

Continual Self-supervised Learning: Towards Universal Multi-modal Medical Data Representation Learning

Yiwen Ye¹ Yutong Xie^{2†} Jianpeng Zhang³ Ziyang Chen¹ Qi Wu² Yong Xia^{1,4,5†}

¹ School of Computer Science and Engineering, Northwestern Polytechnical University, China

² Australian Institute for Machine Learning (AIML), The University of Adelaide, Australia

³ College of Computer Science and Technology, Zhejiang University, China

⁴ Research & Development Institute of Northwestern Polytechnical University in Shenzhen, China

⁵ Ningbo Institute of Northwestern Polytechnical University, China

ywye@mail.nwpu.edu.cn, yutong.xie678@gmail.com, jianpeng.zhang0@gmail.com

zychen@mail.nwpu.edu.cn, qi.wu01@adelaide.edu.au, yxia@nwpu.edu.cn

Abstract

Self-supervised learning (SSL) is an efficient pre-training method for medical image analysis. However, current research is mostly confined to certain modalities, consuming considerable time and resources without achieving universality across different modalities. A straightforward solution is combining all modality data for joint SSL, which poses practical challenges. Firstly, our experiments reveal conflicts in representation learning as the number of modalities increases. Secondly, multi-modal data collected in advance cannot cover all real-world scenarios. In this paper, we reconsider versatile SSL from the perspective of continual learning and propose MedCoSS, a continuous SSL approach for multi-modal medical data. Different from joint representation learning, MedCoSS assigns varying data modalities to separate training stages, creating a multi-stage pre-training process. We propose a rehearsal-based continual learning approach to manage modal conflicts and prevent catastrophic forgetting. Specifically, we use the k -means sampling to retain and rehearse previous modality data during new modality learning. Moreover, we apply feature distillation and intra-modal mixup on buffer data for knowledge retention, bypassing pretext tasks. We conduct experiments on a large-scale multi-modal unlabeled dataset, including clinical reports, X-rays, CT, MRI, and pathological images. Experimental results demonstrate MedCoSS's exceptional generalization ability across 9 downstream datasets and its significant scalability in integrating new modality data. The code and pre-trained model are available at <https://github.com/yeerwen/MedCoSS>.

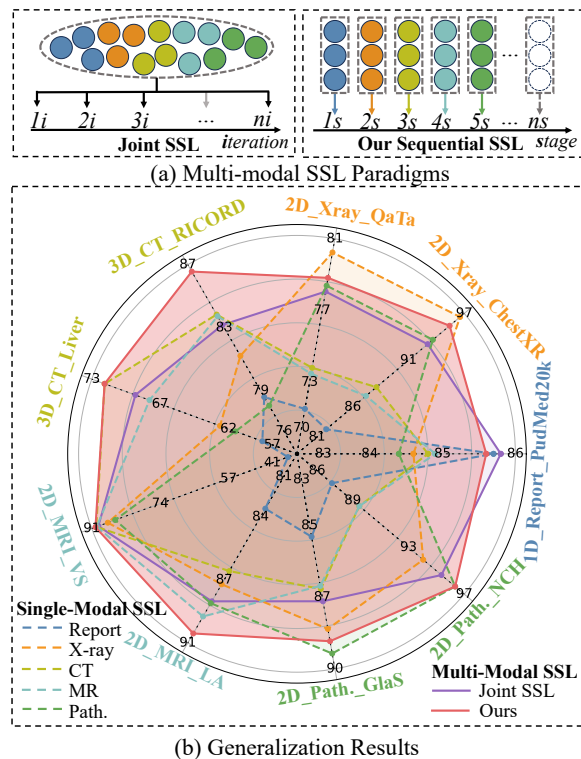


Figure 1. (a) Comparison between joint SSL and our sequential SSL. Different colored balls represent different modalities. The joint SSL suffers from the high costs to integrate knowledge from new data while our sequential SSL shows a distinct advantage in this aspect. (b) Generalization performance of five single-modal pre-trained models (dashed lines) and two multi-modal pre-trained models (solid lines). The label shows the data dimension, modality, and dataset name.

[†] Yutong Xie and Yong Xia are the corresponding authors.

1. Introduction

Medical self-supervised learning (SSL) emerges as a powerful alternative to large-scale supervised learning, showcasing impressive feats in high-quality representation learning. Its essence lies in pre-training strong encoder or encoder-decoder architectures for subsequent downstream tasks by solving a pretext task without human annotations. A review of existing medical SSL approaches highlights a pattern: they primarily focus on a single modality data, such as X-rays [2, 4, 6, 16, 39, 56, 57], computed tomography (CT) scans [16, 19, 41, 50, 52, 53, 56–58], magnetic resonance image (MRI) scans [6, 22, 33, 55], pathological images [23, 44, 51], fundus images [27], multivariate cardiac signals [26], and ultrasound (US) images [8], or only a handful of paired modality data like radiology reports paired with X-rays [9, 29, 30, 48, 54], MRI modality pairs [40], and images paired with genomics [20]. This narrow focus constrains them from extending to the universal medical pre-training since a large-scale paired multi-modal dataset is often impractical. Despite a few pioneering attempts in unpaired multi-modal pre-training [7, 15, 46], the scope of these efforts remains notably constrained. Specifically, these efforts are either limited to only two modalities [46], only for a specific field [7], or restricted to a specific dimension [15]. These limitations hinder their ability to delve deep into the challenges of universal multi-modal SSL in the face of more modalities.

We used the plain ViT/B [13] as the backbone and masked modeling [12, 17] as the pretext task. The results are the average values of accuracy, AUC, and F1 for classification tasks or Dice for segmentation tasks. We observed a consistent trend: single-modal pre-trained models excel in downstream tasks with identical modality but falter considerably when dealing with the data from another modal (see Fig. 1(b)). This observation underscores the urgency to craft a universal pre-training model proficiently handling varied medical multi-modal data. A seemingly intuitive solution to this problem is joint SSL (see Fig. 1(a)), where all accessible medical data (from various modalities) are completely collected and involved in pre-training executed as single-modal per mini-batch considering the diverse dimensions inherent in medical imaging [7, 46]. However, compared to single-modal SSL, where the models are pre-trained and fine-tuned on an identical modality, joint SSL falls short of expectations, despite its richer information base. This is evident from Fig. 1(b), which shows a significant performance drop in most downstream tasks. We attribute this performance degradation to the phenomenon, which we term as *modal data collision*, that the representation learning on two modalities conflicts with each other owing to the vast disparities between those modalities. Moreover, a joint training paradigm is also inadvisable, due to the high costs associated with incorporating the knowl-

edge from both the seen and newly introduced modal data, necessitating a full reiteration of the pre-training process.

