

# ConSlide: Asynchronous Hierarchical Interaction Transformer with Breakup-Reorganize Rehearsal for Continual Whole Slide Image Analysis

Yanyan Huang<sup>1,2\*</sup>, Weiqin Zhao<sup>1\*</sup>, Shujun Wang<sup>3</sup>, Yu Fu<sup>2</sup>, Yuming Jiang<sup>4</sup>, Lequan Yu<sup>1†</sup>

<sup>1</sup>The University of Hong Kong

<sup>2</sup>Zhejiang University

<sup>3</sup>The Hong Kong Polytechnic University

<sup>4</sup>Stanford University

{yanyanh, wqzhao98}@connect.hku.hk, shu-jun.wang@polyu.edu.hk,  
yufu1994@zju.edu.cn, ymjiang2@stanford.edu, lqyu@hku.hk

## Abstract

Whole slide image (WSI) analysis has become increasingly important in the medical imaging community, enabling automated and objective diagnosis, prognosis, and therapeutic-response prediction. However, in clinical practice, the ever-evolving environment hamper the utility of WSI analysis models. In this paper, we propose the *FIRST* continual learning framework for WSI analysis, named **ConSlide**, to tackle the challenges of enormous image size, utilization of hierarchical structure, and catastrophic forgetting by progressive model updating on multiple sequential datasets. Our framework contains three key components. The Hierarchical Interaction Transformer (HIT) is proposed to model and utilize the hierarchical structural knowledge of WSI. The Breakup-Reorganize (BuRo) rehearsal method is developed for WSI data replay with efficient region storing buffer and WSI reorganizing operation. The asynchronous updating mechanism is devised to encourage the network to learn generic and specific knowledge respectively during the replay stage, based on a nested cross-scale similarity learning (CSSL) module. We evaluated the proposed ConSlide on four public WSI datasets from TCGA projects. It performs best over other state-of-the-art methods with a fair WSI-based continual learning setting and achieves a better trade-off of the overall performance and forgetting on previous tasks.

## 1. Introduction

Whole slide images (WSIs) contain rich histopathological information of the tissue sections and are routinely

\*These authors contributed equally to this work.

†Corresponding Author.

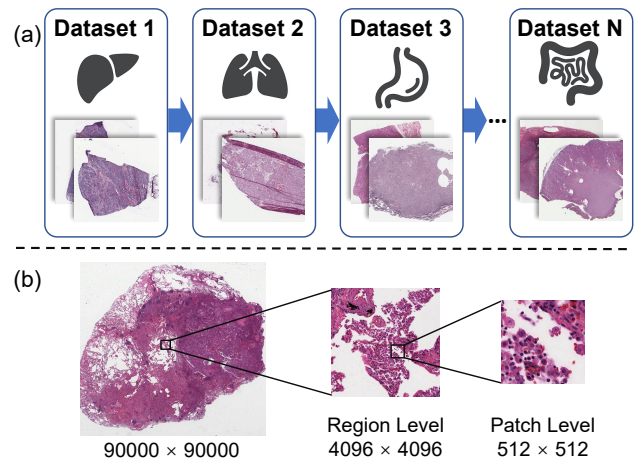


Figure 1. Illustration of concept and challenges of continual WSI analysis. (a) Continual WSI learning aims to alleviate *catastrophic forgetting* while exploiting knowledge from *sequentially incoming* datasets from the different tasks/domains. (b) The hierarchical structure of WSIs raises challenges for network architecture and learning strategy design.

used for clinical practice. With recent advances in deep learning, computational WSI have attracted widespread attention in the medical image analysis community, which can provide automated and objective diagnosis, prognosis, and therapeutic-response prediction [18, 46, 57, 60]. However, the huge size and expensive pixel-level annotations of WSI bring computational challenges in deep model architecture design [47, 28, 31]. Therefore, multiple instance learning (MIL)-based approaches are proposed for weakly-supervised WSI analysis [13, 46, 59], where they first divide WSI into a set of patches, conduct analysis for each patch, and then aggregate them together for slide-level prediction.

Although encouraging results are achieved, these ap-

proaches typically adopt the *static model learning* setting [40, 57, 65, 17, 66, 26]. However, the WSI imaging technology and staining protocols are not static [58], which constrains the model performance in new data, and they are subject to an ever-evolving environment in which WSI analysis methods have to adapt in order to remain effectiveness [37, 53, 61, 21, 32]. Although retrain the model every time new data are available is a intuitive solution, but this process can incur high computational and storage costs, especially for WSIs with huge size, while the number of WSIs for different tasks is growing fast [6]. Another candidate solution is to train a new model for each new dataset. However, it may be difficult to collect enough data to train a model for each domain (*e.g.*, different stain of WSIs). Fine-tuning the pre-trained model on the newly arrived datasets is also a candidate solution. However, during the fine-tuning process, the network is prone to concentrate too much on adapting to the feature domain of the current dataset and disrupting the knowledge acquired from previous datasets, which leads to bad performance on previous datasets. This phenomenon is referred as *catastrophic forgetting* [9, 19, 39].

*Continual Learning* (CL) was recently proposed [45, 19] to overcome the limitations of static model learning and alleviate catastrophic forgetting. The aim of CL is to train models with adaptability when meeting datasets from new tasks without catastrophically forgetting the knowledge learned from previous tasks, which can make deep learning models much more versatile to the constant growth of medical datasets. CL has received much attention in recent years, and it can be achieved either by parameter regularization [34, 56, 36], knowledge distillation [42, 24, 8], apposite architecture design [44, 54, 23], or by data rehearsal-based strategies [55, 11, 14]. Among these strategies, rehearsal-based methods achieved good performance by replaying a subset of the training data stored in a memory buffer. However, the characteristics of WSI pose unique challenges for designing continual WSI analysis frameworks and we are not aware of such frameworks in the existing literature.

WSIs are usually stored at different resolutions, resulting in a hierarchical structure with containing different pathological information, as shown in Figure 1 (b). For example, patch-level images encompass fine-grained cells and tumor cellularity information [25, 51, 3]. Region-level images mainly characterize the tissue information, such as the extent of tumor-immune localization [1, 17, 10], while slide-level images depict the overall intra-tumoral features [5, 30, 48]. Modeling and utilizing this hierarchical characteristic of WSI is critical for accurate WSI analysis [27, 52], while it is quite challenging to handle this hierarchical structure when designing a continual WSI analysis framework. Our preliminary experiments revealed that directly adapting the current CL approaches to hierarchi-

cal WSI models will lead to drastic knowledge forgetting of previous datasets (see Section 5.4). Moreover, WSIs are gigapixel images with only slide-level labels, which brings storage and computational challenges for rehearsal-based CL, as it is impractical to store and replay the representative WSI in the limited memory buffer.

To tackle these challenges, we develop a novel WSI continual analysis framework, named *ConSlide* to enable progressive update of a hierarchical WSI analysis architecture by sequentially utilizing the heterogeneous WSI datasets. To achieve that, we store a representative region set of past datasets, and then regularly reorganize and replay them during the current model update with an asynchronous updating mechanism. Specifically, we first design a novel *Hierarchical Interaction Transformer* (HIT) as the backbone to efficiently model and utilize the hierarchical characteristic of WSI. HIT is possible to aggregate both fine-grained and coarse-grained WSI features for more comprehensive WSI analysis via its bidirectional interaction within the hierarchical structure. Further, to enable the continual update of the designed hierarchical architecture, we follow the rehearsal strategy but develop a novel *Breakup-Reorganize* (BuRo) rehearsal module to tackle the unique challenges of WSI data replay. Particularly, the BuRo rehearsal module utilizes a random sampling strategy to select and store WSI regions of old tasks in an efficient manner, and then reorganize augmented WSIs of old tasks to improve the knowledge fidelity of old tasks in the replay step. Based on the augmented old task WSI data, we devise a new asynchronous updating mechanism with the inspiration of *Complementary Learning System* (CLS) theory [35], to encourage the patch-level and region-level blocks to learn generic and task-specific knowledge, respectively, by conducting a nested *Cross-Scale Similarity Learning* (CSSL) task from both old and current WSI data.

