



Self-supervised learning for classifying paranasal anomalies in the maxillary sinus

Debayan Bhattacharya^{1,2} · Finn Behrendt¹ · Benjamin Tobias Becker² · Lennart Maack¹ · Dirk Beyersdorff³ · Elina Petersen⁴ · Marvin Petersen⁵ · Bastian Cheng⁵ · Dennis Eggert² · Christian Betz² · Anna Sophie Hoffmann² · Alexander Schlaefer¹

Received: 6 December 2023 / Accepted: 1 May 2024 / Published online: 8 June 2024
© The Author(s) 2024

Abstract

Purpose Paranasal anomalies, frequently identified in routine radiological screenings, exhibit diverse morphological characteristics. Due to the diversity of anomalies, supervised learning methods require large labelled dataset exhibiting diverse anomaly morphology. Self-supervised learning (SSL) can be used to learn representations from unlabelled data. However, there are no SSL methods designed for the downstream task of classifying paranasal anomalies in the maxillary sinus (MS).

Methods Our approach uses a 3D convolutional autoencoder (CAE) trained in an unsupervised anomaly detection (UAD) framework. Initially, we train the 3D CAE to reduce reconstruction errors when reconstructing normal maxillary sinus (MS) image. Then, this CAE is applied to an unlabelled dataset to generate coarse anomaly locations by creating residual MS images. Following this, a 3D convolutional neural network (CNN) reconstructs these residual images, which forms our SSL task. Lastly, we fine-tune the encoder part of the 3D CNN on a labelled dataset of normal and anomalous MS images.

Results The proposed SSL technique exhibits superior performance compared to existing generic self-supervised methods, especially in scenarios with limited annotated data. When trained on just 10% of the annotated dataset, our method achieves an area under the precision-recall curve (AUPRC) of 0.79 for the downstream classification task. This performance surpasses other methods, with BYOL attaining an AUPRC of 0.75, SimSiam at 0.74, SimCLR at 0.73 and masked autoencoding using SparK at 0.75.

Conclusion A self-supervised learning approach that inherently focuses on localizing paranasal anomalies proves to be advantageous, particularly when the subsequent task involves differentiating normal from anomalous maxillary sinuses. Access our code at <https://github.com/mtec-tuhh/self-supervised-paranasal-anomaly>.

Keywords Paranasal anomaly · Self-supervised learning · Maxillary sinus · CNN · Classification

Anna Sophie Hoffmann and Alexander Schlaefer have contributed equally.

✉ Debayan Bhattacharya
debayan.bhattacharya@tuhh.de

¹ Institute of Medical Technology and Intelligent Systems, Technische Universitaet Hamburg, Hamburg, Germany

² Department of Otorhinolaryngology, Head and Neck Surgery and Oncology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

³ Clinic and Polyclinic for Diagnostic and Interventional Radiology and Nuclear Medicine, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

Introduction

The paranasal sinuses, air-filled spaces within the craniofacial complex, vary significantly and include the maxillary, frontal, sphenoid, and ethmoid sinuses [1]. Common pathologies like retention cysts, polyps, and mucosal thickening are identifiable through radiological screenings [2–4]. However, their diagnosis is challenging due to their incidental nature and the variability in sinus appearance [5]. Research

⁴ Population Health Research Department, University Heart and Vascular Center, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

⁵ Clinic and Polyclinic for Neurology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

underscores their prevalence and the importance of accurate diagnosis in patient care [6]. 3D imaging from computed tomography (CT) and magnetic resonance images (MRI) is vital for precise diagnosis, and misdiagnosis can lead to patient distress and increased healthcare costs [7, 8]. The anatomical variability of the sinuses [9] necessitates careful application of deep learning for reliable diagnoses.

Convolutional neural networks (CNNs) are recognized for diagnosing paranasal pathologies, evidenced in sinusitis classification [10, 11], differentiating inverted papilloma from carcinomas [12], and detecting MS fungal ball and chronic rhinosinusitis in CT scans [13]. Prior studies have explored contrastive learning and cross-entropy loss for MS anomaly classification [14], and MS extraction techniques from MRI [15]. However, all of the aforementioned methods use supervised learning. Given the difficulty in obtaining well-labelled datasets in clinical settings [16], and the relative ease of acquiring unlabelled data, self-supervised learning (SSL), which learns representations from unlabelled data to improve the downstream task, has not yet been explored for paranasal anomaly classification. SSL efficiently utilizes unlabelled data through tasks like nonlinear compression [17, 18], denoising [19], feature alignment from augmented images [20–22] and inpainting masked regions of images [23]. However, these methods are designed to improve the performance of models exposed to 2D natural images. Hence, they lack a specific focus on enhancing MS anomaly classification from 3D MRI. Our aim is to design an SSL task that enables the models trained on it to achieve maximum data efficiency in classifying paranasal anomalies. We hypothesize anomaly segmentation within MS could be a good SSL task. Without ground truth segmentation masks, we use a UAD framework, applied in brain [24, 25] and paranasal anomaly detection [26], to localize MS anomalies. A 3D convolutional autoencoder (CAE) trained on a labelled *normal* dataset is used to reconstruct MS volumes and localize anomalies in an unlabelled dataset by failing to reconstruct anomalies leading to reconstruction errors. These errors, serving as pseudo segmentation masks are used in the SSL task to localize anomalies. We investigate if a 3D CNN, predicting these errors as SSL task, can improve feature discrimination between anomalous and normal MS in our labelled dataset. Our SSL task leverages available normal MS data, essential for supervised downstream task training.

Overall, our main contributions can be summed up as follows:

- We present a self-supervised method that improves the downstream classification of normal vs anomalous MS. Our self-supervision task explicitly learns to coarsely localize anomalies by reconstructing the residual volumes generated through the UAD-trained autoencoder. This distinguishes our approach from the compared

methods, where anomaly localization is not a primary focus for the self-supervision task.

- Our self-supervised method effectively utilizes labelled healthy MS data reserved for downstream tasks. Hence, we explore how varying the CAE training set impacts downstream classification performance.
- We investigate post-processing strategies and loss function used in the self-supervision task for learning better transferable features for the downstream task.