To circumvent the obstacle of *modal data collision* and recover the scalability of the pre-trained model, we propose the **Medical Continual Self-Supervised (MedCoSS)** paradigm for multi-modal representation learning. MedCoSS diverges from traditional joint pre-training by adopting a sequential pre-training paradigm [35], assigning each stage to a specific modality. To avoid catastrophic forgetting, we design a rehearsal-based continual learning technique that retains a subset (*e.g.*, 5%) of previous modal data within a rehearsal buffer, using the *k*-means sampling strategy over random sampling. Specific to previous knowledge retention, we deploy a feature distillation strategy and intra-modal mixup strategy during subsequent pre-training stages. According to our MedCoSS paradigm, a model undergoes pre-training on the data from five modalities: Report, X-ray, CT, MRI, and Pathological imaging. Evaluation on nine downstream datasets encompassing all seen modalities at the pre-training phase to showcase the superior generalization capability of our model. The model from MedCoSS consistently outperforms the models developed through single-modal pre-training, joint pre-training, and other multi-modal pre-training paradigms, while maintaining cost-effective scalability to include new knowledge, paving the way to build the multi-modal pre-trained medical universal model. In summary, our contributions can be encapsulated as:

- We identify and mitigate the *modal data collision* issue and innovate the MedCoSS paradigm. By shifting from joint to sequential training and introducing continual learning, we alleviate the collisions and cost-effectively scale new knowledge without forgetting the old.
- We conduct an in-depth exploration into unpaired multi-modal SSL, expanding the modality and data dimension. We integrate five prevalent modalities, including the Report, X-ray, CT, MRI, and Pathological imaging, spanning three dimensions (1D, 2D, and 3D) to pre-train a universal model using the proposed MedCoSS.
- The model developed through our MedCoSS achieves state-of-the-art generalization performance on a broad range of downstream tasks, indicating a potential direction to develop the multi-modal pre-trained medical universal model.

2. Related Work

2.1. Medical self-supervised learning

Medical SSL has been increasingly used to improve the performance of medical image analysis, achieving success in various clinical applications, with its success largely dependent on large-scale pre-training data. The modalities of data directly influence the performance of models in target tasks.

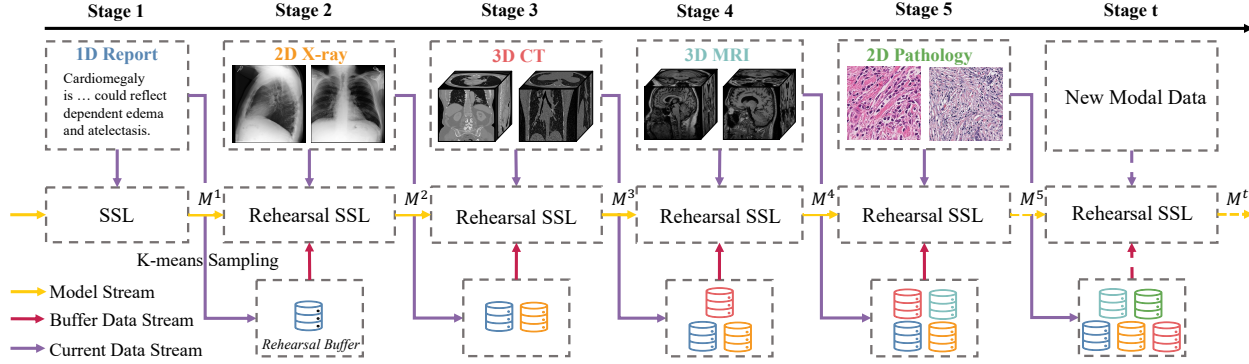


Figure 2. Pipeline of proposed MedCoSS paradigm. MedCoSS leverages a sequential pre-training manner to produce a multi-modal pre-trained model (M^t) through multi-stage learning, each assigned to specific modal data. To alleviate the issue of catastrophic forgetting, a rehearsal buffer is adopted and maintained for previous modal data rehearsal in the subsequent pre-training. After each stage, the model M_t , which consists of three tokenizers and a Transformer-based encoder, will be used to select samples from current data using a k -means sampling strategy and then transferred to the next stage.

Generally, popular medical SSL can be categorized into three primary categories: single-modal SSL, paired multi-modal SSL, and unpaired multi-modal SSL. **Single-modal SSL**: this approach leverages only one modality for pre-training. Typical modalities in this category encompass X-ray [2, 4, 6, 16, 39, 56, 57], CT [16, 19, 41, 50, 52, 53, 56–58], MRI [6, 22, 33, 55], Pathological imaging [23, 44, 51], fundus image [27], multivariate cardiac signals [26], and ultrasound [8]. Despite the versatility of single-modal SSL, our observations in Fig. 1(b) indicate that models trained on one modality often struggle to generalize across other modalities, particularly when these modalities differ in dimensionality. **Paired multi-modal SSL**: this approach leverages multiple modalities with an inherent relationship for pre-training. The design purpose behind the corresponding pretext tasks is to encourage the model to be aware of these relationships. Existing SSL paradigms employ various data pair types, such as clinical reports paired with images [9, 29, 30, 45, 48, 54], captions paired with images [45], or MRI modality pairs [40]. However, collecting such paired multi-modal datasets is resource-intensive and impractical, even for widely used modalities. **Unpaired multi-modal SSL**: this approach presents a more economical alternative to paired multi-modal SSL by relaxing data constraints, allowing for easier scaling. Pioneering efforts have delved into this field. Ghesu *et al.* [15] collected a 2D mixed modality dataset comprising X-rays, CT slices, MRI slices, and US images. Their approach, focusing on multi-modal clustering, enabled distinct online clustering for each modality using individual prototypes. Cai *et al.* [7] devoted to ophthalmology, creating a multi-modal dataset and leveraging masked image modeling (MIM) for effective representation learning. Xie *et al.* [46] introduced an iterative training scheme for joint training of CT and X-ray data, incorporating volume-slice consistency as a form

of cross-dimension regularization during CT pre-training. These studies, each addressing different aspects of unpaired multi-modal SSL, collectively enrich our understanding of this field. However, as modalities and data dimensions increase, these paradigms often fall short in fine-tuning compared to single-modal pre-training, hindered by the modal data collision in joint training. Crucially, the pre-trained models obtained from these paradigms exhibit limited scalability when they come to incorporating knowledge from new data sources. Therefore, we introduce continual learning technology in the SSL paradigm to mitigate these issues, proposing a more practical and efficient approach to multi-modal SSL.