With the above careful designs, our framework can preserve the knowledge of previous WSI datasets to mitigate catastrophic forgetting and improve the generalization of the WSI model for more accurate analysis. We evaluated our framework on four public WSI datasets from TCGA projects. The extensive ablation analysis shows the effectiveness of each proposed module. Importantly, our *Conslide* achieves more improvement than other compared methods and better trade-off of the overall performance and forgetting on previous datasets under a fair WSI CL setting. Our code is available at <https://github.com/HKU-MedAI/ConSlide>.

## 2. Related Work

**Multiple Instance Learning for WSI Analysis.** Multiple Instance Learning (MIL) is widely adopted for WSI analysis. The conventional MIL approach considers handcrafted

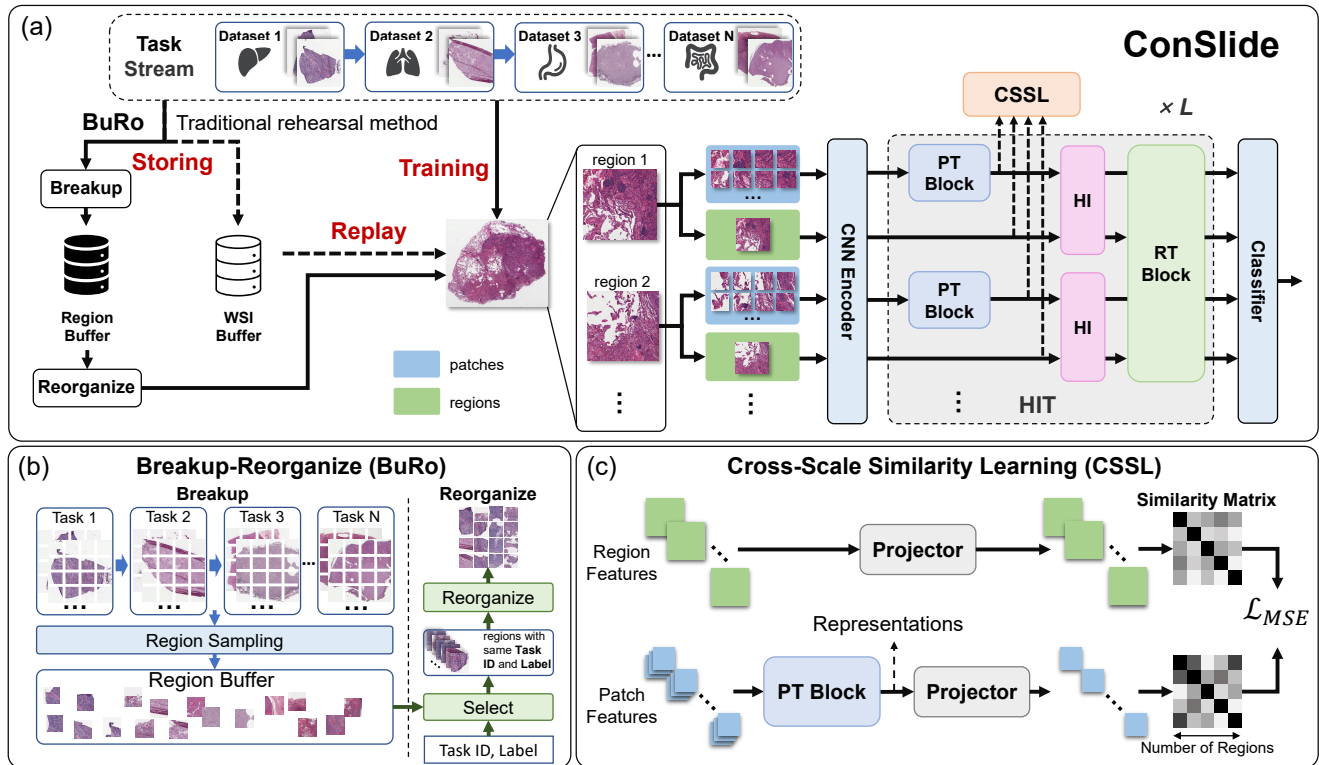


Figure 2. Illustration of our proposed continual WSI analysis framework. (a) The architecture and whole process of our proposed *ConSlide* framework, which includes *Training*, *Storing*, and *Replay*. (b) The details of our proposed Breakup-Reorganize (BuRo) rehearsal method. (c) Illustration of the proposed Cross-Scale Similarity Learning (CSSL) module.

aggregators, such as mean-pooling or max-pooling. Recently, attention-based aggregators have shown great potential. For example, Ilse *et al.* [29] proposed an attention-based MIL model and obtained the bag embedding according to the contribution of each instance. Campanella *et al.* [13] further extended this strategy to WSI analysis. Lu *et al.* [46] present an attention-based global pooling operator for weakly-supervised WSI classification by using ResNet-50 as the instance-level feature extraction. Moreover, Li *et al.* [40] extend the attention-based aggregators for WSI analysis by combining the max pooling operator and attention-based aggregators with a simple summation.

**Hierarchical Structure in WSI.** To utilize the pyramid structure information of WSI, Li *et al.* [40] proposed a multi-scale MIL model that extracts patches at two different resolutions and concatenates features from different resolutions in a pyramidal way. Hou *et al.* [27] proposed a hierarchical graph neural network to model the pyramid structure of WSI. Recently, Chen *et al.* [17] proposed a hierarchical Transformer (*i.e.*, HIPT) to leverage the natural hierarchical structure in WSI. However, HIPT uses fine-grained features to get the coarse-grained features, and there is no interaction between them. Instead, our proposed HIT model implements a novel bidirectional information interaction be-

tween patch- and region-level features so that it can model WSI in a more effective and accurate manner.

**Continual Learning for Medical Image Analysis.** Recently, continual learning has been studied in machine learning and computer vision community. There are different scenarios for continual learning, *e.g.*, class-incremental learning, task-incremental learning, and domain-incremental learning. The continual learning strategies can be divided into several categories: regularization-based [4, 34, 42, 64, 56, 36], distillation-based [42, 11, 24, 8], architecture design-based [41, 44, 62, 54, 23], and rehearsal-based [16, 55, 63, 11, 14]. Readers can refer to [19] for a detailed survey. Recently, there have been some works studying the CL setting in medical image analysis, such as MRI detection and classification [7, 53, 49], X-ray classification [38]. Particularly, some recent works discuss the benchmark of CL in pathology [21, 32, 6]. However, these works are limited in patch-level WSI analysis due to obstacles brought by the gigapixel nature of WSI. As far as we know, this work is the first one to discuss the challenges and solutions of continual slide-level pathology analysis.

### 3. Problem Formulation

**WSI Preprocessing.** WSI can be formed in a hierarchical structure, *e.g.*, WSI-, region-, and patch-level, respectively. For an original WSI  $\mathcal{I}$ , we split it into  $M$  non-overlapping regions  $\mathcal{R}$  with size of  $4,096 \times 4,096$ , and each region is further split into  $N$  non-overlapping patches  $\mathcal{P}$  with size of  $512 \times 512$ , where  $\mathcal{I} = [\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_M]$ , and  $\mathcal{R}_i = [\mathcal{P}_{i1}, \mathcal{P}_{i2}, \dots, \mathcal{P}_{iN}]$ ,  $1 \leq i \leq M$ . The regions  $\mathcal{R}$  are stored in a downsampling version in which we re-scale the size from  $4,096 \times 4,096$  to  $512 \times 512$  for downstream processing. To ease the training, we extract the *patch-level* features  $\mathbf{P}$  from  $\mathcal{P}$  and the *region-level* features  $\mathbf{R}$  from  $\mathcal{R}$  with a pre-trained feature extractor, *e.g.*, ConvNeXt [43]. Therefore, we can obtain the region-level and patch-level features with the following format  $\mathbf{R} = [\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_M]$ ,  $\mathbf{P}_i = [\mathbf{P}_{i1}, \mathbf{P}_{i2}, \dots, \mathbf{P}_{iN}]$ , and  $\mathbf{P}_{ij}, \mathbf{R}_i \in \mathbb{R}^{1 \times C}$ , where  $C$  means feature channels.