Methods

Description of dataset

As part of the Hamburg City Health Study (HCHS) [27], cranial MRI scans were obtained from individuals aged 45–74 years to evaluate neuroradiological parameters. The scans were acquired using fluid attenuated inversion recovery (FLAIR) sequences in the NIfTI format at the University Medical Center Hamburg-Eppendorf. The MRI scans had a resolution of 173 mm x 319 mm x 319 mm. The labelled dataset consisted of 1067 patients. Among the patients, 489 exhibited no pathologies in their left and right MS, while 578 had at least one MS presenting polyp, cyst or mucosal thickening pathology. All these anomalies were grouped into the "anomaly" class. Our unlabelled dataset consists of 1559 patient MRIs. The diagnoses were established by two ENT specialists and one radiologist specialized in ENT. Figure 1b shows coronal slices highlighting the diverse set of anomalies that are present in our dataset.

Dataset preprocessing

In our dataset preprocessing, as outlined in previous work [14, 15], we first align MRIs with a fixed sample from our dataset. Centroid locations of left and right MS regions were recorded for 20 patients, guiding the extraction of MS volumes from larger cranial MRIs. This step isolates the relevant MS volumes for our task of classifying healthy and anomalous MS. We then used the mean centroid location from these 20 recordings to extract left and right MS volumes from all cranial MRIs in our dataset. The extracted volumes, sized 64 mm x 64 mm x 64 mm, cover the entire MS. Figure 1a illustrates this extraction process.

Each cranial MRI yielded one left and one right MS volume. To enhance symmetry, right MS volumes were horizontally flipped to match the left ones. All volumes were normalized to an intensity range of 0 to 1. We employed five-fold cross-validation for evaluation, ensuring diverse labelled datasets (10%, 20%, 40%, 60%, 80%) maintain the anomaly-to-normal ratio. The separation of training, validation, and test sets was strictly maintained, with left or right MS vol-