2.2. Continual learning for knowledge retention

Continual Learning (CL) techniques have been instrumental in aiding models to retain knowledge over time, especially when sequentially exposed to diverse data or tasks. These CL techniques can be summarized in five main categories, including regularization-based approach [25], rehearsal-based approach [5, 35], optimization-based approach [28], representation-based approach [14], and architecture-based approach [32]. The rehearsal-based approach, notably successful in retaining knowledge when equipped with a rehearsal buffer, usually maintains some previous data [35] or features [5] and replays them during subsequent learning. A standout method within this category is DER [5], which retains previous training samples and their associated logits, focusing on logit-matching rather than label-matching. In this study, we adjust DER with three customized modifications, including: (1) we replaced reservoir sampling with a k -means sampling strategy, considering sample diversity; (2) instead of directly preserving logits of samples, we only preserved samples and utilized a freeze encoder to extract logits, catering to the random masking operation for inputs;

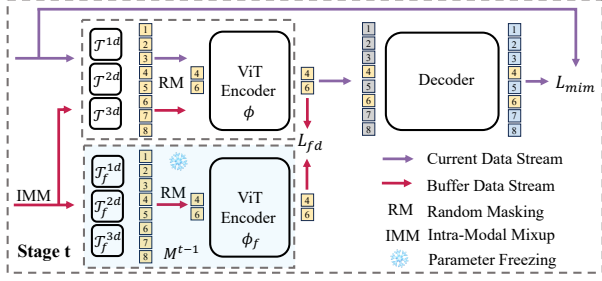


Figure 3. Illustration of proposed rehearsal SSL paradigm. The masked modeling pretext task [12, 17] is used for representation learning on current modal data corresponding to stage t . The feature distillation strategy is used for knowledge retention by building a freeze model M^{t-1} and rehearsal buffer.

(3) we introduced an intra-modal mixup strategy to augment data from the rehearsal buffer.

3. Approach

3.1. Overview

The proposed MedCoSS paradigm is designed as a two-step SSL, which comprises an unsupervised pre-training phase and a fully-supervised fine-tuning phase. During pre-training, we employ masked image/language modeling as the pretext task to extract generalized representations from an integrated set of multi-modal data, specifically clinical reports, X-rays, CT scans, MRI scans, and pathological images. To circumvent the obstacle of *modal data collision* caused by the joint multi-modal pre-training, we introduce a sequential pre-training scheme, in which each stage is tailored for training on a specific modality of data. The potential risk of knowledge forgetting during this sequential process is countered with continual learning techniques. During fine-tuning, the pre-trained encoder is paired with a randomly initialized task-specific head for each downstream task. A comprehensive visualization of our MedCoSS paradigm was illustrated in Fig. 2. We now delve into the details of each part.

3.2. Universal architecture for multi-modal SSL

We craft a universal architecture [59] in line with the aim of universal multi-modal SSL. The medical data from various modalities could be either 1D (e.g., clinical reports), 2D (e.g., X-rays and pathological images), or 3D (e.g., CT and MRI scans). To accommodate these data seamlessly, a dimension-free backbone is imperative. For this study, we employ three dimension-specific tokenizers to convert 1D, 2D, and 3D medical data into token sequences, respectively, and use a plain ViT/B [13] as the encoder for representation learning in a sequence-to-sequence manner regardless of the dimension of the medical data. Specifically, we deploy the

byte pair encoding (BPE) tokenizer [36] for text and 2D/3D image patch tokenizer [13] for 2D/3D visual data. For text, following BERT [12], we mask out 15% words randomly. The model predicts each masked word according to the visible words, using the cross-entropy loss as a constraint. For visual data, following MAE [17], token sequences are randomly masked at a ratio of 75% and only those unmasked token sequences are fed to the encoder. Subsequently, the encoded visible token sequences, together with learnable mask tokens, serve as the inputs of the Transformer-based decoder, aiming to reconstruct the previously masked tokens. The mean squared error (MSE) loss is employed to ensure a high consistency between the original and reconstructed images in masked regions.

3.3. Rehearsal-based continual pre-training

The typical approach to joint pre-training on multi-modal data often encounters the issues of *modal data collision* and high cost of integrating new knowledge. To address both issues, we advocate a shift towards a sequential pre-training paradigm, effectively assigning each stage to a specific imaging modality. Formally, considering T unlabeled subsets of data $D = \{D_1, D_2, \dots, D_T\}$, each being acquired from a unique modality. The correspondence between modalities and stages is random. Unlike the standard practice of directly pre-training a model M on D , our MedCoSS paradigm pre-trains M sequentially on each subset D_t in the t^{th} stage, where the intermediate pre-trained model is denoted by M^t . This paradigm strategically circumvents *modal data collision* by isolating different modalities during pre-training while at the risk of catastrophic forgetting. We counter this risk by integrating rehearsal-based CL techniques to preserve previous knowledge. When stage t unfolds, the pre-training not only focuses on the MIM pretext task using D_t , but also engages an auxiliary feature distillation task to prevent catastrophic forgetting.

Learn from current modality. We utilize M_{t-1} , which includes an encoder ϕ and three tokenizers $\mathcal{T}^{1d}, \mathcal{T}^{2d}, \mathcal{T}^{3d}$, along with a randomly initialized decoder to continually learn new knowledge from current modal data on a masked modeling pretext task.

Learn from previous modalities. We establish a rehearsal buffer $B = \{x_i^j \mid x_i^j \in D_j, j \in (1, 2, \dots, t-1)\}$ to retain a fraction of training data from each of previous stages (see Section 3.3.1 for details). Moreover, we copy a frozen version of M_{t-1} , where an encoder and three tokenizers are denoted as $\phi_f, \mathcal{T}_f^{1d}, \mathcal{T}_f^{2d}, \mathcal{T}_f^{3d}$, respectively. For each sample x from B , the intra-modal mixup (IMM) strategy is utilized (see Section 3.3.2 for details) for augmentation. The augmented sample is processed by the learnable network (i.e., $\mathcal{T}^{1d}, \mathcal{T}^{2d}, \mathcal{T}^{3d}$, and ϕ) and frozen network (i.e., $\mathcal{T}_f^{1d}, \mathcal{T}_f^{2d}, \mathcal{T}_f^{3d}$, and ϕ_f), respectively. The embeddings produced by both networks are then encouraged to be similar

Algorithm 1: MedCoSS’s pre-training algorithm.