**Continual Learning for WSI Analysis.** CL for WSIs is defined as training a model on non-stationary WSI datasets with a series of tasks/classes/domains. We define a sequence of datasets  $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_T\}$ , where the  $t$ -th dataset  $\mathcal{D}_t = \{\mathbf{x}, y, t\}_k$  contains tuples of the WSI sample  $\mathbf{x}$ , its corresponding label  $y$ , and optional datasets identifier  $t$ . The subscript  $k$  means there are  $k$  samples for dataset  $\mathcal{D}_t$ . In the class-incremental scenario, the model needs to predict the classification label  $y$  with WSI  $\mathbf{x}$  as the input, while in the task-incremental scenario, the inputs are the combination of WSI  $\mathbf{x}$  and its corresponding task identifier  $t$ .

### 4. Method

Our rehearsal-based continual WSI analysis framework is shown in Figure 2 (a), which includes three steps: **Training**, **Storing**, and **Replay**. 1) We train the proposed *Hierarchical Interaction Transformer (HIT)* model with the current dataset (and replay data) in the **Training** step. 2) We adopt the proposed *Breakup-Reorganize (BuRo)* scheme to store a representative set of WSI regions from current datasets to buffer in the **Storing** step. 3) We **Replay** the stored WSI data and employ the proposed *asynchronous updating* paradigm to remind the model of the previous dataset knowledge in the next training step. The whole pipeline is conducted in a closed-loop manner.

#### 4.1. Hierarchical Interaction Transformer (HIT)

Different from natural images, WSIs are gigapixel images with a hierarchical structure. Therefore, we design a powerful Hierarchical Interaction Transformer model to abstract the hierarchical structural information of WSI. As shown in Figure 2 (a), the proposed HIT model consists of  $L$  layers, where each layer contains three components: Patch-level Transformer (PT) block, Region-level Transformer (RT) block, and Hierarchical Interaction (HI) block.

The PT and RT blocks are used to aggregate the patch-level and region-level features, while the HI block is employed to conduct the interaction between these two level blocks.

**Patch-level Transformer Block.** Given the patch-level feature set  $\mathbf{P}_i^{l-1}$  of the  $i$ -th region  $\mathcal{R}_i$ , the PT block learns the relationship of different patch features in this region and then calculates the output feature

$$\hat{\mathbf{P}}_i^l = \text{PT} \left( \mathbf{P}_i^{l-1} \right), \quad (1)$$

where  $l = 1, 2, \dots, L$  is the index of the HIT block. The dimensionality of  $\hat{\mathbf{P}}_i^l$  is  $[B, M, N, C]$ , where  $B$  is the batch size,  $M$  is the number of regions,  $N$  is the number of patches in on region, and  $C$  is the feature dimension.  $\text{PT}(\cdot)$  is one layer canonical Transformer [22], including Multi-head Self-Attention (MSA), Multi-Layer Perceptron (MLP), and Layer Normalization (LN). The PT blocks share the same parameters among all patches in the same layer.

**Hierarchical Interaction Block.** We design a HI block to fuse the patch-level features and coarse region-level features of the same region in a bidirectional manner. Specifically, for region  $\mathcal{R}_i$  in the  $l$ -th HIT block, we have its region-level feature  $\mathbf{R}_i^{l-1}$  and the output feature of the PT block  $\hat{\mathbf{P}}_i^l$  in this layer. To get the interacted patch-level feature  $\mathbf{P}_i^l$ , we add  $\mathbf{R}_i^{l-1}$  to  $\hat{\mathbf{P}}_i^l$ . To get the interacted region-level feature  $\hat{\mathbf{R}}_i^l$ , we first process the patch-level features  $\hat{\mathbf{P}}_i^l$  with a convolution and a max pooling operator and then add the processed patch-level feature to the region-level feature  $\mathbf{R}_i^{l-1}$ . The procedure can be represented as

$$\begin{aligned} \mathbf{P}_i^l &= \hat{\mathbf{P}}_i^l + \mathbf{R}_i^{l-1}, \\ \hat{\mathbf{R}}_i^l &= \mathbf{R}_i^{l-1} + \text{Max} \left( \text{Conv} \left( \hat{\mathbf{P}}_i^l \right) \right). \end{aligned} \quad (2)$$

Since  $\hat{\mathbf{P}}_i^l$  is a set of patches feature vector, we add  $\mathbf{R}_i^{l-1}$  vector to each vector of set  $\hat{\mathbf{P}}_i^l$ . The max pooling operation in equation 2 is element-wise. Similar to the PT block, the HI blocks also share the same parameters among all patches in the same layer.

**Region-level Transformer Block.** The final process in our HIT model is to learn the relationship of different interacted region-level features generated from HI blocks. Specifically, the RT block handles the interacted region-level feature set of  $M$  regions  $\hat{\mathbf{R}}^l = [\hat{\mathbf{R}}_1^l, \hat{\mathbf{R}}_2^l, \dots, \hat{\mathbf{R}}_M^l]$ , and calculates the output feature

$$\mathbf{R}_i^l = \text{RT} \left( \hat{\mathbf{R}}_i^l \right). \quad (3)$$

Similar to the PT block,  $\text{RT}(\cdot)$  is also one layer canonical Transformer. We take the first token of the RT block output

---

**Algorithm 1** The process of proposed BuRo Method

---

**Initialize:** dataset  $\mathcal{D}$ , model  $f$ , loss function  $\ell$ , scalar  $\alpha$

```
1:  $\mathcal{M} \leftarrow \{\}$ 
2: for  $\mathcal{D}_t$  in  $\mathcal{D}$  do
3:   for  $(x_k, t_k, y_k)$  in  $\mathcal{D}_t$  do
4:      $(\mathcal{R}', t', y) \leftarrow \text{Select}(\mathcal{M})$ 
5:      $x' \leftarrow \text{Reorganize}(\mathcal{R}')$ 
6:      $\text{loss} = \ell(y_k, f(x_k)) + \alpha \ell(y', x')$ 
7:      $\mathcal{R}'' = \text{Sample}(\text{Breakup}(x_k))$ 
8:      $\mathcal{M} \leftarrow \text{Reservoir}(\mathcal{R}'', t_k, y_k)$ 
9:   end for
10: end for
```

---

at the last layer as the slide-level representation, and then use a linear layer to generate the final slide-level prediction.

**Discussion.** Our proposed HIT model is different from previous hierarchical transformer [17]. We design a novel bidirectional information interaction block between the patch- and region-level features so that our network can model WSI in a more effective and accurate manner.

## 4.2. Breakup-Reorganize Rehearsal Method

Existing rehearsal-based CL methods usually save the whole natural images to the buffer, which is not suitable for gigapixel WSIs. As WSI can be split into non-overlapping regions, we propose a novel rehearsal method, named *Breakup-Reorganize* (RuRo), as shown in Figure 2 (b) and Algorithm 1.

**Breakup.** For each WSI, we first sample some regions randomly, and store these regions in the buffer. We highlight that, instead of the original image, the buffer stores region-level features  $\mathbf{R}_i$  and their corresponding patch-level features  $\mathbf{P}_i$ .