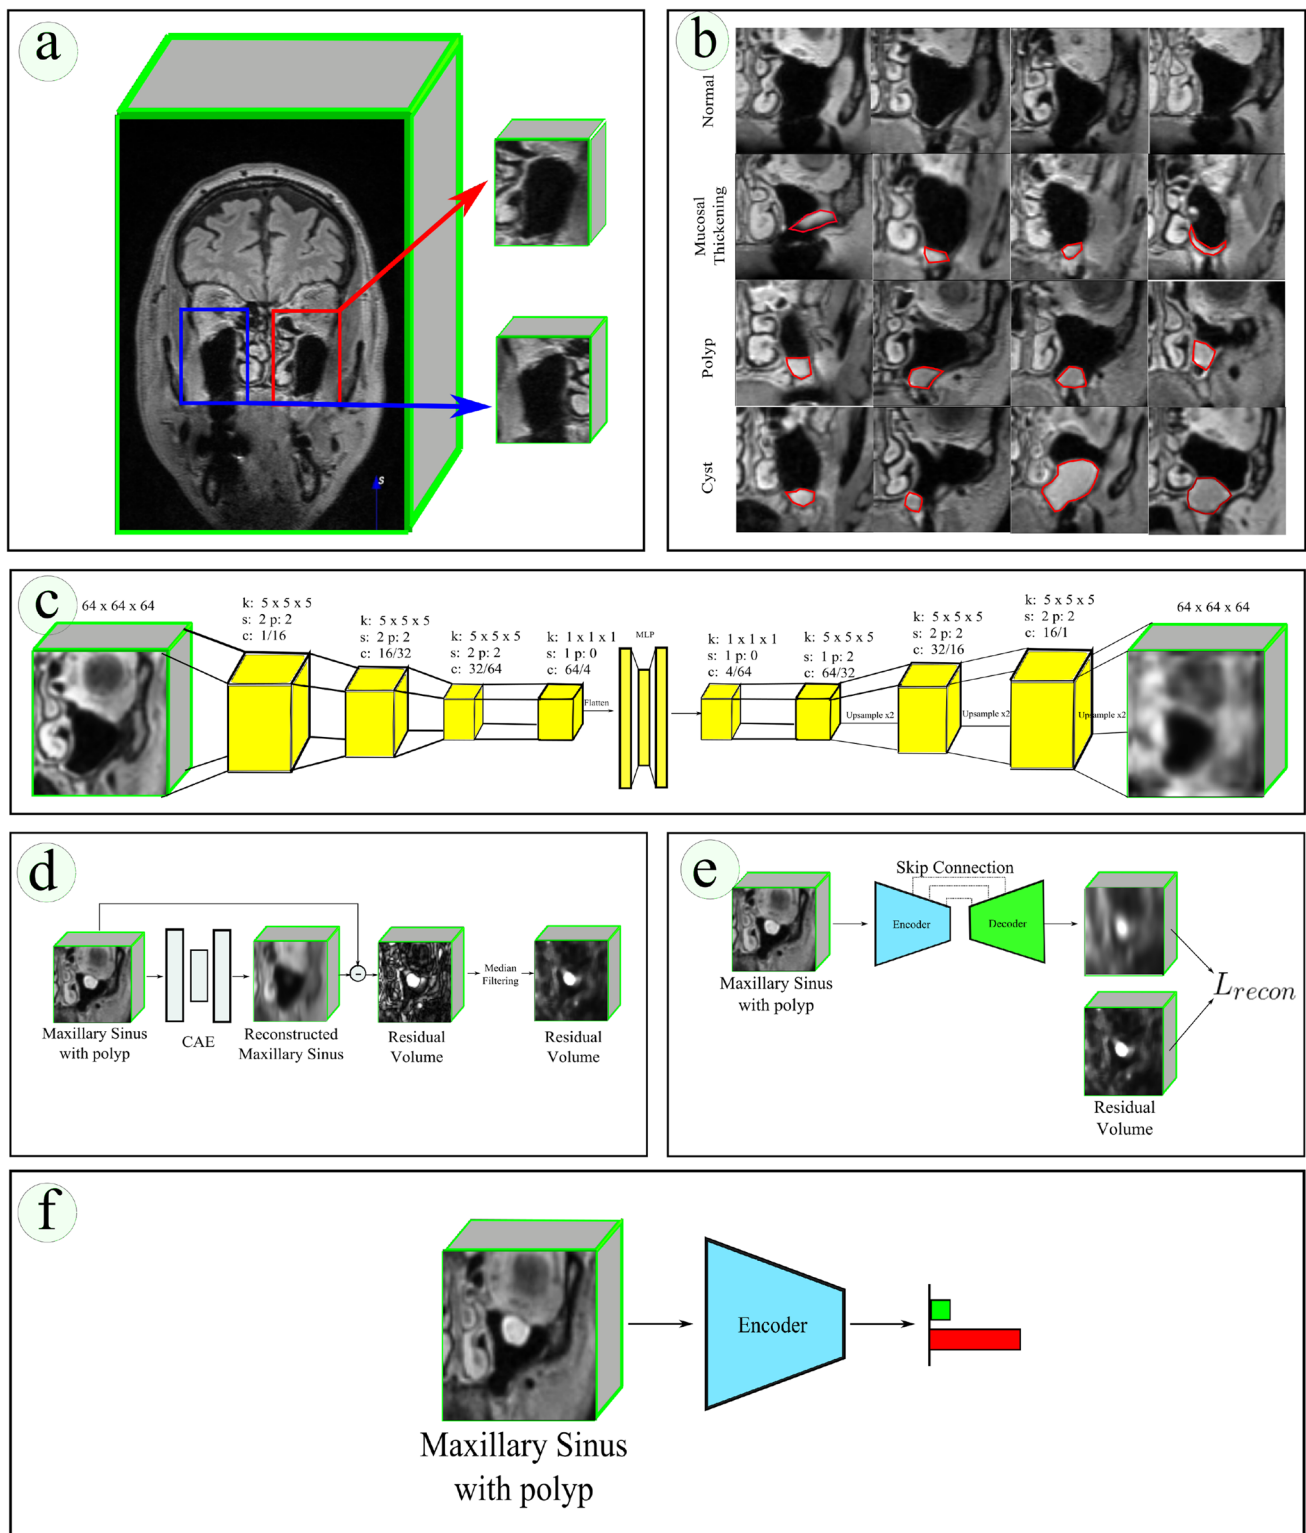


Fig. 1 **a** Extraction of MS volumes from cranial MRI, **b** Exemplary coronal images of normal MS volume and MS with mucosal thickening, polyp and cyst anomaly, **c** Our CAE architecture. Here, k refers to kernel size, s refers to stride, p refers to padding, c refers to channel where, for example, $1/16$ refers to input channel of 1 and output channel of 16. Each stage of the encoder and decoder is formed using 3D con-

volution followed by batch normalization and leaky ReLU. Upsample refers to trilinear upsampling. **d** Generation of residual volume required for the self-supervision task using our CAE, **e** Our self-supervision task where the encoder and decoder is trained to reconstruct the residual volume, **f** Downstream task where the self-supervision trained encoder is trained to classify between normal and anomalous MS

Table 1 Statistics of our labelled dataset D_l

Class	Training set	Validation set	Test set
# Normal MS	708	176	380
# Anomalous MS	487	122	261

umes from the same patient assigned to only one set. Table 1 details our dataset division across these sets.

Architecture

Our CAE, depicted in Fig. 1c, uses 3D convolutional operations with a latent bottleneck dimension of 512. The CNN architecture is U-Net inspired, featuring a 3D ResNet18 encoder $E(\cdot)$ [28] with four stages and channel dimensions of 64, 128, 256, and 512. The decoder $D(\cdot)$ mirrors the encoder, with reverse channel dimensions and trilinear upsampling. Skip connections are used to pass encoder features to the decoder. For Bootstrap your own latent (BYOL), SimSiam, and SimCLR training, only the encoder $E(\cdot)$ is used, with an MLP attached to project the final layer features to a dimension of 512.

Autoencoder training and inference on unlabelled dataset

Consider D_l to be our labelled dataset containing normal and anomalous MS and D_u to be our unlabelled dataset. Further, let $D_l^n \subset D_l$ be a dataset consisting of only normal MS volumes. Let $x \in \mathbb{R}^{64 \times 64 \times 64}$ be an MS volume in D_l . Let the autoencoder be represented as $A(\cdot)$ such that $x' = A(x)$ represents the reconstructed MS volume. We train the autoencoder using L1 reconstruction loss which may be written as $\|x - x'\|$ on D_l^n . Once trained, we use the autoencoder $A(\cdot)$ to generate residual volumes on D_u . Figure 1d illustrates our residual volume generation method.

Transfer learning

Since transfer learning (TL) is a method to achieve data efficiency, we also trained our models initialized with transfer learning weights. However, since our downstream task involves MRI and is in 3D domain, ImageNet [29] weights may not be appropriate. Hence, the model weights we utilized as initial weights were obtained through training on eight diverse public 3D segmentation datasets, covering both MRI and CT modalities. We believe these weights are more suitable than those derived from natural image training and therefore employed them as the basis for our 3D CNN. For further information on the transfer learning model, please see the GitHub repository.¹

¹ <https://github.com/Tencent/MedicalNet>.

Self-supervised training

With the residual volumes generated for D_u , we train $E(\cdot)$ and $D(\cdot)$ to reconstruct the residual volumes again. This, in effect, makes the encoder and decoder learn features relevant for anomaly localization within the unlabelled dataset D_u . We train $E(\cdot)$ and $D(\cdot)$ using L_{recon} which in our case is binary cross-entropy (BCE) loss. Figure 1e illustrates our self-supervised training task. We evaluated our self-supervised learning method against autoencoder (AE), denoising autoencoder (DAE), BYOL, SimSiam, SimCLR and sparse masked modelling with hierarchy (SparK). These methods use similar encoders $E(\cdot)$ and decoders $D(\cdot)$, with BYOL, SimSiam, and SimCLR employing an additional MLP for feature projection. Pretraining with the SparK framework requires sparse encoder $E'(\cdot)$ and a special light decoder which contains 3 convolutional blocks and 3 upsampling blocks [23]. Patch size $8 \times 8 \times 8$ and masking ratio of 60% was used during pretraining. Detailed description and implementation details of our state-of-the-art (SOTA) SSL methods is provided in the supplementary material section 1-7. More details about the other masking ratios and patch sizes tested for SparK can be found in the supplementary material section 11.