Input: T datasets $\{D_1, \dots, D_T\}$ with different modalities, rehearsal buffer B , tokenizers \mathcal{T} , encoder ϕ , model-specific decoders $\psi_i, (i \in (1, \dots, T))$, sampling operation $Sample(\cdot)$

Output: ϕ and \mathcal{T}

- 1: \triangleright **Stage 1:**
 - 2: Training dataset $D \leftarrow D_1$
 - 3: Update ϕ, \mathcal{T} , and ψ_1 by L_{mim}
 - 4: $B \leftarrow Sample(D_1)$
 - 5: \triangleright **Stage 2, 3, ...:**
 - 6: **for** $i = 2, 3, \dots$ **do**
 - 7: Training dataset $D \leftarrow D_i \cup B$
 - 8: **for** $iteration = 1, 2, \dots$ **do**
 - 9: Sample a batch of unlabeled data x from D
 - 10: **if** $x \in D_i$ **then**
 - 11: Update ϕ, \mathcal{T} , and ψ_i by L_{mim}
 - 12: **else**
 - 13: Update ϕ and \mathcal{T} by L_{fd}
 - 14: **end if**
 - 15: **end for**
 - 16: $B \leftarrow B \cup Sample(D_i)$
 - 17: **end for**
-

via minimizing the MSE loss (see L_{fd} in Fig. 3).

The pre-training process of our MedCoSS paradigm is summarized in Algorithm 1.

3.3.1 Rehearsal buffer construction

In our MedCoSS paradigm, the rehearsal buffer B is constructed for subsequent training stages by using the k -means sampling strategy to select representative samples from each subset D_t . Unlike the random sampling that lacks reliability in capturing data diversity, k -means sampling consists of two steps: (1) clustering each subset D_t into C clusters based on the embeddings produced by the pre-trained model (*i.e.*, M_{t-1} in the t^{th} Stage), and (2) selecting K samples, which are nearest to the center, from each cluster. Note that the number of clusters C is empirically set to 1% of the subset size.

3.3.2 IMM Augmentation Strategy

In MedCoSS, the IMM augmentation strategy enhances the variety of the samples drawn from the rehearsal buffer B to deal with a limited buffer size for each modality.

For text data, we denote the batch of text samples as $b \in \mathbb{R}^{N \times L}$, where N represents the batch size and L means the length of text sequences. To perform augmentation, we initiate by copying the batch b to produce b' , which is sub-

Table 1. Overview of the datasets used in this study.

	Dataset	Modality	Task	#Train	#Val	#Test
SSL	MIMIC-CXR	Report	MLM/MIM	227,323	-	-
	MIMIC-CXR	X-ray		356,309	-	-
	DeepLesion	CT		10,594	-	-
	ADNI	MRI		3,789	-	-
	TCGA	Path.		217452	-	-
DS	PudMed20k	Report	Cls	180,040	30,212	30,135
	ChestXR	X-ray	Cls	14,365	3,593	3,432
	QaTa	X-ray	Seg	5,716	1,429	2,113
	RICORD	CT	Cls	194	52	84
	LiTS	CT	Seg	104	-	27
	VS	MRI	Seg	193	-	49
	LA	MRI	Seg	80	-	20
	NCH	Path.	Cls	79,994	20,006	7,180
	GlaS	Path.	Seg	67	18	80

sequently shuffled to create a new sequence order. A binary mixup strategy is then employed to synthesize a new batch b^{mix} , defined as:

$$b^{mix} = \lambda_b \odot b + (1 - \lambda_b) \odot b', \quad (1)$$

where \odot denotes element-wise multiplication, and λ_b represents a binary mask array matching the shape of b . This mask’s values, either 0 or 1, are derived from an array generated by uniformly sampling numbers with a random threshold, also sourced from a uniform distribution within the range of $[0, 1)$.

For visual data, the augmentation process begins similarly by duplicating and shuffling the image batch b to obtain b' . The difference is that we incorporate a continual mixup strategy here, expressed as:

$$b^{mix} = \lambda_c \odot b + (1 - \lambda_c) \odot b', \quad (2)$$

where the continual mask vector λ_c is an array of length N with values randomly chosen between $[0, 1)$. Each number in this vector corresponds to a specific sample, determining how much each image in batch b blends with its counterpart in batch b' .

3.4. Fine-tuning on downstream tasks

After pre-training, we adopt the pre-trained encoder and customize it for various downstream tasks, including the classification and segmentation of all seen modalities. Given a task, a corresponding prediction head is designed according to the dimensionality of input data and the task type. We use the Multi-Layer Perceptron (MLP) head for classification tasks and employ a convolutional-based decoder with a segmentation head for segmentation tasks. More details are in Appendix A.

4. Experiment

4.1. Datasets

The datasets used in this study can be divided into upstream unlabeled and downstream labeled datasets. The summary

of these datasets is shown in Table 1.

Upstream datasets: Our large-scale multi-modal unlabeled dataset encompasses five modalities: Report, X-ray, CT, MRI, and Pathological imaging (Path.). The clinical reports and X-rays are sourced from the JPG version of the MIMIC-CXR 2.0.0 dataset [21]. CT scans are obtained from the DeepLesion dataset [49]. MRI scans are drawn from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database, including ADNI-1, ADNI-2, and ADNI-GO datasets [18]. Pathological images are collected from 7 projects: TCGA-THYM, TCGA-THCA, TCGA-BRCA, TCGA-UCEC, TCGA-UVM, TCGA-OV, TCGA-MESO within the cancer genome atlas (TCGA).

Detailed preprocessing steps for these upstream datasets are provided in Appendix B.1.

Downstream (DS) datasets: Nine datasets covering five different modalities are used for downstream evaluation. These include (1) Report: sentence classification (PudMed20k [11]); (2) X-ray: multi-class COVID-19 classification (ChestXR [1]) and COVID-19 infected region segmentation (QaTa-COV19-v2) [10]; (3) CT: COVID-19 classification (RICORD [43]) and liver and liver tumor segmentation (LiTS [3]); (4) MRI: vestibular schwannoma segmentation (VS [37]) and left atrium segmentation (LA [47]); (5) Pathological imaging: colorectal cancer and healthy tissue classification (NCH [24]) and gland segmentation (GlaS [38]). More details about these downstream datasets are in Appendix B.2.

4.2. Implementation details

The default pre-training modality order in our pre-training setup is Report, X-ray, CT, MRI, and Pathological imaging. We set the size of inputs as follows: the length of 1D vectors is set to 112, the size of 2D patches is set to 224×224 , and the size of 3D patches is set to $16 \times 192 \times 192$. Augmentation techniques such as random crop, resize, and flip were applied to 2D images, while mirror augmentation was used for 3D images. Following MAE [17], the AdamW optimizer [31] was employed, setting the batch size to 512 and the maximum number of epochs per modality at 300. For each pre-training stage, a warm-up strategy was employed during the initial 40 epochs to gradually increase the learning rate from 0 to $1.5e-4$, which was then decreased to 0 in the subsequent training according to the cosine schedule. The sampling number K for each cluster, as detailed in Section 3.3.1, was fixed at 5, indicating that we preserve 5% of the previous data for rehearsal purposes. During the fine-tuning phase, we continued with the AdamW optimizer but tailored the hyperparameters to suit the specific downstream tasks. Extensive details on these configurations were provided in Appendix B.2. Performance in classification tasks was evaluated based on the area under the receiver operator curve (AUC), accuracy (ACC), and F1 score. For

segmentation tasks, we employed the Dice similarity coefficient (Dice) and the 95% Hausdorff distance (HD) as metrics. To ensure the robustness of our conclusions, we averaged the results obtained using three different random seed values, including 0, 10, and 100.