**Reorganize.** During the replaying step, according to a given pair of task id and label each time, we will randomly sample  $n$  regions which have the same task id and label from the reservoir  $\mathcal{R}'$ , and regard these regions as a new augmented WSI  $x'$  (since we don't use position information of regions). Then we reorganize the selected regions to form a new augmented WSI.

**Discussion.** We argue that our proposed BuRo method has three advantages. 1) BuRo is resource efficient in continual WSI analysis scenarios, especially when buffer size is limited, as it can store regions from a bigger number of WSI sources; 2) BuRo can boost the diversity of the buffered data by preserving diverse knowledge from previous tasks; and 3) *Reorganize* operation promotes the generalization ability of the model, as it generates new WSIs with a large number of potential combinations, which means that BuRo is able to “increase” the sample size without any extra buffer size.

## 4.3. Asynchronous Updating with Cross-Scale Similarity Learning

Inspired by CLS theory and other CLS-inspired continual learning models [50, 33, 20, 54], a continual learning system can be composed of a slow-learner and a fast-learner to learn generic and task-specific features, respectively. In the context of our proposed HIT architecture, we encourage the PT blocks to update *slowly* and extract low-level generic features, while we encourage the RT blocks to learn *quickly* on new tasks and extract high-level task-specific features. Particularly, we design a Cross-Scale Similarity Learning (CSSL) scheme under self-supervised training to facilitate asynchronous updating of the PT block and RT block parameters.

Formally, given a WSI image, we use a linear projector to project its region-level features into region-level projections, and then calculate the cosine similarity among the  $M$  region-level projections to form a similarity matrix  $\mathcal{C}^r \in \mathbb{R}^{M \times M}$ . Meanwhile, for the patch-level features of the same WSI image, we take the average of the PT block output features within the same region as the “second” region-level features. We then use the same linear projector to project them into projections and calculate the second cosine similarity matrix  $\mathcal{C}^p \in \mathbb{R}^{M \times M}$  with acquired “second” region-level projections. The objective of the CSSL is to minimize the Mean Squared Error (MSE) loss between these two similarity matrices from different levels:

$$\mathcal{L}_{CSSL} = \sum_i \sum_j (\mathcal{C}_{ij}^r - \mathcal{C}_{ij}^p)^2. \quad (4)$$

During the training phase, we conduct the proposed CSSL on the replayed and current WSI data and utilize this supervision to update the PT block parameters. Also, we calculate the final slide-level prediction loss on the current WSI data and utilize this supervision to update the network parameters. In this case, the PT blocks need to simultaneously perform well on replayed data and current data, so that they are encouraged to retain the generic low-level knowledge. Meanwhile, the RT blocks need to perform well on the current dataset, so that they are encouraged to adapt their parameters to the new task.

**Discussion.** The CSSL module constrains the consistency between region- and patch-level matrices calculated by region- and patch-level features respectively, instead of constraining the consistency between region- and patch-level features themselves. Although our asynchronous updating paradigm is inspired by the classic complementary learning system [11, 54], we specifically take advantage of the natural hierarchical structure to design our updating paradigm, which avoids explicitly adopting two steps/models and further improves the learning efficiency.

Dataset	Tumor type	Cases
NSCLC	Lung adenocarcinoma (LUAD)	492
	Lung squamous cell carcinoma (LUSC)	466
BRCA	Invasive ductal (IDC)	726
	Invasive lobular carcinoma (ILC)	149
RCC	Clear cell renal cell carcinoma (CCRCC)	498
	Papillary renal cell carcinoma (PRCC)	289
ESCA	Esophageal adenocarcinoma (ESAD)	65
	Esophageal squamous cell carcinoma (ESCC)	89

Table 1. The statistics of continual WSI analysis benchmark.

Model	Patch size	AUC	ACC
CLAM-SB [46]	256	$0.972 \pm 0.008$	$0.770 \pm 0.037$
DS-MIL [40]		$0.973 \pm 0.005$	$0.765 \pm 0.010$
TransMIL [57]		$0.976 \pm 0.008$	$0.783 \pm 0.067$
HIPT [17]		$0.977 \pm 0.005$	$0.746 \pm 0.055$
<b>HIT (Ours)</b>		<b><math>0.983 \pm 0.004</math></b>	<b><math>0.833 \pm 0.026</math></b>
CLAM-SB [46]	512	$0.967 \pm 0.006$	$0.750 \pm 0.042$
DS-MIL [40]		$0.967 \pm 0.004$	$0.760 \pm 0.026$
TransMIL [57]		$0.974 \pm 0.003$	$0.765 \pm 0.072$
HIPT [17]		$0.978 \pm 0.005$	$0.766 \pm 0.029$
<b>HIT (Ours)</b>		<b><math>0.984 \pm 0.003</math></b>	<b><math>0.831 \pm 0.018</math></b>

Table 2. Model architecture comparison on merged datasets.

## 5. Experiments

### 5.1. Datasets

As we are the first to study WSI continual learning framework, we setup a WSI continual learning benchmark as shown in Table 1. It contains a sequence of WSI-based tumor subtype classification tasks, which includes four public WSI datasets from The Cancer Genome Atlas (TCGA) repository: non-small cell lung carcinoma (NSCLC), invasive breast carcinoma (BRCA), renal cell carcinoma (RCC), and esophageal carcinoma (ESCA). Note that we simply choose two subtypes of TCGA-RCC for consistency with other datasets (both two subtypes).

### 5.2. Experimental Setting

We mainly adopt the class-incremental setting to compare our framework with other methods. We define the arrival sequence for four datasets as NSCLC, BRCA, RCC, and ESCA to conduct continual learning. In our designed benchmark, each dataset contains two categories. For the model development of the first dataset, it is defined as a two-class learning task. While for each new arrival dataset, we increase the class number by two at each time and conduct multi-class learning. The following results are reported by ten-fold cross-validation.

**Evaluation Metrics.** We assess the model’s final per-

formance with Area Under the Curve (AUC), Accuracy (ACC), and Masked Accuracy (Masked ACC) of all the previous and current datasets after continual learning. Note that the Masked ACC can reflect the performance of CL on the task-incremental scenario and it is calculated by masking irrelevant categories from different datasets. Besides the final performance, we assess the overall performance over the entire continual learning time span by *Backward Transfer (BWT)* and *Forgetting* following previous works [9, 24] for reference only.

### 5.3. Implementation Details

Following the WSI pre-processing in [46], we first segment the foreground tissue and then crop  $4,096 \times 4,096$  tiles from the segmented tissue at 20X magnification as region-level images. Each cropped region is further divided into 64 non-overlapping  $512 \times 512$  patches as the patch-level images. We adopted ConvNeXt [43] as the pre-trained feature extractor to extract both the patch-level features and region-level features from the corresponding images. For a fair comparison, we adopt the same patch-level and region-level features as the model input for all other methods.

### 5.4. Experimental Results

**Comparison with Other WSI Analysis Approaches.** We first compare our proposed HIT backbone with several state-of-the-art WSI analysis approaches to show the effectiveness of utilizing hierarchical structure information for WSI analysis: 1) single-attention-branch CLAM-SB [46], 2) non-local attention pooling-based DS-MIL [40], 3) correlated MIL-based TransMIL [57], and 4) a recent hierarchical structure-based HIPT [17]. We merge the mentioned four TCGA datasets together and formulate an eight-class classification task to predict the tumor type of each WSI.

The average per epoch training time under JointTrain setting with a single RTX 3090 GPU card are 1.32min, 1.97min, 2.87min, 1.72min, and 1.90min for CLAM-SB, DS-MIL, TransMIL, HIPT, and our HIT, respectively. Table 2 shows the comparison results with the above approaches under two different patch size settings of 256 and 512. It is clear that our proposed HIT achieves the best AUC and ACC score under both settings. Specifically, HIT outperforms the SOTA approaches with 0.6% and 5.0% improvement on AUC and ACC respectively when patch size is  $256 \times 256$ , and also acquires a performance improvement of 0.6% AUC and 6.5% ACC respectively when patch size is  $512 \times 512$ . HIPT [17] also utilizes the hierarchical information of WSI for analysis and thus outperforms other WSI analysis methods on the AUC metric. Benefiting from the proposed bidirectional hierarchical interaction scheme, HIT achieves better results than HIPT where no interaction among different level features is utilized, which further shows the effectiveness of the proposed HIT block.