Finetuning

Having successfully trained the $E(\cdot)$ and $D(\cdot)$ using self-supervision, we move onto the finetuning phase. We discard $D(\cdot)$ and focus on training $E(\cdot)$ by leveraging samples from the labelled dataset D_l . For TL models, we initialize $E(\cdot)$ with transfer learning weights. Next, we introduce a MLP as an additional component, responsible for projecting the encoder features from their original dimension of 512 to an intermediate dimension of 256. Subsequently, the MLP maps these features to a final dimension of 2, corresponding to the number of classes. We finetune $E(\cdot)$ using BCE loss.

Figure 2 illustrates the data processing pipeline and elucidates how the different components fit into our overall method.

Implementation details

Our PyTorch and PyTorch Lightning-based code accommodates a maximum batch size of 256 on NVIDIA A6000 with 48GB VRAM for self-supervised pretraining. We optimize models using LARS [30] with a learning rate of 0.2 across 500 epochs, incorporating a 20-epoch linear warmup and cosine annealing. For finetuning, AdamW [31] is employed with a constant rate of $1e-4$ for 100 epochs at a batch size of 16. Models yielding the lowest validation loss are preserved for final evaluation with the test set. The CAE was trained on 708 normal MS volume samples without augmentation. For

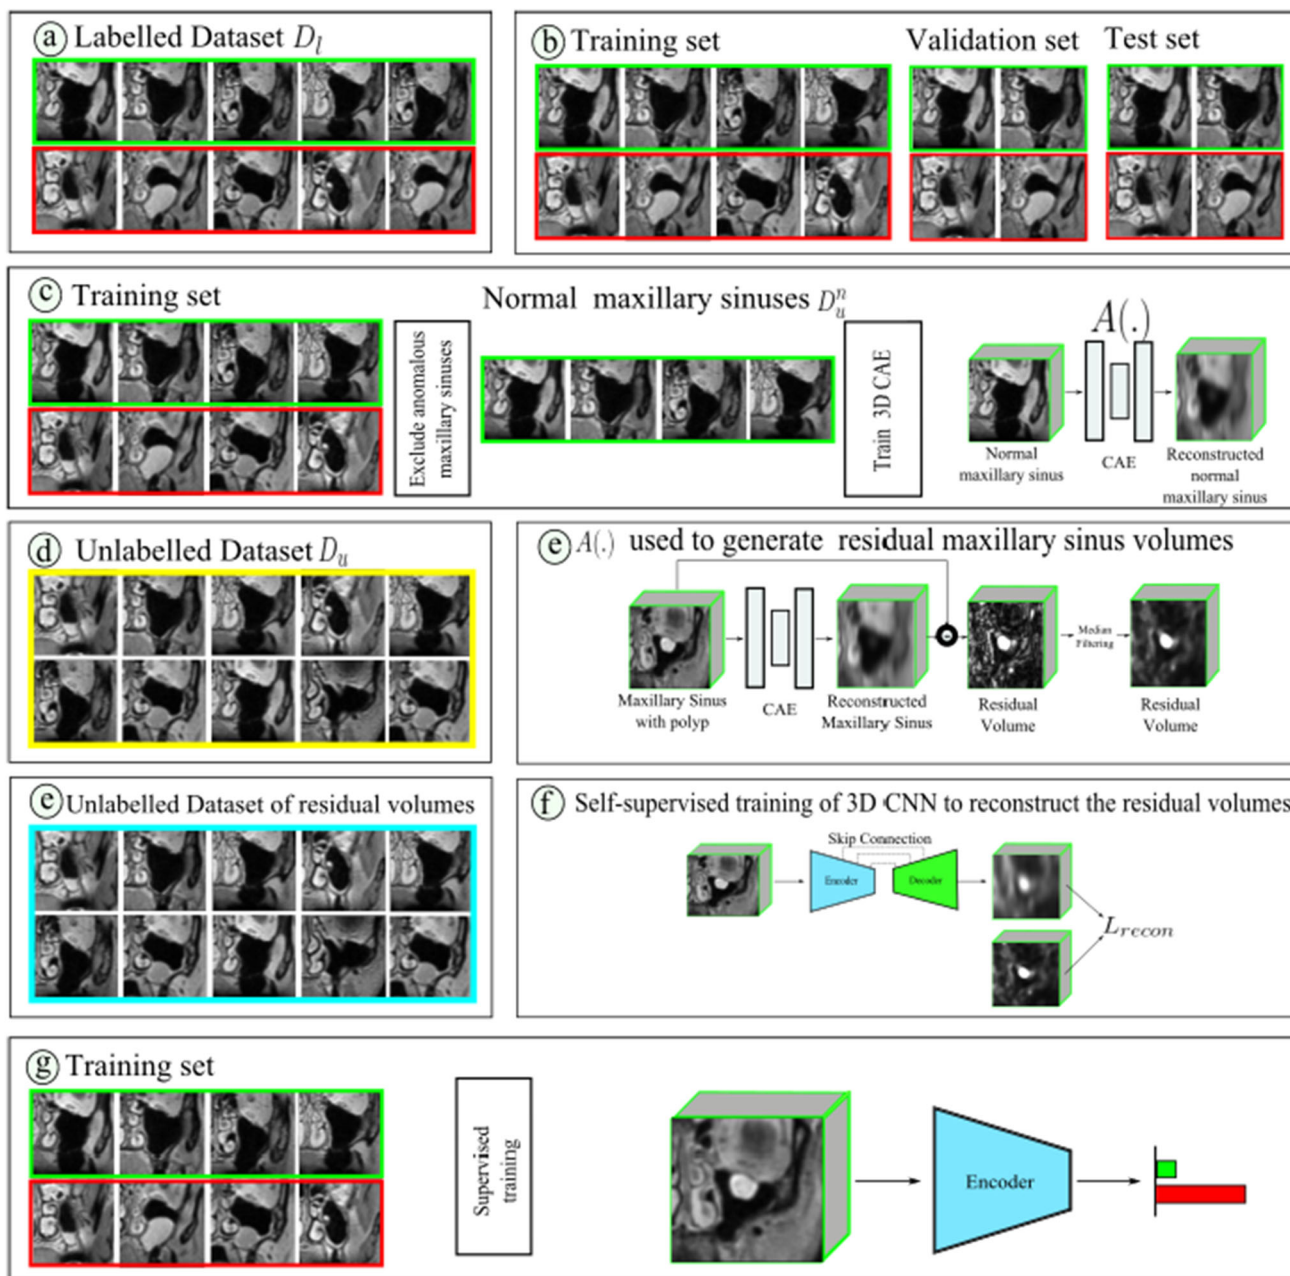


Fig. 2 Our data processing pipeline comprises several steps: **a** The labelled dataset D_l , **b** Splitting D_l into training, validation, and test subsets for downstream classification of normal versus anomalous MS. **c** Normal MS samples from the labelled training set form D_u^n , used to train the 3D CAE $A(\cdot)$ within the UAD framework. **d** Unlabelled dataset D_u , **e** This trained 3D CAE $A(\cdot)$ generates residual volumes from the