4.3. Comparing to state-of-the-art SSL

We comprehensively compared the proposed MedCoSS over other pre-training paradigms, including five single-modal pre-training paradigms (each trained on one of five modalities) and six multi-modal pre-training paradigms, across nine downstream datasets. The multi-modal paradigms are categorized into joint pre-training (including joint SSL with dimension-shared decoders and joint SSL with modal-specific decoders) and continual pre-training (including EWC [25], ER [35], PackNet [32], and CaSSLe [14]). For all SSL paradigms, we employed masked modeling as the pretext task and ViT/B as the backbone with identical training epochs. The performance of all those paradigms is displayed in Table 2. It first reveals that a well-pre-trained model achieves significant performance improvement over training from scratch (TFS).

Comparing to single-modal pre-training paradigms:

While single-modal pre-trained models demonstrate excellence on downstream datasets with identical modalities, their performance significantly drops on datasets with different modalities, particularly those of different dimensions, indicating limited generalization. For instance, the model pre-trained on X-rays outperforms others on ChesXR and QaTa datasets but falls behind the model pre-trained on pathological images by 2.36% in ACC and 3.44% in F1 on the NCH dataset and the model pre-trained on CT scans by 3.58% in ACC and 2.81% in F1 on the RICORD dataset. In contrast, our model from MedCoSS exhibits more generalization performance, achieving the best or second-best results on eight out of nine datasets, with a marginal underperformance on the PudMed20k dataset compared to the model pre-trained on clinical reports.

Comparing to multi-modal pre-training paradigms:

We experimented with two joint SSL paradigms, *i.e.*, joint SSL with dimension-shared decoders (maintaining three decoders for masking modeling) and joint SSL with modal-specific decoders (maintaining five decoders for masking modeling). Despite showing better generalization performance than single-modal pre-trained models, the models from both joint SSL paradigms often underperform compared to single-modal models when fine-tuning on identical modalities and our MedCoSS. Furthermore, they exhibit limited scalability in integrating new knowledge from stream-like data. Among the other continual learning-based paradigms, including EWC, ER, PackNet, and CaSSLe, MedCoSS consistently achieves superior performance across eight datasets in all evaluated metrics.

Table 2. Generalization Performance of our MedCoSS, TFS, five single-modal pre-training models, and six multi-modal pre-training on nine datasets. The best and second-best results in each column are highlighted in red and blue, respectively. *: Joint SSL with dimension-shared decoders. †: Joint SSL with modal-specific decoders.

Method	PudMed20k (Report)			ChestXR (X-ray)			QaTa (X-ray)		RICORD (CT)			LiTS (CT)		VS (MRI)		LA (MRI)		NCH (Path.)			GlaS (Path.)	
	ACC	AUC	F1	ACC	AUC	F1	Dice	HD	ACC	AUC	F1	Dice	HD	Dice	HD	Dice	HD	ACC	AUC	F1	Dice	HD
TFS	82.41	94.82	76.45	82.36	94.32	81.55	73.94	38.38	73.02	79.69	81.25	60.77	67.40	84.27	58.15	85.94	16.62	84.93	97.37	79.15	86.20	51.88
<i>Single-Modal Pre-training</i>																						
Report	83.79	95.33	77.97	79.73	92.67	79.04	72.31	40.29	74.21	81.42	82.36	59.71	63.41	43.17	151.83	83.92	17.45	86.31	97.65	81.07	85.73	56.12
X-ray	82.40	94.80	76.55	95.55	99.21	95.08	80.27	27.47	76.98	84.77	83.28	62.96	55.79	87.07	26.86	87.90	10.85	93.26	99.04	90.57	88.65	47.82
CT	82.59	94.91	76.84	85.37	95.63	84.69	74.40	36.10	80.56	85.74	86.09	71.98	33.79	89.90	10.91	87.21	16.01	88.25	98.31	83.79	87.42	50.81
MRI	82.82	94.93	76.96	84.14	95.14	83.47	74.08	36.22	79.37	87.60	85.11	68.54	43.79	89.83	12.07	89.64	10.24	88.29	98.40	83.72	87.33	52.13
Path.	82.15	94.71	76.22	92.13	98.35	91.45	78.56	30.66	73.41	82.30	80.68	61.78	57.97	85.26	46.74	88.93	11.23	95.62	99.54	94.01	89.51	46.79
<i>Multi-Modal Pre-training</i>																						
Joint SSL*	84.10	95.52	78.30	92.24	98.33	91.59	78.61	30.29	77.78	85.89	84.13	70.59	38.24	89.73	15.33	89.41	13.84	94.63	99.27	92.93	87.86	51.20
Joint SSL†	83.94	95.33	78.13	91.54	98.02	90.85	78.29	30.27	78.97	86.65	84.90	69.60	44.17	89.41	15.83	88.84	14.94	94.48	99.30	92.69	87.83	52.08
EWC [25]	83.39	95.17	77.66	90.48	97.77	89.72	77.44	31.67	75.79	82.18	82.59	64.59	53.66	88.15	23.87	88.00	12.55	94.50	99.37	92.61	88.16	50.81
ER [35]	83.75	95.43	78.15	85.32	95.12	84.43	75.03	36.59	75.79	83.51	82.03	69.53	41.86	88.61	25.47	87.37	14.42	92.21	99.14	89.54	87.61	53.10
PackNet [32]	82.88	95.01	77.06	86.14	95.98	85.28	75.69	35.98	71.43	83.75	79.63	61.89	62.12	83.92	54.07	88.87	12.55	92.42	99.13	90.00	86.96	52.81
CaSSLLe [14]	82.93	95.12	77.04	91.66	98.07	90.99	78.19	31.23	76.59	82.85	84.32	67.04	47.17	87.94	21.47	89.66	10.67	95.01	99.44	93.06	89.12	47.93
MedCoSS	83.59	95.38	77.87	94.31	98.83	93.77	78.98	29.43	83.33	88.74	87.87	72.01	36.50	90.12	7.80	90.46	9.55	95.76	99.51	94.01	89.13	46.69

Table 3. Ablation studies of proposed feature distillation (FD), k -means sampling, and intra-modal mixup (IMM) strategies on five datasets. The baseline is the vanilla ER paradigm with a random sampling strategy. The best result in each column is highlighted with bold.