CL Type	Method	Buffer size	AUC ( $\uparrow$ )	ACC ( $\uparrow$ )	Masked ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	Forgetting ( $\downarrow$ )	
Baselines	JointTrain (Upper)	-	0.984 $\pm$ 0.003	0.831 $\pm$ 0.018	0.902 $\pm$ 0.011	-	-	
	Finetune (Lower)	-	0.700 $\pm$ 0.025	0.280 $\pm$ 0.043	0.733 $\pm$ 0.040	-0.224 $\pm$ 0.062	0.224 $\pm$ 0.062	
Regularization based	LwF [42]	-	0.721 $\pm$ 0.030 ***	0.225 $\pm$ 0.013 ***	0.780 $\pm$ 0.031 *	-0.160 $\pm$ 0.056	0.161 $\pm$ 0.055	
	EWC [34]	-	0.727 $\pm$ 0.007 ***	0.273 $\pm$ 0.044 ***	0.724 $\pm$ 0.042 **	-0.226 $\pm$ 0.055	0.226 $\pm$ 0.055	
Rehearsal based	GDumb [55]	5 WSIs	0.545 $\pm$ 0.034 ***	0.201 $\pm$ 0.055 ***	0.563 $\pm$ 0.040 ***	-	-	
	ER-ACE [12]		0.723 $\pm$ 0.035 ***	0.271 $\pm$ 0.056 ***	0.680 $\pm$ 0.031 ***	-0.084 $\pm$ 0.020	0.103 $\pm$ 0.047	
	A-GEM [15]		0.803 $\pm$ 0.030 ***	0.337 $\pm$ 0.057 ***	0.753 $\pm$ 0.058 **	-0.213 $\pm$ 0.065	0.213 $\pm$ 0.065	
	DER++ [11]		0.781 $\pm$ 0.023 ***	0.339 $\pm$ 0.034 ***	0.766 $\pm$ 0.045 **	-0.169 $\pm$ 0.051	0.171 $\pm$ 0.051	
	ConSlide w/o BuRo	1100 regions ( $\approx$ 5 WSIs)	0.852 $\pm$ 0.013 ***	0.413 $\pm$ 0.009 ***	0.823 $\pm$ 0.016	-0.126 $\pm$ 0.022	0.131 $\pm$ 0.024	
	ConSlide		<b>0.915 <math>\pm</math> 0.015</b>	<b>0.553 <math>\pm</math> 0.033</b>	<b>0.835 <math>\pm</math> 0.032</b>	-0.066 $\pm$ 0.023	0.069 $\pm$ 0.021	
	GDumb [55]		10 WSIs	0.618 $\pm$ 0.064 ***	0.188 $\pm$ 0.064 ***	0.516 $\pm$ 0.069 ***	-	-
	ER-ACE [12]			0.762 $\pm$ 0.024 ***	0.360 $\pm$ 0.060 ***	0.716 $\pm$ 0.054 **	-0.066 $\pm$ 0.066	0.102 $\pm$ 0.075
	A-GEM [15]	0.828 $\pm$ 0.028 ***		0.322 $\pm$ 0.047 ***	0.814 $\pm$ 0.030	-0.122 $\pm$ 0.014	0.122 $\pm$ 0.014	
	DER++ [11]	0.860 $\pm$ 0.025 ***		0.440 $\pm$ 0.056 ***	0.809 $\pm$ 0.048	-0.112 $\pm$ 0.051	0.113 $\pm$ 0.050	
	ConSlide w/o BuRo	2200 regions ( $\approx$ 10 WSIs)	0.869 $\pm$ 0.017 ***	0.464 $\pm$ 0.034 ***	0.822 $\pm$ 0.038	-0.102 $\pm$ 0.018	0.104 $\pm$ 0.016	
	ConSlide		<b>0.931 <math>\pm</math> 0.014</b>	<b>0.594 <math>\pm</math> 0.053</b>	<b>0.837 <math>\pm</math> 0.034</b>	-0.092 $\pm$ 0.026	0.094 $\pm$ 0.023	
GDumb [55]	30 WSIs		0.661 $\pm$ 0.040 ***	0.233 $\pm$ 0.062 ***	0.609 $\pm$ 0.066 ***	-	-	
ER-ACE [12]			0.844 $\pm$ 0.032 ***	0.469 $\pm$ 0.079 ***	0.756 $\pm$ 0.020 *	-0.019 $\pm$ 0.014	0.044 $\pm$ 0.021	
A-GEM [15]		0.855 $\pm$ 0.023 ***	0.353 $\pm$ 0.119 ***	0.800 $\pm$ 0.050	-0.140 $\pm$ 0.067	0.144 $\pm$ 0.064		
DER++ [11]		0.900 $\pm$ 0.024 **	0.597 $\pm$ 0.065 *	0.840 $\pm$ 0.037	-0.078 $\pm$ 0.020	0.082 $\pm$ 0.023		
ConSlide w/o BuRo	6600 regions ( $\approx$ 30 WSIs)	0.940 $\pm$ 0.011	<b>0.668 <math>\pm</math> 0.040</b>	<b>0.866 <math>\pm</math> 0.036</b>	-0.051 $\pm$ 0.028	0.055 $\pm$ 0.024		
ConSlide		<b>0.943 <math>\pm</math> 0.007</b>	0.659 $\pm$ 0.022	0.861 $\pm$ 0.017	-0.075 $\pm$ 0.030	0.076 $\pm$ 0.030		

Table 3. Comparison results among different continual learning methods. The best performances are highlighted as **bold**. \*/\*\*/\*\* denote there are significant different (paired t-test  $p$ -value  $<$  0.05/0.01/0.001) between best performances with other comparisons.

Method	BuRo	CSSL	AUC ( $\uparrow$ )	ACC ( $\uparrow$ )	Masked ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	Forgetting ( $\downarrow$ )
ConSlide			0.860 $\pm$ 0.025 ***	0.440 $\pm$ 0.056 ***	0.809 $\pm$ 0.048	-0.112 $\pm$ 0.051	0.113 $\pm$ 0.050
		✓	0.869 $\pm$ 0.017 ***	0.464 $\pm$ 0.034 ***	0.822 $\pm$ 0.038	-0.102 $\pm$ 0.018	0.104 $\pm$ 0.016
	✓		0.903 $\pm$ 0.013 **	0.509 $\pm$ 0.034 ***	0.833 $\pm$ 0.010	-0.099 $\pm$ 0.021	0.101 $\pm$ 0.023
	✓	✓	<b>0.931 <math>\pm</math> 0.014</b>	<b>0.594 <math>\pm</math> 0.053</b>	<b>0.837 <math>\pm</math> 0.034</b>	-0.092 $\pm$ 0.026	0.094 $\pm$ 0.023

Table 4. Ablation study of RuRo method and asynchronous updating mechanism in ConSlide when buffer size is 10 WSIs (2200 regions).

**Comparison with Other Continual Learning Approaches.** Before comparison to continual learning methods, we jointly train a model with the merged datasets (JointTrain), which can be regarded as upper-bound, and train these four datasets on-by-one with the fine-tuning scheme (Finetune), which can be regarded as lower-bound. We then compare with several state-of-the-art continual learning approaches, including regularization-based methods LwF [42] and EWC [34], and rehearsal-based methods GDumb [55], ER-ACE [12], A-GEM [15], and DER++ [11], under different buffer size settings (*i.e.*, 5, 10, and 30 WSIs). We re-implement these continual learning approaches with the proposed HIT as the backbone for a fair comparison. Table 3 shows the comparison results.