unlabelled dataset D_u , **e** Unlabelled dataset of residual volumes, **f** The 3D CNN undergoes self-supervised training to reconstruct these residual volumes. **g** The 3D CNN’s encoder is initialized with weights from the SSL task and then undergoes supervised training for the final task of classifying normal versus anomalous MS, using the training set created in step (a)

self-supervised methods and MS anomaly classification, we applied data augmentations such as random affine transformations, flipping, and Gaussian noise. The DAE specifically used Gaussian noise with a mean of 0 and standard deviation of 0.6 at 100% probability, while other augmentations were applied 50% of the time. Supplementary material offers comprehensive descriptions and visualizations of SOTA SSL methods.

Results

Comparison to state of the art

Results in Table 2 show our method outperforming others in AUROC, AUPRC, and F1 scores across different labelled dataset scenarios (10%, 20%, 100% of D_l). Our method

Table 2 The table displays the mean and 95% confidence intervals of metrics evaluating model performance in the downstream classification task

Method	Training set percentage D_l	AUROC	AUPRC	F1
No pretraining	10%	0.74 (0.64–0.84)	0.69 (0.56–0.82)	0.64 (0.59–0.69)
Transfer Learning	10%	0.77 (0.72–0.82)	0.73 (0.66–0.79)	0.63 (0.57–0.69)
AE	10%	0.73 (0.68–0.79)	0.68 (0.62–0.74)	0.55 (0.43–0.67)
DAE	10%	0.74 (0.73–0.76)	0.68 (0.66–0.69)	0.62 (0.60–0.64)
BYOL	10%	0.79 (0.76–0.81)	0.75 (0.70–0.79)	0.63 (0.59–0.69)
SimSiam	10%	0.77 (0.72–0.83)	0.74 (0.68–0.79)	0.62 (0.53–0.72)
SimCLR	10%	0.78 (0.74–0.81)	0.73 (0.68–0.78)	0.63 (0.59–0.68)
SparK MAE	10%	0.78 (0.77–0.80)	0.75 (0.73–0.76)	0.65 (0.63–0.67)
Ours	10%	0.81 (0.74–0.88)	0.79 (0.71–0.87)	0.67 (0.58–0.77)
No pretraining	20%	0.81 (0.79–0.82)	0.78 (0.76–0.79)	0.67 (0.65–0.69)
Transfer Learning	20%	0.84 (0.79–0.88)	0.81 (0.74–0.88)	0.68 (0.62–0.75)
AE	20%	0.81 (0.76–0.86)	0.78 (0.72–0.83)	0.67 (0.60–0.74)
DAE	20%	0.79 (0.77–0.81)	0.74 (0.70–0.79)	0.67 (0.64–0.70)
BYOL	20%	0.82 (0.80–0.84)	0.79 (0.77–0.82)	0.70 (0.68–0.71)
SimSiam	20%	0.84 (0.82–0.86)	0.81 (0.78–0.84)	0.70 (0.67–0.74)
SimCLR	20%	0.81 (0.79–0.83)	0.77 (0.74–0.81)	0.68 (0.67–0.69)
SparK MAE	20%	0.80 (0.78–0.82)	0.76 (0.73–0.79)	0.67 (0.65–0.68)
Ours	20%	0.85 (0.83–0.87)	0.82 (0.81–0.83)	0.72 (0.70–0.75)
No pretraining	100%	0.90 (0.89–0.91)	0.89 (0.88–0.90)	0.80 (0.78–0.82)
Transfer Learning	100%	0.92 (0.91–0.93)	0.91 (0.90–0.93)	0.82 (0.80–0.83)
AE	100%	0.92 (0.91–0.93)	0.91 (0.90–0.93)	0.82 (0.80–0.84)
DAE	100%	0.90 (0.88–0.92)	0.89 (0.88–0.91)	0.79 (0.77–0.82)
BYOL	100%	0.89 (0.89–0.90)	0.88 (0.87–0.89)	0.78 (0.76–0.81)
SimSiam	100%	0.92 (0.91–0.93)	0.91 (0.90–0.92)	0.81 (0.79–0.83)
SimCLR	100%	0.90 (0.88–0.91)	0.89 (0.87–0.91)	0.79 (0.77–0.80)
SparK MAE	100%	0.87 (0.85–0.88)	0.86 (0.84–0.87)	0.75 (0.73–0.76)
Ours	100%	0.93 (0.91–0.94)	0.92 (0.90–0.93)	0.83 (0.80–0.86)

These models, trained with varying portions of D_l , were initialized using different SSL methods before supervised training
The bold values signify the best/highest values for the given metrics (AUROC, AUPRC, F1)

demonstrated notable improvements in AUROC (3.34% and 4.93% over SimSiam) and AUPRC (5.33% over BYOL and 5.12% over AE) for 10% and 20% dataset scenarios, respectively. SparK trained models perform generally poorer compared to the other SSL and TL methods with the performance gap between SparK MAE and our method widening with increased training set percentage. Our method had AUPRC 8.21% higher than the TL method when finetuned on a 10% training set. Pretraining models using our method significantly boosted AUPRC by 14.49% and AUROC by 9.45% compared to no pretraining when trained on a 10% training dataset. At 100% dataset finetuning, our method achieved the highest scores, with AE and SimSiam showing similar performance. Compared to no pretraining, our method improved AUPRC by 3.33%. Figure 3 illustrates AUPRC and AUROC trends with increasing training set percentages, respectively. Our method excels in settings with 40% or less training data but aligns with SOTA performance beyond that.