FD	k -means	IMM	PudMed20k (Report)			ChestXR (X-ray)			LiTS (CT)		VS (MR)		NCH (Path.)			Average (Cls)			Average (Seg)	
			ACC	AUC	F1	ACC	AUC	F1	Dice	HD	Dice	HD	ACC	AUC	F1	ACC	AUC	F1	Dice	HD
			83.75	95.43	78.15	85.32	95.12	84.43	69.53	41.86	88.61	25.47	92.21	99.14	89.54	87.10	96.57	84.04	79.07	33.67
✓			83.53	95.25	77.68	92.16	98.27	91.43	71.20	37.56	89.89	11.70	95.17	99.42	93.16	90.29	97.65	87.43	80.54	24.63
	✓		83.65	95.36	77.94	84.72	95.22	83.85	70.61	38.31	88.12	31.46	92.43	99.15	89.96	86.93	96.58	83.92	79.36	34.88
✓	✓		83.50	95.35	77.77	93.76	98.76	93.21	71.56	38.50	89.95	10.56	95.47	99.57	93.88	90.91	97.89	88.28	80.76	24.53
✓	✓	✓	83.59	95.38	77.87	94.31	98.83	93.77	72.01	36.50	90.12	7.80	95.76	99.51	94.01	91.22	97.90	88.55	81.07	22.15

4.4. Ablation studies

To validate the effectiveness of the proposed feature distillation (FD) strategy, k -means sampling strategy, and intra-modal mixup (IMM) strategy, ablation studies were conducted on five downstream datasets, each representing a different modality. We calculated the average values of each metric for segmentation and classification tasks, respectively, to provide a holistic view of each strategy’s effect. The results are displayed in Table 3, where the first row reports the performance of the ER paradigm with a random sampling strategy [35]. Based on these results, we have three conclusions: (1) adopting the FD strategy instead of the defined pretext task markedly enhances knowledge retention and maintains the model’s plasticity to new knowledge. As a result, the FD strategy improves average ACC, AUC, F1, Dice, and HD by 3.19%, 1.08%, 3.39%, 1.47%, and 9.04, respectively; (2) while the k -means strategy alone shows an unclear impact, its combination with the FD strategy yields gains in generalization across all datasets; (3) the integration of the IMM strategy furthers this improvement.

4.5. Impact of buffer size on model performance

In supervised continual learning, a large buffer size usually equates to improved performance due to reduced forgetting [5, 34, 42]. For the multi-modal pre-training, a larger buffer

will bring the risk of intensified *modal data collision* and computational cost, hindering current modality representation learning. To find an optimal balance, we experimented with ER and MedCoSS paradigms using sampling ratios of 1%, 5%, and 10% across five datasets, as shown in Fig. 4. For the ER paradigm, a notable degradation in performance is observed when increasing the sampling ratio from 1% to 5%. On the contrary, the model from our MedCoSS demonstrates significant improvements across all datasets with the sampling ratio from 1% to 5% but shows performance degradation on PudMed20k, LiTS, VS, and NCH datasets except for the ChestXR dataset when the sampling ratio is raised to 10%. Finally, we selected a 5% sampling ratio for the rehearsal buffer in MedCoSS.

4.6. Knowledge retention in sequential pre-training

While successfully mitigating the issue of the *modal data collision*, the sequential pre-training paradigm inherently faces the risk of catastrophic forgetting. To address this, our MedCoSS introduces three CL techniques. Fig. 5 illustrates the performance changes of both MedCoSS and a sequential SSL, where no knowledge retention strategies are employed, during the multi-stage pre-training process. This comparison distinctly shows that MedCoSS effectively preserves previous knowledge, exhibiting minor performance degradation over successive stages. In contrast, sequen-

Table 4. Performance of different pre-training orders on five datasets.

Order	PudMed20k (Report)			ChestXR (X-ray)			LiTS (CT)		VS (MRI)		NCH (Path.)			Average (Cls)			Average (Seg)	
	ACC	AUC	F1	ACC	AUC	F1	Dice	HD95	Dice	HD95	ACC	AUC	F1	ACC	AUC	F1	Dice	HD95
Report, X-ray, CT, MRI, Path. (Default order)	83.59	95.38	77.87	94.31	98.83	93.77	72.01	36.50	90.12	7.80	95.76	99.51	94.01	91.22	97.90	88.55	81.07	22.15
Path., MRI, Report, CT, X-ray	83.63	95.27	77.79	95.70	99.25	95.24	71.43	36.58	89.29	11.93	93.75	99.30	91.19	91.03	97.94	88.07	80.36	24.26




Figure 4. Comparison of our MedCoSS and ER with 1%, 5%, and 10% buffer sampling ratio on five datasets.




Figure 5. Comparison of our MedCoSS and sequential SSL for knowledge retention on the PudMed20k and ChestXR datasets. The sequential SSL means sequential pre-training without any knowledge retention strategies.

tial SSL is observed to have significant sensitivity to catastrophic forgetting. For instance, after the CT pre-training stage, sequential SSL experiences a stark drop, dropping from 95.0% to 89.7% in terms of ACC. MedCoSS, on the other hand, maintains a more stable performance, with accuracy slightly decreasing from 95.3% to 94.4%.

4.7. Pre-training order of modality

We experimented with another set of modality pre-training orders: Pathological imaging, MRI, Report, CT, and X-ray. The results in Table 4 demonstrate that MedCoSS still achieves high generalization performance after the order change. Moreover, the modality placed last usually achieves better downstream performance than other positions. Due to the mitigated risk of knowledge forgetting in our MedCoSS, the performance difference between being last and other positions is insubstantial.

5. Conclusion

In this paper, we conduct an in-depth exploration of multi-modal SSL, specifically focusing on the practical but challenging unpaired data. We rethink the limitations in the prevalent joint SSL paradigm. Our experiments observed a

so-called *modal data collision*, which seriously hinders the representation learning across modalities. To alleviate this issue and ensure the scalability of the pre-trained model, we propose MedCoSS, a sequential pre-training paradigm that approaches multi-modal SSL through a multi-stage manner, assigning each stage to a specific modality. MedCoSS introduces a customized continual learning technique to alleviate inherent catastrophic forgetting within the sequential training scheme. We conduct experiments on a large-scale multi-modal unlabeled dataset, including the modality of Report, X-ray, CT, MRI, and Pathological imaging, and demonstrate the superiority of our MedCoSS on nine downstream datasets. In our future work, we plan to delve deeper into multi-modal medical pre-training within data streams, where each stage needs to handle specific multi-modal data. **Clinical impact:** Although MedCoSS may not currently achieve state-of-the-art performance across all tasks, its significant potential in advancing the multi-modal pre-trained medical universal model should be emphasized. The unique advantage of MedCoSS lies in its cost-effective scalability, which facilitates the seamless integration of new knowledge derived from real-world, stream-like data sources. This makes MedCoSS become a dynamic and adaptable paradigm, well-suited to keeping pace with the rapidly expanding variety of clinical multi-modal data.

6. Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 62171377, Shenzhen Science and Technology Program under Grants JCYJ20220530161616036, Ningbo Clinical Research Center for Medical Imaging under Grant 2021L003 (Project 2022LYKFZD06), and Innovation Foundation for Doctor Dissertation of Northwestern Polytechnical University. Yutong Xie was supported by the Centre for Augmented Reasoning (CAR) project.

References

- [1] Moulay A. Akhloufi and Mohamed Chetoui. Chest XR COVID-19 detection. <https://cxr-covid19.grand-challenge.org/>, 2021. Online; accessed September 2021. 6
- [2] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, et al. Big self-supervised models advance medical image classification. In *ICCV*, pages 3478–3488, 2021. 2, 3
- [3] Patrick Bilic, Patrick Christ, Hongwei Bran Li, Eugene Vorontsov, Avi Ben-Cohen, Georgios Kaissis, Adi Szeskin, Colin Jacobs, Gabriel Efrain Humpire Mamani, Gabriel Chartrand, et al. The liver tumor segmentation benchmark (lits). *MIA*, 84:102680, 2023. 6
- [4] Behzad Bozorgtabar, Dwarikanath Mahapatra, Guillaume Vray, and Jean-Philippe Thiran. Salad: Self-supervised aggregation learning for anomaly detection on x-rays. In *MICCAI*, pages 468–478. Springer, 2020. 2, 3
- [5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *NeurIPS*, 33: 15920–15930, 2020. 3, 7
- [6] Yu Cai, Hao Chen, Xin Yang, Yu Zhou, and Kwang-Ting Cheng. Dual-distribution discrepancy with self-supervised refinement for anomaly detection in medical images. *MIA*, 86:102794, 2023. 2, 3
- [7] Zhiyuan Cai, Li Lin, Huaqing He, and Xiaoying Tang. Uni4eye: Unified 2d and 3d self-supervised pre-training via masked image modeling transformer for ophthalmic image classification. In *MICCAI*, pages 88–98. Springer, 2022. 2, 3
- [8] Yixiong Chen, Chunhui Zhang, Chris HQ Ding, and Li Liu. Generating and weighting semantically consistent sample pairs for ultrasound contrastive learning. *IEEE TMI*, 2022. 2, 3
- [9] Zhihong Chen, Yuhao Du, Jinpeng Hu, Yang Liu, Guanbin Li, Xiang Wan, and Tsung-Hui Chang. Multi-modal masked autoencoders for medical vision-and-language pre-training. In *MICCAI*, pages 679–689. Springer, 2022. 2, 3
- [10] Aysen Degerli, Serkan Kiranyaz, Muhammad EH Chowdhury, and Moncef Gabbouj. Osegnet: Operational segmentation network for covid-19 detection using chest x-ray images. In *ICIP*, pages 2306–2310. IEEE, 2022. 6
- [11] Franck Dernoncourt and Ji Young Lee. Pubmed 200k rct: a dataset for sequential sentence classification in medical abstracts. *arXiv preprint arXiv:1710.06071*, 2017. 6
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2, 4
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 2, 4
- [14] Enrico Fini, Victor G Turrisi Da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *CVPR*, pages 9621–9630, 2022. 3, 6, 7
- [15] Florin C Ghesu, Bogdan Georgescu, Awais Mansoor, Youngjin Yoo, Dominik Neumann, Pragneshkumar Patel, RS Vishwanath, James M Balter, Yue Cao, Sasa Grbic, et al. Self-supervised learning from 100 million medical images. *arXiv preprint arXiv:2201.01283*, 2022. 2, 3
- [16] Fatemeh Haghighi, Mohammad Reza Hosseinzadeh Taher, Michael B Gotway, and Jianming Liang. Dira: Discriminative, restorative, and adversarial learning for self-supervised medical image analysis. In *CVPR*, pages 20824–20834, 2022. 2, 3
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 2, 4, 6
- [18] Clifford R Jack Jr, Matt A Bernstein, Nick C Fox, Paul Thompson, Gene Alexander, Danielle Harvey, Bret Borowski, Paula J Britson, Jennifer L. Whitwell, Chadwick Ward, et al. The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *J. Magn. Reson. Imaging*, 27(4): 685–691, 2008. 6
- [19] Yankai Jiang, Mingze Sun, Heng Guo, Xiaoyu Bai, Ke Yan, Le Lu, and Minfeng Xu. Anatomical invariance modeling and semantic alignment for self-supervised learning in 3d medical image analysis. In *ICCV*, pages 15859–15869, 2023. 2, 3
- [20] Ting Jin, Xingran Xie, Renjie Wan, Qingli Li, and Yan Wang. Gene-induced multimodal pre-training for image-omic classification. In *MICCAI*, pages 508–517. Springer, 2023. 2
- [21] Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Roger G Mark, and Steven Horng. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317, 2019. 6
- [22] András Kalapos and Bálint Gyires-Tóth. Self-supervised pretraining for 2d medical image segmentation. In *ECCV*, pages 472–484. Springer, 2022. 2, 3
- [23] Mingu Kang, Heon Song, Seonwook Park, Donggeun Yoo, and Sérgio Pereira. Benchmarking self-supervised learning on diverse pathology datasets. In *CVPR*, pages 3344–3354, 2023. 2, 3
- [24] Jakob Nikolas Kather, Niels Halama, and Alexander Marx. 100,000 histological images of human colorectal cancer and healthy tissue. *Zenodo10*, 5281, 2018. 6
- [25] 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. *Proc. Natl. Acad. Sci.*, 114(13):3521–3526, 2017. 3, 6, 7
- [26] Xiang Lan, Dianwen Ng, Shenda Hong, and Mengling Feng. Intra-inter subject self-supervised learning for multivariate cardiac signals. In *AAAI*, pages 4532–4540, 2022. 2, 3
- [27] Xiaomeng Li, Xiaowei Hu, Xiaojuan Qi, Lequan Yu, Wei Zhao, Pheng-Ann Heng, and Lei Xing. Rotation-oriented