The performance of JointTrain is significantly higher than Finetune (improved 28.3% in AUC, 52.7% in ACC, 15.6% in Masked ACC), confirming that the naive fine-tuning approach suffers from catastrophic forgetting when meeting new WSI datasets. Moreover, it is observed that regularization-based methods perform worse than rehearsal-based methods and only achieve a slight improve-

ment compared with the lower-bound baseline in AUC metric, which is consistent with natural image observations [11]. By comparing rehearsal-based methods under different buffer sizes, it is observed that GDumb [55] performs worse than the others, although GDumb has good performance in the natural image-based continual learning setting. Particularly, the reason may be that the buffered WSI samples cannot accurately represent the task-specific feature distribution, thus making the model difficult to preserve enough old task knowledge during the replay step. This phenomenon validates that it is non-trivial to extend the natural image CL approaches to WSI due to the unique challenges of analyzing gigapixel WSI. The rehearsal-based ER-ACE [12], A-GEM [15], and DER++ [11] can achieve better performance than GDumb due to advanced learning strategies. Compared with DER++ under different buffer size situations, the proposed ConSlide w/o BuRo already consistently produces 7.1%, 0.9%, and 4.0% improvements in AUC, 7.4%, 2.4%, and 7.1% improvements in ACC, and 5.7%, 1.3%, and 2.6% improvements in Masked ACC, validating the effectiveness of the proposed hierarchical-based

Method	Buffer size	AUC ( $\uparrow$ )	ACC ( $\uparrow$ )	Masked ACC ( $\uparrow$ )	BWT ( $\uparrow$ )	Forgetting ( $\downarrow$ )
DER++ [11]	5 WSIs	$0.787 \pm 0.017$	$0.349 \pm 0.018$	$0.799 \pm 0.034$	$-0.100 \pm 0.022$	$0.109 \pm 0.030$
ConSlide w/o BuRo		$0.822 \pm 0.009$	$0.336 \pm 0.022$	$0.805 \pm 0.028$	$-0.090 \pm 0.046$	$0.112 \pm 0.049$
ConSlide	1100 regions ( $\approx 5$ WSIs)	<b><math>0.893 \pm 0.025</math></b>	<b><math>0.426 \pm 0.047</math></b>	<b><math>0.839 \pm 0.039</math></b>	$-0.005 \pm 0.054$	$0.045 \pm 0.024$
DER++ [11]	10 WSIs	$0.819 \pm 0.015$	$0.390 \pm 0.018$	$0.787 \pm 0.031$	$-0.111 \pm 0.046$	$0.120 \pm 0.051$
ConSlide w/o BuRo		$0.867 \pm 0.029$	$0.427 \pm 0.017$	<b><math>0.839 \pm 0.026</math></b>	$-0.033 \pm 0.040$	$0.052 \pm 0.020$
ConSlide	2200 regions ( $\approx 10$ WSIs)	<b><math>0.896 \pm 0.026</math></b>	<b><math>0.453 \pm 0.052</math></b>	$0.829 \pm 0.040$	$-0.040 \pm 0.055$	$0.076 \pm 0.035$
DER++ [11]	30 WSIs	$0.883 \pm 0.032$	$0.527 \pm 0.037$	$0.852 \pm 0.027$	$-0.036 \pm 0.041$	$0.064 \pm 0.024$
ConSlide w/o BuRo		$0.914 \pm 0.004$	<b><math>0.543 \pm 0.032</math></b>	$0.849 \pm 0.036$	$-0.045 \pm 0.032$	$0.058 \pm 0.035$
ConSlide	6600 regions ( $\approx 30$ WSIs)	<b><math>0.918 \pm 0.008</math></b>	$0.499 \pm 0.025$	<b><math>0.854 \pm 0.039</math></b>	$-0.021 \pm 0.039$	$0.058 \pm 0.032$

Table 5. Continual learning results of different methods when the sequence of datasets is *reversed*.

PL	RL	HI module	AUC	ACC
✓			$0.979 \pm 0.003$ **	$0.805 \pm 0.018$ *
	✓		$0.977 \pm 0.006$ **	$0.784 \pm 0.050$ **
✓	✓		$0.981 \pm 0.006$ *	$0.805 \pm 0.033$ *
✓	✓	✓	<b><math>0.984 \pm 0.003</math></b>	<b><math>0.831 \pm 0.018</math></b>

Table 6. Ablation study on in HIT model. \*/\*\*/\*\* denote there are significant different (paired t-test  $p$ -value  $< 0.05/0.01/0.001$ ) between best performances with other comparisons.

asynchronous updating continual learning paradigm. By incorporating the proposed BuRo module, ConSlide can further boost the continual learning performance. Note that the BWT and Forgetting metrics of our proposed ConSlide are not the best among these methods, since there is a trade-off between absolute performance and forgetting resistance, and these two metrics cannot comprehensively evaluate the effectiveness of our framework.\* Following previous works [9, 24], we list BWT and Forgetting metrics as reference only.

## 5.5. Analysis of Our Framework

**Additional Ablation Analysis of HIT Model** We further conduct an ablation study for the proposed HIT module to investigate the effectiveness of each component. Table 6 lists the ablation results. It is observed that the performance will degrade if we drop either region-level or patch-level features and the performance degradation is large if we drop patch-level features (comparing the first three rows). By incorporating the hierarchical interaction module, the performance is further boosted (comparing the last two rows).

We first conducted experiments on the HIT model without region-level features, and the performance degrades ob-

\*If the absolute performance is low, it is easy to maintain the similar performance during the model updating and thus have a large BWT and low Forgetting value.

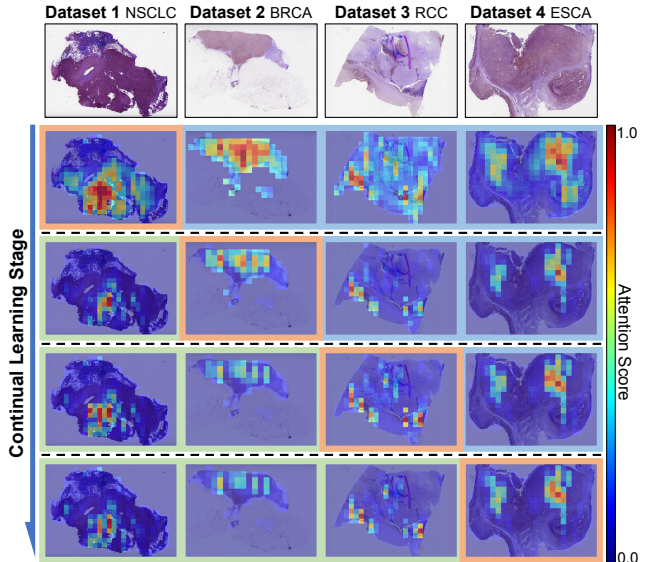


Figure 3. The region-level attention maps of four samples from four consecutive tasks in four different continual learning stages. The figures with orange/green/blue box are the attention maps of current/past/unseen tasks.

viously, while the degradation become larger for HIT model without patch-level features. This pinpoints that the fine-grained information of patch-level features and the coarse-grained information of region-level features are both vital for WSIs analysis. After that, we conducted experiments on HIT model without HI module, *i.e.*, replacing the convolution operation and max pooling in HI module with simple mean pooling, and it indicate the the combination of patch- and region-level features can further promote the performance. However, such simple combination is still less representative than the proposed HI module, which reveals that the HI module is effective for fusing features of different scales.