Effect of varying the CAE training set

The effectiveness of our self-supervised task is contingent on the CAE's proficiency in reconstructing healthy MS volumes. Inaccurate reconstructions yield unreliable residuals, affecting self-supervision. To assess the impact of training set size, the CAE was trained with different proportions (20%, 40%, 60%, 80%, 100%) of the healthy MS dataset D_l^n . After training, the CAE processed dataset D_u to produce residual volumes, which were refined using a median filter with a kernel size of 5. Subsequent supervised training utilized 10% of our labelled dataset D_l . Table 3 presents improvements in the downstream task metrics correlating with increased healthy MS training set sizes, suggesting that larger *normal* dataset D_l^n enhance normal MS representation learning and improve anomaly localization.

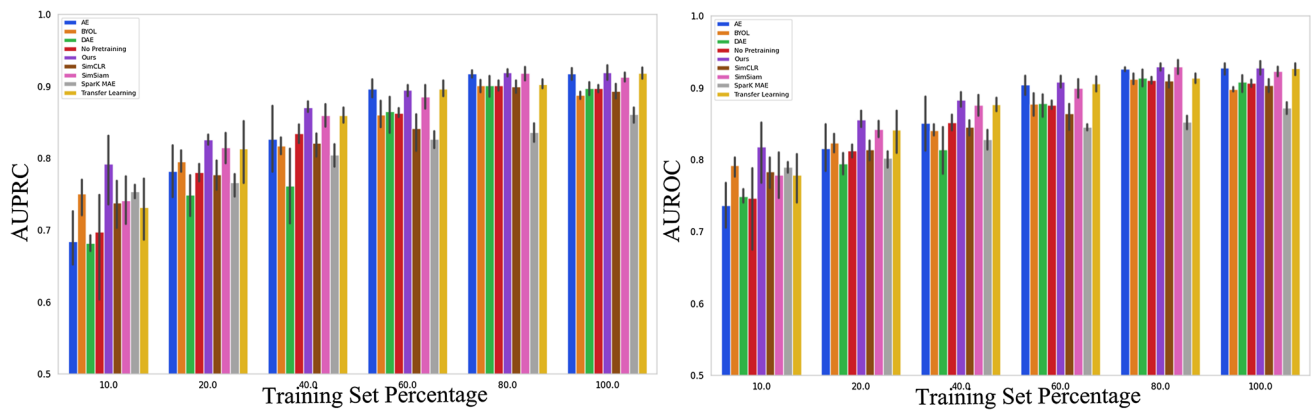


Fig. 3 (LEFT) AUPRC trend vs training set percentage (RIGHT) AUROC trend vs training set percentage

Table 3 The table shows the mean and 95% confidence intervals of metrics for evaluating model performance in downstream classification

Training set percentage D_1^n	AUROC	AUPRC	F1
20%	0.76 (0.70–0.81)	0.72 (0.67–0.77)	0.60 (0.50–0.69)
40%	0.77 (0.73–0.80)	0.72 (0.66–0.78)	0.63 (0.57–0.68)
60%	0.78 (0.75–0.80)	0.74 (0.71–0.77)	0.65 (0.62–0.68)
80%	0.80 (0.76–0.84)	0.76 (0.72–0.81)	0.67 (0.63–0.72)
100%	0.81 (0.74–0.88)	0.79 (0.71–0.87)	0.67 (0.58–0.77)

The CAE was trained on varying proportions of the normal MS volumes dataset (D_1^n) and then used to generate residual volumes from the unlabelled dataset (D_u). Each model was initialized using our proposed SSL method

The bold values signify the best/highest values for the given metrics (AUROC, AUPRC, F1)

Discussion

Tailoring SSL tasks to specific downstream tasks offers distinct advantages [32]. Current SOTA SSL methods [20–22], primarily developed for 2D image classification on datasets like ImageNet, do not address the unique challenges of 3D MRI modalities and the specifics of paranasal anomalies. Our SSL task is specifically tailored to address the challenges associated with 3D environments, MRI modality, and the classification of paranasal anomalies.

We conjecture that segmentation of anomalies as a SSL task, requiring knowledge of anomaly locations, enhances the learning of class-discriminative features for distinguishing normal and anomalous MS. Our SSL task is a segmentation task therefore, it requires segmentation masks highlighting anomalies. To avoid the high costs of annotation, we use a CAE trained in the UAD framework for generating approximate annotations, effective in localizing paranasal anomalies [26]. This CAE training utilizes labelled *normal* datasets, typically accessible in supervised settings. Unlike generic SOTA SSL methods, which do not prioritize anomaly localization, our approach demonstrates improved AUROC and AUPRC (as shown in Table 2), suggesting that effective anomaly localization can enhance classification performance, even with limited labelled data. Methods like BYOL and SimSiam, which aim to maximize agreement

between augmented views, are less effective for paranasal anomaly classification. SimCLR’s performance shortfall is likely due to smaller batch sizes, a necessity given the impracticality of large batches in 3D settings, despite SimCLR’s recommendation of 4096 [33]. Our method is more suited for such constrained computational resources. AE and DAE, focusing on compression-decompression and denoising, do not guarantee discriminative feature learning for downstream classification [34], and were found less effective in our context. When the entire training set is used, our method, AE, and SimSiam yield comparable results, with ours marginally outperforming. We also explored MAE-style pretraining using SparK. However, the results suggest that fine-tuning performance is notably weaker, particularly when fine-tuning with a training set percentage 40% and above. These findings imply that generating masked regions contributes to representation learning; however, the acquired representations do not appear to enhance downstream classification. It is noteworthy that the SparK framework was initially developed and evaluated for 2D natural images. Although we adapted the framework for 3D applications, our findings underscore the necessity for further methodological advancements to effectively support tasks in the 3D domain. Further, TL models exhibit comparable performance to SSL methods when fine-tuning on training sets exceeding 20%. This suggests that transfer learning methods remain viable for paranasal

anomaly classification given an ample supply of labelled samples. However, in the scenario of an extremely limited labelled dataset, such as 10%, our method outperforms TL, indicating that the representations acquired by our approach are especially advantageous in low-data environments. Overall, compared to approaches without pretraining, our tailored SSL task consistently shows superior downstream classification performance, underlining its efficacy.