- collaborative self-supervised learning for retinal disease diagnosis. *IEEE TMI*, 40(9):2284–2294, 2021. 2, 3
- [28] Guoliang Lin, Hanlu Chu, and Hanjiang Lai. Towards better plasticity-stability trade-off in incremental learning: A simple linear connector. In *CVPR*, pages 89–98, 2022. 3
- [29] Bo Liu, Donghuan Lu, Dong Wei, Xian Wu, Yan Wang, Yu Zhang, and Yefeng Zheng. Improving medical vision-language contrastive pretraining with semantics-aware triage. *IEEE TMI*, 2023. 2, 3
- [30] Che Liu, Sibao Cheng, Chen Chen, Mengyun Qiao, Weitong Zhang, Anand Shah, Wenjia Bai, and Rossella Arcucci. M-flag: Medical vision-language pre-training with frozen language models and latent space geometry optimization. In *MICCAI*, pages 637–647. Springer, 2023. 2, 3
- [31] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. 2018. 6
- [32] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *CVPR*, pages 7765–7773, 2018. 3, 6, 7
- [33] Jiahong Ouyang, Qingyu Zhao, Ehsan Adeli, Greg Zaharchuk, and Kilian M Pohl. Self-supervised learning of neighborhood embedding for longitudinal mri. *MIA*, 82:102571, 2022. 2, 3
- [34] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, pages 524–540. Springer, 2020. 7
- [35] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *ICLR*, 2019. 2, 3, 6, 7
- [36] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *ACL*, 2015. 4
- [37] Jonathan Shapey, Aaron Kujawa, Reuben Dorent, Guotai Wang, Alexis Dimitriadis, Diana Grishchuk, Ian Paddock, Neil Kitchen, Robert Bradford, Shakeel R Saeed, et al. Segmentation of vestibular schwannoma from mri, an open annotated dataset and baseline algorithm. *Scientific Data*, 8(1):286, 2021. 6
- [38] Korsuk Sirinukunwattana, Josien PW Pluim, Hao Chen, Xiaojuan Qi, Pheng-Ann Heng, Yun Bo Guo, Li Yang Wang, Bogdan J Matuszewski, Elia Bruni, Urko Sanchez, et al. Gland segmentation in colon histology images: The glas challenge contest. *MIA*, 35:489–502, 2017. 6
- [39] Mohammad Reza Hosseinzadeh Taher, Fatemeh Haghighi, Michael B Gotway, and Jianming Liang. Caid: Context-aware instance discrimination for self-supervised learning in medical imaging. In *MIDL*, pages 535–551. PMLR, 2022. 2, 3
- [40] Aiham Taleb, Christoph Lippert, Tassilo Klein, and Moin Nabi. Multimodal self-supervised learning for medical image analysis. In *IPMI*, pages 661–673. Springer, 2021. 2, 3
- [41] Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *CVPR*, pages 20730–20740, 2022. 2, 3
- [42] Rishabh Tiwari, Krishnateja Killamsetty, Rishabh Iyer, and Pradeep Shenoy. Gcr: Gradient coreset based replay buffer selection for continual learning. In *CVPR*, pages 99–108, 2022. 7
- [43] Emily B Tsai, Scott Simpson, Matthew P Lungren, Michelle Hershman, Leonid Roshkovan, Errol Colak, Bradley J Erickson, George Shih, Anouk Stein, Jayashree Kalpathy-Cramer, et al. The rsna international covid-19 open radiology database (ricord). *Radiology*, 299(1):E204–E213, 2021. 6
- [44] Xiyue Wang, Sen Yang, Jun Zhang, Minghui Wang, Jing Zhang, Junzhou Huang, Wei Yang, and Xiao Han. Transpath: Transformer-based self-supervised learning for histopathological image classification. In *MICCAI*, pages 186–195. Springer, 2021. 2, 3
- [45] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Towards generalist foundation model for radiology. *arXiv preprint arXiv:2308.02463*, 2023. 3
- [46] Yutong Xie, Jianpeng Zhang, Yong Xia, and Qi Wu. Unimiss: Universal medical self-supervised learning via breaking dimensionality barrier. In *ECCV*, pages 558–575. Springer, 2022. 2, 3
- [47] Zhaohan Xiong, Qing Xia, Zhiqiang Hu, Ning Huang, Cheng Bian, Yefeng Zheng, Sulaiman Vesal, Nishant Ravikumar, Andreas Maier, Xin Yang, et al. A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *MIA*, 67:101832, 2021. 6
- [48] Bin Yan and Mingtao Pei. Clinical-bert: Vision-language pre-training for radiograph diagnosis and reports generation. In *AAAI*, pages 2982–2990, 2022. 2, 3
- [49] Ke Yan, Xiaosong Wang, Le Lu, and Ronald M Summers. Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *J. Med. Imaging*, 5(3):036501–036501, 2018. 6
- [50] Xiangyi Yan, Junayed Naushad, Chenyu You, Hao Tang, Shanlin Sun, Kun Han, Haoyu Ma, James S Duncan, and Xiaohui Xie. Localized region contrast for enhancing self-supervised learning in medical image segmentation. In *MICCAI*, pages 468–478. Springer, 2023. 2, 3
- [51] Pengshuai Yang, Zhiwei Hong, Xiaoxu Yin, Chengzhan Zhu, and Rui Jiang. Self-supervised visual representation learning for histopathological images. In *MICCAI*, pages 47–57. Springer, 2021. 2, 3
- [52] Yiwen Ye, Jianpeng Zhang, Ziyang Chen, and Yong Xia. Desd: Self-supervised learning with deep self-distillation for 3d medical image segmentation. In *MICCAI*, pages 545–555. Springer, 2022. 2, 3
- [53] Chuyan Zhang, Hao Zheng, and Yun Gu. Dive into the details of self-supervised learning for medical image analysis. *MIA*, 89:102879, 2023. 2, 3
- [54] Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Weidi Xie, and Yanfeng Wang. Knowledge-enhanced visual-language pre-training on chest radiology images. *Nature Communications*, 14(1):4542, 2023. 2, 3

- [55] Bo Zhou, Jo Schlemper, Neel Dey, Seyed Sadegh Mohseni Salehi, Kevin Sheth, Chi Liu, James S Duncan, and Michal Sofka. Dual-domain self-supervised learning for accelerated non-cartesian mri reconstruction. *MIA*, 81:102538, 2022. [2](#), [3](#)
- [56] Hong-Yu Zhou, Chixiang Lu, Sibeï Yang, Xiaoguang Han, and Yizhou Yu. Preservational learning improves self-supervised medical image models by reconstructing diverse contexts. In *ICCV*, pages 3499–3509, 2021. [2](#), [3](#)
- [57] Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B Gotway, and Jianming Liang. Models genesis. *MIA*, 67:101840, 2021. [2](#), [3](#)
- [58] Jiuwen Zhu, Yuexiang Li, Yifan Hu, Kai Ma, S Kevin Zhou, and Yefeng Zheng. Rubik’s cube+: A self-supervised feature learning framework for 3d medical image analysis. *MIA*, 64: 101746, 2020. [2](#), [3](#)
- [59] Xizhou Zhu, Jinguo Zhu, Hao Li, Xiaoshi Wu, Hongsheng Li, Xiaohua Wang, and Jifeng Dai. Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. In *CVPR*, pages 16804–16815, 2022. [4](#)