**Effectiveness of BuRo Module.** To figure out the effectiveness of the proposed Breakup-Reorganize rehearsal method and asynchronous updating mechanism, we did more ablation study about the BuRo and CSSL module of ConSlide framework. The results are summarized in Table 4. The comparison between ConSlide w/o BuRo and ConSlide w/ BuRo in different buffer size also reported in Table 3, which shows the effectiveness of the proposed Breakup-Reorganize Rehearsal method. Note that for a fair comparison, we set up the buffer size of BuRo in an approximate way (*i.e.*, using the average region numbers of each WSI). By storing a representative set of regions, not the whole WSIs, BuRo can save diverse regions from more WSI sources with a fixed buffer size. By reorganizing the regions to generate more new replay WSI samples, BuRo also promotes the generalization ability of the model. The performance is thus benefited from the above advantages (*i.e.*, 6.3%, 6.2%, 0.3% AUC improvement under 5, 10, and 30 WSI buffer sizes). We can observe that the marginal benefit of BuRo is saturated with the increase of buffer size. Therefore, BuRo is much more effective in CL scenarios on WSIs, especially when buffer size is limited.

**Analysis of Asynchronous Updating.** The comparison between ConSlide w/o CSSL and ConSlide w/ CSSL in Table 4 shows the effectiveness of the proposed asynchronous updating mechanism. Compared with ConSlide w/o CSSL, our asynchronous updating mechanism based on the CSSL consistently achieved better performance in all metrics no matter whether there is BuRo module, demonstrating the prominence of the asynchronous updating mechanism.

**Impact of Changing Task Orders.** We conducted experiments to verify the generalization of our framework. By reversing the sequence of the datasets in Section 5.2, we compare the performance of our framework with the strongest CL baseline DER++ and show the results in Table 5. We can observe that the proposed ConSlide w/o BuRo still achieved obvious performance increases in AUC, ACC and Masked ACC than DER++ under all settings. Besides, ConSlide with BuRo also shows further performance improvement in the reversed datasets, especially under limited buffer size (*i.e.*, 5 WSIs). These results together demonstrate the robustness of our proposed ConSlide framework.

**Attention Map Visualization.** We further visualized the attention heads of the RT block in the HIT model by using “Attention Rollout” [2] algorithm, and get the region-level attention maps of four samples from four consecutive tasks in different continual learning stages in Figure 3. First, we average the attention weights of RT blocks across all heads and then recursively multiplied the weight matrices of all layers to get the attention map between all tokens (including region-level tokens and one class token). After that, we use the  $M$  dimension attention score of the class token

to represent the attention map with single heatmap colour. The different stages are separated by dash lines, and the figures with orange/green/blue boxes are the attention maps of current/past/unseen tasks in each continual learning stage. From the maps with green boxes, it can be observed that the learned importance of tumor regions for the same sample is maintained during different stages, showing that the knowledge of previous tasks can be preserved in the proposed ConSlide framework. Besides, from the maps with blue boxes, we can find that ConSlide can assign high attention scores to tumor regions even for unseen tasks, indicating that the model trained in previous tasks can provide some prior knowledge for subsequent tasks.

## 6. Conclusion

In this paper, we propose ConSlide, an asynchronous hierarchical interaction transformer with breakup-reorganize rehearsal for continual WSI analysis. Our ConSlide sheds light on future WSI-based continual learning, due to its carefully detailed key components to deal with the WSI challenges of catastrophic forgetting, the huge size of the image, and efficient utilization of hierarchical structure. The proposed BuRo rehearsal module is specifically designed for WSI data replay with efficient region storing buffer and WSI reorganizing operation. As we adopt a transformer-based backbone and rely on a low inductive bias of spatial structure, BuRo would not significantly influence the capability of ConSlide to preserve the previous task information, although the reorganization operation in BuRo will interrupt the spatial structure of WSI. The extensive experiments on four WSI datasets of different subtype classification tasks demonstrate the superiority of ConSlide in the effectiveness and better trade-off of overall performance and forgetting on previous tasks. Moreover, the ConSlide is capable of managing survival prediction tasks by incorporating a Cox regression module. We will investigate this potential application in the future.

## 7. Acknowledgments

The work described in this paper was partially supported by grants from the National Natural Science Foundation of China (No. 62201483) and the Research Grants Council of the Hong Kong Special Administrative Region, China (T45-401/22-N).