Our analysis regarding the impact of the CAE training set size shown in Table 3 has demonstrated that the inclusion of a substantial cohort of normal MS volumes yields notable benefits for both the self-supervision task and the subsequent downstream task suggesting that better anomaly localization by the CAE and thereby better representation learning by the CNN in the self-supervision task. We also analysed the influence of the loss function and post-processing used in the self-supervision task which can be found in the supplementary material section 8 and 9.

Our study has limitations that require further investigation. It is based on a single-centre MRI-only study, so multi-centre studies with varied imaging modalities are needed for generalizability. Our methods rely on a cohort of healthy MS volumes, unlike other self-supervised tasks. We focused on convolutional autoencoders, not exploring models like variational autoencoders generative adversarial networks, or transformer-based architectures and diffusion models, which might offer better anomaly localization. We compared L1, L2, and BCE loss functions but not others like the Structural Similarity Index or perceptual loss. Future research should examine these aspects and apply this self-supervision approach to other domains, like brain anomaly detection.

Conclusion

We developed a novel self-supervision task that focuses on anomaly localization to better classify paranasal anomalies in the maxillary sinus, addressing the lack of methods that effectively use unlabelled datasets to learn discriminative features for this purpose. Our approach uses an autoencoder trained on healthy MS volumes to generate residual volumes from an unlabelled dataset. These residuals serve as coarse segmentation masks for localizing anomalies. By training a CNN to reconstruct these volumes, it implicitly learns anomaly localization, thereby developing transferable features for the downstream classification task. Our method outperforms existing self-supervision techniques, proving its effectiveness in this specific domain.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11548-024-03172-5>.

Acknowledgements This work has not been submitted for publication anywhere else. This work is funded partially by the i3 initiative of the Hamburg University of Technology. The authors also acknowledge the partial funding by the Free and Hanseatic City of Hamburg (Interdisciplinary Graduate School) from University Medical Center Hamburg-Eppendorf. This work was partially funded by Grant Number KK5208101KS0 (Zentrales Innovationsprogramm Mittelstand, Arbeitsgemeinschaft industrieller Forschungsvereinigungen). Publishing fees supported by Funding Programme Open Access Publishing of Hamburg University of Technology (TUHH).

Funding Open Access funding enabled and organized by Projekt DEAL. Funding was provided by i3 initiative Hamburg University of Technology, Interdisciplinary Graduate School University Medical Center Hamburg-Eppendorf, Zentrales Innovationsprogramm Mittelstand, Arbeitsgemeinschaft industrieller Forschungsvereinigungen (Grant number: K5208101KS0).

Declarations

Conflict of interest The authors declare no conflict of interest.

Ethical approval The study protocol received approval from the local ethics committee (Landesärztekammer Hamburg, PV5131) and was approved by the Data Protection Commissioners for the University Medical Center of the University Hamburg-Eppendorf and the Free and Hanseatic City of Hamburg. It is registered on ClinicalTrials.gov (NCT03934957) and adheres to Good Clinical Practice, Good Epidemiological Practice, and ethical principles outlined in the Declaration of Helsinki.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Marieb EN (1991) *Essentials of Human Anatomy & Physiology*. Third edition. Redwood City, Calif., Benjamin/Cummings Pub. Co., 1991. <https://search.library.wisc.edu/catalog/9910059601802121>
2. Bal M, Berkiten G, Uyamk E (2014) Mucous retention cysts of the paranasal sinuses. *Hippokratia* 18(4):379
3. Varshney H, Varshney J, Biswas S, Ghosh SK (2015) Importance of CT scan of paranasal sinuses in the evaluation of the anatomical findings in patients suffering from sinonasal polyposis. *Indian J Otolaryngol Head Neck Surg* 68(2):167–172
4. Van Dis ML, Miles DA (1994) Disorders of the maxillary sinus. *Dent Clin North Am* 38(1):155–166
5. Hansen AG, Helvik A-S, Nordgård S, Bugten V, Stovner LJ, Håberg AK, Gårseth M, Eggesbø HB (2014) Incidental findings in MRI of the paranasal sinuses in adults: a population-based study (HUNT

