical image analysis*, 2015. 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. 2, 5
- [7] Davide Chicco and Giuseppe Jurman. The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation. *BMC genomics*, 21:1–13, 2020. 5
- [8] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. 1
- [9] Mohammad Mahdi Derakhshani, Ivona Najdenkoska, Tom van Sonsbeek, Xiantong Zhen, Dwarikanath Mahapatra, Marcel Worring, and Cees GM Snoek. Lifelong: A benchmark for continual disease classification. -, 2022. 1
- [10] Natalia Díaz-Rodríguez, Vincenzo Lomonaco, David Filliat, and Davide Maltoni. Don't forget, there is more than forgetting: new metrics for continual learning. In *Workshop on Continual Learning, NeurIPS 2018 (Neural Information Processing Systems, 2018)*. 4
- [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. 2
- [12] European Commission. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), 2016. 2
- [13] Tarun Kanti Ghosh, Md Kamrul Hasan, Shidhartho Roy, Md Ashraful Alam, Eklas Hossain, and Mohiuddin Ahmad. Multi-class probabilistic atlas-based whole heart segmentation method in cardiac ct and mri. *IEEE Access*, 9:66948–66964, 2021. 3
- [14] Camila Gonzalez, Amin Ranem, Daniel Pinto dos Santos, Ahmed Othman, and Anirban Mukhopadhyay. Lifelong nnunet: a framework for standardized medical continual learning. -, 2022. 2
- [15] Camila González, Amin Ranem, Ahmed Othman, and Anirban Mukhopadhyay. Task-agnostic continual hippocampus segmentation for smooth population shifts. In *MICCAI Workshop on Domain Adaptation and Representation Transfer*, pages 108–118. Springer, 2022. 1, 2
- [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] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [18] Ivana Isgum, Marius Staring, Annemarieke Rutten, Mathias Prokop, Max A Viergever, and Bram Van Ginneken. Multi-atlas-based segmentation with local decision fusion—application to cardiac and aortic segmentation in ct scans. *IEEE transactions on medical imaging*, 28(7):1000–1010, 2009. 3
- [19] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the*

- national academy of sciences*, 114(13):3521–3526, 2017. 1, 2, 5
- [20] Grace Lee, Aytakin Oto, and Mihai Giurcanu. Prostate mri: Is endorectal coil necessary?—a review. *Life*, 12(4):569, 2022. 2
- [21] Guillaume Lemaître, Robert Martí, Jordi Freixenet, Joan C Vilanova, Paul M Walker, and Fabrice Meriaudeau. Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric mri: a review. *Computers in biology and medicine*, 60:8–31, 2015. 4
- [22] Geert Litjens, Robert Toth, Wendy Van De Ven, Caroline Hoeks, Sjoerd Kerkstra, Bram van Ginneken, Graham Vincent, Gwenael Guillard, Neil Birbeck, Jindang Zhang, et al. Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *Medical image analysis*, 18(2):359–373, 2014. 4
- [23] Quande Liu, Qi Dou, and Pheng Ann Heng. Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2020. 4
- [24] Quande Liu, Qi Dou, Lequan Yu, and Pheng Ann Heng. Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data. *IEEE Transactions on Medical Imaging*, 2020. 4
- [25] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 4
- [26] Maria Lorenzo-Valdés, Gerardo I Sanchez-Ortiz, Raad Mohiaddin, and Daniel Rueckert. Atlas-based segmentation and tracking of 3d cardiac mr images using non-rigid registration. In *International conference on medical image computing and computer-assisted intervention*, pages 642–650. Springer, 2002. 1, 3
- [27] Bradley C Lowekamp, David T Chen, Luis Ibáñez, and Daniel Blezek. The design of simpleitk. *Frontiers in neuroinformatics*, 7:45, 2013. 3
- [28] Brian W Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451, 1975. 5
- [29] 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. 2
- [30] Umberto Michieli and Pietro Zanuttigh. Incremental learning techniques for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019. 5
- [31] 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. 2
- [32] Amin Ranem, Camila González, and Anirban Mukhopadhyay. Continual hippocampus segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3711–3720, 2022. 2
- [33] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 2
- [34] Torsten Rohlfing, Robert Brandt, Randolph Menzel, Daniel B Russakoff, and Calvin R Maurer. Quo vadis, atlas-based segmentation? In *Handbook of biomedical image analysis*, pages 435–486. Springer, 2005. 3
- [35] Clare MC Tempny, Peter R Carroll, Michael S Leapman, W Robert Lee, and Diane MF Savarese. The role of magnetic resonance imaging in prostate cancer. *UpToDate. Waltham (MA): UpToDate*, 2018. 2
- [36] Robert Trevethan. Sensitivity, specificity, and predictive values: foundations, pliabilities, and pitfalls in research and practice. *Frontiers in public health*, 5:307, 2017. 5
- [37] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. 2, 5
- [38] Ziv Yaniv, Bradley C Lowekamp, Hans J Johnson, and Richard Beare. Simpleitk image-analysis notebooks: a collaborative environment for education and reproducible research. *Journal of digital imaging*, 31(3):290–303, 2018. 3