## References

- [1] Khalid AbdulJabbar, Shan E Ahmed Raza, Rachel Rosenthal, Mariam Jamal-Hanjani, Selvaraju Veeriah, Ayse Akarca, Tom Lund, David A Moore, Roberto Salgado, Maise Al Bakir, et al. Geospatial immune variability illuminates differential evolution of lung adenocarcinoma. *Nature medicine*, 26(7):1054–1062, 2020. [2](#)
- [2] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*, 2020. [9](#)
- [3] Shahira Abousamra, David Belinsky, John S. Van Arnam, Felicia Allard, Eric Yee, Rajarsi Gupta, Tahsin M. Kurç, Dimitris Samaras, Joel H. Saltz, and Chao Chen. Multi-class cell detection using spatial context representation. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 3985–3994. IEEE, 2021. [2](#)
- [4] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. [3](#)
- [5] Frances R Balkwill, Melania Capasso, and Thorsten Hagemann. The tumor microenvironment at a glance. *Journal of cell science*, 125(23):5591–5596, 2012. [2](#)
- [6] Péter Bándi, Maschenka Balkenhol, Marcory van Dijk, Michel Kok, Bram van Ginneken, Jeroen van der Laak, and Geert Litjens. Continual learning strategies for cancer-independent detection of lymph node metastases. *Medical Image Analysis*, 85:102755, 2023. [2](#), [3](#)
- [7] Chaitanya Baweja, Ben Glocker, and Konstantinos Kamnitsas. Towards continual learning in medical imaging. In *Medical Imaging meets NIPS Workshop*, 2018. [3](#)
- [8] Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-incremental continual learning into the extended der-verse. *IEEE transactions on pattern analysis and machine intelligence*, 2022. [2](#), [3](#)
- [9] Matteo Boschini, Lorenzo Bonicelli, Angelo Porrello, Giovanni Bellitto, Matteo Pennisi, Simone Palazzo, Concetto Spampinato, and Simone Calderara. Transfer without forgetting. In *European Conference on Computer Vision*, 2022. [2](#), [6](#), [8](#)
- [10] Nadia Brancati, Anna Maria Anniciello, Pushpak Pati, Daniel Riccio, Giosuè Scognamiglio, Guillaume Jaume, Giuseppe De Pietro, Maurizio Di Bonito, Antonio Foncubiarta, Gerardo Botti, et al. Bracs: A dataset for breast carcinoma subtyping in h&e histology images. *Database*, 2022, 2022. [2](#)
- [11] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020. [2](#), [3](#), [5](#), [7](#), [8](#)
- [12] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New insights on reducing abrupt representation change in online continual learning. In *International Conference on Learning Representations*, 2022. [7](#)
- [13] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Mirafior, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019. [1](#), [3](#)
- [14] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9516–9525, 2021. [2](#), [3](#)
- [15] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2019. [7](#)
- [16] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaisyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. [3](#)
- [17] Richard J Chen, Chengkuan Chen, Yicong Li, Tiffany Y Chen, Andrew D Trister, Rahul G Krishnan, and Faisal Mahmood. Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16144–16155, 2022. [2](#), [3](#), [5](#), [6](#)
- [18] Toby C Cornish, Ryan E Swapp, and Keith J Kaplan. Whole-slide imaging: routine pathologic diagnosis. *Advances in anatomic pathology*, 19(3):152–159, 2012. [1](#)
- [19] 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. [2](#), [3](#)
- [20] 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, 2022. [5](#)
- [21] Mohammad Mahdi Derakhshani, Ivona Najdenkoska, Tom van Sonsbeek, Xiantong Zhen, Dwarikanath Mahapatra, Marcel Worring, and Cees G. M. Snoek. Lifelong: A benchmark for continual disease classification. In *Medical Image Computing and Computer Assisted Intervention*, volume 13432, pages 314–324. Springer, 2022. [2](#), [3](#)
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*. [4](#)
- [23] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9285–9295, 2022. 2, 3
- [24] Enrico Fini, Victor G Turrisi da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022. 2, 3, 6, 8
- [25] Simon Graham, Quoc Dang Vu, Shan e Ahmed Raza, Ayesha Azam, Yee-Wah Tsang, Jin Tae Kwak, and Nasir M. Rajpoot. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Anal.*, 58, 2019. 2
- [26] Yonghang Guan, Jun Zhang, Kuan Tian, Sen Yang, Pei Dong, Jinxi Xiang, Wei Yang, Junzhou Huang, Yuyao Zhang, and Xiao Han. Node-aligned graph convolutional network for whole-slide image representation and classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18813–18823, 2022. 2
- [27] Wentai Hou, Lequan Yu, Chengxuan Lin, Helong Huang, Rongshan Yu, Jing Qin, and Liansheng Wang. H<sup>2</sup>-mil: Exploring hierarchical representation with heterogeneous multiple instance learning for whole slide image analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 933–941, 2022. 2, 3
- [28] Shih-Chiang Huang, Chi-Chung Chen, Jui Lan, Tsan-Yu Hsieh, Huei-Chieh Chuang, Meng-Yao Chien, Tao-Sheng Ou, Kuang-Hua Chen, Ren-Chin Wu, Yu-Jen Liu, et al. Deep neural network trained on gigapixel images improves lymph node metastasis detection in clinical settings. *Nature Communications*, 13(1):1–14, 2022. 1
- [29] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. 3
- [30] Sajid Javed, Arif Mahmood, Muhammad Moazam Fraz, Navid Alemi Koohbanani, Ksenija Benes, Yee-Wah Tsang, Katherine Hewitt, David Epstein, David Snead, and Nasir Rajpoot. Cellular community detection for tissue phenotyping in colorectal cancer histology images. *Medical image analysis*, 63:101696, 2020. 2
- [31] Sajid Javed, Arif Mahmood, Naoufel Werghi, Ksenija Benes, and Nasir Rajpoot. Multiplex cellular communities in multi-gigapixel colorectal cancer histology images for tissue phenotyping. *IEEE Transactions on Image Processing*, 29:9204–9219, 2020. 1
- [32] Veena Kaustaban, Qinle Ba, Ipshita Bhattacharya, Nahil Sobh, Satarupa Mukherjee, Jim Martin, Mohammad Saleh Miri, Christoph Guetter, and Amal Chaturvedi. Characterizing continual learning scenarios for tumor classification in histopathology images. In *International Workshop on Medical Optical Imaging and Virtual Microscopy Image Analysis*, pages 177–187. Springer, 2022. 2, 3
- [33] Ronald Kemker and Christopher Kanan. Fearnnet: Brain-inspired model for incremental learning. In *International Conference on Learning Representations*, 2018. 5
- [34] 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. 2, 3, 7
- [35] Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534, 2016. 2
- [36] Richard Kurle, Botond Cseke, Alexej Klushyn, Patrick Van Der Smagt, and Stephan Günnemann. Continual learning with bayesian neural networks for non-stationary data. In *International Conference on Learning Representations*, 2019. 2, 3
- [37] Cecilia S Lee and Aaron Y Lee. Clinical applications of continual learning machine learning. *The Lancet Digital Health*, 2(6):e279–e281, 2020. 2
- [38] Matthias Lenga, Heinrich Schulz, and Axel Saalbach. Continual learning for domain adaptation in chest x-ray classification. In *Medical Imaging with Deep Learning*, pages 413–423. PMLR, 2020. 3
- [39] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion*, 58:52–68, 2020. 2
- [40] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2021. 2, 3, 6
- [41] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR, 2019. 3
- [42] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 2, 3, 7
- [43] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 4, 6
- [44] Noel Loo, Siddharth Swaroop, and Richard E. Turner. Generalized variational continual learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 2, 3
- [45] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 2
- [46] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature biomedical engineering*, 5(6):555–570, 2021. 1, 3, 6
- [47] Wenqi Lu, Simon Graham, Mohsin Bilal, Nasir Rajpoot, and Fayyaz Minhas. Capturing cellular topology in

- multi-gigapixel pathology images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 260–261, 2020. 1
- [48] Andriy Marusyk, Vanessa Almendro, and Kornelia Polyak. Intra-tumour heterogeneity: a looking glass for cancer? *Nature reviews cancer*, 12(5):323–334, 2012. 2
- [49] 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. 3
- [50] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019. 5
- [51] Pushpak Pati, Guillaume Jaume, Lauren Alisha Fernandes, Antonio Foncubierta-Rodríguez, Florinda Feroce, Anna Maria Anniciello, Giosue Scognamiglio, Nadia Brancati, Daniel Riccio, Maurizio Di Bonito, et al. Hactnet: A hierarchical cell-to-tissue graph neural network for histopathological image classification. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis*, pages 208–219. Springer, 2020. 2
- [52] Pushpak Pati, Guillaume Jaume, Antonio Foncubierta-Rodríguez, Florinda Feroce, Anna Maria Anniciello, Giosue Scognamiglio, Nadia Brancati, Maryse Fiche, Estelle Dubruc, Daniel Riccio, Maurizio Di Bonito, Giuseppe De Pietro, Gerardo Botti, Jean-Philippe Thiran, Maria Frucci, Orcun Goksel, and Maria Gabrani. Hierarchical graph representations in digital pathology. *Medical Image Anal.*, 75:102264, 2022. 2
- [53] Matthias Perkonigg, Johannes Hofmanninger, Christian J Herold, James A Brink, Oleg Pianykh, Helmut Prosch, and Georg Langs. Dynamic memory to alleviate catastrophic forgetting in continual learning with medical imaging. *Nature Communications*, 12(1):1–12, 2021. 2, 3
- [54] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34:16131–16144, 2021. 2, 3, 5
- [55] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *European conference on computer vision*, pages 524–540. Springer, 2020. 2, 3, 7
- [56] 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, 3
- [57] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in Neural Information Processing Systems*, 34:2136–2147, 2021. 1, 2, 6
- [58] Yiqing Shen, Arcot Sowmya, Yulin Luo, Xiaoyao Liang, Dinggang Shen, and Jing Ke. A federated learning system for histopathology image analysis with an orchestral stain-normalization gan. *IEEE Transactions on Medical Imaging*, pages 1–1, 2022. 2
- [59] Artem Shmatko, Narmin Ghaffari Laleh, Moritz Gerstung, and Jakob Nikolas Kather. Artificial intelligence in histopathology: enhancing cancer research and clinical oncology. *Nature Cancer*, 3(9):1026–1038, 2022. 1
- [60] Junko Tanizaki, Hidetoshi Hayashi, Masatomo Kimura, Kaoru Tanaka, Masayuki Takeda, Shigeki Shimizu, Akihiko Ito, and Kazuhiko Nakagawa. Report of two cases of pseudoprogression in patients with non-small cell lung cancer treated with nivolumab—including histological analysis of one case after tumor regression. *Lung Cancer*, 102:44–48, 2016. 1
- [61] Jeroen Van der Laak, Geert Litjens, and Francesco Ciompi. Deep learning in histopathology: the path to the clinic. *Nature medicine*, 27(5):775–784, 2021. 2
- [62] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *Advances in Neural Information Processing Systems*, 33:15173–15184, 2020. 3
- [63] 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. 3
- [64] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 3
- [65] Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, Sarah E Coupland, and Yalin Zheng. Dtdfmil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18802–18812, 2022. 2
- [66] Yushan Zheng, Jun Li, Jun Shi, Fengying Xie, and Zhiguo Jiang. Kernel attention transformer (kat) for histopathology whole slide image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 283–292. Springer, 2022. 2