- MRI). *BMC Ear Nose Throat Disord* 14(1):13. <https://doi.org/10.1186/1472-6815-14-13>
6. Tarp B, Fiirgaard B, Christensen T, Jensen JJ, Black FT (2000) The prevalence and significance of incidental paranasal sinus abnormalities on MRI. *Rhinology* 38(1):33–38
 7. Brierley J, Gospodarowicz MK, Wittekind C (eds) (2017) TNM classification of malignant tumours. Eighth edn. John Wiley & Sons Inc, Chichester West Sussex UK and Hoboken NJ
 8. Gutmann A (2013) Ethics. The bioethics commission on incidental findings. *Science* 342(6164):1321–1323. <https://doi.org/10.1126/science.1248764>
 9. Papadopoulou A-M, Chrysikos D, Samolis A, Tsakotos G, Troupis T (2021) Anatomical variations of the nasal cavities and paranasal sinuses: a systematic review. *Cureus* 13(1):12727
 10. Jeon Y, Lee K, Sunwoo L, Choi D, Oh DY, Lee KJ, Kim Y, Kim J-W, Cho SJ, Baik SH, Yoo R-E, Bae YJ, Choi BS, Jung C, Kim JH (2021) Deep learning for diagnosis of paranasal sinusitis using multi-view radiographs. *Diagnostics*. <https://doi.org/10.3390/diagnostics11020250>
 11. Kim Y, Lee KJ, Sunwoo L, Choi D, Nam C-M, Cho J, Kim J, Bae YJ, Yoo R-E, Choi BS, Jung C, Kim JH (2019) Deep learning in diagnosis of maxillary sinusitis using conventional radiography. *Investig Radiol* 54(1):7–15. <https://doi.org/10.1097/RLI.0000000000000503>
 12. Liu GS, Yang A, Kim D, Hojel A, Voevodsky D, Wang J, Tong CCL, Ungerer H, Palmer JN, Kohanski MA, Nayak JV, Hwang PH, Adappa ND, Patel ZM (2022) Deep learning classification of inverted papilloma malignant transformation using 3d convolutional neural networks and magnetic resonance imaging. *Int Forum Allergy Rhinol*. <https://doi.org/10.1002/alr.22958>
 13. Kim K-S, Kim BK, Chung MJ, Cho HB, Cho BH, Jung YG (2022) Detection of maxillary sinus fungal ball via 3-D CNN-based artificial intelligence: Fully automated system and clinical validation. *PLoS ONE* 17(2):1–19. <https://doi.org/10.1371/journal.pone.0263125>
 14. Bhattacharya D, Becker BT, Behrendt F, Bengs M, Beyersdorff D, Eggert D, Petersen E, Jansen F, Petersen M, Cheng B, Betz C, Schlaefer A, Hoffmann AS (2022) Supervised contrastive learning to classify paranasal anomalies in the maxillary sinus. In: Wang L, Dou Q, Fletcher PT, Speidel S, Li S (eds) *Medical image computing and computer assisted intervention-MICCAI 2022*. Springer, Cham, pp 429–438
 15. Bhattacharya D, Behrendt F, Becker BT, Beyersdorff D, Petersen E, Petersen M, Cheng B, Eggert D, Betz C, Hoffmann AS, Schlaefer A (2023) Multiple instance ensembling for paranasal anomaly classification in the maxillary sinus. *Int J Comput Assist Radiol Surg* 19(2):223–231
 16. Pang G, Shen C, Cao L, Hengel AVD (2021) Deep learning for anomaly detection: a review. *ACM Comput Surv*. <https://doi.org/10.1145/3439950>
 17. Pihlgren G, Sandin F, Liwicki M (2021) Pretraining image encoders without reconstruction via feature prediction loss. In: 2020 25th international conference on pattern recognition (ICPR), pp 4105–4111. IEEE Computer Society, Los Alamitos, CA, USA. <https://doi.org/10.1109/ICPR48806.2021.9412239>
 18. Xie Y, Thurey N (2023) Reviving autoencoder pretraining. *Neural Comput Appl* 35(6):4587–4619. <https://doi.org/10.1007/s00521-022-07892-0>
 19. Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol P-A (2010) Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J Mach Learn Res* 11:3371–3408
 20. Grill J-B, Strub F, Altché F, Tallec C, Richemond P, Buchatskaya E, Doersch C, Avila Pires B, Guo Z, Gheshlaghi Azar M, Piot B, kavukcuoglu k, Munos R, Valko M (2020) Bootstrap your own latent—a new approach to self-supervised learning. In: Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H (eds.) *Advances in neural information processing systems*, vol. 33, pp 21271–21284. Curran Associates, Inc., . https://proceedings.neurips.cc/paper_files/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf
 21. Chen X, He K (2021) Exploring simple siamese representation learning. In: 2021 IEEE/CVF Conference on computer vision and pattern recognition (CVPR), pp 15745–15753 . <https://doi.org/10.1109/CVPR46437.2021.01549>
 22. Huang S-C, Pareek A, Jensen M, Lungren MP, Yeung S, Chaudhari AS (2023) Self-supervised learning for medical image classification: a systematic review and implementation guidelines. *NPJ Digit Med* 6(1):74. <https://doi.org/10.1038/s41746-023-00811-0>
 23. Tian K, Jiang Y, qishuai diao, Lin C, Wang L, Yuan Z (2023) Designing BERT for convolutional networks: sparse and hierarchical masked modeling. In: The eleventh international conference on learning representations. <https://openreview.net/forum?id=NRxydtWup1S>
 24. Baur C, Denner S, Wiestler B, Navab N, Albarqouni S (2021) Autoencoders for unsupervised anomaly segmentation in brain MR images: a comparative study. *Med Image Anal* 69:101952
 25. Behrendt F, Bengs M, Rogge F, Krüger J, Opfer R, Schlaefer A (2022) Unsupervised anomaly detection in 3D brain MRI using deep learning with impured training data. In: 2022 IEEE 19th international symposium on biomedical imaging (ISBI), pp 1–4. <https://doi.org/10.1109/ISBI52829.2022.9761443>
 26. Bhattacharya D, Behrendt F, Becker BT, Beyersdorff D, Petersen E, Petersen M, Cheng B, Eggert D, Betz C, Hoffmann AS, Schlaefer A (2022) Unsupervised anomaly detection of paranasal anomalies in the maxillary sinus. arXiv. <https://doi.org/10.48550/ARXIV.2211.01371>. <https://arxiv.org/abs/2211.01371>
 27. Jagodzinski A (2019) Rationale and design of the Hamburg city health study. *Eur J Epidemiol* 35(2):169–181
 28. Tran D, Wang H, Torresani L, Ray J, LeCun Y, Paluri M (2018) A closer look at spatiotemporal convolutions for action recognition. In: 2018 IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 6450–6459. IEEE Computer Society, Los Alamitos, CA, USA. <https://doi.org/10.1109/CVPR.2018.00675>
 29. Deng J, Dong W, Socher R, Li L.-J, Li K, Fei-Fei L (2009) ImageNet: a large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition, pp 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
 30. Ginsburg B, Gitman I, You Y (2018) Large batch training of convolutional networks with layer-wise adaptive rate scaling. <https://openreview.net/forum?id=rJ4uaX2aW>
 31. Loshchilov I, Hutter F (2019) Decoupled weight decay regularization. In: International conference on learning representations. <https://openreview.net/forum?id=Bkg6RiCqY7>
 32. Ozbulak U, Lee HJ, Boga B, Anzaku ET, Park H-M, Messem AV, Neve WD, Vankerschaver J (2023) Know your self-supervised learning: a survey on image-based generative and discriminative training. *Transactions on Machine Learning Research*. Survey Certification
 33. Chen T, Kornblith S, Norouzi M, Hinton G (2020) A simple framework for contrastive learning of visual representations. In: Proceedings of the 37th international conference on machine learning. ICML. JMLR.org
 34. Raina R, Battle A, Lee H, Packer B, Ng AY (2007) Self-taught learning: transfer learning from unlabeled data. In: Proceedings of the 24th international conference on machine learning. ICML '07, pp. 759–766. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/1273496.1273592>